

Bose-Einstein Condensates  
and  
Degenerate Fermi Gases

A thesis submitted for the degree  
of Doctor of Philosophy of  
the Australian National University

Cameron Stuart Fletcher

January 7, 2005

## Declaration

This thesis contains no material which has been accepted for the award of any other degree or diploma in any university. To the best of the author's knowledge and belief, it contains no material previously published or written by another person, except where due reference is made in the text. This work was conducted under the supervision of Dr. John Close, Dr. Craig Savage, and Prof. Hans Bachor.

A handwritten signature in blue ink that reads "Fletcher". The signature is written in a cursive style with a long horizontal stroke extending to the right.

Cameron Fletcher  
January 7, 2005

For Michelle.

# Acknowledgments

I believe in brilliant, succinct acknowledgments that thank everybody. I used to labour under a similar misapprehension regarding PhDs, as well.

Thank you:

ANU-BEC ...	John, Jess and Ulli. I couldn't have done it without you.
Nick ...	You improved every aspect of the lab, from the BEC machine, to the conversation.
Hans ...	for those occasions where you said the perfect thing, <i>in the perfect way</i> .
Simon, Adele and Jessica Trevena.	I hope you enjoyed working with me as much as I did with you.

To all those who have had the pleasure of working with me ... congratulations!

To all my:

housemates ...	I've learnt at least as much about life as about physics in the last 5 years, and you've all left your mark. You also ended up my best friends. Thank you Alan, Jill, Adam, Caroline, Jessica Trevena, Ben Buchler, Michelle McCaan, Brett and Nicky.
my other friends ...	ditto. Thanks Andrew Lance, Joe, Janet, Sebastian, Beate, Max, Brendan, Laura, Ray, Wez, Liv and Nicole.

To those people that put me in a position to do what I have done ... I hope you're happy with yourselves. I am. Thank you Cyril, Ralph, Mr Knight and Mr. Scarle.

My love and appreciation go to:

Mum and Dad ...	the more I look into myself, the more I see each of you peering back at me. You made me what I am today.
Felicity ...	you're the best sister a guy like me has. I love you to bits.
My second family ...	you have each made my life better. Thank you Jenny, Bek and Daniel.
Michelle ...	what can I say. This thesis is dedicated to you.

and, finally, to all those people that I couldn't possibly afford not to mention by name ...

Thank you.

# Summary of Thesis

*Bose-Einstein condensates and degenerate Fermi gases* is a deceptively succinct, but appropriately portentous moniker with which to emblazon this tome. Part I describes a large volume of experimental work, leading to the production of the first Bose-Einstein condensate in Australia. Part II, in turn, describes the first progress towards the production of a degenerate Fermi gas in Australia.

Creating a machine to generate a Bose-Einstein condensate is a difficult task. It was first achieved in 1995, and for their efforts, the discoverers were rewarded with the 2001 Nobel prize in physics. It is an exciting new field, growing rapidly in terms of the applied workforce, the experimental base, the scope of research, and the promise of things to come.

The work described in part I of this thesis was conducted as part of a group of four people. The entire ANU-BEC machine is summarized, but contemplative emphasis is applied to those aspects of the original machine that formed the primary foci of this thesis. Specifically, the vacuum system, magnetic trap, CCD cameras and computer control systems are elaborated in chapters 4, 6, 8 and 9, to an extent useful for a reader reproducing such an experiment. At all stages the success of the machine is compared and contrasted with the improvements made since its first implementation. This part of the thesis is concluded in chapter 10, with a summary of results garnered from the ANU-BEC machine, including the first BEC results of ANU-BEC-I, and the first atom laser results harvested after its reincarnation as ANU-BEC-II.

Degeneracy in a Fermi gas was first achieved in 1999, but the more ponderous experimental difficulties exhibited by such a system prevented any further realization until 2001. To date only six of the biggest and most experienced experimental groups worldwide can lay claim to owning such an apparatus. The focus of the second part of this thesis was the development of ANU's own research investment toward producing degenerate Fermi gas. Unlike the Bose-Einstein condensation experiment, this research evolved solely from the endeavors of this thesis.

Progress towards construction of a degenerate Fermi gas constitutes the subject of chapters 11 - 14. The hurdles that nature places in the path of the degenerate Fermi gas researcher are introduced in chapters 1, 2 and 11, and the solutions to some of these are outlined in chapters 12, 13 and 14. These generally involve some mix of unique and almost-unique research, dependent on the status of the field when they were broached.

These experiments create new states of matter, ones that purify and amplify a collection of atoms into well described quantum states. More than that, it achieves it in a manner convenient to experiment and theory alike. The explosively dynamic area where these two disciplines meet promise the opening of new vistas of quantum mechanical understanding. These, in turn, foretell of new applications, and new technology, emerging from the fruits of research such as that described herein.

# Contents

<b>Acknowledgments</b>	<b>i</b>
<b>Summary of Thesis</b>	<b>iii</b>
<b>List of Figures</b>	<b>viii</b>
<b>List of Acronyms</b>	<b>xiii</b>
<b>List of Publications</b>	<b>xvii</b>
<b>1 And so it begins ...</b>	<b>1</b>
1.1 Degenerate quantum gases . . . . .	2
1.2 Bose Einstein condensates . . . . .	4
1.3 Degenerate Fermi gases . . . . .	7
1.4 A final motivation . . . . .	10
<b>2 The pursuit of degeneracy</b>	<b>11</b>
2.1 Background . . . . .	11
2.2 Phase space density . . . . .	12
2.3 Practical considerations . . . . .	14
2.4 Thermodynamics . . . . .	17
2.5 Bose-Einstein condensation . . . . .	19
2.6 Degenerate Fermi gases . . . . .	20
2.7 Beyond degeneracy . . . . .	22
<b>I Bose-Einstein condensation</b>	<b>23</b>
<b>3 ANU-BEC: Australia's first Bose-Einstein condensate machine</b>	<b>25</b>
3.1 The ANU-BEC in context . . . . .	26
3.2 Overview of the ANU-BEC-I machine . . . . .	26
3.3 The ANU-BEC-I machine in operation . . . . .	33
<b>4 The vacuum system</b>	<b>37</b>
4.1 Fundamental considerations . . . . .	38
4.2 Practical considerations . . . . .	40
4.3 The ANU-BEC-I vacuum system . . . . .	43
4.4 Improvements to the vacuum system for ANU-BEC-II . . . . .	47
4.5 Future improvements to the vacuum system . . . . .	47
4.6 Characterizing the ANU-BEC vacuum system . . . . .	49

<b>5</b>	<b>The optical system</b>	<b>53</b>
5.1	Practical considerations . . . . .	53
5.2	The ANU-BEC-I optical system . . . . .	56
5.3	Improvements to the optical system in ANU-BEC-II . . . . .	58
<b>6</b>	<b>The magnetic trap</b>	<b>59</b>
6.1	Fundamental considerations . . . . .	61
6.2	Modeling the magnetic trap . . . . .	66
6.3	Practical considerations . . . . .	73
6.4	The ANU-BEC-I magnetic trap . . . . .	75
6.5	Improvements to the magnetic trap for ANU-BEC-II . . . . .	80
6.6	Characterizing the magnetic trap with cold atoms . . . . .	83
<b>7</b>	<b>Evaporative cooling</b>	<b>87</b>
7.1	Theoretical considerations . . . . .	87
7.2	Evaporative cooling in the ANU-BEC-I experiment . . . . .	89
7.3	Improvements to evaporative cooling during ANU-BEC-II . . . . .	91
<b>8</b>	<b>Imaging and CCD cameras</b>	<b>93</b>
8.1	Fundamental considerations . . . . .	94
8.2	Choosing a CCD camera . . . . .	100
8.3	Improvements to the imaging system for ANU-BEC-II . . . . .	105
8.4	Future improvements to the imaging system . . . . .	106
8.5	Future Plans . . . . .	107
<b>9</b>	<b>The control system</b>	<b>109</b>
9.1	Hardware . . . . .	109
9.2	Computer . . . . .	113
9.3	Software . . . . .	114
9.4	The user interface . . . . .	116
9.5	The ANU-BEC-I control system . . . . .	118
9.6	Improvements to the control system during ANU-BEC-II . . . . .	120
9.7	Future improvements to the control system . . . . .	121
<b>10</b>	<b>ANU-BEC results</b>	<b>123</b>
10.1	First results of BEC . . . . .	123
10.2	Fine-tuning evaporative cooling . . . . .	126
10.3	Fine-tuning magnetic trap switching . . . . .	128
10.4	First results on ANU-BEC-II . . . . .	131
10.5	The atom laser . . . . .	132
<b>II</b>	<b>Degenerate Fermi gases</b>	<b>135</b>
<b>11</b>	<b>ANU-DFG</b>	<b>137</b>
11.1	Overview of the DFG machine . . . . .	138

11.2 Progress of the ANU-DFG experiment . . . . .	139
<b>12 External cavity diode lasers</b>	<b>143</b>
12.1 Fundamental considerations . . . . .	145
12.2 Practical considerations . . . . .	152
12.3 The ANU-DFG external cavity diode lasers . . . . .	153
12.4 Considerations for improvement . . . . .	155
12.5 Seeding a tapered amplifier with multiple wavelengths . . . . .	156
<b>13 Extended temperature tuning of external cavity diode lasers</b>	<b>161</b>
13.1 Fundamental and practical considerations . . . . .	162
13.2 A Solid-State, Economical Ultra-Cold Diode Laser . . . . .	171
13.3 Summary . . . . .	177
<b>14 The self locked MOT</b>	<b>179</b>
14.1 Fundamental considerations . . . . .	180
14.2 Practical considerations . . . . .	181
14.3 The self-locked magneto-optic trap . . . . .	182
<b>15 Thinking back and looking forward</b>	<b>187</b>
<b>Bibliography</b>	<b>190</b>



# List of Figures

1.1	Reaching degeneracy. . . . .	3
2.1	The path to degeneracy. . . . .	15
3.1	The main ANU-BEC-I table. . . . .	26
3.2	The ANU-BEC-I control systems. . . . .	27
3.3	The vacuum chamber. . . . .	28
3.4	The optical setup. . . . .	29
3.5	The magnetic trap coils. . . . .	30
3.6	The imaging optics and CCD camera. . . . .	32
4.1	Photo of the ANU-BEC-I vacuum system. . . . .	38
4.2	Photo of the ANU-BEC-I collection chamber. . . . .	44
4.3	The ANU-BEC-I science chamber. . . . .	45
4.4	Photo of the ANU-BEC-III vacuum system. . . . .	48
4.5	Atoms remaining after variable hold times. . . . .	50
4.6	Lifetime of atoms in the magnetic trap. . . . .	50
5.1	Diagram of the optical layout. . . . .	54
6.1	Diagram of the magnetic trap. . . . .	60
6.2	Photo of the magnetic trap. . . . .	61
6.3	Form of a quadrupole magnetic trap. . . . .	63
6.4	Form of a QUIC magnetic trap. . . . .	65
6.5	The simple BiotSavart model. . . . .	66
6.6	The complex BiotSavart model. . . . .	67
6.7	The Mathematica model. . . . .	68
6.8	The magnetic fields due to the trap. . . . .	69
6.9	Diagram of the magnetic trap coils. . . . .	76
6.10	The constant current sink. . . . .	78
6.11	The ANU-BEC switching circuits. . . . .	79
6.12	Diagram of the ANU-BEC-II magnetic trap. . . . .	81
6.13	Photo of the ANU-BEC-II magnetic trap. . . . .	82
6.14	The evolution of the QUIC trap. . . . .	84
7.1	Evaporation to BEC. . . . .	90
8.1	Absorption imaging. . . . .	94
8.2	Diagram of a CCD camera. . . . .	100
8.3	Data collected by the ANU-BEC-I imaging system. . . . .	104
8.4	Data collected by the ANU-BEC-II imaging system. . . . .	106

9.1	Block diagram of the control software. . . . .	110
9.2	The main code window of the computer program. . . . .	115
9.3	The main control window of the computer program. . . . .	117
10.1	Bose-Einstein condensation! . . . . .	124
10.2	Analysis of condensation. . . . .	125
10.3	Logarithmic evaporation ramp. . . . .	126
10.4	Linear evaporation ramp. . . . .	127
10.5	Evaporative cooling comparison. . . . .	128
10.6	Switching off the magnetic trap . . . . .	129
10.7	The improved switching circuit. . . . .	130
10.8	Optimized ANU-BEC-II BEC. . . . .	131
10.9	The semi-continuous atom laser. . . . .	133
10.10	The pulsed atom laser. . . . .	134
12.1	An external cavity diode laser. . . . .	144
12.2	A laser diode. . . . .	146
12.3	Detail of the grating mount. . . . .	149
12.4	Detail of a reflection grating. . . . .	150
12.5	External cavity diode laser specifications. . . . .	154
12.6	Injecting a tapered amplifier with multiple wavelengths. . . . .	157
13.1	An ultracold diode laser system. . . . .	162
13.2	Heat flows in a TEC. . . . .	164
13.3	Performance of a TEC lagged to a reservoir. . . . .	166
13.4	Performance of a TEC lagged to a heatsink. . . . .	167
13.5	Multiple TECs stacked in series. . . . .	168
13.6	Typical use of a cold plate . . . . .	170
13.7	The ultracold external cavity diode laser system. . . . .	172
13.8	Temperature dependance of the free-running wavelength. . . . .	173
13.9	Minimum temperature versus $n_{Hot}$ and $G_{Hot}$ . . . . .	175
13.10	Saturated absorption spectrum of potassium at 766.7 nm. . . . .	177
14.1	Saturated absorption spectroscopy. . . . .	180
14.2	Frequency modulation spectroscopy . . . . .	182
14.3	The experimental set up. . . . .	183
14.4	The demodulated FMS signal. . . . .	184
14.5	The lock acquisition plots. . . . .	185

# List of Acronyms

The following acronyms are used throughout this thesis. They are provided here as a ready reference for the reader.

AC	Alternating current
ANU	Australian National University
AOM	Acousto-optic modulator
ASCII	American standard code for information interchange
AWG	American wire gauge
BCS	Bardeen-Cooper-Schrieffer
BEC	Bose-Einstein condensate
BIOS	Basic input / output system
BNC	Bayonet Neill Concelman
CCD	Charge coupled device
CCS	Constant current sink
CPU	Central processing unit
DAQ	Data acquisition
DAVLL	Dichroic atomic vapour laser locking
DC	Direct current
DFG	Degenerate Fermi gas
DMA	Direct memory access
DQG	Degenerate quantum gas
ECDL	External cavity diode laser
EOM	Electro-optic modulator
FET	Field effect transistor
FMS	Frequency modulation spectroscopy
FP	Fabry Perot
FSR	Free spectral range
GPIO	General purpose information bus
GPS	Global positioning system
IEEE	Institute of electrical and electronic engineers
I/O	Input / output
IP	Ioffe-Pritchard
ISA	Industry standard architecture
JL	Jessica Lye
MO	Master oscillator
MOPA	Master oscillator-power amplifier
MOT	Magneto optic trap
NR	Nick Robins
PA	Power amplifier

PBS	Polarizing beam splitter
PCI	Peripheral component interconnect
PDH	Pound-Drever-Hall
PID	Proportional-differential-integral
PXI	PCI extentions for instrumentation
PSD	Phase space density
QUIC	Quadrupole-Ioffe configuration
RAM	Random access memory
RF	Radio frequency
RSPHYSSE	Research school of physical sciences and engineering
RTSI	Real time system integration
SAS	Saturated absorption spectroscopy
SLMOT	Self-locked magneto optic trap
TA	Tapered amplifier
TEC	Thermoelectric cooler
UHV	Ultra-high vacuum
UI	User interface
USB	Universal serial bus

# List of Publications

The following publications were borne of the work conducted in this thesis. Some of the results were collected from the ANU-BEC-I machine, described as the major development of part I of this thesis. Some were collected during the ANU-BEC-II iteration of the BEC machine, which built on the ANU-BEC-I machine, as discussed throughout. Several of the papers apply to technical developments generated during the implementation of the ANU-DFG prototype degenerate Fermi gas machine described in part II of this thesis. Those papers most applicable to the reports provided in this thesis are starred.

- \* [1] J. E. Lye, C.S. Fletcher, U. Kallman, H. A. Bachor, and J. D. Close., Images of evaporation to Bose Einstein condensation, *J. Opt. B, Quantum and Semiclass. Opt.*, **4** 57 (2002).
  
- \* [2] C. S. Fletcher, J. E. Lye, N. P. Robbins, and J. D. Close, A self locked magneto-optic trap, *Opt. Comm.*, **212** 85 (2002).
  
- \* [3] C. S. Fletcher and J. D. Close, Extended temperature tuning of an external cavity diode laser. *Appl. Phys. B*, **78** 305 (2004).
  
- [4] N.P. Robins, J. E. Lye, C. S. Fletcher, S. A. Haine, J. Dugue, C. Breme, J. J. Hope, and J. D. Close, Dynamical effects of backcoupling on an atom laser, In P. Hannaford, A. Sidorov, H. Bachor, and K. Baldwin, editors, *Laser Spectroscopy, Proc. XVI Int. Conf.* World Scientific, (2004).
  
- [5] N. P. Robins, C. M. Savage, J. J. Hope, J. E. Lye, C. S. Fletcher, S. A. Haine, and J. D. Close, Classical noise and flux: the limits of multi-state atom lasers, *Phys. Rev. Lett.* (2004).



# 1 And so it begins ...

This thesis is somewhat unique in the world of degenerate quantum gases in that it covers work towards not one, but two, degenerate quantum gas (DQG) experiments. The PhD covered work towards Australia's first Bose-Einstein condensate (BEC) and went on to lay the groundwork for the production of a degenerate Fermi gas (DFG) at ANU.

Even taken individually, each of these experiments is new and complicated. Dilute gas Bose-Einstein condensation was first achieved in 1995, while the first degenerate Fermi gas was created in 1999. Since then roughly 40 groups have created BECs, and only six have achieved DFG.

Although the last ten years has seen significant advances in the specific implementation and construction of machines to produce degenerate quantum gases, the general principles applied in these experiments have remained the same. Creating a machine to produce a DQG is still a difficult and expensive undertaking that eludes many groups.

The BEC machine constructed as part of this thesis has been recognized as one of the cheapest, most quickly assembled machines in the world. This is doubly impressive given the relative inexperience of the group members when construction began in early 1999.

In stark contrast, there is still very little available with which to compare the ANU-DFG machine. The second group to produce a DFG only succeeded in 2002, despite the first result being published 3 years earlier. To date, several groups have reported the successful creation of a DFG, including one based on the same system proposed for study using the ANU-DFG experiment. The fundamental and technical reasons limiting the production of DFGs are discussed further in chapter 2.

Producing a thesis elucidating some of the finer experimental results of quantum physics following the construction of an entire DQG machine is a horrendously time-consuming task. Making significant parts of two separate machines precludes the possibility entirely. This thesis is unashamedly technical in flavor, predominantly describing the construction of the machines. It does, of course, contain much enabling physics along the way.

The size of these projects necessarily makes them group efforts. The machines described within this thesis have since passed under the overarching control of others in the ANU group. Since their initial implementation, many interesting results have been generated, some of which are mentioned in this thesis, or elucidated further in the theses of other group members [1, 2].

The rapid success in producing the ANU-BEC drove the production of new and interesting results quite prodigiously given the scope of the available resources [3–

14]. Some of these results have played a large part in this thesis. In an effort to relate the bulk of the story behind the ANU BEC experiment many interesting results have been presented, and their application to this thesis has been clearly indicated throughout.

The thesis is separated into two major parts, describing work on the two separate experiments. The first part relates the construction of each part of the ANU-BEC machine, focussing on those that formed a significant part of this thesis. The second part describes progress towards the construction of the degenerate Fermi gas machine. In addition, the introduction, consisting of chapters 1 and 2, gives background on both the field and the physics behind these experiments, and the conclusion, chapter 15, summarizes the key achievements reported in this thesis.

## 1.1 Degenerate quantum gases

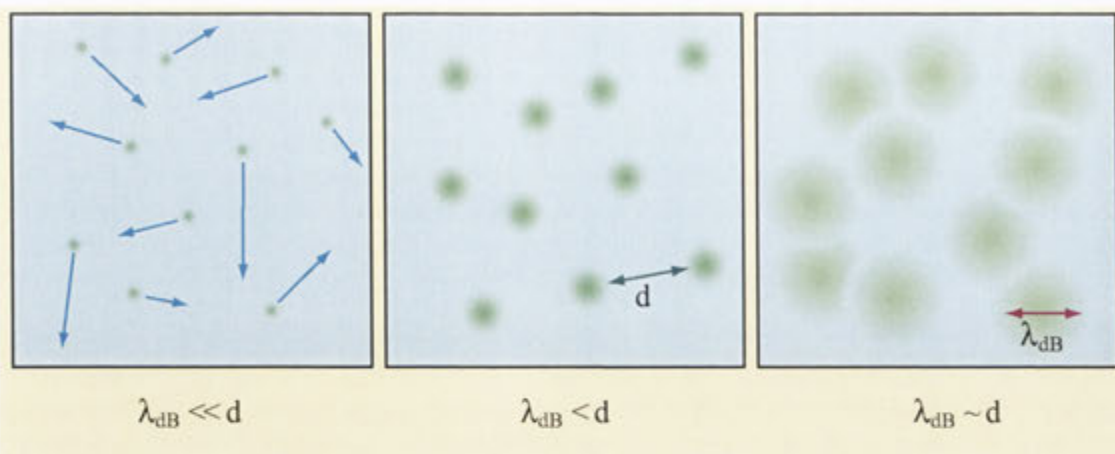
A degenerate quantum gas is a gas that has been cooled and confined so that it has reached degeneracy. In this context, degeneracy refers to the point at which individual atoms or molecules of the gas can no longer be distinguished based on their position. This situation arises from the fundamental “blurriness” of the universe described by quantum mechanics, as shown in figure 1.1. It is a vitally important threshold in the study of quantum systems, because a *polarized* sample of gas (one that has been selected with all atoms in the same internal quantum state), loses its last parameter of distinguishability once the atoms can no longer be identified by position.

This has dramatic effects on the thermodynamics of the gas sample. The Maxwell-Boltzmann statistics describing the traditional ideal gas are necessarily superseded by the appropriate quantum statistics: those described by Bose and Einstein, for integer spin particles [15], or those drawn from the work of Fermi and Dirac, in the case of half-integer spin particles. At the onset of degeneracy, Bosonic particles undergo a distinct phase-change driven purely by the statistical impetus of their quantum nature. Fermionic atoms exhibit no such dramatic change when they first reach degeneracy, but their behavior does alter more subtly as the thermodynamics becomes dominated by quantum, rather than classical, interactions.

The motivations for developing the ability to create degenerate quantum gases can be broadly separated into two categories: fundamental and applied. There are overwhelming reasons for the rapid expansion of this field from both points of view. Quantum mechanics forms one of the fundamental tenets of modern physics, and any method of elucidating it, especially one as promising as degenerate gas experiments, deserves consideration. The relative economy of these table-top experiments provides a refreshing contrast to other methods of understanding the smallest aspects of our universe. Contrasted with the potential gains when products based on “quantum electronics” and “quantum communications” begin entering the market in the not-too-distant future, little further applied justification need be made for investment in this field of physics.

A key interest in the pursuit of degeneracy in dilute gas samples is the de-





**1.1: Reaching degeneracy.** Phase space degeneracy occurs when the average spread of atoms, quantified by the deBroglie wavelength  $\lambda_{dB}$ , exceeds the average interparticle spacing,  $d \propto n^{-1/3}$ , where  $n$  is the standard spatial density of the gas sample. Normal room temperature conditions consist of “billiard ball” dynamics, where atoms are well considered as small, localized spheres, separated by vast interparticle spacings. Laser cooling techniques create samples in which atoms exhibit distinct delocalized properties, but are still separated by many times the average size of the atom. To date, in dilute gases, only combinations of both laser cooling techniques and evaporative cooling, have achieved temperatures and densities such that the delocalization of the confined atoms causes them to “overlap”.

velopment of quantum “laboratories” enabling the direct experimental observation of the multifarious counter-intuitive subtleties of quantum mechanics. Quantum mechanics has been called “the most successful theory of all time” because of its success explaining extremely complicated and bizarre behavior from elegant, relatively straightforward, first principles. However, the elegance of quantum mechanics derives from its inherent power, and the straightforward statement from first principles belies an unprecedented subtlety and complexity that may never be fully understood. A laboratory, containing apparatus to prepare, manipulate and dissect “quantum states”, can only serve to increase our understanding of this subtle and profound aspect of the universe.

Quantum mechanics is a fact of life. To pursue the construction of “quantum laboratories”, when every thing existing around us experiences and illustrates the working of quantum mechanics, may seem like folly. But in the best reductionist tradition, dilute degenerate gases offer the ability to transform the obscured complex interactions of real world quantum mechanics into the relative purity and precision of the laboratory. By working with dilute gases, traditional “classical” interactions

between atoms can be minimized. By creating large samples of coherent degenerate gases, the quantum statistics and interactions can be magnified. By working with atomic species that have accessible optical transitions and magnetic moments, the creation, manipulation and detection of quantum experiments can be performed elegantly and efficiently.

Understanding need not lead to exploitation and subjugation, but it sure helps. Along humanity's chosen path of technological idolatry, quantum control and engineering must be seen as the next step along the electronic-photonic timeline. "Quantum computing" and "quantum communication" are just two of the buzz words indicating the importance of increasing the understanding and control of quantum systems. If they have as large an influence on the world as their classical counterparts, experiments will soon be proceeding as rapidly through the private sector as they currently do in universities.

The development of the ANU-BEC is important, and will become increasingly important in the near future. From both a fundamental and technical viewpoint, the field of dilute degenerate quantum gas experiments appears to be a major new direction. The 2001 Nobel prize in physics was awarded to the experimenters who first realized dilute gas Bose-Einstein condensates in 1995. The Australian government has recently significantly increased its investment in this field. All predictions point towards degenerate quantum gases opening new fields of science, technology and industry in ways similar to laser physics and electronics over the last 50 years. As the first degenerate quantum gas experiment in Australia, the ANU-BEC, and this report of it, may represent the start of a very big part of Australia's future.

## 1.2 Bose Einstein condensates

The possibility of a phase change in an ideal gas of bosonic atoms, driven purely by the fundamental statistics of quantum mechanics, was proposed by Einstein in 1925, as an extension of the work of Bose on the statistics of photons [15]. At the time, this suggestion represented a lateral extension of the rapidly emerging universality of quantum mechanics, rather than a proposal for future experiments. In fact, only by the diligent work of London, Tisza and Landau was the theory revived to explain actual experimental results in liquid helium. They proposed the existence of a Bose-condensed fluid within a sample of "normal" helium to explain the observed superfluidity of ultracold helium. The concept has also been extended to explain the existence and behavior of superconductors.

The creation of a Bose-Einstein condensate in a dilute gas of atoms was proposed by Hecht in 1959 [16]. By the 1980s this led to a series of experiments investigating spin-polarized hydrogen in dilution refrigerators, then using magnetic traps and evaporative cooling [17]. In parallel, laser cooling of alkali and alkali metal atoms was being developed, ultimately leading to the Nobel prize in physics in 1997 [18]. Dilute gas BEC was finally realized in 1995, leading to the Nobel prize in 2001.

Since the first demonstration of dilute gas BEC in 1995, over 40 groups worldwide have achieved condensation of either  $^{87}\text{Rb}$  [19],  $^{23}\text{Na}$  [20],  $^7\text{Li}$  [21],  $^{133}\text{Cs}$  [22],  $^{85}\text{Rb}$

[23],  $^{41}\text{K}$  [24],  $^4\text{He}^*$  [25, 26],  $^2\text{H}$  [27] or  $^{174}\text{Yb}$  [28]. Many other groups are attempting to produce BEC, but it is still a complex and difficult experiment. The design of BEC machines is still evolving, yielding systems for producing larger, faster or more controllable BECs. The work presented in part I of this thesis represents a summary of the first ANU-BEC machine, the design of which borrowed heavily from the best aspects of the state-of-the-art machines operating in 1999.

There is no discernible limit to the future of BEC research. However, in the short time since its discovery, BEC has already begun to be used for tests of fundamental physics and preliminary demonstrations of applied serviceability.

Some of the first experiments with BECs aimed to measure the coherence of the matter waves created in a condensate. In the first of these experiments, Ketterle et. al. showed interference between two separate condensates as they expanded and overlapped from a double-well potential [29]. Later, Cornell et. al. used two pulses of polarized light to couple the constituents of a two-component condensate: the first to initially set the relative phase between the components; the second, after a variable delay, to read out the relative phase evolution between the components [30]. Phillips et. al. showed interference of multiple samples generated from a single condensate by standing-wave pulses operating as diffraction gratings [31].

Phillips et. al. went on to demonstrate the existence of distinct non-linearities in the mixing of packets of matter-waves generated from condensates. Using a very similar setup to that used in [31], they showed that three colliding wave-packets generated from a condensate using a series of Bragg pulses would generate a fourth in a manner analogous to four-wave mixing in non-linear optics [32].

These experiments certainly provided the first fundamental results generated using dilute-gas condensates. However, another vein of experiments was proceeding just as rapidly: those that hoped to generate practical tools for the more precise study of condensates and those that aimed to generate new physical tools employing condensates themselves.

Perhaps the best example of the latter type of experiment was the generation of the first "atom laser" in 1997, little more than a gradually released condensate [33]. However, the demonstration of this concept has since been followed by many experimental and theoretical studies, highlighting its importance. Phillips et. al. [34], Esslinger et. al. [35], as well as the ANU group [13], have all demonstrated early atom laser systems. A continuously pumped atom laser appears to be approaching, with the demonstration of various repetitive BEC generation and loading schemes [36].

Many new methods of probing and controlling condensates have been developed. One of these was the development of optical trapping of condensates in 1998 [37]. This led to the development of multiple-spin state condensates, now known as spinor condensates. These systems allow the study of fundamentally different physics to that explored with spin-polarized condensates.

A new technique for probing condensates was developed by Ketterle in 1999. Bragg spectroscopy was used to elucidate the mean-field energy, momentum uncertainty and coherence of a BEC [38]. The use of such a spectroscopic technique gave new levels of accuracy and precision to studies of BECs.

As new apparatus were developed, and new techniques were demonstrated, further experiments investigating the fundamental physics of condensates were performed. Dalibard et. al. observed arrays of up to 11 quantized vortices in rotating BECs in 2000 [39, 40]. Exquisitely controllable arrays of over 100 quantized vortices were observed by Ketterle et. al. in 2001 [41]. These experiments demonstrated one of the most graphic illustrations of the superfluid nature of BECs.

The extension of the optical trap to lattice-type structures allowed the manipulation of arrays of BECs. These optical-lattice techniques promised to provide a new level of control over the formation and interaction of multiple BECs. Early experiments by Kasevich et. al. demonstrated phenomena analogous to the Josephson effect and band transport of BEC in an optical lattice [42]. Several experiments studying superfluid flow and the Mott-insulator phase have since being conducted [24, 43]. The Mott insulator phase, in particular, promises significant advances in control over the formation of BEC states in preparation for further experiments. It has already been harnessed to reveal the analogy of macroscopic quantum atomic states to quantum light states. Using this technique, Kasevich et. al. claimed squeezing of the number state of a BEC [44].

Enhanced control has also allowed the development of precise studies of BECs of various dimensionalities. Many theoretical studies have investigated the behavior of dilute condensed gas in lower dimensions, leading to interesting predictions as to the behavior of such systems. The Tonks-Girardeau gases [45, 46], and Fermionization of Bosonic gases in low dimensions [47] are but two. Experimentally, optical traps, in particular, have provided new levels of control, opening new regimes of trap shape and trap tightness to be explored [48].

The development of chip-based BEC systems promises both exquisite control and suitability for use in applied and, eventually, industrial scenarios. Complicated and precise systems of magnetic fields, coupled with enhanced optical control, are already yielding experimental results unobtainable using other techniques. The first traps were successfully demonstrated in 2001 [49, 50]. Since then, many other groups have been developing similar traps due to their economy and configurability. Experiments harnessing these advancements are emerging as the field develops [51–53].

A major extension to the field of degenerate Bose gases was the emergence of the first molecular BECs. Much effort has been expended creating ultracold molecules by generating dimers from a standard BEC [54–57] or degenerate Fermi gas [58–61] sample. Recently, the key signatures of condensation have been observed in the clouds of molecules so formed, especially in systems based on degenerate Fermi gas experiments [62–67]. This is an extremely new field, but promises extensibility, and the gradual development of our understanding of quantum mechanics to more complicated systems.

### 1.3 Degenerate Fermi gases

And so the world of Bose-Einstein condensates came into being, allowing superlative control of macroscopic, detectable quantum states. Already, these abilities have borne a most exquisite and interesting fruit. Hope sprung toward the creation of new table-top “laboratories” that would bring experiments and theories together to elucidate the inner workings of quantum mechanics like never before.

But BECs are only half the story. The totality of the statistical description of quantum mechanics is borne of the synergy between two counterpoised counterparts. In addition to the essentially “inclusive” nature of the Bosonic interaction, the “exclusive” Pauli effect dominates the Fermionic description of reality. The somewhat undemonstrative twin of the “condensed” degenerate Bose gas is the “supported” degenerate Fermi gas.

The first experiment to produce a dilute degenerate Fermi gas was reported by Jin et. al. in 1999 [68]. They overcame the limitations inherent in the evaporative cooling of polarized Fermions to extremely low temperatures, discussed further in chapter 2, by creating a mixture of two spin-states of  $^{40}\text{K}$ . They also overcame many other technical limitations, such as inopportune wavelength requirements [69], and the difficulties of working with comparatively rare isotopes of common alkalis [70]. They produced a degenerate Fermi gas of approximately  $7 \times 10^5$   $^{40}\text{K}$  atoms, at approximately  $0.5 T_F$ .

Although mixing spin states was the first method used to produce a DFG, and has continued in its successful application by Jin et. al. since its first implementation, its difficulties and limitations meant that this approach was not used by most of the next groups to produce DFGs in 2001. Almost all of the new systems worked by sympathetically cooling the Fermionic isotope,  $^6\text{Li}$  or  $^{40}\text{K}$ , in the presence of a Bosonic species: either another isotope of the same atom, or a completely different atom [24].

Early in 2001, the second dilute degenerate Fermi gas was born [71]. It used the Bosonic isotope of lithium,  $^7\text{Li}$ , to sympathetically cool the Fermionic isotope,  $^6\text{Li}$ . This was a drastically different approach to that used by Jin et. al. in 1999. Sympathetic cooling had previously been applied in a two-species trapped ion experiment [72], and, more recently, in a two component  $^{87}\text{Rb}$  experiment [73]. Although a relatively new technique, sympathetic cooling represented a powerful generalization of the evaporative cooling process used in quantum degenerate gas experiments. It had the potential to extend evaporative cooling to previously unsuitable atomic species, including Fermionic samples, as well as more general systems, such as molecular samples, vastly increasing the potential scope of degeneracy experiments. It also provided a means of creating large degenerate quantum gases, as the population of the species of interest could be maintained, whilst the “cooling” species was depleted during evaporative cooling [74].

This second demonstration of Fermionic degeneracy was quickly followed by others, many using very similar systems. Salomon et. al. produced a  $^6\text{Li}$  Fermi gas sympathetically cooled by  $^7\text{Li}$  later in 2001. Early in 2002, O’Hara et. al. used a different approach to cool  $^6\text{Li}$  to degeneracy: they used Feshbach resonances to

tune the interactions in a mixture of the lowest two hyperfine states [75]. This allowed them to produce a DFG “all optically” [76]. A variation on this technique was later used by Grimm et. al. to produce degeneracy in  ${}^6\text{Li}$  [58]. By the end of 2002, two groups had also managed to sympathetically cool Fermionic atoms to degeneracy using a completely different cooling species. Inguscio et. al. cooled  ${}^{40}\text{K}$  to degeneracy using  ${}^{87}\text{Rb}$ , the same system that is considered in part II of this thesis [77]. Ketterle et. al. cooled  ${}^6\text{Li}$  to degeneracy using  ${}^{23}\text{Na}$  [78], in the process creating the largest DFG available [74]. Such sympathetic cooling schemes are so beneficial, in fact, that even Jin et. al. have attempted to augment their two-spin state experiment with a  ${}^{40}\text{K} / {}^{87}\text{Rb}$  experiment [79].

The relatively recent “mass” realization of DFG, coupled with the truly massive popularity of BEC experiments, may have suggested that little interesting Fermionic physics has yet been revealed in either experiment or theory. In practice, the development of degenerate Fermi gas studies have often been tied strongly to the development of the corresponding Bose study. However, this does not imply redundancy. It is precisely the differences between these two quantum protagonists that generates interest in these simultaneous studies.

Meystre et. al. has been particularly productive in the theoretical aspect of this field. Many of the experiments performed with BECs since their discovery in 1995 have been recast for the equivalent Fermionic system, yielding significant and testable results. The four-wave mixing results of Phillips et. al. [32] have been reconsidered for Fermionic constituents: very similar results ensue, but for drastically different reasons [80]. In a similar vein, diffraction of Fermions off an “atomic grating” leads to distinctly unique behavior dependent on whether the DFG is a normal fluid, or a Cooper-paired superfluid [81]. The same group have considered the two-component, or “spinor” Fermi gas [82]. Ketterle also considered matter wave amplification in the context of degenerate Fermi gases [83].

There are, of course, many behaviors of DFGs that cannot be related, in any distinct manner, to the antitheses of the corresponding Bosonic behavior. Many of these mannerisms derive from the suppression of interactions between Fermions below degeneracy, due to the Pauli exclusion principle and the absence of available transition states within the Fermi surface of the system, the so-called “Pauli blocking” effect. Examples include the existence of a zero sound mode [84], suppression of light scattering from the Fermi gas sample [41, 85–88], and the emergence of shell structure within the gas sample [89, 90]. Other theoretical studies have focussed on the dynamics of excitations within the DFG, and the ensuing observable changes to be tested by experiment [91–96].

Many of the early applied theory papers significantly preceded the experimental realizations of degeneracy in Fermionic systems. A fraction of these papers aimed to elucidate techniques for detecting the onset of Cooper-pairing in heavily degenerate Fermionic gases [97]. Other papers came soon after the realization of degeneracy, in an attempt to improve the efficiency of DFG production. One example is provided by Jin et. al., with their study of optimum cooling of two-component Fermionic systems [98]. Ketterle et. al. and Thomas et. al. both considered collisions in DFGs: in the zero-temperature limit, [99], and the strongly-interacting unitary

limit [100], respectively.

Similarly, the first experiments revealing physics associated with DFGs, like those conducted on the first BECs, were extensions and characterizations of the experimental method used to create the DQG. Jin et. al. published results on the near-degenerate collision properties of a gas of  $^{40}\text{K}$  four months before degeneracy was finally reported [101]. They went on to publish a variety of physical results on their unique system [102]. They were able to measure Pauli blocking of collisions within the DFG in 2001 [103]. They observed the behavior of their DFG across the transition from the collisionless to the hydrodynamic regime later that year [104]. In 2002, they reported observations of spin excitations in their degenerate Fermi gas [105].

Other groups had also started to produce experimental results from their DFGs. Hulet et. al. collected images illustrating the onset of a “supporting” Fermi pressure as their dilute gas sample reached degeneracy [71]. Thomas et. al. also studied the mechanical stability of their DFG in the strongly interacting unitary limit [106]. Inguscio et. al. reported the collapse of a DFG of  $^{40}\text{K}$  in the presence of a BEC of  $^{87}\text{Rb}$ , due to attractive interparticle interactions [107]. They went on to study dipolar oscillations of their DFG in 2003 [108]. Salomon et. al. studied the variation of the interaction energy inherent in their DFG near a Feshbach resonance [109]. Ketterle applied precision RF spectroscopy to verify the absence of strong interparticle interactions, and the correspondent clock shift that limits the accuracy Boson-based atom clocks [110].

An interesting sideline, especially from the perspective of groups working with sympathetically-cooled DFGs, is the productive physics to be exploited in systems containing both Bosonic and Fermionic, or two component Fermionic, degenerate gases. This has led to many theoretical studies, investigating effects such as stability of two-component Fermi gases [97] and the hydrodynamic description of Bose-Bose, Fermi-Fermi and Fermi-Bose systems [111]. More recently, Meystre et. al. studied two-component systems in further detail, calculating the quasiparticle spectrum of a Bose-Fermi system [112], and a method of tunable Raman coupling of the interactions between components in a two-component system [82]. Lewenstein et. al. have considered mixtures of Bose-Fermi two component systems in optical lattices with controllable interactions [113].

These types of systems have recently become available experimentally. Inguscio et. al., especially, have focussed on the mixtures of degenerate Bosonic  $^{87}\text{Rb}$  and Fermionic  $^{40}\text{K}$  created by their sympathetic cooling apparatus. They have shown DFG collapse as a result of attractive interactions mediated by degenerate Bosons in the mixture [107]. They have analyzed the effect on the stability of Bose-Fermi mixtures with respect to the mean-field energy of the sample [114]. They have also manipulated both Bosons and Fermions using optical lattices, and compared results [115]. They have yet to publish results on lattice-based manipulation of Bose-Fermi mixtures.

Perhaps the most intriguing research direction for degenerate Fermi gases, however, is the study, beyond degeneracy, of Cooper pairing and Bardeen-Cooper-Schrieffer (BCS) transitions [116]. Of particular interest is the crossover region

between BEC of tightly-bound diatomic molecules of Fermionic atoms, and the weakly-bound Cooper pairs of the same. This has elicited major interest in all aspects of the field. Innumerable theoretical papers discuss everything from the mechanisms of pairing [117, 118] to the most likely candidate BCS systems [97], to the methods of detecting the onset of the BCS transition and superfluidity [97].

The most favored technique, however, has been the careful manipulation of interaction potentials within degenerate systems using magnetic fields near Feshbach resonances. A large effort has gone into developing theories of Feshbach resonances in suitable systems [119]. Many groups have reported experimental verification of these types of theoretical studies [109, 120, 121].

The first fruit borne of Feshbach resonance studies was not the direct observation of the BCS transition that had been so anticipated. It has, however, generated its own important and interesting results. Molecular BECs generated from composite Fermions appear to provide many major advantages over molecules generated from Bosonic atoms. The suppressive nature of interactions amongst degenerate Fermions becomes vital, not so much during the formation of molecules, as afterwards, as they thermally equilibrate to degeneracy [122]. The lifetime of dimers generated from BECs is generally so short that it is not clear that the molecules so formed ever Bose condense in terms of their quasi-equilibrium distribution. Dimers formed from Fermions, on the other hand, can exhibit very long lifetimes, and distinct signatures of degeneracy have been observed.

In the rush to observe the BCS transition in a DFG, many experimental groups have produced molecules from a DFG using Feshbach resonance manipulation of scattering lengths [58–61]. Many of these have gone on to observe distinct signatures of condensation [64–67].

One group recently published the first intimation of observation of the BCS crossover. Jin et. al. reported observation of resonance condensation of Fermionic  $^{40}\text{K}$  pairs [123]. Hot on their tail, both Ketterle et. al. and Salomon et. al. have submitted articles detailing their experimental studies of the BEC-BCS crossover region in  $^6\text{Li}$  [124, 125].

## 1.4 A final motivation

Clearly, the field of dilute degenerate quantum gases has exploded into the awareness of modern physics. The theory envelops an amazing mix, from pure quantum mechanics, to heavily applied experimental modeling. The experiments are compact, fast, exciting, and proliferating at a rate with which it is difficult to remain abreast. The potential rewards for such endeavor, both fundamental and applied, are staggering.

The context has been set.



## 2 The pursuit of degeneracy

Producing and exploring degenerate quantum gases offers perhaps the best experimental insight into the bizarre world of quantum mechanics that modern physicists have come to accept as one of their most fundamental theories. Degeneracy purifies and amplifies the effect of the quantum world. It promises the creation of quantum laboratories, where the well-formed reductionist traditions of science can be brought to bear, slowly but surely unraveling the secrets of this subtle but powerful theory.

### 2.1 Background

“Degeneracy” intimates the situation in which two separate states can no longer be distinguished. Fundamentally, identical quantum systems, such as atoms, can not be distinguished. However, in everyday experience this is rarely an issue, as atoms do not occupy the same quantum volume of space, and can be localized based on their position. In this regime, the number of such accessible position states far outweigh atoms to fill them. Particles are adequately described by Maxwell-Boltzmann statistics, based on the statistical mechanics of distinguishable particles interacting with the quantum states of the system. A gas reaches phase-space degeneracy when it is confined in a trap such that the wavefunctions describing the atoms within the trap begin to overlap. In this regime, the indistinguishability of the particles becomes vitally important, and the statistics they exhibit diverge from the classical description. Bosonic particles experience Bose stimulation, tending to fall to the ground state of the trap, while Fermionic particles demonstrate Fermi exclusion, stacking up in the lowest energy states of the trap.

Degeneracy is important because it elucidates the fundamental aspects of quantum mechanics. Although our day-to-day experience of quantum mechanics is diluted through the multifarious interferences of complicated real-world systems, the underlying physical principles are the same. Degenerate systems purify and amplify the bizarre effects of quantum mechanics, enabling unequivocal comparisons with theory. They present situations in which the classical view of the universe can not explain the observed results, and motivate the development of our understanding of quantum mechanics.

Dilute degenerate quantum gases, as opposed to liquids, such as helium, and strongly-interacting gases, such as electron gases in metals, present an ideal system for manipulating and quantifying these effects. Classical interactions between atoms are minimized due to the relatively large separation between essentially non-interacting neutral particles. This is why these systems are seen as miniature quantum laboratories: the independent physical principles behind quantum mechanics

can be isolated, then manipulated and measured using a variety of tools at the experimenters' disposal.

Trapped atoms can be made to overlap by confining them closer together, or by making individual atoms larger. The size of an atom at very low temperatures is determined by its wave-like delocalization, which, as shown in figure 1.1, is quantified by the deBroglie wavelength,  $\lambda_{dB} \sim m/k_B T$ , where  $m$  is the particle mass,  $k_B$  is Boltzmann's constant, and  $T$  is the characteristic temperature of the sample of which the atom is a part. As an atom cloud is cooled, the average deBroglie wavelength of the constituent atoms increases. If this can be achieved while maintaining the density of the cloud, the atoms will eventually begin to overlap, and degeneracy will be achieved. Of course, if the density of the atoms can be increased, the interparticle spacing decreases, and degeneracy will occur at a smaller atom size, or higher temperature.

## 2.2 Phase space density

This dichotomy is quantified by the phase space density (PSD),  $\mathcal{D} = n\lambda_{dB}$ , where  $n$  represents the normal spatial density of the atomic sample. The PSD can be thought of as a measure of how many particles exist per quantum of phase space,  $\hbar^3$ . When the PSD exceeds one, degeneracy is achieved, as, on average, more than one particle cohabits the same quantum region of space. The PSD of a sample of air at standard temperature and pressure is  $\sim 10^{-8}$ . To reach degeneracy at standard pressure, such a sample of gas would need to be cooled to  $\sim 1$  mK. However, even this proves an unrealistic proposal.

Phase-space degeneracy is usually not observed in macroscopic gases because of processes that limit the phase-space density of normal everyday materials. The most serious limitation is interactions between particles. These interactions prevent the indefinite increase of the density of a gas of atoms. As the density increases, and atoms get closer together, interparticle interactions dominate the random thermal distribution of atoms in the system. Eventually, these interactions reorganize the structure of the cloud, typically forming an incompressible liquid or solid structure. A similar reaction occurs as the randomizing effect of thermal motion decreases as an atom cloud is cooled. This explains why cooling a standard sample of air to 1 mK does not create a degenerate gas: at approximately 77 K, the nitrogen becomes a liquid, at 63 K a solid.

Simple extrapolation to absolute zero would predict the solidification of practically every real substance. The exception is helium, which remains a liquid due to its low mass and high zero-point motion. The weak interactions that prevent helium solidifying also manifest themselves in the gas-like behavior of the liquid. At a temperature of 2.17 K,  $^4\text{He}$  reaches degeneracy and undergoes a phase change due to the loss of positional distinguishability of individual particles in the liquid.  $^3\text{He}$  reaches degeneracy at a similar temperature, but does not undergo a rapid phase change because it is a Fermionic particle. However, at a temperature of 2.6 mK, the degenerate  $^3\text{He}$  atoms begin to Cooper pair, forming composite Bosons, and

condensing.

Although it is difficult to store a sample as a gas near absolute zero, it is possible to cool it to very low temperatures for a short time, because the processes that lead to thermal equilibrium operate on different time scales from those that lead to solidification. More specifically, gases rethermalize via energy sharing elastic two-body collisions, while they begin to combine into larger structures via inelastic three-body recombination collisions. Both processes are mediated by collisions: the balance between the “good” two body collisions and the “bad” three body collisions is tunable, to some extent, via control over the density of the gas. The rate for two body collisions is proportional to the density,  $\nu_{2b} \propto \sigma_{2b}n_{av}$ , while the three-body collision rate varies as the square of the density,  $\nu_{3b} \propto \sigma_{3b}n_{av}^2$ .

Clearly, there exists a compromise between increasing density and the elastic collision rate, thereby speeding rethermalization, and minimizing density and the three-body recombination rate. The conditions at which these effects achieve optimum balance depends crucially on the collision cross-sections,  $\sigma_{2b}$  and  $\sigma_{3b}$ . These are essentially atomic parameters, and their ratio determines the suitability of certain atomic species for use in degenerate quantum gas experiments. New implementations of techniques for independently tuning cross-sections, such as the use of Feshbach resonances, are rapidly being developed [22, 109, 120, 121].

The other limitation mediating the precise form of these degenerate quantum gas experiments is technical in nature. The timescales for losses and heating through technical noise in the experiment need to be at least comparable to the timescales defined by the more fundamental considerations above. This places strict requirements on the level of isolation provided by the vacuum systems and magnetic trapping fields used in the last stages of the experiment, as discussed further in chapters 4 and 6.

The balance between all these factors depends crucially on the specific characteristics of the atomic species being studied. In most current degenerate quantum gas experiments, the timescales for factors such as collisions with background gas atoms due to imperfect vacuum are of the scale of one hundred seconds. In such experimental systems, the atoms most commonly condensed, such as  $^{87}\text{Rb}$  and  $^{23}\text{Na}$ , achieve optimum collision ratios between elastic and inelastic collisions at densities of the order of  $10^{13} \text{ cm}^{-3}$ . Higher densities generate disadvantageous three-body recombination rates. Lower densities limit the evaporation rate to time scales longer than the losses due to technical noise in the system.

The key to cooling a gas toward degeneracy is, therefore, the maintenance of suitable parameters, over appropriate timescales, to enable the realization of ultracold temperatures, at sufficient densities, without the loss of the gas. The practical methods, outlined below, can be separated broadly into two approaches: cooling at low densities to prevent the formation of non-gas phases, and cooling at high speeds to temporarily prevent the formation of non-gas phases. In practice, each technique is a combination of both: fighting to cool macroscopic gas samples to unprecedented temperatures over timescales rapid enough to prevent solidification.

At densities of  $10^{13} \text{ cm}^{-3}$ , a PSD of greater than one is only achieved below temperatures of  $\sim 500 \text{ nK}$ . Extraordinary levels of both cooling and environmental

isolation are necessary to reach these temperatures in a laboratory-based experiment. These requirements are discussed further below, and throughout Part I of this thesis. At the very pinnacle of technical implementation, across many simultaneous pieces of apparatus, conditions are maintained just long enough to cool a dilute gas of atoms to temperatures low enough, with densities high enough, that the wavefunctions describing the atoms making up the gas begin to overlap into degeneracy.


## 2.3 Practical considerations

Cooling a gas of atoms from room temperature to 500 nK is a difficult goal. In addition, fine control over the gas must be maintained at the final stages of the experiment to ensure that densities of the order of  $10^{13} \text{ cm}^{-3}$  can be maintained to achieve the best balance of elastic two-body rethermalizing collisions to inelastic three-body “re-phasing” collisions. In addition to providing an active cooling mechanism, the gas must be well isolated from the environment to maintain such low temperatures.

Current experiments achieve these goals through a series of techniques that have been developed over the past 20 years. Although many different techniques have been implemented, on many different atomic species, over this time, only one class of experiment has produced a dilute degenerate quantum gas successfully. This technique optically pre-cools a neutral atom gas, before evaporatively cooling it, at a carefully selected density, in a tunable conservative trap. This somewhat convoluted path uses some of the most advanced techniques developed in contemporary atom manipulation experiments, making the journey to degeneracy a potentially difficult process. However, as no other technique has demonstrated the required levels of control and isolation, these techniques are becoming more widely known.

The path through phase space is illustrated in figure 2.1. Neutral atoms are studied because they minimize long distance interactions between particles in the atom cloud. They are isolated in an ultra-high vacuum chamber to prevent interaction with any room temperature gas. The cloud is contained in a magnetic or AC electric field so that the atoms do not contact the wall of the vacuum chamber. To date, two main styles of cooling have been employed on the path to degeneracy: optical techniques and evaporative cooling. Optical techniques include magneto-optic trapping, optical molasses, and polarization gradient cooling. Degeneracy in a dilute gas was finally achieved through an ingenious method of forced evaporative cooling that was implemented in a magnetic trap [126–128].

In many modern atom optics experiments, an alkali gas is generated and isolated for study using a dispenser mounted on an electrical feedthrough in a vacuum system. Dispensers consist of a nichrome or stainless steel boat “filament” containing a salt of the alkali, and a reducing agent in powdered form. When current is passed through the nichrome boat, it heats up to over  $700 \text{ }^\circ\text{C}$ , driving a reduction-oxidation reaction that releases the alkali as a reasonably pure gas, while gettering unwanted species. At this point, the vacuum chamber is filled with approximately



**2.1: The path to degeneracy.** The path through phase space from a room temperature gas to a degenerate quantum gas is long and somewhat convoluted. Magneto-optic trapping bears the bulk of the responsibility, capturing dilute room temperature gas, cooling it to the order of  $200 \text{ } \mu\text{K}$ , as well as providing some significant compression. More advanced sub-Doppler optical techniques can cool the gas a little further. Adiabatic compression directly trades this extra cooling for higher position-space densities, without incurring a degradation of phase-space density. This prepares an atom sample at densities exhibiting a fortuitous balance of elastic (energy sharing) and inelastic (loss promoting) collision rates, allowing efficient evaporative cooling into the degenerate regime.

room temperature gas, typically at a density of  $10^{-9}$  torr, or  $10^9 \text{ cm}^{-3}$ . This yields a phase space density of  $10^{-18}$ , and represents the starting point on the journey to degeneracy.

A magneto-optic trap (MOT) can increase the PSD of such a gas by up to eleven orders of magnitude. It achieves this by both cooling and confining a cloud of gas from the sample liberated into the experimental chamber. The limitations on the temperature and density accessible in a MOT arise from the quantized interaction of the light with the atom cloud. The average velocity of the atoms in the MOT is limited to about one hundred times the “recoil” imparted by a single photon during the cooling process. The density is limited to values such that a photon emitted from an atom will not be re-absorbed by another atom in the cloud. These factors limit the lowest temperatures in the ANU-BEC MOTs to somewhat less than  $1 \text{ mK}$ , and the highest densities to  $10^{10} \text{ cm}^{-3}$ . These values yield a PSD of approximately

$10^{-7}$ , a massive increase, but still seven orders of magnitude away from degeneracy.

The MOT is vital for confining and cooling atoms from a room temperature sample. However, the configuration that yields the best collection efficiency from the background gas in the experimental chamber does not necessarily generate the highest PSD possible with optical techniques. Typically, the atoms from the collection MOT are transferred to a separate MOT in a ultra-high vacuum (UHV) chamber. This isolates the atoms to be cooled to degeneracy from the room-temperature vapor from which they were collected. Optical molasses and polarization gradient cooling are then implemented. These techniques can cool the atom gas almost to the recoil limit, an order of magnitude lower than the temperatures achievable in a MOT. As neither of these techniques confines the atoms, they are performed over a short period, during which the atom sample falls under gravity, and expands due to its residual kinetic energy. However, after only  $\sim 5$  ms of these sub-Doppler cooling techniques, the cloud has lost little density, but is an order of magnitude colder, yielding a phase space density of  $\sim 10^{-6}$ . Optical cooling techniques have great trouble increasing the phase space density further because of the recoil momentum associated with an absorption-emission cycle.

This problem is overcome in *conservative* traps. There are two main types in use in this field: far off-resonant optical traps, and magnetic traps. The far detuned optical traps use the AC electric fields generated by an intense laser beam to interact with dipole moments of the neutral atoms. The magnetic traps interact directly with the net magnetic moment of the alkali atoms. Both traps exert only conservative forces on the atoms, and do not perform any cooling per sé. They simply hold the atoms in a *tunable* trap, that can be tightened or loosened to adiabatically move along contours of constant PSD in figure 2.1. This is vital for controlling the density of the gas, balancing the rethermalizing two-body collisions with the depleting three-body collisions. The balance that is implemented depends on the time scales that are required for the experiment. Optical traps are generally tighter, and denser, yielding condensates more quickly. Magnetic traps are usually looser, and produce condensates less rapidly, but because the density of a magnetically trapped cloud is lower, three-body losses do not limit the lifetime so stringently. In either case, the atoms out of the MOT are generally at a density too low to optimize the collision rate balance. After catching the atoms, the tunable conservative trap adiabatically compresses them to their optimum density, raising the temperature, but maintaining phase space density.

The atoms are cooled to degeneracy by forced evaporative cooling. This process is analogous to the well-known natural process of evaporative cooling. Evaporative cooling involves selectively removing the most energetic atoms from a sample of gas in thermal equilibrium, reducing the average energy per particle, then allowing the remaining gas to re-equilibrate to a lower temperature. It can occur in any system in nature that demonstrates an energy-selective removal process, and a energy sharing process for the re-thermalization of the gas sample. This is the case, for instance, in a hot cup of coffee, where only the most energetic atoms can overcome the binding energy at the liquid-air surface of the coffee, to escape from the system. The atoms that remain have a lower average energy, and after re-distributing this energy via

collisions, will again reach thermal equilibrium, but at a lower temperature. The process reaches a consistent steady state when the coffee becomes thermalized with the surrounding environment.

Forced evaporative cooling involves implementing a tunable energy-selective removal process. This allows the extension of the concept of evaporative cooling to progressively lower temperatures. By setting the escape energy to remove the most energetic atoms from a trap, allowing the remaining atoms to rethermalize, then lowering the energy at which evaporation takes place, more and more energy can be removed from the trap, leaving fewer and fewer atoms. In fact, the primary limits on the temperatures achievable using this process are the initial number of atoms and the efficiency with which energy is carried away for each particle that is lost from the trap. This powerful system was the first, and to date, the only, method to directly produce a dilute degenerate quantum gas. Many of the practical considerations of evaporative cooling are considered further in chapter 7.

Efficient forced evaporative cooling represents a subtle interplay between the properties of a system defined by nature, and the ability of the experimenter to work within these properties to produce the desired result. At the most fundamental level, evaporation is mediated by the rethermalizing collisions that redistribute energy throughout the cloud after the most energetic atoms are removed. Effective rethermalizing collisions are two-body and elastic in nature. Competing collision pathways include inelastic three-body recombination processes. As described above, the balance between these factors determines the required specifications of every aspect of the final stages of a degenerate quantum gas experiment, from the speed of evaporation, to the tightness of the trap confining the atoms, to the required vacuum specification in the experimental chamber.

## 2.4 Thermodynamics

The specific nature of the rethermalization process is well understood for an ideal gas of distinguishable particles. Atoms are reasonably well described by Maxwell-Boltzmann statistics, apparently colliding as billiard balls, to redistribute energy throughout a sample. From the beginning of the experiment to the end of optical precooling, the conditions in the atomic gas sample are such that a cloud left to its own devices would be suitably described by this sort of system. The sample is said to be “thermal”, and, given time, equilibrates to share energy throughout the sample as described by equation 2.1, where  $n_i$  is the population of the  $i^{\text{th}}$  energy level, with energy  $\epsilon_i$ ,  $\mu$  is the chemical potential,  $k_B$  Boltzmann’s constant, and  $T$  the temperature of the gas sample.

$$n_i = \frac{1}{e^{-(\epsilon_i - \mu)/k_B T}} \quad (2.1)$$

The internal structure and quantum nature of the atoms is not overly important at this stage of the experiment. Although neglecting the quantum mechanical characteristics of the particles is a gross simplification of reality, practically, the energy

sharing behavior is well described by assuming that colliding atoms are distinguishable. Collision rates can be summarized as discussed in section 2.2, above. At the quantum mechanical level, the two-particle system involved in a collision can be described by equation 2.3, where  $\Psi$  is the total wavefunction of the particle,  $\psi_a$  the position wavefunction of the  $a^{\text{th}}$  particle, and  $\xi$  the spin wavefunction.

$$\text{For separable position and spin coordinates:} \quad \Psi = \psi(\mathbf{r})\xi(\mathbf{s}) \quad (2.2)$$

$$\text{Spatial part of two-species wavefunction:} \quad \psi(\mathbf{r}_1, \mathbf{r}_2) = \psi_a(\mathbf{r}_1)\psi_b(\mathbf{r}_2) \quad (2.3)$$

This representation fundamentally implies distinguishability, as it assigns distinct quantum states (a and b) to particles 1 and 2, respectively. One method of extending this description of the collision into the calculable domain is using partial wave representations. The Lippmann-Schwinger representation describes the incoming and outgoing partial waves under the interaction of the collision, where  $P_l$  is the Legendre polynomial of order  $l$ ,  $\theta$  is the angle between the incident particles as they collide,  $i = \sqrt{-1}$ ,  $k$  is the momentum of the collision in the center-of-mass frame, and  $f_l$  [129].

$$\langle \mathbf{x} | \Psi \rangle \propto \sum_l P_l(\cos\theta) \left( [1 + 2ikf_l(k)] \frac{e^{ikr}}{r} - \frac{e^{-i(kr-l\pi)}}{r} \right) |\xi \rangle \quad (2.4)$$

The partial wave representation of the collision is particularly useful at the low temperatures demonstrated in a degenerate quantum gas experiment, because only the first few partial wave modes are accessible to atoms with such low kinetic energy. An analysis of the contribution of various partial waves to collisions at various temperatures shows that, for instance, the P and D partial wave contributions “freeze out” of the collision dynamics for a  $^{40}\text{K}$  sample at temperatures of  $100 \mu\text{K}$  and  $500 \mu\text{K}$  respectively [69]. Below  $100 \mu\text{K}$ , only the s-wave partial wave mode contributes significantly to collision processes. Rubidium exhibits similar properties at slightly lower temperatures:  $30 \mu\text{K}$  and  $150 \mu\text{K}$  respectively.

Distinguishable particles imply no further regulation on this description of the collision. However, a similar description of indistinguishable quantum particles necessitates consideration of the symmetries of the system: the wave function for Bosonic particles must maintain an overall symmetry under particle interchange, while the Fermion wave function must exhibit anti-symmetry.

As the gas cools towards degeneracy, these Maxwell-Boltzmann statistics of distinguishable particles become less and less appropriate. The quantum nature of the particles being cooled becomes a dominant effect. As the gas is cooled, the average constituent atom is left with only enough energy to probe the lowest collision mode partial wave: s-wave scattering. This, in turn, leaves the spatial part of the collision wavefunction represented in equation 2.4 symmetric, yielding significant consequences for the dynamics of colliding indistinguishable quantum particles.

Collisions within the magnetically trapped, spin-polarized gas samples, are, by definition, spin-symmetric under particle interchange. As the gas cools, the spatial dynamics becomes limited to symmetric s-wave interactions. Bosonic particles



begin to experience Bose stimulation. Fermionic particles, which require overall anti-symmetry in the two-particle wavefunction, simply stop interacting. At this point, the quantum nature of the particles becomes vitally important to the collision dynamics that mediate the efficacy of evaporative cooling. These effects dominate the evaporative cooling stage of the experiment, and the approach to degeneracy.

## 2.5 Bose-Einstein condensation

Bosons are integer spin particles, exhibiting a symmetric position state under particle interchange. This, of course, has important implications for the consideration of collisional dynamics, evaporative cooling, and the journey towards degeneracy in a Bosonic gas. The spatial part of the wavefunction describing the interaction between Bosonic particles during a collision is modified, from equation 2.3, to ensure symmetry is maintained:

$$\psi(\mathbf{r}_1, \mathbf{r}_2) = \frac{1}{\sqrt{2\pi}} [\psi_a(\mathbf{r}_1)\psi_b(\mathbf{r}_2) + \psi_b(\mathbf{r}_1)\psi_a(\mathbf{r}_2)] \quad (2.5)$$

Considering the collision process at very low energies, it is apparent that the symmetry of the spatial part of the s-wave component of the partial wave description of the collision in equation 2.4, suits the Bosonic requirement. That is, as the s-wave description of the collision is symmetric under particle interchange, and the collision involves *polarized* Bosons in the same spin state, the entire collision wavefunction is symmetric. Collisions between spin-polarized Bosons at low temperatures tend to satisfy the quantum mechanical symmetry requirements, unlike the equivalent analysis for Fermions, which is discussed further below.

This behavior leads to a modification of the expected scattering rates. Bosonic stimulation leads to increased elastic scattering rates as a gas of Bosons is cooled closer to degeneracy. These increased scattering rates increase the efficiency of evaporation at the very end of a cooling cycle towards BEC.

The other significant modification is to classical thermodynamic statistics near degeneracy. Equation 2.6 describes the distribution of energies within a Bosonic sample in thermal equilibrium. Far from degeneracy, the chemical potential is large and negative, making the -1 a small perturbation to the classical energy distribution of equation 2.1. Near degeneracy, however, this perturbation gains importance as the magnitude of the chemical potential decreases, eventually becoming positive.

$$n_i = \frac{1}{e^{-(\epsilon_i - \mu)/k_B T} - 1} \quad (2.6)$$

As a gas of Bosons nears, then passes degeneracy, the distribution elucidated by equation 2.7 describes an enhancement of ground state occupation. The maximum occupation possible in each state of the system is revealed by re-expressing the distribution of equation 2.6, using the maximum value of the chemical potential: the ground state energy.

$$n_{i_{max}} = \frac{1}{e^{-(\epsilon_i - \epsilon_0)/k_B T} - 1} \quad (2.7)$$

The most obvious feature of the availability of states throughout such a system is the fact that the ground state ( $\epsilon_i = \epsilon_0$ ) can maintain an essentially limitless occupancy. As more and more thermal energy is removed from a gas of Bosons, the constituent atoms begin to aggregate in the ground state of the system. The distinct onset of these statistics, as degeneracy is achieved, is the process of Bose-Einstein condensation.

The temperature at which these effects begin to occur is referred to as the critical temperature,  $T_{crit}$ . It is the temperature of this sharp phase transition associated with Bose-Einstein condensation, where  $m$  is the mass of the particles being condensed,  $n$  is the density of the sample, and  $\zeta$  is the Riemann-Zeta function.

$$T_{crit} = \frac{2\pi\hbar^2}{k_B m} \left[ \frac{n}{\zeta(3/2)} \right]^{2/3} \quad (2.8)$$

In a three dimensional harmonic trap, the density is related to the temperature, and the transition temperature can be expressed purely in terms of the total number of atoms,  $N$ , and the trapping frequencies of the confining trap, with geometric mean  $\omega$ , as shown in equation 2.9.

$$T_{crit} = \frac{\hbar\omega}{k_B} \frac{N^{1/3}}{[\zeta(3)]^{1/3}} \quad (2.9)$$

Although a clear phase transition to BEC occurs at  $T_{crit}$ , the number of atoms in the ground state builds rapidly, but continuously, as the temperature is lowered beyond  $T_{crit}$ . The number of Bosons in the ground state,  $N_0$ , at temperature  $T$  is given by equation 2.10

$$N_0 = N[1 - (T/T_C)^{3/2}] \quad (2.10)$$

A key benefit of systems of dilute gas BEC are that they contain a macroscopically occupied, amplified quantum sample, with only small residual classical interparticle interactions relative to other systems in which BEC is observed. This is most important because it allows the construction of useful theoretical models, even though these systems can contain up to  $10^6$  individual atoms. This is a consequence of the fact that the BEC can be well described by a single wavefunction constructed from a series of identical individual atom wavefunctions.

## 2.6 Degenerate Fermi gases

Fermions are half-integer spin particles. They support an anti-symmetric wavefunction under particle interchange, as manifested in the wavefunction describing the spatial part of a collision between two spin-polarized Fermions, shown in equation 2.11. A sample of Fermions polarized into identical internal quantum states

cannot support inter-particle collisions if only symmetric s-wave channels are available.

$$\psi(\mathbf{r}_1, \mathbf{r}_2) = \frac{1}{\sqrt{2\pi}}[\psi_a(\mathbf{r}_1)\psi_b(\mathbf{r}_2) - \psi_b(\mathbf{r}_1)\psi_a(\mathbf{r}_2)] \quad (2.11)$$

This is of drastic importance to the production of degenerate Fermi gases, because no collisions implies no evaporative cooling. There is little to say beyond this: the laws of physics are immutable, and Fermions in identical quantum states simply cannot collide in a symmetric manner.

The solution is to provide other, allowed, collision modes to mediate the rethermalization of the gas sample during evaporation. This is achieved by providing atoms in *other* quantum states, mixed throughout the Fermionic sample, with which the polarized Fermions can collide. These other atoms simply need to be in another quantum state: they can be the same species in a different, magnetically-trappable, internal quantum state; they can be a different isotope of the same atom; or they can be another species of atom entirely. The only requirements, as with all evaporative cooling, are a suitable manner of holding both species in a single region of space, and good elastic collision cross-sections between the species to facilitate energy sharing and rethermalization. These requirements are by no means trivial, however, and already several systems in each of the classes mentioned have been studied in the pursuit of degeneracy.

In the same way that Bosons experience amplified collision rates as they near degeneracy, Fermions experience a suppression of collision rates, even with other species, as they near degeneracy. It is a direct result of Fermi statistics, illustrated in equation 2.12, and Pauli blocking.

$$n_i = \frac{1}{e^{-(\epsilon_i - \mu)/k_B T} + 1} \quad (2.12)$$

The defining behavior of a Fermi gas as it reaches degeneracy is the stacking into available trap states, with no two particles ever occupying the same state. The result is an apparent “Fermi pressure” supporting the cold atom cloud against further compression, in stark contrast to the “collapse” associated with a Bose-Einstein condensate. The other effect, alluded to above, is the suppression of collisions and interaction between trapped Fermions, as the population of energetically available states becomes saturated. The net suppression of interactions promises to be one of the most interesting aspects of degenerate Fermionic experiments: already molecular quantum gases are harnessing these characteristics for the benefit of other systems. This behavior, again, contrasts with BEC behavior, in which the interaction energy in the condensate is typically many times larger than the residual kinetic energy of the constituent particles.

The Fermi energy,  $\epsilon_F$ , shown in equation 2.13, represents the maximum energy state occupied in the system, and is the value of the chemical potential at zero temperature.

$$\epsilon_F = \frac{(6\pi^2)^{2/3} \hbar^2}{2m} n^{2/3} \quad (2.13)$$

The onset of degeneracy occurs at the Fermi temperature,  $T_F = \epsilon_F/k_B$ , as shown in equation 2.14. This is similar to the transition temperature of a BEC, within a factor of the order of unity.

$$T_F = \frac{(6\pi^2)^{2/3} \hbar^2}{2 m k_B} n^{2/3} \quad (2.14)$$

For a three-dimensional harmonic trap, this reduces to:

$$T_F = \frac{\hbar\omega}{k_B} (6N)^{1/3} \quad (2.15)$$

## 2.7 Beyond degeneracy

This briefest summary of the journey and the physics of along the path to degeneracy, aims but to familiarize the reader with the challenges involved in the construction of the apparatus outlined throughout the rest of the thesis. Like chapter 1, this chapter furnishes field-wide context for the story that follows.

**Part I**


**Bose-Einstein condensation**

### 3 ANU-BEC: Australia's first Bose-Einstein condensate machine

The initial 18 months of this thesis were employed working to produce the first Bose-Einstein condensate in Australia. These endeavors were aided by the assistance of three co-workers: John Close, the supervisor of this thesis and leader of the group; Ulrich Kallmann, a postdoctoral fellow from Germany who worked with the group for the first year; and Jessica Lye, a fellow postgraduate student. The team worked towards the common goal, at times working on independent aspects of the machine, at others, combining efforts to push toward BEC.

Part I of the thesis aims to present a full account of the ANU-BEC apparatus. The main focus is the first BEC machine created at ANU: ANU-BEC-I. However, much of the work of this thesis that lead towards the creation of ANU-BEC-I was built-on and improved to create ANU-BEC-II, the second iteration of the ANU-BEC machine. Brief summaries of the improvements that constituted ANU-BEC-II are included throughout each chapter, to contrast those aspects of ANU-BEC-I that worked well, with those that were refined. As this thesis is being written, a new machine, ANU-BEC-III, is being assembled. As this machine further builds on lessons learnt during ANU-BEC-I and ANU-BEC-II, some brief summaries of its design are included. However, as ANU-BEC-III has not yet produced BEC it is not emphasized strongly in the text because it is difficult to make solid comparisons as to performance.

Many people have worked on the ANU-BEC-I, ANU-BEC-II and ANU-BEC-III machines described in this part of the thesis. The small size of the ANU-BEC group has ensured that everybody has contributed through discussion and advice to almost all aspects of each machine. However, inarguably, each part of the machine has developed under the watchful eye of one or two main protagonists. Creating a useful report summarizing these machines inevitably involves reporting on work primarily conducted by some of these people. Those aspects of the machines which constituted a significant investment and return on the work completed for this thesis are emphasized. Those parts of the machine that were contributed to mainly through group discussion and isolated problem solving, are simply summarized. In the case of the ANU-BEC-I machine, their full disclosure is left to Jessica Lye, hereafter JL, in her thesis, *Dynamic non-destructive detection of Bose-Einstein condensates and atom lasers* [1]. The ANU-BEC-II machine is fully summarized in Nick Robins' (NR) thesis, *Bose-Einstein condensation and the atom laser* [2].



**3.1: The main ANU-BEC-I table.** The photo includes the vacuum system, optics, magnetic trap and the imaging system. It does not show the laser systems, which are on another optical table to the left.


### 3.1 The ANU-BEC in context

The ANU-BEC was a significant achievement. Assembly of the apparatus began in March 1999, and BEC was first observed late one evening in April 2001. This made it the first BEC in Australia. It was also one of the fastest realizations of BEC in the world, made doubly impressive by the fact that none of the original team had significant experience in the field of laser cooling and trapping before the experiment began.

Such success can probably be attributed, in part, to the careful selection of the best of the technologies that were beginning to emerge from the main BEC groups in the world in 1999, as they built their second generation machines. This allowed the ANU-BEC team to build a remarkably elegant, state-of-the-art device on the first attempt, that to this day represents the epitome of design for a BEC machine based on double magneto-optic traps and a macroscopic magnetic trap.

### 3.2 Overview of the ANU-BEC-I machine

Photos of most of the complete ANU-BEC-I machine is shown in figures 3.1 and 3.2. The following overview summarizes each important aspect of the machine, to provide context for the more complete descriptions in the chapters following.



**3.2: The ANU-BEC-I control systems.** The photo includes the computer control system and much of the supporting electronics of the experiment.

In section 3.3 below, a typical experimental “run” of the ANU-BEC-I machine is presented to illustrate the integration of each part of the experiment to create BEC.


ANU-BEC-I is a Bose-Einstein condensate of  $^{87}\text{Rb}$  atoms. It can be produced as a dual-component cloud containing up to  $10^5$  atoms in the BEC, or a pure BEC of approximately  $10^4$  atoms.

The ANU-BEC-I machine is based on the well known double MOT design first introduced by Myatt et. al. [130]. The far right hand side of figure 3.3 shows the basic configuration of the double MOT chamber. This design supplies atoms to the ultrahigh vacuum “science cell” MOT from a much higher pressure “collection MOT”.

A key digression of the ANU machine from Myatt’s design is that the MOTs are oriented vertically, the collection MOT positioned approximately 50 cm above the science MOT. This allows the use of only a simple resonant push beam to transfer atoms from one trap to the other, without the need for hexapole magnetic guiding.

The key features of the vacuum system shown in figure 3.3 are the glass “science chamber” and stainless steel “collection” chamber at one end of the machine, and the ion pump and titanium sublimation pump at the other. The pumps are separated from the science end of the machine to minimize the effect of stray magnetic fields from the ion pump on the atoms during transfer between the collection MOT and science MOT, and in the magnetic trap itself. The design and construction of the vacuum system are detailed thoroughly in chapter 4.






**3.3: The vacuum chamber.** The double MOT chambers are shown on the right. The stainless steel collection MOT chamber is at the top, and the glass science chamber is positioned below. The ion pump and titanium sublimation pump are shown on the left. The tubes connecting the pumps to the experimental chambers are kept large (4") to ensure good conductance and pumping speeds.

The ANU-BEC-I used a comprehensive system of lasers and optics to generate light for manipulation, cooling and trapping of atoms throughout the experiment. Figure 3.4 shows detail of the optical setup. The most important task performed by the optical setup was the generation and preparation of light for: locking the laser frequency; the collection MOT; the science MOT, “pushing” atoms between the two MOTs; polarization gradient cooling; “pumping” atoms into the correct  $m_F$  state before magnetic trapping; and imaging.

Light for the magneto-optic traps was generated by external cavity diode lasers: one at the  $5^2S_{1/2}F = 2, 5^2P_{3/2}F' = 1-3$   $^{87}\text{Rb}$  crossover, for the cooling light, and one at the  $5^2S_{1/2}F = 1, 5^2P_{3/2}F' = 2$  transition for the repumping light. The cooling light was amplified using a tapered amplifier, to provide approximately 300 mW of light. This light was split into two major components, for the collection MOT and science MOT, respectively, and one minor component, for locking the laser and other diagnostics. The lasers were locked with a compact saturated absorption setup using Zeeman modulation to generate the locking signal.

The cooling light for each magneto-optic trap was double passed through a 110 MHz AOM to bring it close to the  $5^2S_{1/2}F = 2, 5^2P_{3/2}F' = 3$  cooling transition. The




**3.4: The optical setup.** Shown is a snapshot from the main experimental table. It include much of the optics for splitting and preparing the MOT beams, as well aa beams for optical pumping and imaging. Not shown is the laser system, locking optics, the AOM circuits or the optic fibers that connect the main cooling laser to the experiment. They are positioned on a separate table to ensure maximum stability.

cooling light for each MOT could be thus be attenuated and detuned independently. The cooling light was coupled to the rest of the experiment through single-mode polarization maintaining fibers, yielding  $TEM_{00}$  MOT beams. The light was divided into individual MOT beams using half wave plates and polarizing beam splitter cubes, yielding linearly polarized light. The light was rendered circularly polarized using quarter wave plates set to produce  $\sigma^+$  and  $\sigma^-$  beams as required. The repump light was mixed with the cooling light on a spare port of one of the mixing polarizing beam splitter cubes, as its polarization and intensity balance were not critical.

The retro-reflected collection MOT was run 15 MHz detuned, with a total power at the MOT of approximately 20 mW in the three, 15 mm diameter, beams. The six-beam science MOT was run 15 MHz detuned, with approximately 45 mW in the six, 15 mm diameter, beams.

Optical molasses was implemented over 3 ms in the science MOT by turning off the MOT coils, leaving a very small residual field of the order of 10 mG. Polarization gradient cooling was implemented by further attenuating and detuning the laser 25 MHz away from the  $5^2S_{1/2}F = 2, 5^2P_{3/2}F' = 3$  resonance over 3 ms. Shutters blocked all but one of the MOT beams, and a small coil colinear with the remaining beam



**3.5: The magnetic trap coils.** The coil in the foreground, and the coil on the opposite side of the glass cell, are the “quadrupole” coils, that are used for creating the science MOT, catching the atoms in the initial quadrupole trap, and finally, in the full Ioffe-Pritchard trap. The third coil, oriented perpendicular to the quadrupole coils, is the smaller conical Ioffe coil. It transforms the quadrupole trap into a full IP trap as current is ramped from zero to the full operating specification. Also seen at the back of the closest quadrupole coil are the central hole for the MOT beams, the wiring harness, and the water cooling pipes, and, behind them, the slit in the copper former.

was turned on for 1 ms, generating  $\sigma^+$  light to optically pump the atoms to the  $F = 2$ ,  $m_F = +2$  ground state prior to magnetic trapping.

The imaging beam was derived from the collection MOT light. Control over these two parts of the experiment was separated using mechanical shutters. The properties of the imaging beam were thus controlled by the collection MOT AOM. The length and offset of the pulse used for imaging were finely controlled using an external pulse generator triggered by the computer control system. The detuning, attenuation and pulse characteristics were varied to produce good optical densities and signal-to-noise ratios on the images collected. The optical system is summarized in chapter 5.

The magnetic trap was based on the QUIC trap design of Esslinger et. al. [131]. A photo of the trap is shown in figure 3.5. This was a new design, just published at the start of 1999, when it was chosen for the ANU-BEC-I machine. However, it has proved to be a most fortuitous choice, as many groups have moved towards this

less flexible but much simpler and more efficient design since.

The QUIC trap is a simply implemented, low current, three coil magnetic trap that produces an Ioffe-Pritchard type magnetic field. What it trades for its structural simplicity is a lack of dynamic tunability, and a consequent emphasis on initial design.

The important stages of operation of the trap comprised: loading of atoms into the quadrupole trap; transporting atoms into the full Ioffe-Pritchard magnetic trap; tuning the trap parameters to optimize evaporation; and rapid switch-off of the trap for imaging.

The atoms are loaded into the magnetic trap by switching the quadrupole coils on at 12 A, creating a field gradient of approximately  $150 \text{ G cm}^{-1}$ . These are the same coils used in the science MOT, running at higher current. Thus, the magnetic fields of the science MOT and quadrupole trap are perfectly aligned. The current at switch-on is set to optimize the capture of atoms into the magnetic trap. At 12 A up to 75 % of atoms from the science MOT are captured in the quadrupole trap. The current is then ramped up over 2 s to the full operating current of the trap, 35 A, creating field gradients of approximately  $450 \text{ G cm}^{-1}$ . This compresses the cloud adiabatically, trading increased density for a higher temperature.

The trap is morphed from a linear quadrupole trap to a harmonic Ioffe-Pritchard (IP) trap using a third coil, the Ioffe coil, oriented perpendicular to the quadrupole coils. This is by ramping the current in the Ioffe coil from 0 - 35 A over 1 s. This yields radial trapping gradients of  $200 \text{ G cm}^{-1}$ , and an axial field curvature of  $220 \text{ G cm}^{-2}$ . Again, this process is roughly adiabatic. As well as changing the form of the trap, this process physically moves the atoms 9 mm towards the Ioffe coil. This transfer can be implemented with up to 80 % efficiency.

The defining features of the IP trap are its harmonic shape and its lack of a field zero in the trapping region. Its most important specifications are the harmonic trapping frequencies of the trap, and the stability of the bias field at the bottom of the trap. One of the key trade-offs in the QUIC implementation of the IP trap is the lack of independent control over the bias field and confinement of the trap. The interaction between the bias field and the trapping frequencies must be designed into the relative geometries of each coil. The ANU-BEC-I machine generated an IP trap with trapping frequencies of approximately  $200 \times 200 \times 20 \text{ Hz}$ , with a bias of 1 G.

The electronics running the magnetic trap were a mix of commercial and home-made parts. The low current (35 A) of the trap allowed the use of commercial power supplies and semi-conductor switches. These were integrated into a control system using other, home-made, electronics and circuits. In addition, the ANU-BEC-I used a home-made "constant current sink", a low power, current-controlled dissipative supply, to provide some small tunability of the trap. These electronics combined to control each aspect of the trap, from the shape of the potentials formed, to the transition between potentials, to the switch off of the magnetic field.

The design and implementation of this trap formed a large portion of this thesis, and it is discussed extensively in chapter 6.

Atoms were cooled across the BEC transition using radio-frequency evapora-



**3.6: The imaging optics and CCD camera.** The CCD camera can be seen in the background. Much of the imaging optics is ensconced within the general laser cooling optics.

tion from the magnetic trap. This simple and elegant technique has been used on most machines since it was first implemented in 1995 [132]. Despite its amazing performance, evaporative cooling is technically quite easy to implement. The bulk of the execution of an evaporative cooling system is achieved via efficient optical pre-cooling and a well designed magnetic trap. The RF-induced evaporation itself involves a low-RF signal generator (30 MHz) and a single-loop antenna. The indistinctly visible antenna is located between the far quadrupole coil and the glass cell in figure 3.5.

The application of RF evaporation is largely a matter of fine-tuning the parameters used in the RF frequency sweep. A short summary of the procedure is presented in chapters 7 and 10.

The primary method by which data is collected from the ANU-BEC-I machine is via absorption imaging using a CCD camera. There are two main aspects to the imaging process: the pulse of light used to illuminate the atoms, and the system for collecting the light after it has interacted with the cold atom cloud. The pulse of light is defined by its frequency, intensity, duration and the time allowed for expansion of the cloud before it is applied. The collection of data is implemented using imaging optics and a CCD camera.

The imaging light is generated from the upper MOT optical setup, as explained above. The imaging in the ANU-BEC-I experiment was performed relative to the  $F = 2, m_F = 2$  to  $F' = 3, m_{F'} = 3$  closed transition, using  $\sigma^+$  polarized light. Typical values for the detuning, intensity, delay allowing expansion, and duration of the imaging pulse when imaging a BEC, were  $\sim 5$  MHz, somewhat less than 1 mW, tuned relative to the detuning to provide appropriate optical depth, 10 ms and 100  $\mu$ s, respectively.

The CCD camera used in the ANU-BEC-I experiment was a Princeton Instru-

ments model RTE-768-K. This CCD camera had  $768 \times 512$ ,  $9 \mu\text{m}$  square pixels. The pixels had a well depth of  $\sim 80,000 e^-$ , and the camera had a read noise of approximately  $12 e^-$ . It used a 12-bit digitizer to yield true 12-bit information. The imaging system used was a single lens 2-f system yielding one-to-one magnification.

The first of these systems, the image pulse generation and manipulation, is summarized in chapter 8. The design of an optimized imaging system and selection of the CCD camera is considered in great detail throughout the rest of that chapter.

### 3.3 The ANU-BEC-I machine in operation

Table 3.3 shows a succinct breakdown of the timing at various stages of the ANU-BEC-I experiment. The timing information in table 3.3 is incorporated into the experiment via the computer control software, described briefly above, and more completely in chapter 9.

The first stage of the experiment involves repeatedly creating a magneto-optic trap (MOT) in the collection chamber, and transferring atoms from the collection MOT into the science MOT. This process forms the first part of table 3.3. The most important points to note include: the collection MOT is repeatedly loaded for 200 ms; all optical and magnetic fields comprising the collection MOT, as well as the Rb dispenser, are turned off during the push; and the laser is tuned on to resonance during the push. Not indicated is the fact that the dispensers producing rubidium are generally run continuously at all times other than when they are turned off for the push. Note that the loading process is repeated 250 times over 50 s to capture a large number of atoms in the science MOT. Each parameter indicated in the table has been optimized to produce the largest possible number of atoms in the science MOT by the end of this stage of the experiment.

During the loading process, the science MOT is simply run continuously: atoms pushed from the collection MOT are captured as they pass through it. When the loading process has completed, approximately  $10^9$  atoms are stored in the science MOT, at a density of approximately  $10^{10} \text{ cm}^{-3}$  and a temperature of  $200 \mu\text{K}$ . This corresponds to a phase space density of roughly  $10^{-7}$ .

At this stage, the experiment moves to the second part of the table. All magnetic fields are turned off, and for 3 ms the atoms undergo optical molasses cooling. Over the following 3 ms, the MOT beams are detuned from 15 MHz to 25 MHz. This implements  $\sigma^+/\sigma^-$  polarization gradient cooling. It decreases the temperature of the atoms to approximately  $30 \mu\text{K}$ , increasing the phase space density to  $10^{-7}$ . During this stage of the experiment the atoms are not trapped, so they begin to fall under gravity.

The atoms are recaptured in a magnetic trap. However, before the trap is switched on, the atoms are optically “pumped” into the  $F = 2$ ,  $m_F = +2$  ground state of  $^{87}\text{Rb}$ . This requires a short pulse of  $\sigma^+$  polarized light. This is implemented by pulsing one MOT beam on for 1 ms using a combination of mechanical shutters and an AOM, and by pulsing a magnetic field coil aligned with the beam.

The quadrupole coils of the magnetic trap are then turned on as quickly as

Event Description	Start Time	Duration	Atom Numbers
<b>1. Load atoms into the lower MOT</b>			
The upper and lower MOT trapping beams are red detuned 15 MHz.	0.000 s		
The upper MOT coils and beams are switched on, and the trap fills from the surrounding Rb vapor.	0.040	200 ms	$N = 8 \times 10^8$ $T = 300\mu K$
The upper MOT coils and beams are switched off, and a circularly polarized, $2mWcm^{-2}$ , resonant beam pushes the atoms to the lower, UHV chamber through a 50 mm by 5 mm tube.	0.240	5 ms	
Atoms are transferred from the upper to lower MOT 250 times.	0.245	55 s	
<b>2. Molasses and Polarization Gradient Cooling</b>			
The lower MOT coils are switched off, and the Earth's magnetic field is compensated leaving a residual field on the order of 10 mG.	55.000	600 $\mu s$	
The lower MOT beams are left on at detuning of 15 MHz for optical molasses.	55.000	3 ms	$N = 8 \times 10^8$ $T = 30\mu K$
The lower MOT beams are red detuned 25 MHz for polarization gradient cooling.	55.004	3 ms	
<b>3. Pumping all atoms into the <math>F = 2, m_F = +2</math> state</b>			
The vertical Earth's magnetic field compensation coil is switched off, leaving the 730 mG vertical component of the Earth's field.	55.007	< 1 ms	
A circularly polarized, $2mWcm^{-2}$ , resonant vertical beam optically pumps the atoms into the $5^2S_{1/2}F = 2, m_F = +2$ Zeeman level.	55.007	1 ms	$N = 8 \times 10^8$ $T = 30\mu K$
<b>4. Magnetic Trapping</b>			
The lower MOT coils (quadrupole) are switched on at 12 A, then the current is ramped up to 35 A.	55.008	2 s	$N = 6 \times 10^8$ $T = 200\mu K$
A constant current sink in parallel with the quadrupole coils is switched on.	57.009		
A third coil (Ioffe), which provides a non-zero trap minimum of 1 G, has its current ramped up from 0 - 35 A.	58.009	1.5 s	$N = 4 \times 10^8$ $T = 200\mu K$ $n_{pk} = 10^{12} cm^{-3}$
<b>5. Evaporative Cooling</b>			
RF frequency is ramped from 20 - 1.86 MHz.	61.000	25 s	$N_C = 3 \times 10^5$ $T = 400nK$
<b>6. Imaging</b>			
The magnetic coils are switched off.	61.000	600 $\mu s$	
The cloud expands for $x$ ms, and an absorption image is taken with a detuned $\sigma^+$ beam.	100 $\mu s$		

**Table 3.1:** A recipe for BEC.

possible. They are turned on at a current intermediate between the current used during operation of the science MOT, and the current used in the full magnetic trap. This current was optimized to allow capture of the largest fraction of atoms from the optical cooling stage of the experiment. The currents are gradually ramped up over 2 s to the full operating current of the trap, compressing the cold atom cloud. Just under  $10^9$  atoms are held in the quadrupole magnetic trap.

At this point, the current through the Ioffe coil is ramped from 0 A to full operating current, 35 A, over 1.5 s. This pulls the trap from the linear, zero field minimum quadrupole trap to the harmonic, non-zero field minimum Ioffe Pritchard trap. It also physically pulls the atom cloud 9 mm towards the Ioffe coil. The transfer process appears to maintain up to 80 % of the atoms originally magnetically trapped.

The atoms in the full IP trap experience a roughly harmonic potential of frequencies  $20 \times 200 \times 200$  Hz. Approximately  $5 \times 10^9$  atoms are trapped at a temperature of  $200 \mu\text{K}$  and a density of  $10^{12} \text{ cm}^{-3}$ , yielding a phase space density of  $10^{-6}$ .

Part 5 of table 3.3 shows the command entered to evaporate atoms from the trap. The most important information about the process is the specific frequency at which the RF ramp ends. Only a rough value is included in the table because it is a variable that is adjusted to produce a cold atom cloud, a dual-component cloud containing both a cold atom cloud and a BEC, or a pure BEC.

The experiment is completed by acquisition of information about the BEC produced via absorption imaging. This is implemented by a short pulse of  $\sigma^+$  polarized light from a MOT beam. The very short pulse time is implemented outside the computer control program using a pulse generator with adjustable delay and pulse width. The image is collected on a CCD camera interfaced to a second experimental computer. Images of both a second imaging beam (sans BEC), and of the effected dark field (sans BEC and imaging beam), are taken each experimental run to calibrate the data.

The data collected is converted to calibrated information by analysis of these image sequences on the imaging computer. The information is fitted with models of atom distribution to determine total number of atoms, temperature, condensate fraction, and local column density fluctuations. This information is used to decide what changes to this sequence of events to implement on the next experimental run.





## 4 The vacuum system

The vacuum system is perhaps the part of a BEC machine that benefits most from careful initial design and implementation. It needs to support an “ultra-high vacuum” (UHV) of at least  $10^{-11}$  torr. This is an exceptionally good vacuum, and requires specialist equipment and techniques to successfully implement.

### 4.0.1 Vacuum systems for BECs


BEC experiments are somewhat different to other research requiring UHV, such as surface science experiments, because they can have the benefit of containing very little equipment actually inside the vacuum, making them extremely small. This has another great benefit: the vacuum system, once closed, can remain closed indefinitely. The vacuum system on the current BEC experiment was sealed in September of 1999, and has never been opened since. This means that a well-constructed system is a worthwhile investment, and some great expense, especially in time, can be justified in the initial construction.

The aspect of BEC vacuum system design that requires some innovation is the allowance for optical access and extremely close positioning of the the magnetic coils to the region in which the atoms are trapped. In the ANU-BEC experiments, the solution to both problems was to implement an optical grade quartz cell grafted to a metal conflat seal.

### 4.0.2 The vacuum system in context

The design of the ANU-BEC vacuum system was very compact and efficient for its time. In BEC laboratories around the world, many lessons were being quickly learnt as atom optics groups who had previously required vacuums operating at  $10^{-8}$  or  $10^{-9}$  torr attempted to build systems capable of operation at  $10^{-11}$  torr. Experiments and vacuum systems evolved into large, complicated beasts in the rush to produce BEC. The ANU-BEC design incorporated the best features of the systems of the day, producing a relatively small, mostly off-the-shelf, reliable vacuum system design.

The ANU-BEC vacuum system design also benefited from the development of new methods of loading high numbers of atoms into a UHV MOT. Specifically, the ANU-BEC was one of the earliest groups to adopt the double MOT design pioneered by Myatt et. al., using rubidium dispenser filaments instead of a separate oven apparatus [130].



**4.1: Photo of the ANU-BEC-I vacuum system.** The double MOT chambers are shown on the right. The ion pump and titanium sublimation pump are shown on the left. Not shown is the small ion pump attached to the collection chamber. Note the dimensions of the various tubes used throughout the system. The tubes connecting the pumps to the experimental chambers are kept large (6") to ensure good conductance and pumping speeds.

#### 4.0.3 The vacuum system in brief

The ANU-BEC-I vacuum system is shown in figure 4.1. It consists of vertically-aligned dual chambers, connected via a constricted tube to enable differential pumping. The design uses ion pumps on each chamber, and a titanium sublimation pump on the UHV science chamber. The upper (collection) chamber was a complicated geometric stainless steel structure. The lower (science) chamber was a quartz cell, attached via a long, high quality, glass-to-metal seal. The pumps were connected to the science chamber by long, large diameter vacuum tubes to ensure low magnetic field leakage whilst maintaining high pumping rate.

### 4.1 Fundamental considerations

It is generally regarded as necessary to maintain a vacuum of at least  $10^{-11}$  torr in the region of magnetically trapped atoms to be evaporatively cooled [133–135]. This is because dilute gas BECs rely predominantly on cooling the gas to increase its phase space density. The dilute gas stipulation is necessary because,

at densities significantly above  $10^{13} \text{ cm}^{-3}$ , a trapped cloud of  $^{87}\text{Rb}$  atoms starts to undergo significant three-body recombination losses. At these densities, the atoms must be cooled to temperatures of the order of 100 nK before the average de Broglie wavelength exceeds the mean inter-particle spacing, and degeneracy is achieved.

In order to cool a system in thermal equilibrium to such low temperatures, the isolation from the surrounding environment must be carefully considered. One of the main ways this is achieved is by performing the experiment in an excellent vacuum to ensure the trapped atoms are not heated by the surroundings.

The key issue is reducing the number of room temperature background gas atoms in the experimental chamber in which the BEC is produced. This is vitally important for extremely cold dilute gas BECs, because background gas atoms efficiently heat the trapped gas, especially via grazing-angle collisions [134, 136], as well as knocking them from the trap via more direct collisions. The second of these effects is often the dominant technical limitation of the lifetime of atoms stored in a magnetic trap in the vacuum, and, consequently, its minimization is very important.

It is possible to quantitatively link the background gas pressure to the lifetime in the magnetic trap. The lifetime of the magnetic trap is an easily measurable quantity, and can be used, conversely, to calculate the background pressure in the experimental chamber. This is important because there are few readily available vacuum gauges that measure pressures as low as are needed to produce BEC. Fortunately, this test can be conducted with magnetically trapped cold atoms rather than a BEC, so it is possible for experimental groups to verify the pressure in their vacuum chamber before pushing to BEC. This measurement and calculation are documented for the current apparatus in section 4.6, below.

The use of macroscopic magnetic traps with coils located outside the vacuum system implies relatively low trapping frequencies in many DQG experiments. The low densities and elastic collision rates accessible in this sort of trap lead to a long rethermalization rate and a necessarily slow evaporation cycle (typically 20 - 40 seconds). If the lifetime of atoms in the trap is significantly less than the time taken for an evaporation ramp, background collisions will reduce the number of atoms available for evaporation significantly over an experimental run. This can be enough to reduce the number of atoms below the threshold condition for efficient RF evaporation, preventing the realization of BEC.

Experimentally, the collision rate is directly determined by the density of the trapped gas, which is, in turn, determined by the trapping frequencies generated by the magnetic trap, as detailed in chapters 2 and 7. A simple relationship between the elastic collision rate (determined, ultimately, by the trap geometry and current), and the trap lifetime (determined by the vacuum system), has been published by Ketterle et. al. [137]. To ensure the realization of efficient runaway evaporation,  $\tau_{lifetime}/\tau_{elastic}$  should exceed 400.

It is interesting to note that with the advent of magnetic microtraps and optical trap evaporation techniques, both yielding significantly higher trapping frequencies, and faster evaporation ramps, pressure limitations for BEC experiments have relaxed somewhat as shorter lifetimes become more acceptable.

## 4.2 Practical considerations

Designing any piece of experimental apparatus is always a matter of balancing speed, economy and reliability. The vacuum system for a BEC experiment is no exception. Much of the vacuum system can be made from standard off-the-shelf equipment designed for ultrahigh vacuum use. This is usually fast, cheap, and dependable. Often, some parts need to be custom made. Keeping custom made parts away from critical areas of the vacuum can minimize difficulties in construction and reliability. The other conflict generated by the design of a BEC vacuum chamber is the balance between specialization and generalization. The ANU-BEC-I machine leaned strongly to the side of specialization: yielding excellent long term performance and reliability over the last 3 years. The ANU-BEC-III machine is designed with more general experiments in mind, such as the ability to insert equipment inside the vacuum chamber if necessary.

The essence of most current BEC vacuum systems is the industry-standard stainless steel “conflat” connector. These connections are not overly economic, but they are reliable, and present one of the few standard fitting systems for operation at UHV. They come in various sizes and are available on a range of standard tubes, t-pieces, cross-pieces, and cubes. They can incorporate optical quality windows, and are the standard connection to UHV pumps and gauges. The only significant perturbation from this method of construction, is used by some few groups, such as that of Ketterle at MIT, who have the (envious) resources to make large portions of their vacuum systems entirely from glass. However, even they must default to conflat connections at pumps and gauges, and between glass chambers.

The main benefits of glass chambers over stainless steel are their optical access, lack of magnetization, and their small size for complicated shapes. The first is vital for laser cooling and trapping. The second and third are vital for magnetic trapping and evaporation.

The trade-off in using a glass chamber in a BEC experiment is the necessity of at some point connecting to a conflat-based stainless steel system. The connection between the steel and glass forms a focal point for stresses, making it difficult to create a join good enough to cope with the mechanical shocks typical in assembling such a vacuum system.

A much bigger problem for UHV glass-to-metal seals, however, is the need to bake them to high temperatures to drive impurities out of the surface layers of metal inside the vacuum system. The coefficients of expansion of optical quality quartz and stainless steel are very different, and only a tiny differential is necessary to “open up” such a joint.

The most common solution is to use a series of glasses of different thermal expansion coefficients to bridge the gap between the optical quality quartz of the glass cell and the stainless steel conflat connection. A disadvantage of this type of glass-to-metal seal is its extra length, which increases the size of the vacuum system. This is especially important in a configuration where atoms need to be transported out of a glass cell, as in a double-MOT configuration, when the extra length may lead to low transfer efficiency.

A well designed UHV chamber that is rarely opened to atmosphere actually has very modest pumping requirements. It does require continuous pumping, however, due to effects such as the slow diffusion of gas molecules out of the walls of the chamber itself. The key requirements of a pump used in such a system are very clean, low pressure, long term, low vibration and low maintenance operation. The other aspect of pumping a UHV chamber is the initial pump down from atmosphere. This generally requires a separate pump to that specified above because few pump designs work well at both high and low pressure. The “pump down” pump must be capable of pumping a relatively large volume of gas from an atmospheric chamber down to at least  $10^{-4}$  torr, where the UHV pumps can be switched on.

Ion pumps are the “workhorse” UHV pump used in most BEC experiments. They work by ionizing atoms and molecules that enter a high voltage discharge region of the pump, via collisions with high energy electrons, then accelerating the resultant ions with strong electric and magnetic fields through a high voltage mesh so that they are embedded in the wall of the pump. In sputter-ion pumps, this process is combined with a small, but regulated and continuous, sputtering of titanium over the wall surfaces of the pump, coating and securing the pumped ions. Ion pumps are designed to work in the “molecular” regime of fluid flow: they don’t “suck” atoms out of the chamber; they just prevent any atom that enters the pump from leaving it. This allows operation at the extremely low pressures of UHV. They are extremely clean and UHV compatible as they have no moving parts, so require no lubricating oil, and require only standard UHV electrical feedthroughs into the vacuum chamber. They can operate continuously at low pressures, needing only a high voltage supply of insignificant current. At high pressures ( $> 10^{-4}$  torr), however, the current drawn through the pump electrodes can heat them to destruction. The magnetic field can be problematic if positioned near trapped atoms, but this can be minimized by positioning ion pumps away from trapped atoms and making the chamber from 316-grade non-magnetic stainless steel.

Most early BEC experiments, including ANU-BEC-I, implemented a titanium sublimation pump. A titanium sublimation pump is an electrode of an alloy of titanium and molybdenum, through which a controlled current is passed to sublimate the titanium, coating the walls of the surrounding chamber. As titanium is highly chemically active, reactive background gas colliding with a coated wall will stick, and be effectively removed from the system. This process only works while there is fresh titanium available on the walls, so after a monolayer of pumped material forms, fresh titanium must be evaporated. The pumping speed of this type of pump is limited by the area that can be coated, and the conductance between the area being pumped and the titanium coated surfaces. Despite their apparent high pumping speeds, titanium sublimation pumps are unsuitable for pumping unreactive gases. Most importantly, they cannot effectively pump any of the inert gases or methane, and so, in practice, must be coupled with another style of pump to provide the full pumping requirements of a BEC system. In practice, those groups that have titanium sublimation pumps have found they are unnecessary, and groups building new machines are omitting them to minimize size and cost of the vacuum system.

Pumping down a chamber from atmospheric pressure requires a very different

pump to UHV operation. There are two main differences: the pump must be able to remove a large amount of gas, typically venting it to atmosphere; and the pump must operate effectively in the laminar regime. This is the regime where pressure differentials across the vacuum system “suck” the gas out of the chamber. In reality most pumps can operate in a combination of both regimes to some extent, but few so well as a turbo-molecular pump. A turbo pump consists of a multi-bladed turbine that sucks gas out of the vacuum system in the laminar regime, and acts like a one-way turnstile for gas atoms that wander into its path in the molecular regime. Combined with a rotary roughing pump, a good turbo can pump all the way from atmosphere to  $10^{-9}$  torr. Its weaknesses for use in a BEC experiment are the unavoidable vibration associated with the rapidly moving turbine mechanism, and, to a lesser extent, the associated limited service life and lubrication requirements. In 1999 the ANU-BEC was pumped down with an oil-free turbo pump to pressures sufficiently low to start the ion pumps that have run to this day.

Another trade-off in constructing the vacuum system in a BEC apparatus is the investment in pressure diagnostics. No cheap, readily available, gauges exist that can measure the pressures at which BEC experiments take place. A gauge attached to a small BEC system becomes a significant extra volume and one extra leak-prone joint. Once the permanently closed vacuum system is pumped down, the gauge is essentially useless unless something goes wrong. The other side of the argument, however, is that as the vacuum is so important some small increase in cost and complexity is justifiable even if only to have corroborating proof that everything is OK. To this end, many BEC machines include an ion gauge.

An ion gauge works on very similar principles to an ion pump: it uses high voltages to ionize atoms that enter its gauge volume, then measures the ensuing currents. The limitation comes when you ionize so few atoms that you cannot accurately measure the current, in a typical ion gauge at  $10^{-9}$  torr. There are various incarnations of this design that measure pressure better: the new ANU-BEC-III machine uses a cold cathode gauge that can read below  $10^{-10}$  torr. It is important to note, that due to the similarity in their operation, most ion pump controllers provide a roughly calibrated estimate of the pressure inside the pump, inferred from the current being drawn through their discharge electrodes. It is possible to use this rough pressure measurement quite successfully as a measure of the vacuum chamber pressure, especially when corroboration is available via other means, such as the lifetime measurements discussed in section 4.6.

The final important aspect of UHV system design is the bake-out, which is performed in stages as the machine is assembled. There are generally two or three stages of bake-out. The benefits of baking are maximized at very high temperatures. However, inevitably, some components of the system have lower damage thresholds than others. Baking the most rugged components at  $\sim 350$  °C increases the diffusion rate of molecules out of the surface layers of the inner walls of the vacuum chamber. More sensitive components can often be baked at  $\sim 150$  °C to remove all traces of water from any surfaces inside the vacuum chamber. In some cases it is necessary to perform a low temperature bake-out (50 - 80 °C ) if there are parts inside the vacuum chamber that should not be exposed to higher temperatures.

Ideally, the entire vacuum system could be baked out to 400 °C for several weeks. Metal left in atmosphere acquires a steady-state population of atmospheric gases diffused into its crystal matrix. If a vacuum is applied to such a metal surface, the steady-state population in the metal is much lower. Gas begins to diffuse out of the walls into the vacuum chamber, acting as a “virtual leak”. In UHV systems this process can easily limit the ultimate vacuum attainable in a chamber.

By baking at elevated temperatures, and under vacuum, the natural diffusion of gas particles out of the walls inside the vacuum chamber is accelerated dramatically. This reduces the number of molecules stuck in the surface layer of the walls, and thereafter, the virtual leak load due to this effect. The hotter the bake is performed, the faster the diffusion occurs: practical limits are set by damage thresholds of parts inside the vacuum. Most parts designed for UHV can be baked in some form to 350 - 400 °C . Ion pumps must typically have the magnets removed, and sensitive electronics for vacuum gauges are usually removable.

A high temperature bake-out is extremely important for optimum operation of a UHV system, but it is also a significant effort to bake to such high temperatures for so long. Unless a vacuum system is heavily contaminated, it usually only needs one high-temperature bake-out. Even briefly opening the system to atmosphere is unlikely to be problematic, as it takes a significant time for room temperature atmospheric gases to diffuse deep into the metal matrix. However, gases do become adsorbed to the surface of the walls very quickly. For this reason, a UHV chamber usually needs to be baked out to a medium temperature of 150 °C , until the pressure stabilizes, each time it is opened. This temperature ensures that all water is removed from surfaces inside the system.

Lower temperature bake-outs are rarely needed except where temperature-sensitive parts must be added to the chamber after the higher temperature bake-outs to avoid damaging them. This situation is hardly ideal, as temperatures of 150 °C would be recommended to ensure all water has been removed from the system. If it is not possible to avoid this situation, it is better to position the effected areas away from those that need the best vacuum.

### 4.3 The ANU-BEC-I vacuum system

Figure 4.1 shows the detailed layout of the vacuum system used in the initial BEC experiment at ANU. It was designed, and assembled, on the principles described above, as well as with consideration of how the rest of the experiment needed to fit around it.

The system is designed to hold a double MOT in a vertical alignment, with the collection MOT positioned above the science MOT [130]. The science MOT is created in a square-section glass chamber, attached to the rest of the system with a glass-to-metal seal to a conflat connection. The collection MOT is created in a custom made stainless steel chamber with optical-quality windows sealed in place with indium seals. The two MOT chambers are separated by a 5 mm diameter, 50 mm long tube, that allows a pressure differential to be maintained between the






**4.2: Photo of the ANU-BEC-I collection chamber.** The chamber is a truncated cube: the six large faces hold indium-sealed optical quality windows for the MOT beams; the eight tubes provide 1 1/3" conflat connections used for: connecting the chamber to the rest of the vacuum system; an electrical feedthrough driving the Rb dispensers; and windows for the push beam and general observation.

UHV science MOT and the lower vacuum collection MOT. The science MOT is pumped continuously by a 40 l/s ion pump and a titanium sublimation pump. These pumps are removed from the science MOT chamber by approximately one metre to minimize the effect of stray magnetic field from the ion pump. They are connected to the science MOT by large diameter (6") tubes, exhibiting significant conductance, to ensure pumping speed is maintained despite their separation. The collection MOT is pumped by a small 2 l/s ion pump.

Figure 4.2 shows a photo of the collection MOT chamber. It contains six MOT-beam windows mounted with lock-rings on indium seals. It contains eight 1 1/3" conflat ports for attachment to the differential pumping impedance tube, the small collection-chamber ion pump, the electrical feedthrough supporting the rubidium dispensers, and the mounting of four additional windows for the push beam and diagnostic access.

The collection MOT chamber was custom built in the physics department mechanical workshop from a single piece of stainless steel. It is arranged in a complicated but perfect truncated cube. It is extremely small for its complication. This design produced good optical access for large MOT beams, and multiple connection options. Its small size reduces the pumping load and positions the MOT coils close to the atoms for maximum efficiency.

In hindsight, this level of construction was completely unnecessary. Worse, it entailed a compromise: to fit the optical windows for the MOT compactly, indium seals were used. Indium is a very malleable metal often used in high vacuum applications. Unfortunately, it is not UHV compatible because it can only be baked



**4.3: The ANU-BEC-I science chamber.** The square section of the chamber is formed of 3 mm thick optical-quality quartz, the outside surface of which is anti-reflection coated and 36 mm to a side. The cylindrical section interfaces this experimental region to the metal conflat flange. It consists of over thirty samples of glass of various compositions, gradually matching the thermal expansion of the quartz to the metal.

to approximately 70 °C . This prevented full high temperature baking of the entire machine, and required a third, low temperature bake out, as described below.

The “impedence” between the collection MOT and science MOT is an off-the-shelf 50 mm long, 5 mm diameter conflat-mounted tube. The 5 mm diameter is large enough to transmit atoms pushed from the collection MOT to the science MOT. The tube has a conductance of 0.3 l/s. This conductance limits the load on the UHV science chamber from the higher pressures in the collection MOT chamber. The high capacity science chamber pumps can pump this load sufficiently fast that the pressure remains below  $10^{-11}$  torr.

The glass construction of the science chamber, shown in figure 4.3, achieves excellent optical access with small size. The cell is anti-reflection coated on the outer surfaces. The experimental region is square section, constructed from 3 mm thick quartz, with outside dimension 36 mm. This is fused to a 10 cm diameter glass tube. Along this tube, a gradient of over thirty samples of glasses gradually matches the thermal expansion of the quartz cell to the stainless steel conflat flange to which it is attached. The conflat flange is bolted to a standard 4 1/2” six-way conflat cube. This cube interfaces to the rest of the vacuum system: the impedance tube directly above, and the massive pump-tube to one side.

This cell was donated, in the interests of collaboration, from the Hannover BEC group. It is an excellent piece of apparatus. Its main limitation is the extent of the thermal expansion matching glass gradient. This extra length is not ideal for maximizing transfer efficiency between the collection and science MOTs.

The collection chamber is pumped with a Physical Electronics 2 l/s model pump,

via a 1 1/3" conflat connection. The science chamber is pumped with a Physical Electronics 40 l/s model pump, via a 2 3/4" conflat connection, and approximately 80 cm of 6" conflat tubing. Attached to a conflat cross, near the science chamber ion pump, is a 50 cm long, 10" diameter tube. Inside this can, four titanium electrodes from a titanium sublimation pump are mounted on an eight-way, high current electric feedthrough. The pumping speed of this pump is limited by the inside area of the can. In practice, the conductance of the 6" diameter tubes connecting the pump to the experimental chambers, and the 10 cm diameter gradient section of the glass cell, is a much stricter limit on the pumping speed at the science chamber.

The pumps were distanced from the science chamber by the 6" conflat tube to minimize stray fields from the ion pump. The length of the connection tubes required the large diameter to maintain a reasonable conductance. Since this system was implemented, many groups have shown that ion pumps can be placed as close as 20 cm to the science MOT and magnetic trap, as long as the stray fields are carefully cancelled with compensation coils. The much reduced length allows standard 2-3/4" conflat tubes to be used whilst maintaining conductance. This greatly reduces the size of the vacuum system, simplifying construction and improving performance.

The vacuum chamber was pumped from atmosphere, and roughed during baking, using a oil-free turbo-molecular pump, backed by an oil-free mechanical diaphragm pump.

The vacuum system included a Physical Electronics ion gauge near the science chamber pumps. It measured pressures between  $10^{-5}$  and  $10^{-9}$  torr. In addition, the science chamber ion pump had a pressure display feature on the controller. This displayed pressure in the ion pump as a derivative of the current drawn in the pump.

The bake out of the ANU-BEC-I machine was made problematic by the use of indium seals on the collection chamber windows, and the failure of a commercial feedthrough in the titanium sublimation pump. Its relatively small size eased the construction of an oven for bake-out. However, the delicate glass-to-metal seal required care during baking to ensure its safety.

The ANU-BEC-I machine was baked in an oven constructed around it from aluminum foil coated ceramic walls. The oven was powered by four Elstein FSR 1000W 240V ceramic radiative heaters [138], and was fan-forced from the top of the chamber by a commercial oven fan. The radiative heaters were shielded from direct view of the machine by reflective metal shields. The glass cell was additionally shielded with foil sheaths. These precautions attempted to heat the entire machine at a relatively even rate to minimize thermal gradient induced strains. Thermal gradients were monitored with a series of seven thermocouples attached around the machine and near the glass cell. Of particular worry was the glass-to-metal interface, which was monitored carefully. The temperature of the machine was raised extremely slowly to ensure less than 10 °C temperature differential across all measured thermistors.

The first bake-out was conducted up to 360 °C . It took a day to gently raise the temperature, was held at maximum temperature for 2 days, and was cooled over a further several day period. Before opening to atmosphere, the chamber was vented with dry nitrogen. This bake-out included the entire conflat backbone of

the chamber, the ion pump body, the titanium sublimation pump can, the six-way conflat cube, with windows, and the glass cell. It did not include the collection chamber, the ion pump magnets or the titanium sublimation pump electrodes.

The chamber was opened to install the titanium sublimation pump electrodes. The second bake out was then performed at 150 °C . This was hot enough to drive off water vapor and impurities adsorbed to the inner surface of the chamber, and to bake out the titanium sublimation pump surfaces. It was performed over a similar time scale to the hot bake out.

The chamber was opened again to install the collection chamber. Once installed, the indium seals used on the collection chamber windows prevented baking above 70 °C . Given the low temperature, the last bake out was performed for three days to allow time for the water vapor to reduce significantly.

After the hot bake-out, the pressure of the vacuum system did not initially pump down below the sensitivity of the ion gauge. After various simplistic attempts to isolate the cause of the high pressure, a helium leak tester was borrowed from another research group at the ANU. It was attached in place of the turbo-molecular pump, and a dispenser attached to the unit was used to spray helium at each joint over the machine.


The leak was found in the electrical feedthrough supporting the titanium sublimation pump. This was a commercial part, operated within specifications. It was replaced, and after further high temperature baking the vacuum system attained ultimate pressures below the ion gauge sensitivity.

## 4.4 Improvements to the vacuum system for ANU-BEC-II

The vacuum chamber was not upgraded for the ANU-BEC-II iteration of the experiment. The only notable difference between operation of the ANU-BEC-II machine and the design of the ANU-BEC-I machine, was the discovery that parts of the system were redundant. Both the 2 l/s ion pump attached to the collection chamber, and the titanium sublimation pump, were not operated during this phase of the experiment. The collection chamber pump had been accidentally disconnected at some time after the initial construction of the vacuum system. The system worked well, so it was never reconnected. Although small improvements in trap lifetime were noticed after operating the titanium sublimation pump, other factors were found to limit this lifetime more significantly. As the pressure was, by all means, sufficient for producing BEC, periodic reactivation of the titanium sublimation pump was discontinued.

## 4.5 Future improvements to the vacuum system

The new vacuum system for the ANU-BEC-III machine, shown in figure 4.4, builds on the success of the ANU-BEC vacuum system. It incorporates those design features found useful in ANU-BEC-I, whilst removing superfluous features. It



**4.4: Photo of the ANU-BEC-III vacuum system.** The key design features of the ANU-BEC-III vacuum system are: the use of a glass cell for the collection chamber as well as the science chamber; the omission of a titanium sublimation pump, and the close positioning of the ion pump to the experimental region. These factors combine to make the new vacuum system both significantly smaller, and much more dependent on commercially available components.

also incorporates new BEC machine designs developed since 1999, allowing further reduction of the size and complexity of the system.

A key lesson learnt from ANU-BEC-I was the benefit of using off-the-shelf components where possible. The ANU-BEC-III machine replaced the complicated collection MOT design of ANU-BEC-I with a glass cell similar to that used for the science chamber. The connection to the science chamber, and the connection for the feedthrough for the rubidium dispenser, are implemented via a 2-3/4" six-way conflat cube.

The glass cells for this experiment were purchased from Starna. This company specialized in constructing glass cells for spectroscopy. They have much experience making high quality optical cells, but less making UHV compatible components. The cells have a significantly shorter gradient section connecting the square-section glass region to the metal conflat connector. They have similar outer dimensions to the ANU-BEC-I cell (36 mm square). They were purchased in two lengths: 20 cm long (similar to the ANU-BEC-I cell), and 10 cm long. They were not anti-reflection coated on their outside surfaces, due to cost.

The simpler glass gradient of the cells does not seem to effect their bake-out

performance. However, of six cells purchased, one broke under the strain of attaching it to the system. The lack of anti-reflection coating definitely effects the transmission of light and the operation of a retro-reflected MOT. However, this is a common configuration, and these defects can be compensated outside the cell. The main benefit of the cells is their short length, yielding a compact system that will hopefully enhance transfer efficiency between the collection and science MOTs.

The ANU-BEC-III vacuum system is significantly smaller than the ANU-BEC-I system. This is due to the omission of the titanium sublimation pump and collection chamber ion pump, and the placement of the science chamber ion pump much closer to the chamber. This allowed the tube connecting the ion pump to the rest of the chamber to use the standard 2-3/4" diameter used throughout the system, simplifying joints for reliable performance. It utilizes a cold-cathode gauge, an advanced ion gauge design that allows measurement of pressures as low as  $5 \times 10^{-11}$  torr.

The one digression from simple, off-the-shelf design has been the implementation of an impedance tube and mirror *inside* the vacuum system [139]. This is to allow operation of a 2D collection MOT to generate high atom flux to the science MOT. The tube is mounted in a commercial solid copper gasket and is contained completely within the vacuum system. This simplifies construction because no home-made component is required to maintain the full pressure differential between UHV and atmosphere.


The ANU-BEC-III vacuum system is much smaller than the ANU-BEC-I system. It is made of standard UHV components. The size and construction allows the whole chamber to be easily baked at a high temperature of 350 °C . This makes the chamber as a whole more flexible than the ANU-BEC-I chamber, as it can be opened and re-implemented during the life of the machine with only moderate difficulty.

## 4.6 Characterizing the ANU-BEC vacuum system


There are several common methods by which the background gas pressure may be ascertained in an atom optics experiment, using the atoms themselves as the gauge. The most obvious, perhaps, is the use of "lifetime" measurements, in cases where the lifetime is predominantly limited by the background gas pressure.

In this situation, the process of determining the pressure becomes as simple as observing the decay of atoms from a sample held for increasing amounts of time in the region of interest. Figure 4.5 shows a typical data run from the ANU-BEC-I experiment. The data is collected as a discrete series of measurements, due to the incumbent experimental setup. This does not significantly affect the determination of the trap lifetime however: only a few data points are required to fit a reasonable exponential decay curve to the data, as illustrated by figure 4.6.

The data pictures in figure 4.5 are collected using absorption imaging of laser-cooled atoms held in the magnetic trap for variable hold times. The magnetic trap is populated by first loading the science MOT 250 times from a collection MOT run for 200 ms, typical parameters of operation discussed in chapter 3. The atoms undergo



**4.5: Atoms remaining after variable hold times.** A series of images of cold clouds of atoms after different hold times in the magnetic trap. The reduction in atom number is assumed to be from losses due to collisions with background gas atoms. Simple analysis of the loss rate can yield a value for the background pressure in the experimental chamber.



**4.6: Lifetime of atoms in the magnetic trap.** Plot of the relative number of atoms in each of the images in figure 4.5 , relative to the hold time for each image.

optical molasses and polarization gradient cooling, similar to an experimental run which is to produce a BEC. They are optically pumped, and loaded, first into a quadrupole magnetic trap, and consequently, into the full IP trap. This is the “hold condition”, and no further activity occurs until, after some variable delay, the cloud is imaged using a typical absorption imaging sequence, as discussed in chapter 8.

The primary limit on the accuracy of the data collected is the repeatability with which the magnetic trap can be prepared with a given number of atoms. This factor of performance has been progressively improved through each iteration of the ANU-BEC experiment. Typically, as long as the variation is not too bad, this limitation can be managed by combinations of careful machine preparation and data averaging. The data collected here are for the first configuration of the BEC machine, ANU-BEC-I. These same techniques have been implemented on each iteration of the ANU-BEC machine to enable comparisons and quantify improvements.

Figure 4.6 shows the number of atoms remaining in the trap after a variable hold time. The relative number of atoms is determined directly from the images in figure 4.5, using the standard absorption imaging algorithm based on self-calibrated integrated differential absorption, developed by JL. The solid line is an exponential fit, performed in a mathematical software package, yielding a lifetime of between 30 - 50 s.

For data well-fit by a single exponential time constant, it seems reasonable to assume that there is one clearly preeminent process, causing exponential decay, dominating the losses from the trap. Moreover, by temporarily improving the vacuum, by, for instance, running a titanium sublimation pump, and observing the improved trap lifetime, the background gas pressure may be identified as this dominant loss mechanism. Other observations, discussed below, have indicated that this somewhat simplistic analysis may be misleading. Neglecting these issues temporarily, the exponential time constant apparent in figure 4.6 can be directly related to the background gas pressure in the chamber.

Quantitatively relating this lifetime to the pressure in the vacuum chamber can be achieved extremely simplistically, or somewhat more accurately. Specifically, equation 4.1 shows a simple relationship, defined by Metcalf, between background pressure and trap lifetime, at typical pressures of interest in BEC experiments [140]. Although this formula may appear simplistic, the accuracy of the data in figures 4.5 and 4.6, limited by repeatability of the BEC apparatus and statistical variation, is already a significant restriction on the accuracy of the whole process.

$$\tau \sim \frac{10^{-8} \text{ torr}}{P} \quad (4.1)$$

This simple relation yields a pressure of approximately  $10^{-10}$  torr for a lifetime of 100 s.

Conversely, however, Schmeidmayer et. al., relate a condition for the loss rate per atom, as indicated in equation 4.2, in which  $n_{bg}$  is the spatial density of background gas atoms,  $\bar{v}_{bg}$  is the average velocity of the background gas atoms,  $\sigma$  is the collision cross-section between background gas atoms and those stored in the trap, and  $p_{bg}$  is the pressure of the background gas [141].



$$\gamma = n_{bg}\bar{v}_{bg}\sigma = 5.1 \times 10^{-3}\text{s}^{-1} \frac{p_{bg}}{10^{-10}\text{torr}} \frac{\sigma}{1\text{nm}^2} \quad (4.2)$$

The collision cross section between typical ultra-high vacuum species has been reported by Bali et. al., as approximately  $3500 \text{ \AA}^2$  for Rb-Rb collisions,  $198 \text{ \AA}^2$  for Rb-He collisions, and  $295 \text{ \AA}^2$  for Rb-H<sub>2</sub> collisions [142]. With a background gas pressure of  $10^{-10}$  torr, these figures yield typical lifetimes of the order of 10 s when the background gas is dominated by hydrogen or helium, or only 1 s if it is dominated by Rb gas lingering in the chamber from previous experimental runs. Assuming the former case, the  $\sim 40$  s lifetime apparent in figure 4.6, implies a background pressure of the order of  $3 \times 10^{-11}$  torr. This represents a significant divergence from the simple relation in equation 4.1 above. It does agree more closely with the pressure observations in the ANU-BEC laboratory, which read off-scale on the ion gauge and calibrated ion pump controller (somewhat below  $10^{-10}$  torr).

This experimental verification of vacuum performance has been used many times, on each iteration of the BEC machine. As the consistency of each machine has improved, the spread of the data have improved also. In particular, NR details similar measurements for the ANU-BEC-II apparatus in his thesis [2], generating trap lifetimes of approximately 150 s, with a reasonable degree of accuracy.

NR was also able to determine trap lifetime using a second technique, fluorescence detection of the background collision rate on a small MOT [136]. This technique is immune to some types of loss experienced by held magnetic trap lifetime measurements. Specifically, it is less sensitive to losses due to background resonant light, due to the higher trap depth of the MOT. A measured lifetime of over 250 s was achieved using this technique.

The discrepancy between these two results, along with general experience of operating the BEC machine and variability in reported lifetime requirements in the literature [135, 139], has lead to a new thesis regarding trap lifetime. Resonant optical light, scattered off equipment and the walls of laboratory, is believed to be a significant and serious limitation of many improperly baffled BEC apparatus, especially in the final stages leading to BEC. This belief is backed by multiple private communications detailing limitations in other groups' BEC machines, but is not widely reported in the literature.

# 5 The optical system

Laser cooling and trapping provides the bulk of the cooling in most dilute gas BEC experiments. In addition, optical manipulation and probing are vital for performing experiments in this new breed of “quantum laboratory”. These tools are implemented in a BEC machine via a system of lasers and optics. This chapter is merely a summary of the pertinent aspects of the optical system of the ANU-BEC machine.

## 5.0.1 Overview of the ANU-BEC optical system

The mainstays of the optical system in ANU-BEC-I are the two magneto-optic traps (MOTs). They are arranged in fairly standard configurations: the collection MOT as a high power, large beam, retroreflected MOT; the science MOT a high power, independent six-beam MOT. The light for the MOTs, as well as push beams, optical pumping and imaging, is provided by external cavity diode laser systems (ECDLs). The power of the main cooling beams is increased using a commercial master oscillator power amplifier (MOPA) configuration with a tapered amplifier chip. The frequency of the lasers is directly stabilized to naturally broadened transitions of  $^{87}\text{Rb}$  using saturated absorption spectroscopy locking. The frequency and intensity of light throughout the experiment is manipulated using acousto-optic modulators. Light switching is provided by mechanical shutters for the extremely sensitive magnetic trapping stage of the experiment. Spatial mode quality of the beams is guaranteed using single mode optic fibers.

## 5.1 Practical considerations

To create a MOT, the trapping laser light must have a linewidth that is at least as narrow as the natural linewidth of the atomic transition used to cool the gas. Lasers used to trap alkali metals typically have linewidths of 1 MHz or less. Extremely cheap, 780 nm single mode laser diodes, useful for cooling rubidium, are available commercially, due to their use in compact disc writing drives for computers. Variations in carrier density along the small length of the laser diode cavity (typically  $100\mu\text{m}$ ) leads to large linewidths of the order of 100 MHz. Several techniques have been developed to narrow the linewidth of such laser diodes: in one such scheme they are mounted in an external cavity, approximately 3 cm long, that reflects about 30 % of the light back to the laser diode [143]. This narrows the linewidth to approximately 500 kHz, and provides a convenient method for tuning the wavelength of the diode laser. The development of external cavity diode lasers, such as these, for the degenerate Fermi gas experiment, is discussed extensively in chapter 12.



The ability of the collection MOT to collect a large number of atoms is strongly dependent on the size and power of the MOT beams used. More power than is available from standard diode lasers is generally required. A (relatively) cheap and elegant solution is to use a tapered amplifier diode in conjunction with an ECDL in a master oscillator - power amplifier configuration. A tapered amplifier is simply a large semiconductor laser diode anti-reflection coated for single-pass operation. It is designed with a tapered gain medium to allow light to build up with constant power density, and can produce powers of over 600 mW. Its large size would yield bad frequency and spatial mode characteristics if it were free-run, but injecting the small input facet with  $\sim 10$  mW from a stable ECDL, generates an amplified output beam with the frequency and spatial qualities of the injected beam. A new commercial design implements a tapered amplifier directly in an external cavity, omitting the requirement of a master oscillator and the supporting electronics, to produce a cheaper system with similar performance characteristics.

The external cavity configuration narrows the linewidth of a laser diode to 500 kHz on the short time scales determined by fluctuations in carrier density. Over longer time scales (minutes to hours), the frequency of an ECDL can drift due to laboratory temperature variations and drifts in the driving electronics. To guarantee the absolute frequency stability of the lasers in a BEC experiment they are "locked" to an atomic transition. The locking circuit needs to provide stability of the order of 1 MHz, so the locking signals are typically derived from a naturally broadened atomic transition. This is usually achieved using saturated absorption spectroscopy [144]. Some more advanced locking schemes were developed as part of this thesis, and are discussed in chapter 14.

Locking the lasers to a naturally broadened atomic transition is vital for excellent absolute frequency stability of the light used in the experiment. However, most parts of the experiment actually require light slightly detuned from resonance. This light is generated from the locked lasers using acousto-optic modulators (AOMs). A limitation of AOMs is the loss of intensity, which is generally  $\sim 25$  % per pass, yielding less than 60 % of the input light as useful trapping light for the experiment.

The benefits of AOMs far outweigh this cost, however. They allow easy analog voltage control of both the intensity (0.1 % to 100 %) and frequency ( $110\text{MHz} \pm 20\text{MHz}$ ) of the laser light. This level of frequency control is usually sufficient, but the intensity control is not good enough to exclusively employ AOMs as shutters for the laser light into the experimental region. In the magnetic trap, and as the atoms are cooled towards BEC, any stray resonant light is disastrous, so further precautions need to be implemented.

Mechanical shutters are the best way to completely eliminate transmitted light into the science chamber during evaporative cooling and BEC experiments. Their main limitation is their speed. The fastest shutters close in  $100\mu\text{s}$ , but suffer from a variable trigger lag of the order of 1 ms. The practical solution is to combine the fast almost-off switch of the AOM with the slower completely-off mechanical shutters.

Optimum operation of a MOT relies on beams with good spatial profiles to guarantee repeatability and maximum size. Good beam quality is generated using

a spatial filter: either a free-space filter, or a single-mode optic fiber. Typically 40 % of the input light from an ECDL is lost during this spatial filtering process. This inefficiency is a direct consequence of the bad mode quality generated by laser diodes.

The detailed operation of MOTs and polarization gradient cooling, especially regarding their implementation in the ANU-BEC experiments, have previously been reported by JL and NR. MOTs require six red-detuned beams, arranged in three mutually orthogonal counter-propagating pairs. Standard 3D MOTs can be roughly separated into two groups: those that retroreflect three independent beams and those that use six independent beams.

A MOT cools and traps atoms very effectively, but a further 1 - 2 orders of magnitude cooling can be realized by more subtle laser cooling techniques. After atoms are collected in the science MOT, the magnetic fields are switched off and the lasers detuned further away from resonance to implement polarization gradient cooling.

MOTs and polarization gradient cooling create a cold, relatively dense sample of atoms spread fairly indiscriminately across the magnetic sublevels available. Magnetic traps only trap atoms in magnetic sublevels that seek weak magnetic fields, so the cold atoms must be “pumped” to the required  $m_F$  level before the magnetic trap is switched on. This is achieved using a pulse of  $\sigma^+$  polarized light with a small on-axis magnetic field created by a coil.

## 5.2 The ANU-BEC-I optical system

The ANU-BEC-I experiment used a TUI Laser (now Toptica) TA100 master oscillator-power amplifier laser as the main source of light for the entire experiment. This laser was capable of generating up to 500 mW of single mode light with a linewidth less than 1 MHz. It exhibited poor mode quality, typical of diode lasers, limiting the efficiency of spatial filtering. The high power of the laser allowed light to be split up for use in different parts of the experiment. Using a single laser for all cooling, imaging and pumping light simplified the locking and stabilization requirements of the experiment.

The main laser was augmented by light from an external cavity diode laser made for the experiment by collaborators in the Research School of Physical Science and Engineering at the ANU. This laser generated approximately 10 mW of light for repumping the MOT. This laser was also locked to a naturally broadened transition of rubidium, but was not manipulated with AOMs or spatially filtered.

The optical layout of the locking circuits is shown in Figure 5.1. The locking circuits are a compact implementation of saturated absorption spectroscopy, based on retroreflection using a polarizing beam splitter cube. The rubidium sample is provided in a room-temperature vapor cell. The dither for lock-in detection is applied via magnetic coils around the vapor cell, which modulate the Zeeman splitting of the rubidium transitions. This modulation exists on the saturated absorption signal recorded at the photodiode, and can be demodulated for phase-sensitive detection

by a lock-in amplifier. The locking signal thus generated is fed back to the grating of the external cavity diode laser to maintain frequency stability. This technique avoids modulating the light used in the experiment directly.

The frequency of the cooling, imaging, and push light was controlled in the ANU-BEC-I experiment using acousto-optic modulators. Due to considerations of cost and inefficiency, only two AOMs were implemented. One was used to generate the cooling light for the collection MOT, as well as the push beam and imaging light. The other generated the cooling light for the science MOT, and the light for optical pumping.

The AOMs were configured in a double-pass configuration to minimize beam steering at different modulation frequencies, as shown in Figure 5.1. This configuration also allows the main cooling laser to be locked to the very strong  $5^2S_{1/2}F = 2 \rightarrow 5^2P_{3/2}F' = 1, 3$  crossover from the saturated absorption signal, then shifted 210 MHz to the  $5^2S_{1/2}F = 2 \rightarrow 5^2P_{3/2}F' = 3$  transition used for laser cooling, with a common 110 MHz AOM. The repump laser, which is not controlled using AOMs, is locked directly to the  $5^2S_{1/2}F = 1 \rightarrow 5^2P_{3/2}F' = 2$  transition.

A combination of fast AOM switching and slow mechanical shuttering was used to control the intensity of light throughout the experiment. The acoustic shock generated by the mechanical shutters tended to unlock the lasers. The long-term solution to this problem was to move the lasers to a separate optics table and use optic fibres to transmit the light to the BEC table.

The ANU-BEC-I experiment created good beam shape using single mode optic fibers as a spatial filter. The quality out of such a fiber is a pure  $TEM_{00}$  beam. The other benefit of such a system (over a pinhole spatial filter) is that the fiber can be used to transport light to a remote location, such as another optic table. It also generates some independence between stages of the experiment, because if something is misaligned before a fiber, only those components before the fiber need be realigned. The problematic trade-off, however, is polarization instability. Even “polarization preserving” fibers exhibit significant time-varying output polarizations, which become intensity imbalances in the MOT beams when they are split on a polarizing beam splitter cube. A pinhole spatial filter should not exhibit this problem.

The collection MOT was operated as a retroreflected MOT. The science MOT was operated as a six-beam MOT due to losses expected in the transfer through the glass cell and in the dense MOT itself. The six-beam configuration required careful optimization to ensure good intensity balance between counter-propagating beams.

The collection MOT used three  $\sim 7$  mW beams of 15 mm diameter, typically detuned 15 MHz below resonance. It collected  $10^8$  atoms in 200 ms. The atoms were pushed to the science MOT with a pulse of resonant light. The science MOT used six 7 mW beams of 15 mm diameter, typically detuned 15 MHz below resonance. It caught atoms transferred in this way with approximately 30 % efficiency. The collection MOT was pulsed 250 times in 50 s. The science MOT collected a few times  $10^9$  atoms at a temperature of approximately 1 mK. These MOT configurations were changed significantly during the ANU-BEC-II iteration of the experiment. Most notably, the science MOT was re-implemented as a three-beam retro-reflected MOT, yielding better stability.

The ANU-BEC-I experiment used 3 ms of optical molasses and 3 ms of polarization gradient cooling, after magneto-optic trapping, to cool the atoms to approximately  $30 \mu\text{K}$ . This was implemented by turning off the magnetic fields of the MOT and, after 3 ms, gradually detuning the laser light up to 25 MHz away from resonance, while reducing its intensity. A key component of polarization gradient cooling was careful canceling of the Earth’s magnetic field, to within 10 mG, using compensation coils around the science chamber.

Before the magnetic trap was switched on, atoms were pumped into the  $5^2S_{1/2}F = 2, m_F = +2$  state. This was achieved using a pulse of  $\sigma^+$  circularly polarized light from one of the MOT beams. The light was polarized relative to a small magnetic field created by a “pumping coil” along the axis of the pumping beam. This process strongly effected the efficiency of transferring atoms to the magnetic trap.

### 5.3 Improvements to the optical system in ANU-BEC-II

The fundamental structure of the ANU-BEC-I optical system was very successful. The laser systems, frequency stabilization, frequency control, intensity control and beam shape stabilization methods were optimized, but not improved in any significant manner on the ANU-BEC-II machine. However, several very important improvements were made to the more advanced parts of the optical system during the ANU-BEC-II iteration of the experiment. The magneto-optic trap and optical pumping methods were significantly refined, yielding larger, more reliable samples of laser cooled atoms for magnetic trapping.

The biggest improvement to the MOT systems in ANU-BEC-II was the implementation of the science MOT as a retroreflected MOT. This was a practical, rather than fundamental, improvement. By removing the requirement for precisely balanced counter-propagating MOT beams, it increased the day-to-day stability of the MOT. This allowed better optimization on average, yielding larger MOTs. It also generated significantly better optical access around the science cell due to the reduced amount of optics.

The other major improvement in MOT implementation was the use of permanent quantitative diagnostics on both the collection and science MOTs. Although ANU-BEC-I used closed-circuit TV cameras to qualitatively estimate the quality of each MOT, the calibrated photodiodes implemented in ANU-BEC-II provided a much better measure of performance. This information is vital because it determines the total number of atoms magnetically trapped at the start of evaporation, which in turn determines whether the conditions for runaway evaporation are ever achieved. Access to this information during the experimental run allowed further optimization of the ANU-BEC-II machine to produce larger, more reliable MOTs.

## 6 The magnetic trap

The design and implementation of the magnetic trap forms a large part of this thesis. It involved extensive modeling, mechanical design, prototyping and characterization. This magnetic trap design yielded BEC in the first generation ANU-BEC machine, and was built on extensively in the second and third generation machines.

### 6.0.1 Magnetic traps for BECs

The implementation of magnetic trapping for laser cooled alkali atoms was a key development in the realization of the first dilute gas BECs. Magnetic traps are *conservative* for cold neutral atoms. They simply store the atom cloud, isolating it from surrounding environment. Any further cooling is provided by other techniques, most commonly forced evaporation.

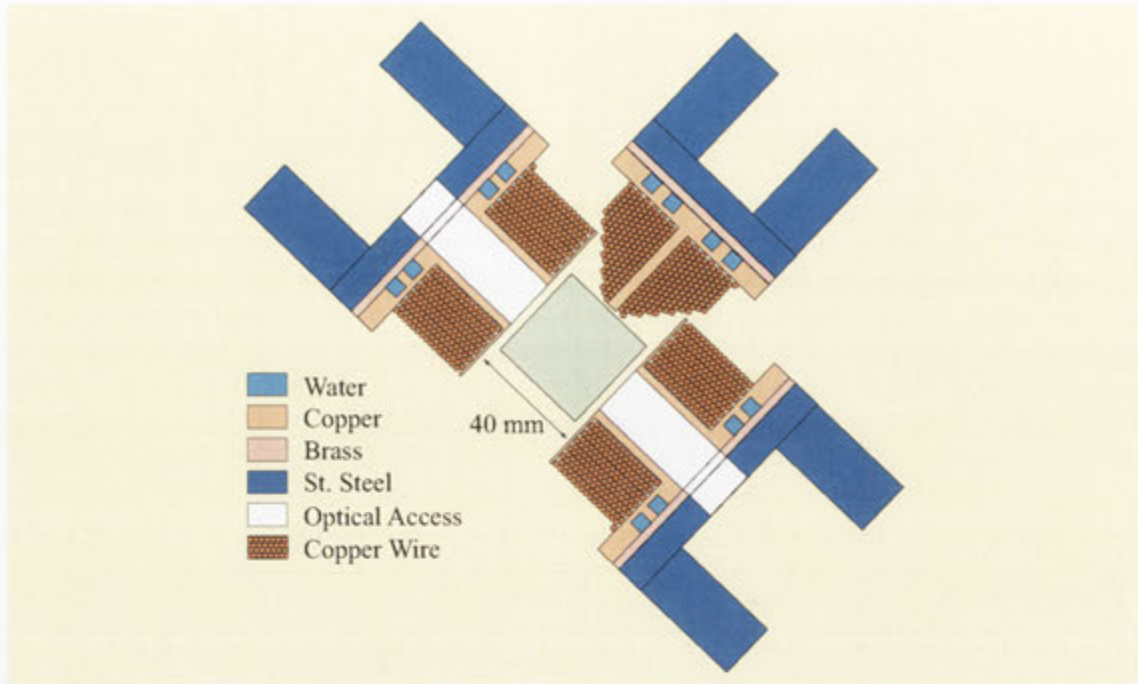
Storing atoms in a conservative trap is important because it minimizes interactions that may limit the accessible temperatures achievable in the trap, such as photon recoil in magneto-optic traps. Neutral atoms have no net electric charge. However, they can exhibit a magnetic moment, and may be contained by an external magnetic field.

The difficulty with trapping atoms magnetically lies in the small magnetic moments exhibited by alkali atoms. Although all sub-atomic particles have moments associated with their intrinsic spin, stable atomic structure leads to paired spins that reduce the overall magnetic moment of the atom. Magnetic dipoles due to orbiting electrons also arrange to minimize the total moment of the atom, both with other orbit-induced dipoles, and the intrinsic spins of the atom through Russell-Saunders coupling. These considerations are alleviated, to some small extent, for alkali atoms, with a single unpaired electron in the valance shell.

The net result of the small moments typical in neutral atoms is that large magnetic field gradients are needed to make effective traps. Conversely, at easily accessible fields and gradients, magnetic traps are not very deep. Typically, traps used in degenerate gas experiments can contain atoms with a temperature of the order of 1 mK. Thus, magnetic traps are complemented perfectly by laser pre-cooling methods that cool small samples of gas from room temperature to below 1 mK.

Despite the seemingly passive role played by the magnetic trap during the production of a BEC, its precise implementation is of vital importance in determining the size and density of the stored atoms. This becomes increasingly important as forced evaporation is implemented to cool atoms to BEC. Tuning the magnetic trap parameters is one of the most important ways of optimizing evaporation, and therefore one of the key aspects of producing a BEC.






**6.1: Diagram of the magnetic trap.** Details include an accurate representation of the cross-section of each coil, the copper coil formers, including channels for water cooling and the silver-soldered brass back plate, and top view of the stainless steel coil mounts.

### 6.0.2 The magnetic trap in context

When the magnetic trap for ANU-BEC was being developed there were several designs predominate worldwide. The time orbiting potential (TOP) trap was used to create one of the first BECs, and its small size and low currents made it a favorite amongst many groups [145]. The cloverleaf trap was a newer design that provided the benefit of a static Ioffe-Pritchard (IP) magnetic field, at the cost of extremely high currents ( $\sim 300$  A), and the associated difficulties that went along with high current electronics [146]. As the choice for the ANU-BEC trap was being finalized, a new trap design, the quadrupole-Ioffe configuration (QUIC) trap, was published by Esslinger et. al. [131]. The QUIC trap combined low currents, similar to the TOP trap, with the static IP field of a cloverleaf trap. The costs of this new design were decreased dynamic tuning of trap parameters and less optical access.

The new trap was by no means an obvious choice. All the large successful groups in the world used other trap designs. Only Esslinger's group had demonstrated BEC in a QUIC trap. The design achieved simplicity by reducing not only the complexity, but also the flexibility and tunability. This amplified the importance of the initial design of the trap to ensure the correct parameters were attained. However, every magnetic trap design involved compromises, and the compromises of the QUIC trap suited the ANU-BEC apparatus. The time saved in not dealing with dynamic trap



**6.2: Photo of the magnetic trap.** The foreground coil and its opposed partner constitute the quadrupole coils. They are used during operation of the science MOT, the quadrupole trap and the full quadrupole-Ioffe configuration trap. The smaller, perpendicularly oriented Ioffe coil is gradually ramped up to full operating current to morph the quadrupole trap into the QUIC trap.

potentials and high current circuits was invested in designing and testing a QUIC trap and the associated electronics.

## 6.1 Fundamental considerations

A magnetic trap is made by arranging magnets to create a field that atoms will see as a bowl-shaped potential. This is possible when atoms exhibit a magnetic moment that will interact with a magnetic field.

The intrinsic spin and moving electric charges of subatomic particles in an atom give it a magnetic field that is well described by a simple dipole magnetic moment. When placed in a relatively small external magnetic field this moment precesses about the field, with a quantized projection along the field. The quantized projection is described by the magnetic sub-level quantum number,  $m_F$ , and the precession ensures a minimum uncertainty in the total angular momentum of the atom.

The magnetic sub-level quantum number of a Bosonic atom can take an integer value from  $-F$  to  $+F$ , including  $m_F = 0$ . The  $m_F = 0$  sublevel has zero projection along the magnetic field, and therefore has no energy dependence on the field. Positive or negative  $m_F$  sublevels have an energy dependence that either increases or

decreases with increasing magnetic field, depending on the sign of the gyromagnetic ratio,  $g_F$ , for a particular hyperfine level. The  $F = 1$  ground state of  $^{87}\text{Rb}$  has negative  $g_F$  and the energy of the positive (negative)  $m_F$  state increases (decreases) with increasing field. The  $F = 1, m_F = +1$  state is said to be *strong-field seeking* while the  $F = 1, m_F = -1$  state is said to be *weak-field seeking*. In contrast, the  $F = 2$  ground state of  $^{87}\text{Rb}$  has positive  $g_F$ , yielding weak-field seeking states of the positive  $m_F$  states.

The energy of a magnetic moment in a magnetic field is  $E = -\mu \cdot \vec{B}$ . Atoms in a field will move to minimize their potential energy, so strong-field (weak-field) seeking states will move towards regions of high (low) magnetic field. Gauss' law elegantly proves that a magnetic field may not diverge within a volume of free space. It is easy to see that the field produced in the middle of a region surrounded by magnets (but containing no magnets itself) cannot support a magnetic field maximum, as corroborated more strictly by Wings' theorem [147, 148]. However, it can support a field minimum. It is only possible to create a potential bowl, then, for weak field seeking states that see a magnetic field minimum as a potential minimum.

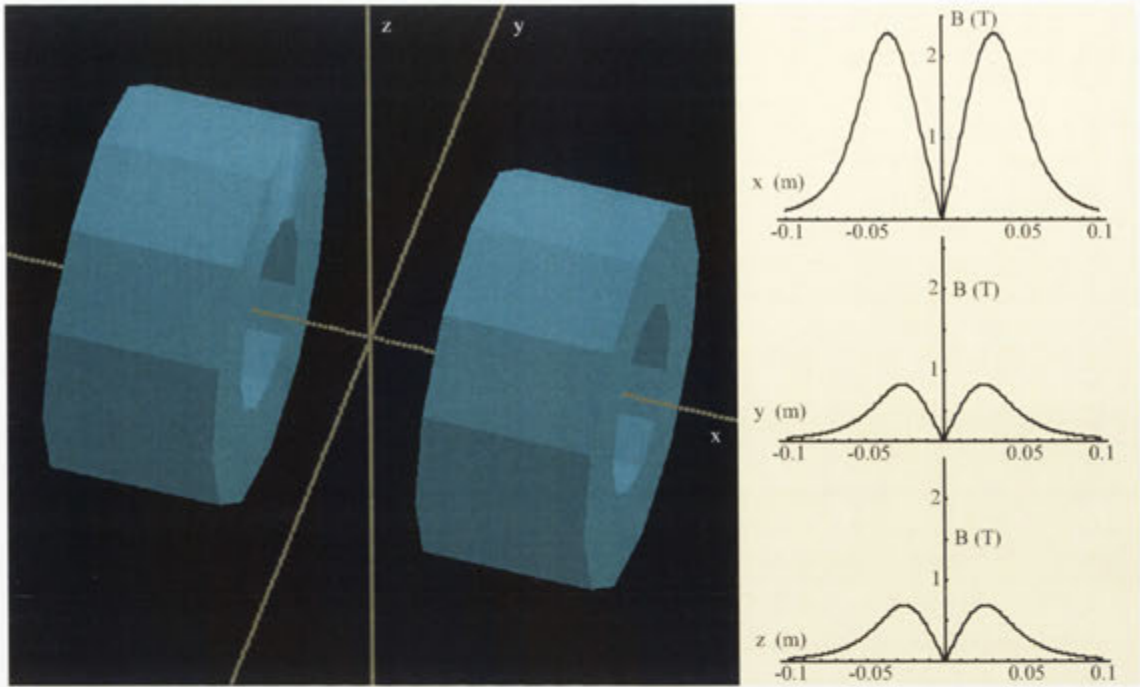
These considerations assume that a trapped atom maintains its magnetic sub-level,  $m_F$ , relative to the magnetic field as it moves around in the trap, even though the field direction changes. This is a plausible assumption because an atom slowly traversing a region of changing magnetic field adiabatically "follows" the field, maintaining the projection of its magnetic moment along the field. That is, it remains in the same quantum state, and interacts with the field in the same way, independent of the laboratory-referenced direction of the field. "Slowly traversing" a region of changing magnetic field, in this case, implies that the field changes direction very little over a single period of precession. Equation 6.1 yields the full condition.

$$\text{In 3D:} \quad |\mathbf{v}_T \cdot \nabla \hat{\mathbf{B}}| \ll \omega_L \quad (6.1)$$

$$\text{In 1D, along } \rho: \quad \sqrt{\frac{8k_B T}{M\pi}} \frac{dB_\rho}{d\rho} \ll \frac{g_F \mu_B B_\rho^2}{\hbar} \quad (6.2)$$

where  $v_T = \sqrt{8k_B T / M\pi}$  is the thermal velocity of a gas of atoms with mass  $M$  and temperature  $T$ ,  $\hat{\mathbf{B}} = \mathbf{B}/|\mathbf{B}|$  is the magnetic field unit vector, and  $\omega_L = g_F \mu_B B_\rho^2 / \hbar$  is the Larmor precession frequency about the magnetic field (in the  $\rho$  direction) of the magnetic moment attributed to an atom in the field. Note that  $g_F$  is the Lande-g factor for the atom being studied, and  $\mu_B$  is the Bohr magneton.

The adiabatic criterion is a general condition elucidating the speed at which an atom moving through a magnetic field gradient will no longer be able to follow the magnetic field lines. Below this speed, the atom will follow the field, even if it reverses direction completely. The direct result of this behavior is that the energy of the atom in the field reduces to  $E = -\mu|\vec{B}|$ , dependent only on the *magnitude* of the field. A simple potential minimum is therefore created by any magnetic field that has a zero-field crossing. The most famous of these is the so-called "anti-Helmholtz" or "quadrupole" configuration using two identical electromagnets of opposite orientation separated by roughly the coil diameter, shown in Figure 6.3.



**6.3: Form of a quadrupole magnetic trap.** A representation of the physical form and magnetic field characteristics of a quadrupole magnetic trap. The trap is typically formed of two classic short solenoids, running counter-propagating currents, separated by the order of their diameter. The subsequent equal magnitude but opposite direction axial field cancels exactly at the center of the trap, midway between the coils. The zero magnitude radial component at the mid-point builds away from this minimum. For atoms “following” the field, the potential is proportional to the magnitude of the field, which forms a zero-field minimum at the center of the trap. In this model, the quadrupole coils are 30 mm long with inner diameter 24 mm and outer diameter 70 mm. They carry a current of 35 A.

The figure shows that near the centre of the trap the magnetic field magnitude approaches zero field along all components. In the weak fields near the center of the field magnitudes shown in figure 6.3, weak-field seeking atoms experience a linear Zeeman shift, and thus experience a potential directly proportional to the field magnitudes - a magnetic trap!

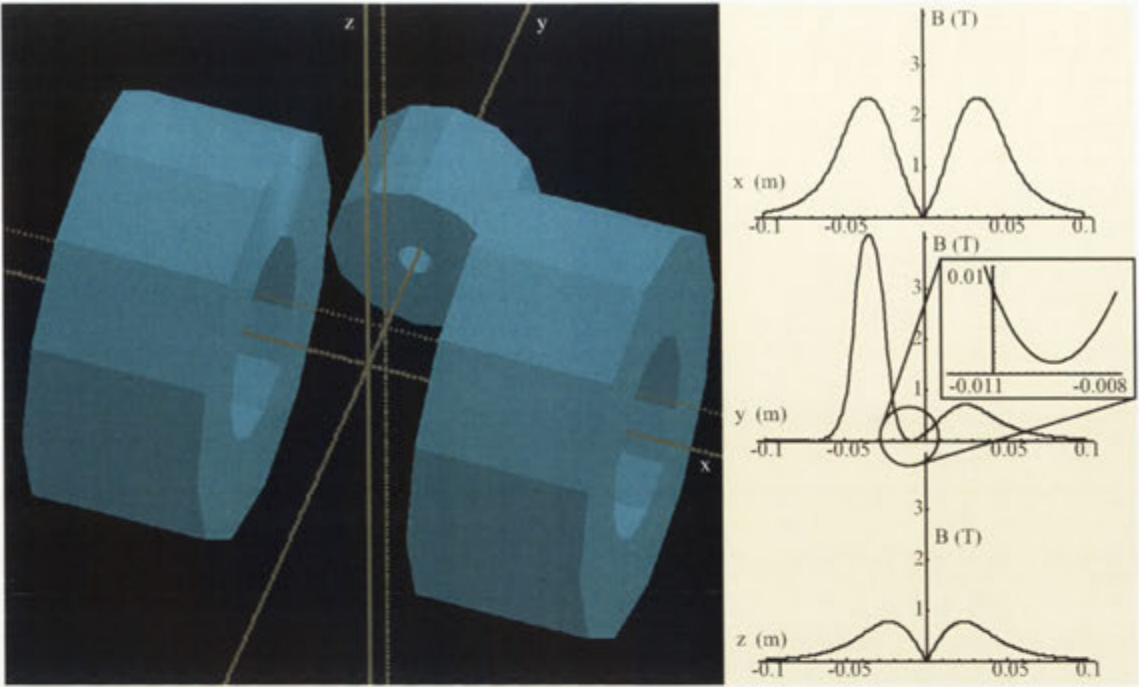
The quadrupole trap was the fundamental design of both machines that are originally credited with discovering BEC. The quadrupole design has very tight linear fields, requires low currents and has excellent optical access. The problem with the quadrupole configuration is the region of zero field at the bottom of the trap. Atoms passing through the region of zero field will never meet the adiabatic criterion of equation 6.1. They will experience Majorana spin-flips into other  $m_F$  states, and be lost from the trap. This becomes more problematic as the atoms get colder, because they relax further towards the bottom of the trap. Laser cooled atoms

( $\sim 100\mu K$ ) are generally hot enough to be contained in typical quadrupole traps used to create BEC. However, as such an atom cloud is evaporatively cooled towards BEC, more of the atoms spend more time near the zero-field region, and the losses become fatal. To reach BEC both original groups had to prevent atoms entering this zero-field region, by plugging the hole with a blue-detuned laser beam [20], or moving it faster than the atoms could respond [145].

As a sample of atoms gets colder, and the average velocity decreases, the minimum field criterion stipulated by equation 6.1 relaxes to prevent significant probability of Majorana flops. However, this does not take into account the distribution of the gas sample within the trap as the temperature drops. A thermal gas obeys a Boltzmann density distribution, with  $n(r) \propto e^{-V(r)/k_B T}$ . In the case of a harmonic oscillator potential,  $V(r) = 1/2 M \omega^2 r^2$ , the distribution is a Gaussian with  $\text{FWHM} \propto \sqrt{k_B T / M \omega^2}$ , where  $\omega$  is the frequency of the harmonic component of the trapping fields. As the gas cools, although the Majorana loss criterion becomes less stringent, more and more atoms occupy the region around the very minimum of the trap. In the extreme limit, the clustering of atoms around the field minimum as the gas is cooled dominates the relaxation of the adiabatic criterion, and any field zero, such as that in the quadrupole trap, will cause catastrophic loss.

A particularly appropriate solution is to create a magnetic trap with no zero-field region in it at all. The most common field configuration currently used was developed for neutral atom trapping by Pritchard [149, 150], and is based on the Ioffe design for trapping plasma [151]. The quadrupole-Ioffe configuration (QUIC) trap achieves this field configuration in a remarkably simple and elegant way. A single ‘‘Ioffe’’ coil is added perpendicular to the ‘‘quadrupole’’ coils, as shown in Figure 6.4. The figure shows how the Ioffe coil ‘‘pulls’’ the potential minimum of the trap from between the quadrupole coils towards the Ioffe coil, eventually raising the minimum to a non-zero ‘‘bias’’ field. Part of the beauty of this configuration is that the quadrupole coils are run at low current to form the science MOT, then at higher current to transfer the atoms into the magnetic trap. This ensures perfect alignment of the MOT and magnetic trap fields, and promotes good transfer efficiency to the magnetic trap. The full IP field is generated by simply ramping the current in the Ioffe coil up to its operating value, which literally ‘‘pulls’’ the atoms towards the Ioffe coil, into the non-zero magnetic trap minimum.

One defining specification of an IP magnetic trap is the trapping frequency. This specification models the very bottom of the trap as a 3D simple harmonic oscillator. This approximation is valid because extremely cold atoms and BECs probe only the very bottom of the trap, where the potential is well described as a harmonic trap. The trapping frequencies determine the distribution of a given ensemble of atoms in the trap and, hence, the density of the atom cloud, as indicated, for one dimension, in equation 6.3, where  $n(r)$  is the density,  $m$  the mass and  $T$  the temperature of the atoms in the sample, and  $\omega$  is the frequency of the harmonic approximation to the trapping fields. The density of the atomic cloud is the key tunable parameter determining the elastic and inelastic collision rates accessible in the trap. These collision rates are the defining parameters for the preparation of efficient evaporative cooling and are therefore of utmost importance in the push to




**6.4: Form of a QUIC magnetic trap.** The addition of a third coil to the quadrupole trap from figure 6.3 creates a potential energy minimum in a region of non-zero field. This becomes increasingly important as atoms are cooled towards BEC, spending relatively more time near this region of the trap. The quadrupole coils have the same dimensions as in figure 6.3, 30 mm long, 24 mm inner diameter and 70 mm outer diameter. The Ioffe coil is conical, 35 mm long and 47 mm in diameter at its maximum extent. A current of 35 A flows through the coil in this model.

produce a BEC. Evaporative cooling is discussed in section 7.

$$n(r) \propto e^{-m\omega^2/k_B T} \quad (6.3)$$

An IP field is the configuration that yields the strongest DC trapping potential with an inherently non-zero bias field. However, even in this field configuration, macroscopic magnetic traps installed outside the vacuum system access only the lower range of acceptable trapping frequencies for evaporative cooling to BEC. A low current QUIC trap is the weakest of the macroscopic IP trap designs. Most of the design process outlined in section 6.2 attempted to maximize trap frequencies within a fairly limited range of operating currents and coil geometries. The simplest way of increasing trapping frequencies for a given QUIC trap is to increase the current flowing through all coils. This is not ideal, however, because it leads to increased heating of the coils, destabilizing the trap, and complicates the driving electronics. In fact, the main modification to the new generation of QUIC traps designed by NR, has been minimization of the power dissipation for a given trap confinement,




**6.5: The simple BiotSavart model.** Note that it is built from the provided pre-compiled composite solenoid models, which affects how BiotSavart calculates the fields. This tends to make it very fast, but maybe less accurate than a more advanced model.

reducing the drive current and improving trap stability. The coil geometries are limited, in practice, by the need for good optical access to the science chamber. This was the main flaw of the initial QUIC trap design, and the newest QUIC trap design has used a quite unorthodox method to avoid this limitation.

## 6.2 Modeling the magnetic trap

Proper optimization of these parameters is achieved by extensive computer based modeling of magnetic trap designs. Detailed modeling was performed as part of this thesis, and NR has extended these models for the new generation of QUIC traps, emphasizing aspects of the trap that are more important for new experiments. The modeling conducted for this thesis was mostly numerical, using a commercial magnetic field program called BiotSavart, produced by a company called Ripplon Software, Inc. [152]. The model was extended to raw Mathematica code when the limitations of BiotSavart became problematic [153].

Modeling the field generated by a single perfect current loop, or even several perfect current loops, is very simple, as the vector fields from each loop can simply be summed at each point in space. Unfortunately, in reality very few perfect current loops exist. Real wire loops have wire “tails” connected to the rest of the driving



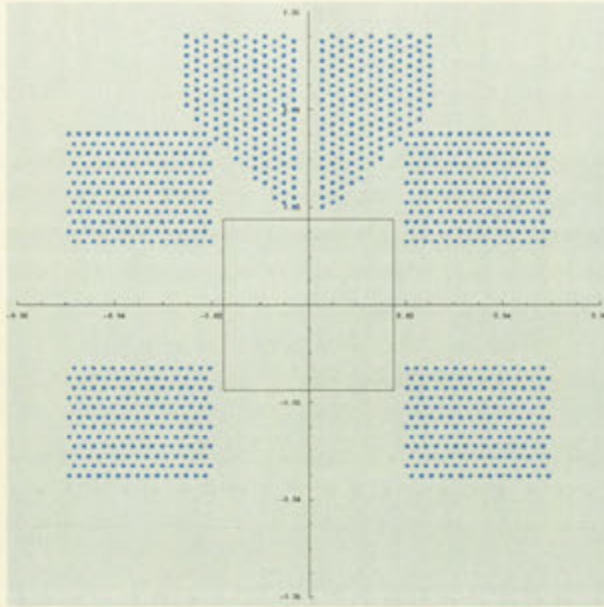
**6.6: The complex BiotSavart model.** Note that it contains individual loops, rather than pre-compiled solenoids, describing each coil. More difficult to appreciate is that the loops are arranged as counter-propagating helices that closely mimic the real position of wires in the coil. The effect of these changes may be perceived in the qualitatively different shape of the field contours, relative to figure 6.5.

circuit. Real magnetic coils, like those used in the QUIC trap, have many wire loops, arranged in complicated counter-propagating helices, rather than individual perfect current loops. While it may be difficult to gauge exactly where the pursuit of the perfect model reaches the point of diminishing returns, it is possible to test some of these limitations relatively quickly with the right tools.

Modeling is vitally important for the design and understanding of the key specifications of the magnetic trap. These specifications include the trap depth, the gradients and curvature of the magnetic fields, and, perhaps most critically, the trapping frequencies at the bottom of the trap where it is well approximated as a simple harmonic oscillator potential.

Within the QUIC trap design, it is the final three parameters that are most important: the gradients of the field in the tight radial direction and the curvature in the loose axial direction determine how laser-cooled atoms interact with the field. They are important for ensuring efficient transfer of atoms from the MOT to the magnetic trap, and at the beginning of evaporation. As the atoms are cooled further, the trap becomes essentially harmonic, and can be specified by trapping frequencies in three dimensions, or the geometric mean of the trapping frequencies,



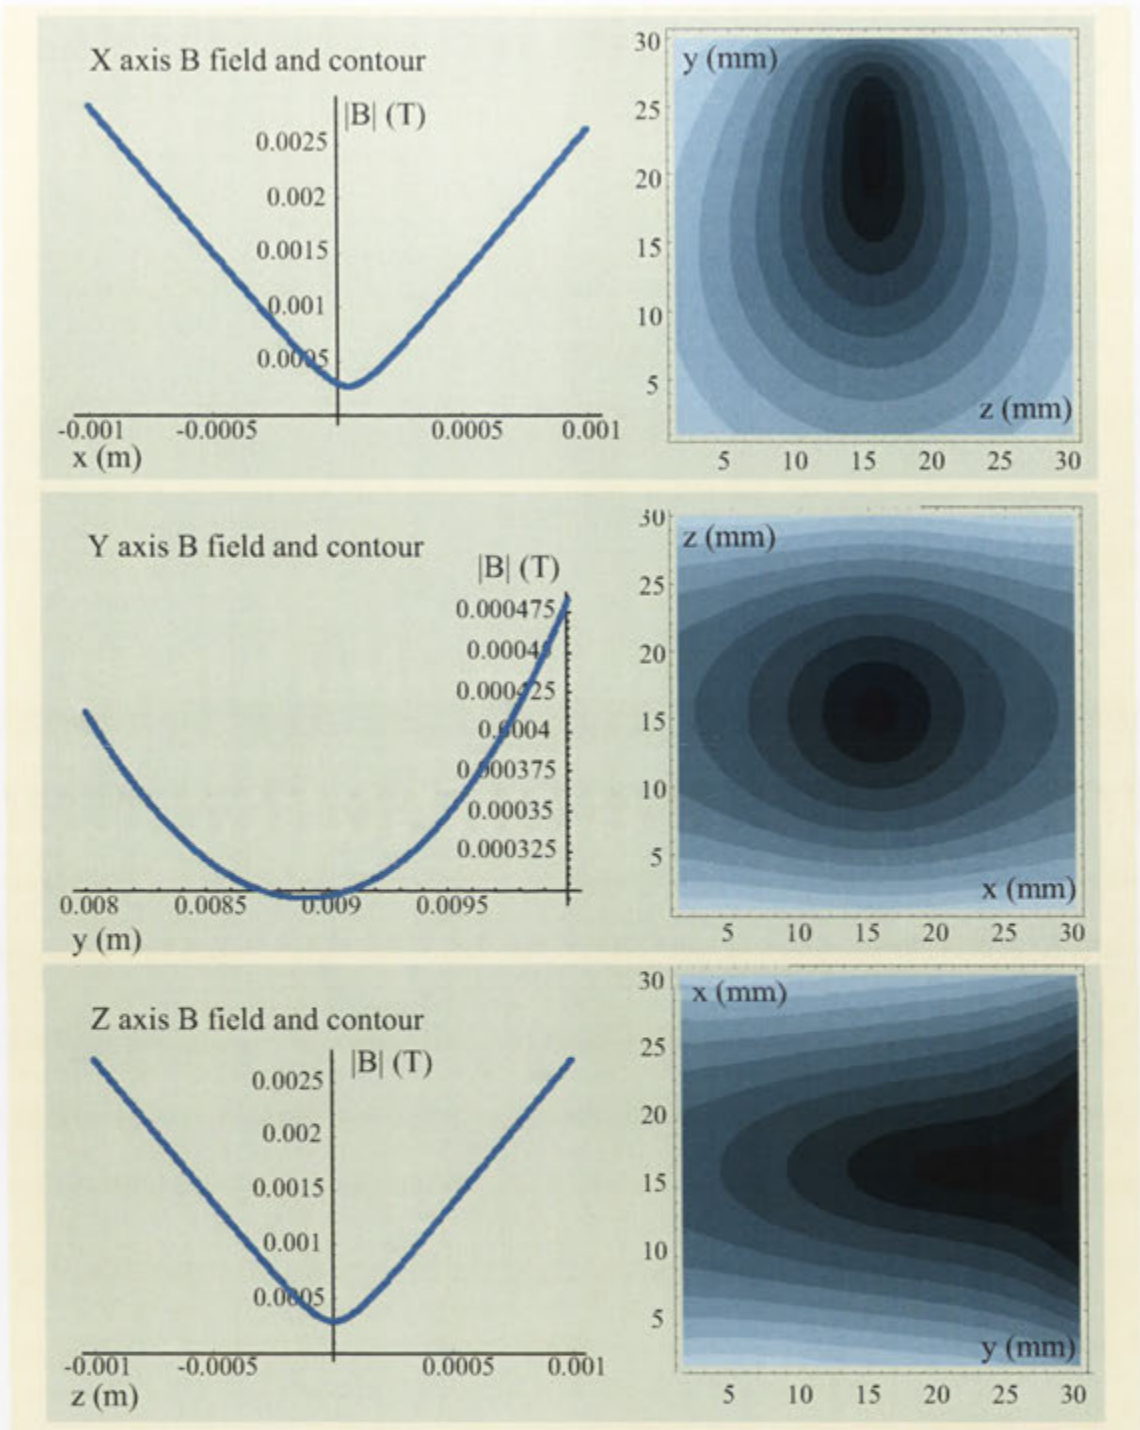


**6.7: The Mathematica model.** Note that it contains individual loops rather than solenoids. As opposed to the complex BiotSavart model these are just loops, rather than a helix, which will effect the accuracy of the model. However, the Mathematic model overcame calculation limitations of Biot-Savart.

$\bar{\omega} = \sqrt[3]{\omega_x \omega_y \omega_z}$ . The trapping frequencies are vitally important because they determine the distribution and densities of atoms within the trap, and hence the collision parameters that characterize evaporative cooling.

These parameters were extracted from the fully numeric models considered as part of this thesis, by fitting the appropriate polynomials using an advanced mathematical package such as Mathematica. Their effect was estimated by comparing the values with specifications published by other successful BEC groups around the world. NR extended these models into the analytic domain, allowing exact calculation of the various orders of the trapping fields. Also considered, from a theoretical viewpoint, was the importance of these quantities to the production of BEC via evaporative cooling.

A large amount of time was invested modeling various QUIC trap configurations in the BiotSavart magnetic field modeling software package, an example of which is shown in figure 6.5. This program has a user friendly, but quite powerful interface, which was very important for quickly analyzing and understanding models and results. The main benefit of BiotSavart over other numerical modeling techniques, such as coding from scratch, is that it intelligently selected computation methods at each datapoint in a calculation. Thus, if a particular datapoint was exactly (nearly)



**6.8: The magnetic fields due to the trap.** Further detail of the magnetic fields due to the magnetic trap. Shown for each axis is a plot of the field magnitude along the axis, and a contour plot of the field in the plane perpendicular to the axis, at the origin. The Ioffe coil is aligned along the  $y$ -axis. The quadrupole coils are aligned along the  $x$ -axis.

on-axis with a coil in the model, an analytic (perturbation) method would be used. If necessary it could perform a full numeric solution. This made BiotSavart very fast for even quite populated models, especially if the built-in composite objects, such as solenoids, were used.

The benefit of these low accuracy models was extremely fast prototyping of general aspects of the trap for broad geometry classes. These features were used to produce the first models of the QUIC trap to be used in the ANU-BEC experiment, verifying the results published in Esslinger’s paper [131]. In particular, these results highlighted the dynamic “morphing” of the trap from a quadrupole to IP field configuration as the current through the Ioffe was ramped up to its operating condition. Having a fast, easily modified, full three dimensional model to tweak, lead to an intuitive understanding of the trap that was vital for the final design.

They also elucidated the most important aspect of magnetic trap and coil design: *the exact location and geometry of the “front” of each coil is of primary importance.* This is obvious given the strong position dependence of magnetic fields, but illustrated in the context of creating a magnetic trap it gains added significance. Much more advanced design followed to optimize this principle to best effect. Results from a model of the final trap used in the ANU-BEC-I experiment are shown in figure 6.8. The key parameters of the plot are the bias field, and the curvature of the trap in three dimensions. These values can be calculated from the plots shown in figure 6.8 to be 1 G,  $\omega_x \sim 2\pi \times 200\text{Hz}$ ,  $\omega_y \sim 2\pi \times 20\text{Hz}$  and  $\omega_z \sim 2\pi \times 200\text{Hz}$ , respectively.

The extreme importance of the precise position and shape of the front of each solenoid-like coil, suggested that finer detail might also be important. Most worrying was the limitation of modeling a real coil as some limit of a perfect solenoid. Each layer in a real coil exists as a helix. Every second layer winds the helix in the opposite direction. There is clearly a complicated interplay between the precise position of an individual turn on the coil, and the specifics of its size, shape and angle relative to the coil axis.

Luckily, BiotSavart is quite capable of maintaining such a model, through judicious combinations of a composite helix model that it supplies. Once a small range of optimum size and shape was ascertained for each coil using the simple solenoid models described above, the model was extended to create each coil as pairs of layers of counter-propagating helices, joined by a straight piece of wire. Figure 6.6 shows the detail in the model. Apart from development time, the main trade-off accepted in extending the model to this level of complication was the decreased application of the “smart calculation” available in BiotSavart. This made most calculations fully numeric, and made calculation times long enough to limit the number of iterations tested. A combination of fast, simple, solenoid models, and slower, more complete, helix models, ensured that many prototype traps were produced and tested, however. The more complicated models did produce noticeably different results to simple solenoid calculations, especially with regard to the front-most turns on the Ioffe coil. The specifics of the front turns of the Ioffe coil were so disproportionately important that much time was invested attempting to add extra turns to the front of the coil by slightly varying wire diameter, coil geometry, and former shape.

The results presented in figure 6.8 are generated from the final complex Biot-Savart model of the ANU-BEC-I trap.

The main practical limitation with BiotSavart was purely functional. It was designed and programmed with very good spatial resolution specifications, completely sufficient for most practical calculations. The most detailed investigations of the ANU-BEC trap, however, involved analyzing the fields to the micrometre level. Although BiotSavart supported calculations to this level, it could not export data at these precisions because it used single precision floating point numbers in its exported data files.

The solution was to develop raw code in the commercial Mathematica package, produced by Wolfram software [153]. This level of modeling was only performed when the QUIC trap design had been finalized to within a very small range of geometries. The strength of coding a dedicated program is the underlying flexibility available to solve any general problem. As soon as the program is coded, however, it becomes less flexible than a general tool like BiotSavart. Consequently, the choice of features to be included in the model were carefully considered prior to coding.

BiotSavart was a very general tool, with good user interface and many useful pre-compiled composite models. In moving to coding the program directly, most of these features were lost. It was deemed necessary to recreate a key aspect of the BiotSavart user interface: the graphic representation of the coils being modeled. This very descriptive “double check” of the underlying code saved unnecessary calculation several times. An illustration of the model is shown in figure 6.7.

The first feature of the advanced BiotSavart model to be lost was the helical coils. By limiting the model to simple perfect loops both the coding, and perhaps more importantly, the calculation run times, were significantly reduced. This was not seen as a fatal flaw given that the final models of the QUIC trap were reproduced at all levels discussed in this section: a simple solenoid BiotSavart model; a “precise” helix-based BiotSavart model; and an “accurate” Mathematica model.

An important practical advantage of the Mathematica models over the Biot-Savart models, was the integration with further calculation of parameters required for trap analysis. The Mathematica model was fundamentally slower than the Biot-Savart models, because it did full numerical calculations for even simple loop-based models. However, the entire process of modeling a field, viewing it at different resolutions, and calculating the trap specifications from the results, was quicker and more efficient when executed completely within the Mathematica environment.

The Mathematica models verified the results of the advanced BiotSavart models at the smallest scales, and extrapolated them to scales unavailable in the BiotSavart models. The bias fields and trapping frequencies calculated using these models agreed with those generated from the complex BiotSavart model above.

Many of the key points ascertained through extensive modeling have been mentioned above. As many turns as possible should be created on the front of the Ioffe coil. Within the cone-shaped design, this is typically a trade-off between precise wire size and the precision with which a 45 °cone may be wound. The Ioffe coil should be positioned as close as possible to the science cell. This was achieved by designing an advanced former with a removable front piece so that the front coils

of the Ioffe could be maneuvered within 0.5 mm of the science cell.

In contrast, the quadrupole coils are not very important. They can be made in a classic short solenoid design. The main limitations become purely geometrical: they should not obstruct the optimum design of the Ioffe coil, and they should provide sufficient optical access to the science chamber.

To allow all coils to be run in series off a single power supply, they should be designed to run at the same current when producing an IP field. This typically requires extremely precise positioning of the coils, so needs to be considered in conjunction with coil mount design.

When all these factors had been taken in to account, the final consideration in the original ANU-BEC design was to maximize trapping frequencies, given available power supplies and expected switching performance. This was achieved, but at the cost of significant resistive heating of the coils during an experimental run. The water cooling design described in section 6.4 did not perform as well as expected. Although the trap performed well enough to produce BEC, this became a problem for more the advanced atom laser experiments that have been performed since. When the trap was redesigned for these experiments, more emphasis was placed on minimizing trap heating than further maximizing trapping frequencies.

The other problem unforeseen in the design of the original QUIC trap was the full extent of the limitations on optical access. The quadrupole coils were designed as short solenoids with large central holes (18 mm) to facilitate the MOT beams required during the laser cooling stage of the experiment. They were also expected to facilitate passage of a beam for absorption imaging the BEC. The dynamic behavior of the trap as it “morphed” from the central quadrupole trap to the shifted IP trap had been investigated. The IP trap was expected to be pulled 8-9 mm from the center of the quadrupole coils towards the Ioffe coil. The size of the hole through the quadrupole fields was marginal, but maximized the number of turns on the front of the coil within the expected requirements. In reality, the atom cloud was pulled 9 - 10 mm under the IP configuration, and was not accessible using the holes in the quadrupole coils. This led to a significantly more complicated imaging system. In future trap designs, the awareness of this problem, and the use of bigger MOT beams, made imaging through the quadrupole coils a reality.

NR has been able to extend these original models to the analytic domain. This has yielded the ability to mathematically optimize various QUIC trap designs for different, explicit, operating conditions. The most interesting result is probably the extremely low power designs that optimize trapping frequency at low resistive dissipation to ensure low heating. These designs have reduced power dissipation from over 1500 W in the original trap design, to less than 50 W in the optimized design, while maintaining trapping frequencies. One extremely low power design also utilizes a transport stage to move magnetically trapped atoms from the distant science MOT, to the full QUIC trap, relaxing the geometric optical access limitations and allowing more turns on the Ioffe coil very close to the science chamber. These changes have led to inherently more stable traps, very important for investigating noise in atom lasers.

## 6.3 Practical considerations

A major role of the magnetic trap is to isolate cold atoms from the surrounding environment, so that they can be cooled to temperatures low enough to reach BEC. Macroscopic magnetic traps tend to perform this task admirably, because vibrations in either the driving current or position of the coils tend to excite broad oscillations in the field, that are manifested as common mode oscillations across an entire trapped cloud [134]. These oscillations do not couple efficiently to general heating mechanisms, except in very small bandwidths around harmonics of the trapping frequencies. When a magnetic trap is employed during evaporative cooling, however, a new stability criterion arises. Specifically, variations in the driving current and position of the magnetic trap coils must not perturb the stability of the bias field at the bottom of the trap. This is vital, because the energy position of the RF knife is defined relative to the magnitude of the bias field. The trap coils must exhibit sufficiently stability, both in terms of the current creating the field, and the position of the coils themselves, in order to create reliable and repeatable bias fields during evaporation, yielding reliable and repeatable BECs. In addition, the fields should be precisely dynamically controllable, enabling rapid and reliable switching.

The temperatures required to produce BEC depend explicitly on the number of atoms in the sample, and, therefore, implicitly on the efficiency of evaporation, through the particular atomic collision characteristics and the type and tightness of conservative trap used. In the ANU-BEC experiments, temperatures of the order of 100 nK were typically reached after an evaporation cycle producing BEC. Producing atoms at temperatures this low requires sufficient stability in the magnetic fields creating the bias field that defines the bottom of the magnetic trap. A temperature of 100 nK implies stability of the bias of the trap of the order of  $10^{-3}$  G, or 0.1 %. This, in turn, implies stability of a similar order in the currents that define the magnetic fields created by the coils that make up the magnetic trap. Commercial power supplies exhibiting better than this level of stability are available.

In exactly the same way that current noise effects the stability of the bias field, and the repeatability with which BEC may be created in a given magnetic trap, vibration in the position of individual coils of the magnetic trap will perturb the bias field of the trap, and the accuracy of BEC production. Vibrations in the position of the Ioffe coil along its axis couple to the value of the bias field as in the stringent current criterion considered above. Although it is more difficult to analytically estimate the magnitude of this effect, simple numerical modeling shows that the relative position of the Ioffe coil must be maintained to significantly less than  $1 \mu\text{m}$  to prevent oscillations in the bias field equivalent to 100 nK resolution during the evaporative cut.

Achieving this level of mechanical stability in a BEC experiment usually relies on the art of the mechanical support staff rather than scientific design. To ensure stable magnetic fields the coils must be constructed ruggedly. They must be mounted solidly to the optical table to prevent large-scale vibrations. They must also be bonded solidly together so that individual turns of the coil cannot vibrate. The specific implementation of these general precepts in the ANU-BEC experiments is

considered in section 6.4 below.

The relatively high currents flowing through an operating magnetic coil heat it. As the coils heat, they undergo thermal expansion, even if they are otherwise ruggedly constructed. The best solution to this instability is to minimize current in the coils to reduce resistive heating. Unavoidable residual heating should be minimized by active cooling, such as flowing water through the former. In extreme cases, it may be necessary to wind coils out of small tube, and to pump water through the “wire” itself. Proper temperature stabilization is the most difficult factor in producing a very stable magnetic trap. In a well constructed magnetic trap, using high-quality power supplies, it is often the variation in the magnetic fields caused by thermal expansion of the magnetic coils that limits the overall stability of the magnetic trap.

To have full control over the experiment, it must be possible to manipulate the magnetic fields quickly. The currents in the coils themselves can be manipulated electronically, independent of the inductance of the coils. However, any closed conductive circuits near the magnetic coils will experience eddy currents induced during rapid changes of magnetic field. This is a much more difficult problem to remedy, and is best prevented during the design stage by minimizing highly conductive closed paths near the coil. This can be difficult if a highly thermally conductive former is required, as most such materials intrinsically exhibit high electrical conductivity. In this case, the closed-loop path should be broken by slitting the former completely through at one place to reduce dipole excitation of the eddy current loop around the former as the coils are switched on and off.

In addition to the power supply requirements considered above, currents need to be switched, ramped and delivered to the coils. In general, the ease and economy with which the control circuitry is implemented corresponds to the magnitude of the currents used in the coils. Currents below  $\sim 40$  A are controllable by standard high current semiconductor devices, significantly simplifying construction of the controlling circuitry required by the magnetic trap.

A key job of the electronics supporting the magnetic trap is to switch it off quickly relative to the motion of atoms in the trap. In macroscopic traps with geometric trapping frequencies of  $\sim 200$  Hz, this implies trap switch off times of less than 1 ms.

Switching high current, high turn count electromagnets quickly is a difficult task due to the large inductance of the coils, as shown by equations 6.4 - 6.6.

$$B = \frac{\mu_0 N I}{l} \quad (6.4)$$

$$L = \frac{\mu_0 N^2 A}{l} \quad (6.5)$$

$$I(t) = I_0 e^{-t/\tau}, \text{ where } \tau = L/R \quad (6.6)$$

Equation 6.4 is the magnetic field, in Tesla, inside a solenoid, of length  $l$  meters and number of turns  $N$ , carrying a current of  $I$  amperes. Although it is not the exact field applicable in magnetic trapping, it serves as a reference of proportionality. The

inductance of such a solenoid, with a cross-sectional area of  $A$  meters, is shown in equation 6.5. The key point of note is that the inductance scales as  $N^2$ , while the field only scales as  $N$ . Equation 6.6 shows the inherent time constant associated with such a coil if it has a residual resistance of  $R$  Ohms. The time constant simply scales the current driving the coil in an exponential manner during switch off. To increase the magnetic trapping field, either the number of turns or current need be increased. Either of these options increases the time required for the residual fields to fall to some pre-defined “negligible” value, when the coil is switched off.

It is possible to make  $\tau$  arbitrarily small, independent of the field-producing ability of the coil, by choosing a suitable resistance,  $R$ . Figure 6.11 (a), presented in the discussion of the ANU-BEC-I magnetic trap below, shows a typical circuit to achieve this end. The voltage generated by switching off such a circuit is shown in equation 6.7.

$$V(t) = L \frac{di}{dt} = -\frac{I_0 L}{\tau} e^{-t/\tau} = -I_0 R e^{-t/\tau} \quad (6.7)$$

If the resistance is made high to decrease the switching time, the “backlash” voltage generated can be very large. As it is opposed to the normal direction of current flow in the circuit, semiconductor devices can reach their reverse break down voltage and behave unexpectedly. A better solution can be a circuit that limits the voltage during switch-off, as discussed further in the report of the ANU-BEC-II magnetic trap, and in figure 6.11 (b), below.


## 6.4 The ANU-BEC-I magnetic trap

The evolution of the magnetic trap has probably been the most significant improvement at each iteration of BEC machine design at ANU. This is a result of its importance in producing a “high quality” BEC: one that has well defined parameters and can be re-produced reliably. A description of the magnetic trap used in each stage of the experiment is the best way of illustrating the evolutionary process, as well as the qualities that are most useful in magnetic traps for modern BEC experiments.

The ANU-BEC-I magnetic trap is shown in figures 6.1 and 6.2. It is a quadrupole-Ioffe configuration (QUIC) design based on that published by Esslinger et. al. [131]. The quadrupole coils are run without the Ioffe coil to provide fields for the science MOT, and for the initial capture of atoms in the magnetic trap. The three coils are run off a single supply during full Ioffe-Pritchard trapping and evaporative cooling. The full trap produces magnetic field gradients of roughly 200 G / cm, a curvature of 200 G / cm<sup>2</sup> at a bias field of 1 G, and trapping frequencies of  $\omega_x \sim 2\pi \times 200\text{Hz}$ ,  $\omega_y \sim 2\pi \times 200\text{Hz}$  and  $\omega_z \sim 2\pi \times 20\text{Hz}$ .

The key parameters of the coils are borne out by figure 6.1 and 6.9. They are wound of 2.2 mm round copper wire, as an intuitively chosen compromise between turn-density maximization and resistive heating minimization. The quadrupole coils are classic short cylindrical solenoids, approximately 30 mm long, with an inside diameter of 24 mm and an outside diameter of close to 70 mm. The Ioffe coil





**6.9: Diagram of the magnetic trap coils.** Detail of the quadrupole and Ioffe coils forming the ANU-BEC-I magnetic trap. The dimensions, in millimeters, are clearly marked and commented on within the text. The coil cross-section is positionally accurate to a reasonable degree (no account is taken of the helix nature of each turn). The slit through each former may be seen. Also apparent are the water cooling channels milled into the back of each former. These were sealed with a brass back plate, attached using a paste-applied silver solder technique.

is formed of a filled cylindrical back section and a filled conical nose section, of largest outside diameter 47 mm and overall length 35 mm. The conical nose section allows optical access for two perpendicularly-oriented MOT beams, each 45 ° from horizontal.

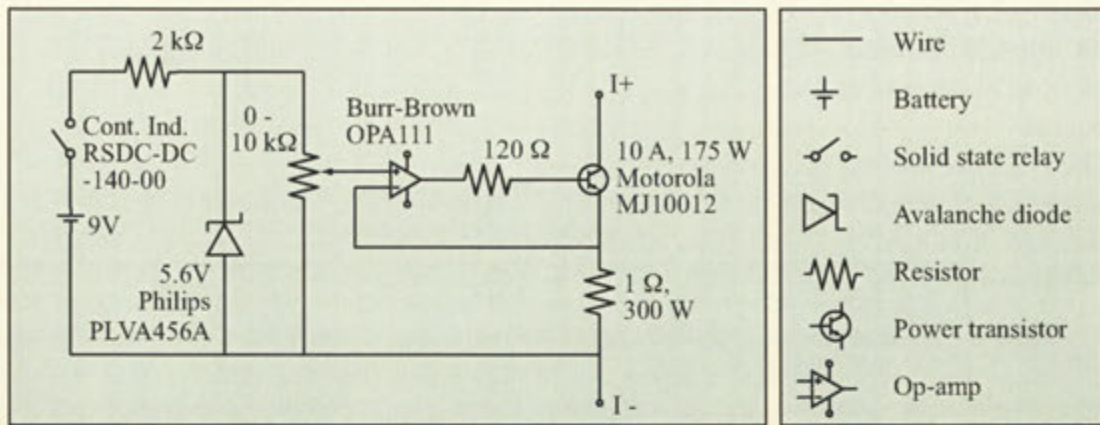
The coils are wound on precisely machined copper formers. Each layer of turns is covered in thermally conductive epoxy, to ensure good thermal conduction to the outside of the coil. The former and epoxy ensure that the magnetic coils are physically stable. The mounts positioning the coils must be solid, yet small to allow good access for optics near the science chamber. The mounts were constructed of non-magnetic stainless steel and kept very small. They appeared not to limit the stability of the trap.

The cooling in the original trap was confined to the back plate of the copper former supporting each coil. The highly conductive copper formers embraced the inside of each coil, and the thermal epoxy conducted heat to the copper efficiently. The back plate was 1 cm thick, and had a 1 cm wide double "C" shaped track milled into it around the slit in the former. A bronze plate with water tube connections was silver soldered over the copper plate to complete the liquid circuit. This formed a very compact, high flow, water cooling mechanism. Unfortunately, the conduction of heat from within the coils to the water cooled former was not sufficient. This led to the front of the coil, the part furthest from the cooling plate and most important in determining the stability of the magnetic trap, heating significantly. Newer trap designs focussed more on reducing the current requirements (and the consequent heating), rather than improving this active cooling design.

Eddy currents were prevented by ensuring no closed circular paths existed in any highly conductive medium near the magnetic coils. As the coil formers were built of copper to maximize thermal conduction, they had to be slit right through at one point to prevent the generation of massive eddy currents. The former was mechanically stabilized by gluing a precisely matched piece of perspex in to the slit. The newer coil designs, using potted coils rather than permanent formers, do not require this manipulation.

The stability of the current in the magnetic trap coils directly effects the stability of the magnetic field. Section 6.3 showed that a fractional power supply current ripple of  $10^{-3}$ , correlated between the coils of the magnetic trap, provides sufficient stability in the bias field to evaporate to the temperatures required for BEC. Fortunately, off-the-shelf power supplies with excellent current noise are available. Unfortunately, ANU-BEC only had access to a relatively cheap Powerbox SM-3540 switch mode power supply until the start of 2001. BEC was finally achieved after this power supply was replaced with a new F.u.G. Elektronik NLN 1400M-35 linear power supply, claiming  $10^{-4}$  fractional stability.

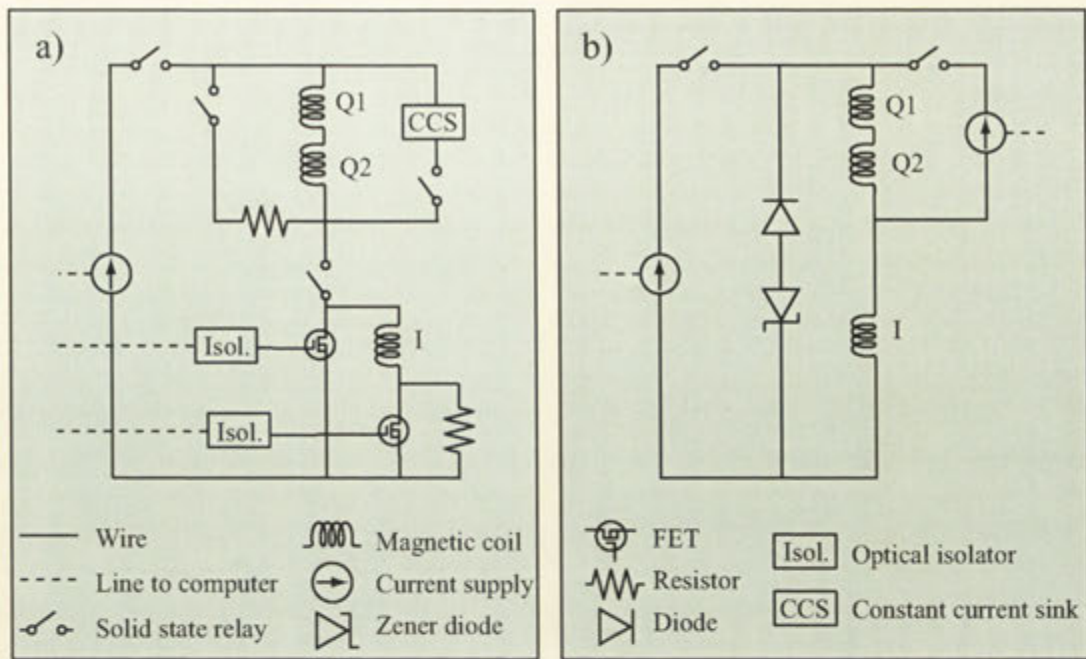
The original design for the ANU-BEC QUIC trap involved dynamic control of the Ioffe coil current relative to the quadrupole current in order to tune the precise nature of the magnetic trap. Initially, this was achieved using two independent power supplies, to drive the Ioffe coil and quadrupole coils respectively. However, driving all three coils, in series, from a single supply, suppresses common-mode oscillations that make the bottom of the trap fluctuate, leaving only less critical



**6.10: The constant current sink.** The constant current sink was an actively stabilized, tunable shunt impedance, that allowed several amps of current to bypass the quadrupole coils. This adjusted the current through the quadrupole coils relative to the current in the Ioffe coil, to maintain the optimum IP trap configuration. This active shunt was vital, because the ANU-BEC-I quadrupole coils exhibited significant resistive heating, and a corresponding impedance variation throughout each run of the experiment.

common-mode fluctuations. This was partially implemented by driving all the coils with a single supply, while shunting a small (few amp) current around the quadrupole coils, using a constant current sink. Although the absolute noise requirements for the constant current sink were identical to that for a second power supply, the fractional noise (the most common specification for power supply noise), was an order of magnitude smaller, because only a fraction of the total driving current was shunted.

The constant current sink is shown in figure 6.10. It is a relatively straightforward current follower design, but it was carefully constructed over several iterations to operate within the required noise specifications. The key components regulating the noise performance of the circuit are those constituting the voltage reference and comparator section of the circuit: the avalanche diode, the high-power 1 Ω sense resistor, and the precision Burr-Brown operational amplifier. The battery-powered voltage reference exhibited extremely low AC noise, but the avalanche diode was vital to prevent slow sag of the reference voltage as the battery discharged. In a similar manner, the 1 Ω sense resistor was over specified to minimize drift in its resistance due to ohmic heating during constant operation. The op-amp continuously compared the voltage of the low-noise reference to the voltage generated across the sense resistor, adjusting the dissipation in the power transistor to maintain a steady current, typically  $\sim 2$  A, through the device, independent of the parallel impedance of the quadrupole coils. The constant current sink exhibited an excellent fractional stability of  $10^{-4}$ , and was a key development in the progress towards producing



**6.11: The ANU-BEC switching circuits.** The circuit on the left is that used to control and switch the ANU-BEC-I magnetic trap. Its key characteristics are the use of a constant current sink parallel to the quadrupole coils, and the use of resistors to control magnetic trap switch off. The circuit on the right is the control circuit implemented for ANU-BEC-II. Its inherent simplicity arises from the use of micrometer stages, adjusting the physical position of each coil, to tune the zero-point operation of the trap. The other great improvement was the use of a zener diode to control the decay of fields during switch off.

BEC. However, it did require regular maintenance, and was a source of day-to-day variation in the performance of the ANU-BEC-I machine.

To avoid implementing a constant current sink, yet still drive all coils in series off a single power supply, their position must be adjusted with a micrometer stage. This was implemented on the new trap design, and has proven to be most superior.

Figure 6.11 (a) shows the type of circuit used to tune the switching specifications of the magnetic trap during the ANU-BEC-I experiment. Several variations of this circuit were implemented, but the key characteristic was the use of resistors to accelerate the decay of the fields due to the magnetic coils. The switching time was generally maintained at somewhat less than 1 ms, by selecting the value of the resistors according to equation 6.6. The large backlash voltages described by equation 6.7 were a difficult problem, due to their tendency to exceed the reverse breakdown voltages of the solid state relays used in the control circuit, causing unpredictable behavior. This was resolved with brute force, by pairing semiconductor

switches with reversed polarities wherever it was a problem. This was not an elegant solution, but allowed sufficient control to produce BECs.

## 6.5 Improvements to the magnetic trap for ANU-BEC-II

The ANU-BEC-II trap, designed by NR, is shown in figures 6.12 and 6.13. They provide a graphic comparison to the ANU-BEC-I illustrations in figures 6.9 and 6.2, above.


The dimensions of the ANU-BEC-II trap contain many subtle improvements. The most apparent is the size of the holes in the quadrupole coils, and the addition of a small hole through the center of the Ioffe coil. The hole in the center of each quadrupole coil was increased to 30 mm diameter, allowing optical access to the entire width of the glass cell forming the science chamber. This was a massive improvement in terms of the possibilities available for imaging the BEC at the conclusion of the experiment. It was a vital step in the development of more advanced experiments using the ANU-BEC, such as the atom laser work considered in chapter 10.

Less obvious is the fact that the wire diameters of the quadrupole and Ioffe coils were reduced to 1.6 mm and 1.4 mm, respectively. This change was a key result of the advanced, analytic, modeling of the new trap. The dimensions of the coils, their position, and the wire diameter, were simultaneously optimized to yield significantly lower drive power, and less resistive heating. In turn, because the coils dissipate less heat, the active cooling mechanism is less vital to repeatable operation of the trap.


The mount is a major improvement over the ANU-BEC-I design. The coils are positioned precisely using micrometer screw adjustments. This negates the need for a constant current sink around the quadrupole coils in the driving circuitry. This allows the entire circuit to be implemented from readily available commercial parts, significantly simplifying construction and increasing reliability. Along with the improvements in imaging, this is probably the most important difference between the ANU-BEC-I and ANU-BEC-II designs.

The ANU-BEC-II coils *are* slit in a manner similar to the ANU-BEC-I coils. However, due to a planning oversight, they have been situated poorly. Figure 6.12 shows that they are placed downwards, so that when attached to the mounts, the circuit is completed across the slit via the mount itself. This effect is exacerbated by the high conductivity aluminum construction of the mounts. ANU-BEC-I used low permittivity 316 stainless steel mounts partly for their low conductance and reduced eddy current generation.

However, switching performance was not compromised, due to the implementation of a zener diode instead of resistors to dissipate energy released from the magnetic coils at switch off. Figure 6.11 (b) shows the circuit. This circuit clamps the voltage, instead of the dissipation, during switch-off, to ensure the semiconductor switches operated properly. Although this system leads to slightly slower initial switch off, it creates an almost linear decay, so actually reduces the currents to use-



**6.12: Diagram of the ANU-BEC-II magnetic trap.** Detail of the ANU-BEC-II coils for direct comparison to figure 6.9. The key differences are the provision of larger holes through the quadrupole coils, a hole through the Ioffe coil, and the use of smaller diameter wire. The extra optical access made the trap significantly more useable than in the ANU-BEC-I experiment. The smaller wire optimized the field for a given level of resistive heating in the trap.



**6.13: Photo of the ANU-BEC-II magnetic trap.** The most apparent difference to the ANU-BEC-I trap is the size and nature of the coil mounts. These provide the platform for the operation of the micrometer stages used to tune the trap.

ful levels much faster than the exponential decay afforded by the resistive switching circuitry of ANU-BEC-I.

### 6.5.1 Future improvements to the magnetic trap

The ANU-BEC-III magnetic trap follows the principles established in the design of the ANU-BEC-I and ANU-BEC-II traps to their natural conclusion. The defining difference in the design of the ANU-BEC-III magnetic trap is the Ioffe coil. Making the Ioffe coil conical to allow access for MOT beams, as in ANU-BEC-I and ANU-BEC-II, represented a major compromise to the performance of the magnetic trap as a whole. The Ioffe coil can be wound cylindrically if a mechanical translation stage is used to physically move the MOT coils from the position used to make the science MOT, to a remote location in which the magnetic trap is formed. This maximizes performance of the trap during evaporative cooling, at the cost of increased complexity earlier in the experiment.

Releasing the limitation on the Ioffe coil to provide MOT beam access to the science chamber opened a whole new parameter space for the dimensions of the QUIC trap coils. NR implemented an analytic model of the trap, coupled with numerical modeling in BiotSavart, to optimize the trap for minimum resistive heat-

ing during operation. The result predicted by the new model is spectacular: the ANU-BEC-III trap is expected to dissipate approximately 50 W, as opposed to the 700 W dissipation of ANU-BEC-II, or the 1500 W dissipation of ANU-BEC-I. This is expected to lead to unprecedented stability during operation.

The new paradigm has also yielded significant simplifications of the design. The exceptionally low resistive dissipation in the coils displaces the requirement for active water cooling. In addition, the coils were wound with conductive potting epoxy on removable teflon formers. Although the epoxy is conductive, providing some small thermal dissipation, it is not so conductive as to require an eddy-reduction slit. The more efficient trap allows the use of more economical, lower capacity power supplies, and readily available commercial parts in the control circuitry.

## 6.6 Characterizing the magnetic trap with cold atoms

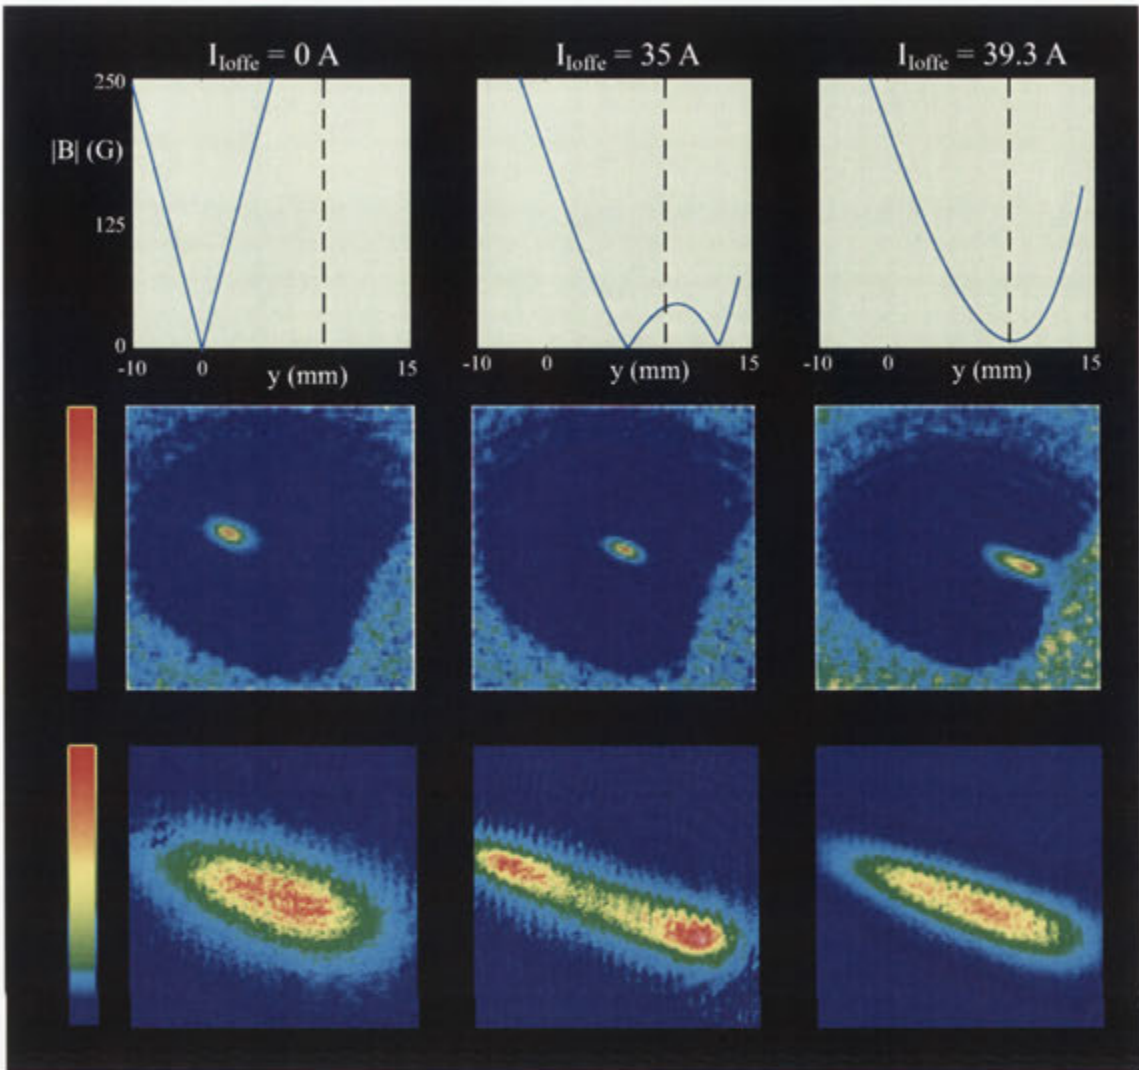
A large part of the benefit of mathematically modeling the magnetic traps during the design and construction phase of the experiment was realized after the magnetic trap was actually constructed. By characterizing various properties of the magnetic trap with cold (but not Bose condensed) atoms, the initial design models were extended to provide a very good, experimental “working model” of the actual trap produced for the experiment. This was important because many of the underlying principles and precise facets of the design elucidated by the initial modeling described above, were extremely difficult to implement and build at similar levels of accuracy.

The most obvious way to characterize the trap experimentally would have been the direct measurement of the magnetic fields produced using a magnetometer. Unfortunately, the available equipment was not sufficiently accurate to provide any useful data for comparison to theory. Since the traps for ANU-BEC-II and ANU-BEC-III have been developed, more advanced magnetometers have been employed to provide good characterization of the traps before they are implemented into experiments.

The key characterizations readily accessible in the original ANU-BEC-I magnetic trap were related to the trap evolution from a quadrupole configuration to an IP configuration, during the ramp up of current through the Ioffe coil. Figure 6.14 (a) shows how the trap evolves. It is very similar to a figure in the original paper published by Esslinger, but is created using the models described above [131]. The ability of the trap to “morph” gently from the quadrupole configuration to IP configuration allows the quadrupole coils be used for the science MOT, initial magnetic quadrupole trap and final IP trap, inherently building in perfect alignment of all magnetic fields in the experiment. The trade-off accepted to use such a configuration, however, is that the quadrupole trap and the IP trap must be in physically different places. This does not appear to effect the transfer efficiency of atoms between the MOT and the QUIC trap: Esslinger’s paper claimed a transfer efficiency of over 90 %.

A benefit of this situation, however, is that it presents a manner of accurately





**6.14: The evolution of the QUIC trap.** A series of plots showing the evolution of the QUIC trap from a quadrupole to an Ioffe-Pritchard configuration.

determining the nature of the magnetic fields created by the magnetic trap, using the precise position of atoms in the trap rather than an accurate magnetometer. The BEC imaging system provides the measurement device, and by recording a series of images during the transfer process, and comparing to the calculated results, the mathematical models can be fine-tuned to iteratively improve the minutiae of the design.

Figure 6.14 (b) and (c) shows the images of atom clouds associated with the models shown in figure 6.14 (a). The agreement is excellent well beyond the level apparent in the images presented. Only very fine tuning was required to adequately adapt the original models to the actual constructed coils.

As the push towards BEC progressed, the excellent understanding of the magnetic trap also allowed the modification of the experiment to achieve configurations unforeseen in the original design. The model was extended for the implementation of the constant current sink QUIC trap described above. It also allowed the modification of the experiment to run at much lower currents after BEC was first achieved, and the medium term thermal stability of the magnetic trap was found to limit the performance of the machine. In both these scenarios, the time invested in developing an accurate model of the magnetic trap paid large dividends in extending the usability of the initial investment into new regimes of operation.

The magnetic trap is arguably the most crucial aspect of a BEC experiment, differentiating it from less stringent laser-cooled cold-atom experiments. The importance of having an excellent understanding of this part of the experiment can not be overstated. It certainly represents a key step along the path to successfully producing a BEC.



# 7 Evaporative cooling

Evaporative cooling is the method by which a dilute gas of Bosonic atoms can finally be cooled far enough that they undergo condensation. Its demonstration was vitally important to the production of dilute gas BECs. The development of the RF-knife method of evaporation from a magnetic trap increased the efficiency and ease with which BEC could be achieved [137]. Counter-intuitively, its experimental implementation seems to be quite robust in terms of care and optimization. This chapter is a summary of the evaporative cooling systems used on the ANU-BEC-I and ANU-BEC-II machines.

## 7.0.1 Evaporative cooling to BEC

The RF-knife method of evaporation from a magnetic trap uses energy selective removal of atoms from the trap to reduce the average energy of atoms left in the trap. This process can be continued to arbitrarily reduce the average energy of atoms left in the trap, which, when rethermalized, corresponds to a lower temperature. The RF-knife allows the energy selective removal of atoms to be performed without loosening the trap, maximizing efficiency by maintaining atom density and collision rate during the evaporation process.

## 7.0.2 Evaporative cooling in context

There are many reports calculating the optimum scheme for evaporative cooling. Groups have invested much time testing different setups for creating, amplifying and transmitting RF signals into the science chamber. It appears that the precise experimental technique employed is just not that important, or, at least, not accomplished with the precision of the calculations. Achieving the required trapping frequencies in the magnetic trap, thereby creating sufficient density in the atom cloud during evaporation, is as important as the RF signal that drives the process.

## 7.1 Theoretical considerations

The RF-knife technique is the most popular method of implementing evaporation in magnetic traps. At a given magnetic field, the RF-knife signal resonantly drives atoms between Zeeman states. Such atoms are driven out of trapped weak-field seeking  $m_F$  sublevels to untrapped (or antitrapped) sublevels, whence they are lost from the trap. By beginning this process at high RF frequencies, the resonant condition for transition between states occurs at high magnetic fields. In the conservative magnetic trap, only atoms with high energy can ever populate these high

magnetic field regions. Thus, the most energetic atoms can be removed from the trap, reducing the average energy of those atoms left in the trap. As the atoms left in the trap undergo elastic collisions, they rethermalize to an equilibrium Maxwell-Boltzmann distribution of energies. If the RF-knife frequency is reduced further, the process occurs again. More and more atoms are lost, and the remaining atoms rethermalize to lower and lower temperatures.

The limiting factors in evaporation are the starting number of atoms and the efficiency of evaporation. The temperatures that can be attained with evaporation are only limited by how many atoms can be lost, and how much “above-average” energy they can carry away from the trap. Normal atoms in thermal equilibrium exhibit a Maxwell-Boltzmann distribution of energies. An RF-knife may be set to cut a few high energy atoms from the trap, or very few, very high energy atoms from the trap. In practice, the optimum operation point is determined by the length of time evaporation is performed over.

Long evaporation times are limited by technical sources of heating and loss in the trapped atom population. The lifetime achievable in UHV magnetic trap systems is typically of the order of 100 s. Evaporation should be performed faster than this to avoid non-cooling losses. These factors have been discussed quite extensively in chapter 4, with reference to the requirements of the BEC vacuum system and magnetic trap lifetimes.

Short times are limited by the rethermalization rate of the atoms in the trap. Evaporation will not be efficient unless the atoms remain roughly thermally equilibrated throughout the cycle. A truncated Maxwell-Boltzmann distribution, created, for example, by a high-energy cut in the magnetic trap with an RF knife, will have its spread of energies re-distributed after approximately two to five collisions per atom, to yield a new thermal distribution [137]. Thus, the rethermalization rate is directly related to the elastic collision rate,  $\gamma = \sigma_{2b} n v_{th}$  [135]. The elastic collision rate is determined by the collision cross-section,  $\sigma_{2b}$ , the density,  $n$ , and the thermal velocity of the atoms,  $v_{th}$ . The collision cross section is fixed for a given species, but the density can be increased by tightening the magnetic trap. In a harmonic trap, the density and thermal velocity can be re-expressed in terms of the total number of atoms and the temperature of the sample, to yield  $\gamma \propto \sigma_{2b} N/T$ .

Runaway evaporation occurs when the rethermalization rate increases as the evaporation proceeds. This is an important definition of efficient evaporation, because it guarantees continued efficiency throughout the entire ramp. Runaway evaporation occurs when the rate of temperature decrease drops faster than the rate at which atoms are lost due to evaporation. This is an interesting criterion, because it differs from the requirement for increasing phase-space density. BEC depends on the density and average delocalization of the constituent atoms in the cloud, according to  $\mathcal{D} = n \lambda_{dB}$ , as discussed in chapter 2. In a harmonic trap, this reduces to the relation  $\mathcal{D} \propto N/T^3$ . Contrasting the number-temperature dependence of the phase-space density to that of the elastic collision rate, it becomes apparent that there are regimes in which the phase-space density will increase, but runaway evaporation will not be achieved. Achieving increasing phase-space density, but not runaway evaporation, at the start of an evaporation run, will not guarantee increas-

ing phase space density throughout the run, because the collision rate will decrease, until, ultimately, the phase-space density, too, ceases improving.

Inelastic collisions that lead to loss from the trap also depend on density. Three-body recombination collisions, in particular, depend on the square of the density. Optimum runaway evaporation proceeds when the density is set to yield the best ratio of rethermalizing elastic to lossy inelastic collisions. These parameters are controlled, almost exclusively, by tuning the specifications of the magnetic trap storing atoms before and during evaporation. This matter has been considered in chapter 6.

A procedure for estimating the collision rate throughout the evaporation ramp, using standard absorption images, was published by Lye, Fletcher, Kallmann and Close in 2002 [6]. Beer's law defines the optical depth as  $O = \ln(I_0/I_t)$ , where  $I_0$  is the light incident on the cold atom cloud, and  $I_t$  is that transmitted through the cloud. Its peak value, measured at the center of a cold atom cloud expanded for time  $t$ , can be calculated as in equation 7.1 [154].

$$O_{pk} = \frac{\lambda^2 m}{4\pi^2 k_B t^2} \frac{N}{T} \quad (7.1)$$


This result is important, because, for a given imaging setup, with a given imaging wavelength and expansion time, changes in  $O_{pk}$  are directly related to changes in  $\frac{N}{T}$ . Thus, the peak optical density of an imaged cloud is directly related to the elastic collision rate of the atoms in the trap. The peak optical density is easily resolved from experimental absorption images of cold atom clouds, and changes in its value throughout an evaporation run directly intimate changes in the elastic collision rate, and whether an experimental cycle achieves runaway evaporation. This result became a key diagnostic in the push towards the first BEC at ANU, as reported in chapter 10.

The RF frequencies necessary for evaporation from a macroscopic magnetic trap begin near 30 MHz, and end near the bias field of the trap: typically  $\sim 1$  MHz. These frequencies are created by an RF signal generator, amplified, and transmitted into the science chamber using a simple loop antenna. The precise requirements for signal purity, intensity and transmission efficiency seem to be very flexible. Many groups report designing an optimized system, changing it significantly and observing no significant reduction in performance.

## 7.2 Evaporative cooling in the ANU-BEC-I experiment

The ANU-BEC used a Stanford Research Systems DS-345 synthesized function generator to create the RF-knife signal. It was transmitted into the science chamber using a simple single loop antenna placed, coaxially, in front of a quadrupole coil. At times, the signal was amplified, but often in an empirical manner, with little consideration to the efficiency of coupling this power into the BEC chamber.

A typical evaporation cycle to BEC consists of a single exponential ramp from 25 MHz to 1 MHz with a time constant of 8 s. The consequent cooling of an atom



**7.1: Evaporation to BEC.** Sequence of images showing the progression of the temperature, density and phase space density of an atom cloud as it is cooled to BEC. These data were recorded on separate experimental runs, due to the destructive nature of the imaging process. This is possible due to the repeatable performance of ANU-BEC-I.

cloud is shown in figure 7.1. The start frequency is not critical. The time constant is important, but not overly sensitive around the optimized value. The end point of the ramp is absolutely crucial in defining the temperature of the atoms remaining after evaporation. Its effective magnitude is defined relative to the bias field at the bottom of the magnetic trap.

The temperature of the atoms remaining after evaporation roughly corresponds to the energy difference between the bottom of the trap and the final frequency reached during the evaporation cycle. This must reach temperature-equivalent energies of 500 nK or less to produce BECs. This implies great accuracy in both the stability of the magnetic trap bias field ( $\sim 10$  mG), as discussed in chapter 6, and the end of the RF ramp ( $\sim 10$  kHz).

The DS-345 function generator has excellent control over the absolute frequency it produces, due to the digital synthesis technique that is used to generate the signal. Unfortunately, it uses an imperfect attenuator, limiting the smallest amplitude field that can be generated. That is, it can not be completely turned “off” using electronic means. It also harnesses its own very general and powerful digital programming scheme for producing the preset frequency ramps used in the BEC experiment. Unfortunately, this technique restricts the quantization of the experimental ramp

in this particular model to between 1500 and 3000 equal time-step frequencies. This makes the last few frequency decrements inordinately step-like.

### 7.3 Improvements to evaporative cooling during ANU-BEC-II

The evaporation setup was modified throughout operation of the ANU-BEC-II to eke out the last few iota of performance and stability in efforts to produce an atom laser. However, the only significant long term changes in the evaporation setup was the addition of an external RF switch outside the Stanford Research Systems DS-345 function generator to augment the units' own attenuation abilities, and the addition of an Agilent 33250A 80 MHz arbitrary waveform generator of similar design. The Agilent generator used similar digital ramp generation, but exhibited better control over the output amplitude. Using combinations of both signal generators, NR was able to evaporatively cool a cloud of  $^{87}\text{Rb}$  to BEC, then carefully outcouple matter waves from this amazingly fragile object, to create various configurations of atom lasers. These results are summarized in chapter 10.






## 8 Imaging and CCD cameras

Recording an image of the cold atom cloud is a vitally important aspect of a BEC experiment, as it is the primary way in which information is gathered. Much effort has been expended worldwide designing and implementing advanced imaging systems [11, 38, 155]. However, the bulk of data generated on current machines is collected using a simple absorption imaging system, shown diagrammatically in figure 8.1.

### 8.0.1 Imaging and CCD cameras in BEC experiments

A typical absorption imaging sequence from the ANU-BEC machines proceeds as follows. On construction of the BEC machine a laser beam is aligned through the region of the experiment where the BEC will form. The beam is made much larger than a typical BEC, and is controlled in both frequency and intensity by an acousto-optic modulator. During BEC production the imaging light is shuttered off from the experimental region. After a BEC has formed in the magnetic trap, and any manipulations or experiments have been performed, it is released into free-fall by rapidly turning off the magnetic fields. During this process, a small bias field is applied along the imaging beam path to encourage the freed atoms to adiabatically realign their moments co-linear with the imaging beam. After a variable delay of some milliseconds, the imaging light, detuned from resonance by some megaHertz, is pulsed on for some hundreds of microseconds. The light travels through the BEC, experiencing dispersion and absorption due to the real and imaginary parts of the refractive index describing the atom cloud. Both the absorption and dispersion information is collected using lenses, and imaged on to a CCD camera, which quantifies the amount of absorption, in two dimensions, inflicted on the beam by the atoms, integrating the third “column” dimension.

This technique, when combined with calibrating procedures and some significant analysis, provides many pieces of information regarding the cold atom cloud or BEC. In short, it yields the two dimensional column density of the atom cloud imaged. In itself, this goes some way to describing the three dimensional shape of the cloud, and is capable of detecting shape oscillations [156], vortices [41], and many other interesting features of BECs. Integrated, this provides information as to the total atom number in the BEC, enabling studies of stability and any process that causes loss. Finally, by varying the expansion time, the temperature of the atom cloud may be gleaned. This is a vitally important diagnostic throughout all BEC experiments.



**8.1: Absorption imaging.** Diagram of a the configuration used during absorption imaging. After formation and manipulation, the BEC is released from the magnetic trap into free-falling expansion. A large, near resonant beam is projected through the region containing the atom cloud. The cloud absorbs photons, both diffracting and refracting the beam. The lenses in the imaging system collect the transmitted beam, as well as light deviated due to the action of the cloud, re-imaging it onto the CCD chip. The reconstructed image information portrays the shadow cast by the BEC on the imaging beam due to absorption of photons from the beam. This information allows calculation of various parameters of the atom cloud, including the total number of atoms and the temperature of the cloud.

## 8.1 Fundamental considerations

Optimal absorption imaging system design in a BEC experiment should fulfill a simple aim. It should produce images containing the most information about the BEC possible: with the highest spatial resolution, at the greatest quantitative “depth”. Specifically, it should maximize the signal-to-noise of both these parameters. In general, there are both fundamental and technical limits on the signal to noise: an optimum imaging system should be limited by fundamental issues only.

Within the description of absorption imaging given above there are several aspects that can be manipulated to change the fundamental signal and noise sources in the experiment. The intensity, duration and detuning of the imaging laser pulse can be manipulated. This effects how much, how long and how efficiently the atoms interact with the imaging light. In addition, the expansion time allowed after switch-

ing off the magnetic trap effects both the size and density of the atom cloud being imaged.

Technical limitations in the experiment are the resolution limit of the imaging optics and the noise added to the atom signal (and noise) by the CCD camera detection and electronics. The choice of optics and electronics can greatly alter the amount of technical noise in the imaging system. There may be other technical noise sources, such as instability of laser light, but they are not considered in the following analysis.

It is not overly difficult to design an imaging system limited by fundamental noise within these parameters. The surprising fact is that many groups over-engineer some aspects of their imaging systems unnecessarily due to misconceptions about where the dominant noise sources occur.

The most obvious, perhaps, is the confusion of the “pixel count” with the “resolution” of a CCD camera. Pixel count is very rarely a factor that limits the information collecting capability of the system as a whole. This is because BECs are very small objects, even when they are expanded, and the fundamental resolution limits imposed by the system are quite large. Standard imaging configurations yield between  $10 \times 10$  and  $100 \times 100$  “pixels” of information at most.

On the other hand, the resolution of a CCD camera system can present a significant technical limitation on the performance of the imaging system. In this context, it is the *size* of the pixels that determine the suitability of a given CCD camera for a given imaging system, or vice-versa. Smaller pixel sizes require less magnification from the imaging system to achieve appropriate spatial resolutions. A simple matter of important note, is the inherent reduction of the apparent resolution due to the spatial “quantization” of the CCD. That is, the impact of the Nyquist effect on CCDs. The precise application of this specific 1D sine-wave theory to the 2D data spread of a CCD is somewhat controversial, but the practical limitation is qualitatively similar. This effectively reduces the spatial resolution of the CCD to the order of twice or three times the size of the pixels. [157]

The sensitivity, or dynamic range, of the CCD camera may also be expected to limit the performance of the imaging system, but yet again, more fundamental factors outweigh its effect. The real limit in most systems is determined by the unavoidable quantum, or shot, noise inherent on the light used to probe the BEC. Typical parameters of an optimized system are derived below, showing that quantum noise may be expected to degrade the signal at up to the 1 % level. A cheap 8-bit CCD camera will resolve this noise, and more expensive cameras with higher sensitivity will resolve it even better. However, they will not extract more information about the BEC.

It turns out that in most well designed imaging systems, the real limits on the quality of images produced are a direct function of the small size and fragile nature of the BEC itself. The calculations below emphasize this “fundamental” limit, comparing it, where necessary, to the more traditional considerations that are often assumed to limit absorption imaging systems.

### 8.1.1 Resolution

The limited resolution of optical systems arises from the finite ability of constituent lens components to collect scattered light from an imaged object. This represents a loss of information, and causes degradation of the reconstructed image. In the language of spatial information processing, the finite size of the optical system imposes a low-pass filter on the scattered optical information, removing high spatial frequencies, or fine spatial detail.

Lenses are fundamentally necessary for accurate imaging, because light passing around any object will be diffracted, scattered away from its incident direction. The established limit quantifying such effects is the diffraction limit, a simple approximation to which is  $\Delta\theta \sim \lambda/2d$ , where an object of size  $d$ , imaged using light of wavelength  $\lambda$ , causes a misdirection of the scattered light through an angle of  $\Delta\theta$ . When  $\lambda$  is significantly smaller than the size of the object being imaged, as is the case for most personal everyday observing situations, the diffraction limit poses but a small perturbation to direct observation. However, as the desire for precision leads to the observation of smaller and smaller entities, the magnitude of  $d$  approaches that of  $\lambda$ , leading to a significant angular perturbation of light leaving the object.

It is the extent of this perturbation that is often used during the design of imaging systems quantifying BECs. A necessary, though not sufficient, condition for designing a “diffraction limited” imaging system is the collection of all light scattered from the object at angles up to  $\Delta\theta$ , ensuring that the spatial information is not degraded due to loss of information from the system. This requirement can be calculated intuitively for the case of the first collection lens in the system: the diameter of the lens must subtend the angle of diffraction,  $\Delta\theta$ . For a 2-f imaging system, yielding unity magnification, the first lens is positioned at twice its focal length from the BEC. Assuming  $\Delta\theta$  is a small angle, the diameter of the lens need only be  $D = 2f\Delta\theta = f\lambda/d$ .

The specific design of the optical system often offers further complications: at all points throughout the collection region the same condition must be met regarding the angular resolution. In BECs it is often it is not the first lens of the system that limits its operation, but the window exiting the vacuum chamber. Other factors include the resolution performance of the device used to record spatial information: the CCD camera, as mentioned above.

These considerations represent a traditional approach to imaging system design. They highlight the importance of diffraction as the fundamental limit to the scattering of imaging light during interaction with the object. Using these calculations, an imaging system can be designed that collects almost all the scattered light from the object, thereby maintaining spatial information, allowing image reconstruction at the optimum resolution. The cut-off frequency of the “low-pass” filter, constituted by the imaging processing system, is set just beyond the maximum frequency of useful spatial information.

However, if some unconsidered effect diverts light more voraciously than the process of diffraction, information will be scattered through angles larger than the diffraction angle calculated above. An imaging system designed only to capture

light deflected by the diffraction angle will disregard some of this more distinctly misdirected light. The cut-off frequency of the imaging system will fall below the highest information frequencies, information will be lost, and the image will be degraded.

Just such an effect may befall imaging systems observing BECs. In addition to diffraction, incident imaging light suffers refraction by the cold atom cloud. The BEC itself “lenses” the imaging beam. This is an effect of the dispersive part of the refractive index describing the BEC. The refractive index,  $n_{ref}$ , is related to the atomic density and light detuning as shown in equation 8.1, below.

$$n_{ref} = 1 + \frac{\sigma_0 n \lambda}{4\pi} \left[ \frac{i}{1 + \delta^2} - \frac{\delta}{1 + \delta^2} \right] \quad (8.1)$$

Equation 8.1 assumes that  $n - 1 \ll 1$ , and the entire probe beam exceeds saturation throughout the cloud. In addition, if the optical transitions are not well described by a single cross-section,  $\sigma_0$  must be adapted to correctly attribute contributions from all absorbing transitions. The detuning,  $\delta = \frac{\omega - \omega_0}{\Gamma/2}$ , is expressed in half linewidths. The equation applies to both BECs and expanded atom clouds, within the stated assumptions, via the explicit dependance on the density,  $n$  of the cloud being imaged.

Lensing of the optical beam by the BEC is due to the dispersive part of the refractive index. The BEC behaves like an indistinctly shaped, variable refractive index lens. The effect on a laser beam can be calculated, although in general it is fairly complicated. Modeling the BEC as a ball lens, of diameter  $d$ , with a uniform refractive index corresponding to the typical density of a real BEC, yields an effective focal length as in equation 8.2.

$$f_{BEC} = f_{ball} = \frac{n_{ref} d}{4(n_{ref} - 1)} \quad (8.2)$$

The phase shift through the ball lens is  $\delta = 2\pi d(n_{ref} - 1)/\lambda$ , which may be used to re-express the focal length as  $f_{BEC} = \pi d^2/2\delta\lambda$ . The refractive angle at which light is scattered away from such a ball lens may be approximated as  $\Delta\theta_{ref} \sim d/2f_{BEC} = \delta\lambda/\pi d$ .

The condition of interest, then, is that when  $\Delta\theta_{ref}$  exceeds the diffraction angle,  $\Delta\theta = \lambda/2d$ . This clearly occurs for lens-induced phase shifts,  $\delta$ , exceeding  $\pi/2$ . In this regime, the refractive scattering of light by the lens-like nature of the BEC will exceed the diffractive scattering due to the small size of the atom cloud. Any imaging system designed with the traditional “diffraction limited” criterion will experience degraded image quality due to the inability to collect all of the light refractively scattered by the BEC.

The conditions at which these limits apply depend on the size and density of the atom cloud. In the case of typical expansion-based absorption imaging, these two quantities are, in turn, related to each other through the expansion of the cloud. This depends on the time allowed for expansion between releasing the magnetic

trap, and pulsing the imaging beam to record the absorption image. The general form of the relationship is illustrated in equation 8.4.

$$n = n_{peak} e^{-\frac{x^2}{\sigma_x^2} - \frac{y^2}{\sigma_y^2} - \frac{z^2}{\sigma_z^2}} \quad (8.3)$$

$$\sigma_i = \sqrt{\frac{2k_B T}{m} \frac{1 + \omega_i^2 t^2}{\omega_i^2}} \quad (8.4)$$

For realistic experimental parameters, it becomes apparent that  $\delta$  approaches  $\pi/2$  only in the case of extremely high densities, typical of those achieved in BECs held within magnetic traps. Imaging such systems involves the refraction of a significant fraction of the probe beam out of the beam direction due to the action of the BEC as a lens [158].

In contrast, for expansion based absorption imaging, the action of the atom cloud as a lens is a small perturbation to the ordinary diffraction limit of the cloud. Given the behavior of such a cloud, described by equation 8.4, an imaging system designed to approach the diffraction limit of a trapped BEC will be able to resolve smaller and smaller fractions of the cloud as it expands.

### 8.1.2 Sensitivity

A firm grasp on the real resolution limitations of the imaging systems allows further investigation of the sensitivity requirements. The fundamental noise imposed by consideration of the interaction of the BEC with the imaging light limits the system far more stringently than more intuitively understood technical limitations. The key criterion for maximizing sensitivity in the imaging system is the maximization of the signal created by the BEC itself. This implies maximizing the number of photons absorbed by the BEC. The limit to this proposal is that the absorption of imaging light should not perturb the atom cloud so much that the image gets blurred. That is, the average momentum imparted to the atoms by the absorption and re-emission of photons should cause them to move less than the resolution limit appropriate to the measurement, as considered above.

The controllable parameters of the imaging procedure are the intensity, duration and detuning of the imaging pulse. The upper limit on useful intensity is determined by saturation of the atomic transition, yielding a photon absorption time  $\tau$ , related to the lifetime of the atomic state being probed. More light than this simply increases shot noise and saturates the CCD camera, degrading the signal. The length of the imaging pulse can then be determined by calculating the consequent motion of atoms and ensuring that they do not move further than the resolved distance of the imaging system.

The motion of an atom in a relatively weak resonant laser beam, such as that posed by a BEC absorption imaging system, represents a well-known momentum-space analog of thermal systems exhibiting Brownian motion. Stimulated absorption of photons from the laser beam constitute a replenished source of energy, providing unlimited opportunity for essentially random spontaneous emission events. Along

the direction of the laser beam, the atom picks up a distinct momentum increase,  $p_{atom}$ , due to the absorption of photons, yielding recoil momentum,  $p_{rec}$ , from the beam,  $p_{atom} = \sum p_{rec}$ . The atom dispels this momentum in a random, almost isotropic manner, leading to both a net momentum imbalance, accelerating the atom along the direction of the beam, coupled with a sort of random walk in momentum space.

The random momentum kicks transverse to the laser beam constitute a heating mechanism [159]. This is manifested as an increase in the average magnitude of the transverse velocity,  $v$ , expressed as a root-mean-square value in equation 8.5.

$$v_{rms}^2 = \frac{\alpha}{3} \left( \frac{p_{rec}}{m} \right)^2 \frac{t}{\tau} \quad (8.5)$$

This measure of the spread of transverse atom velocity arises due to an impulse equivalent to the recoil velocity of the atom,  $p_{rec}/m$ , over a time period,  $\tau$ , related, as above, to the lifetime of the atomic state being addressed by the resonant laser beam. The pre-factors represent a correction due to the dipole emission of the atom,  $\alpha$ , and a denominator due to the dimensionality over which the velocity is spread.

This random velocity spread quickly induces a random position spread, that can be quantified as  $x(t) = \sum \delta v_k(t - t_k)$ . Following Pritchard et. al., this may be related to the mean squared displacement, as shown in equation 8.6, by virtue of the independence of each emission event [160].

$$x_{rms}(t)^2 = \frac{v_{rec}^2}{3} \frac{t^3}{3\tau} = v_{rms}^2 \frac{t^2}{3} \quad (8.6)$$


Equation 8.6 provides a measure of the blurring effect of a resonant imaging pulse on the atom cloud. The most apparent point of observation is the strong dependence of the positional spread on the *time* over which the pulse is implemented. That is, within the limit of operating at the saturation intensity, the image can be blurred less by imaging over a shorter time, yielding a lower signal level. Conversely, for a given acceptable level of blurring, equation 8.6 relates the maximum length of pulse allowable, at which the imaging signal will be maximized.

Taking typical experimental parameters from the ANU-BEC imaging system: an atomic state lifetime of  $\tau \sim 27$  ns, photon recoil of  $v_{rec} \sim 5.9 \times 10^{-3} \text{ms}^{-1}$ , and a desired resolution of  $x_{rms} \sim 10 \mu\text{m}$ ; the optimum pulse length becomes just less than 100  $\mu\text{s}$ . This corresponds well with that observed in the development between ANU-BEC-I and ANU-BEC-II, as the imaging systems were upgraded to provide imaging pulses shorter than 1 ms.

### 8.1.3 Dynamic Range

These calculations of signal strength and resolution allow a much simpler, but just as important, calculation of the dynamic range requirements of typical absorption imaging experiments with BECs. Yet again, the realistic performance of well designed systems prove not to be limited by equipment, but by the fundamental nature of the measurement itself. In the case of dynamic range, it is the quantum





**8.2: Diagram of a CCD camera.** The left image shows a stylized rendering of a charge coupled device (CCD). The main components are the CCD array, consisting of  $10 \times 10$  pixels, a 10 pixel linear transfer buffer constructed of masked pixels, and a massive, masked readout buffer allowing on-chip binning. The right image shows detail of a single CCD pixel. The control electronics (surface components) are enlarged to show detail. The size of each pixel is inversely related to the pixel-count and directly related to the well-depth specification of the CCD chip.

noise, or shot noise, of the imaging process that provides the most stringent limit on the measurement.

A BEC containing  $10^5$  atoms will absorb somewhat less than  $10^9$  photons when illuminated as in the calculation above. This BEC, expanded to 1 mm square, contains  $10^4$  pixels of information. On average, each pixel will be illuminated with  $10^5$  photons, which will yield a shot noise of  $3 \times 10^2$  photons, or 0.3 %. This level of noise is very large. An 8-bit CCD camera is capable of resolving 256 “shades of grey”. This amounts to the ability to quantify intensity variations in an image with 0.35 % precision, which represents an almost perfect match for this imaging system. Most BEC experiments around the world over-specify the resolution of their CCD cameras to 12-bit, or even 16-bit, operation.

## 8.2 Choosing a CCD camera

CCD cameras remain a misunderstood aspect of BEC system design, even though complete explanations of their use, such as the one above, have been published [135]. Many experiments report use of significantly over-specified CCD cameras, that represent a loss of resources that could have been used elsewhere. This is possibly due to the bewildering array of choices faced when purchasing CCD cameras, and the rapid improvement of technology driven by mass-market consumer products in recent years.

Most specifications of a CCD camera are fundamentally linked together. Making the right choices therefore requires a somewhat holistic approach. The key specifications are: pixel count (ambiguously referred to as resolution), pixel size, well depth, and dynamic range. Briefly, CCD chips are usually rectangular, with sides up to 10 mm long, as illustrated in figure 8.2. Within this dimension, higher pixel count implies smaller pixels. The volume of the semiconductor material making up a pixel directly determines the well depth. The well depth, divided by the typical instrumental noise in operating the CCD chip, determines the dynamic range of the camera.

The analysis above leads further into this chicken-or-the-egg dilemma. It affirms that a CCD camera with a resolution of  $100 \times 100$  pixels and dynamic range of 8-bits is sufficient for imaging expanded BECs. Additionally, if pixel size is kept below about  $5\mu\text{m}$ , a simple unity gain imaging system can be used. At one time, few CCD chips of such specification existed. With the expansion of the market driven by consumer products, however, the proliferation of products has been substantial. CCD cameras geared towards ultra-fast capture generally use low pixel count (by modern standards), small pixel size chips, that fit the desired specification perfectly [161].

Another option worth considering is the use of a CCD chip with over-specified pixel-count, such as  $1000 \times 1000$  pixels. This is a valid option because this specification is currently the mainstay of the consumer market, and such chips are inordinately economical. As long as the other parameters of the chip are up to specification it can be used quite successfully, especially if on-chip region-of-interest is enabled, effectively allowing use of just one  $100 \times 100$  region of the CCD. The main difficulty with this approach is ensuring good response in the infrared. Consumers need CCDs that respond to visible wavelengths, and, more specifically, that do not respond to infrared. Such CCDs use an infrared filter to suppress unwanted wavelengths. These filters are either a coating on the chip itself, or a coating on a glass window in front of the chip. Glass window mounted filters can usually be removed without damaging the camera to make it useful for BEC experiments.

Many CCD cameras are cooled, because thermal noise can contribute significantly to signals that are integrated over long times. This situation does not arise with BEC imaging, where images are exposed for hundreds of microseconds. The dark current is not measurable on these time scales, and does not degrade the signal. Cooled CCD cameras are complicated, expensive and add nothing to the performance of a BEC imaging system.

One specification that should be understood is the capture architecture of the CCD chip. Most chips use either full frame or interlaced architectures, and some advanced, or custom-made chips offer frame transfer architecture. These architectures determine the “electronic shutter” performance of the CCD chip. Electronic shuttering refers to the process of: removing all charge from all image pixels; integrating charge on all image pixels for an exposure time; reading information from all image pixels. If the first and third steps could be performed instantaneously CCD cameras would never require a mechanical shutter. Step one can be performed very quickly, because translating charge across pixel wells to the read-out well is fast: typically

50 ns/pixel, so within 25 ms for a 500k pixel CCD chip. Step three, however, is limited by the process of analog-to-digital conversion at the read-out well, which takes a lot longer, typically  $5 \mu\text{s}$  / pixel, taking 2.5 s for the same CCD.

An image nominally integrated for 1 ms using the electronic shutter will indeed expose the cells closest to the read-out well for about 1 ms. Those on the other side of the CCD, however, are exposed for up to 2.501 s, while being shifted across the chip as other pixels are read out, leading to a blurred and overexposed mess instead of an image. Interlaced and frame transfer chips provide CCD pixels covered by an opaque mask, that the image on the unmasked pixels can be shifted to very quickly, after which it can be read-out at leisure without over-exposing or blurring the original image.

This is not a useful feature for BEC absorption imaging, because a very good “electronic shutter” is provided by excellent control over the laser light. Typical optimum pulse lengths calculated above were less than 1 ms, after which there is no light to obscure the image on the CCD in any case. The real use of frame-transfer architecture is for taking multiple images in quick succession with more advanced imaging techniques like phase-contrast [158]. Interlaced CCDs are never a good option, as they cover half the active chip area with masked pixels, reducing the efficiency of image information collection.

Some CCDs are “backlit”, meaning that they are constructed traditionally on a semiconductor substrate, which is then etched away extremely close to the beginning of the back of the CCD sensor. This is useful, because CCDs are created in circuit layers, starting with the actual collection wells on the bottom layer, as shown in figure 8.2. The driving circuitry on layers above the collection wells is typically arranged to obscure pixels as little as possible to light of visible wavelength. Longer wavelength infrared light can experience decreased transmission through these structures, reducing the efficiency of detection. Backlit CCD cameras allow infrared light to access the collection wells almost directly via transmission through the ultra-thin etched substrate. Unfortunately, the well defined crystal structure of the substrate leads to beautifully flat and parallel surfaces which can act as an etalon for laser light.

### 8.2.1 The imaging system in the ANU-BEC-I experiment

The imaging system on the ANU-BEC-I machine used a single lens in a 2-f configuration, yielding unity magnification. It was arranged in a fashion similar to that in the stylized representation in figure 8.1. This somewhat odd orientation was necessary due to the poor optical access provided by the original magnetic trap. The main improvisation was the need to image across a diagonal axis of the magnetic trap by using “flipper” mirrors that rotated into the MOT beam path after the atoms were magnetically trapped. The consequences were difficulty in optimizing optical alignment, placing collection lenses near the experimental chamber, and in analyzing image data.

The CCD camera used in the imaging system was a Princeton Instruments RTE-768-K, with a thermoelectrically cooled CCD exhibiting a quantum efficiency

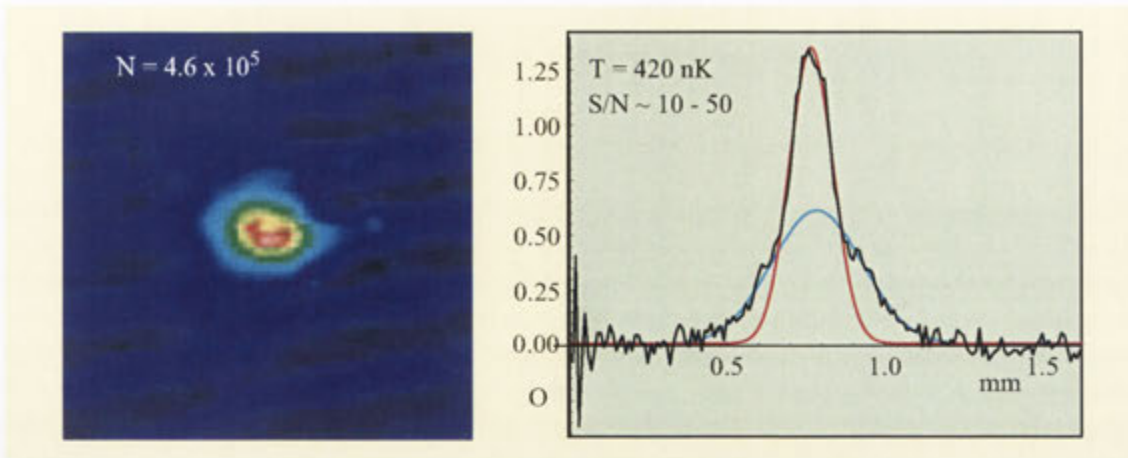
of roughly 30 % at 780 nm, containing  $756 \times 512$ ,  $9 \mu\text{m}$  square pixels. Each pixel had a well depth of  $\sim 80,000 e^-$ , and the read noise of the camera electronics was approximately  $12 e^-$ , which, coupled with the 12-bit analogue to digital convertor, yielded true 12 bit information. It was a legacy unit from a previous experiment. It worked quite well, and was used to detect the first BEC. However, it contained many unnecessary features and was not optimized for this experiment.

The diffraction limited resolution of this imaging system was limited by the size and position of the mechanical flipper mirror required to maneuver the imaging system into place after optical cooling was complete. The 25 mm mirror was placed approximately 25 cm from the location of the BEC, yielding a distance resolution limit of  $10\lambda \sim 8 \mu\text{m}$ . This was closely matched by the size of the CCD pixels  $\sim 9 \mu\text{m}$ . As the atom clouds underwent significant expansion prior to absorption imaging, the resolution of the system would not have been limited by the lensing effect of the BEC considered above.

A typical experimental run expanded a cold atom cloud or BEC, containing  $10^5$  to  $10^7$  atoms of temperatures from 500 nK to 100  $\mu\text{K}$ , for the order of 10 ms. In such a system, a BEC expanded for 10 ms extends approximately  $100 \mu\text{m}$ , yielding  $10 \times 10$  pixels of information. The thermal cloud can exhibit significantly larger extensions. The  $10^5$  atoms of the BEC, spread across  $10^2$  pixels, yields approximately  $10^3$  atoms per pixel. Using the analysis above, a  $50 \mu\text{s}$  pulse of resonant light generates of the order of  $10^3$  photons per atom, or a photon count per pixel of approximately  $10^6$ . However, in this case, the ability of imaging process to collect the signal is limited by the quantum efficiency and well depth of the CCD to approximately  $2 \times 10^4$  photons per pixel. The imaging process could be rescaled to produce signals of this level by changing the expansion time or the length of the imaging pulse. However, in practice, it was the detuning of the laser from resonance, used to reduce the rate of photon scattering per atom, that reduced the signals to prevent saturation of the CCD. In such a system, yielding approximately  $8 \times 10^4 e^-$  information per pixel, the shot noise is of the order of  $3 \times 10^2$ , or,  $\sim 0.4 \%$  of the total signal. The 12-bit dynamic range of the CCD camera ( $\sim 0.025 \%$ ) could accurately reproduce these signals without adding significant instrumental noise.

It may appear that, in some sense, this analysis is back-to-front, as an absorption image records the least amount of light in the regions of highest atomic density. However, the entire absorption imaging process references the shadow of the cloud cast in the main data image to a second image collected without atoms present. It is the difference between these two images, or, more correctly, the logarithm of the ratio of these two images, that constitutes the data analyzed above, so the regions of greatest atomic density do create the regions of highest signal strength. It is also true that the averaging applied in this quantification of performance underestimates the quantum-noise limited dynamic range at the center of the cloud.

Theoretically, this imaging process could be improved by avoiding saturating the signal pixel by either increasing the expansion time, or, perhaps, increasing the well-capacity of the CCD itself. Increasing the expansion time would decrease the number of atoms imaged by each pixel, so that the number of photons necessary to fully elucidate the imaging process would be reduced. It also increases the size of the



**8.3: Data collected by the ANU-BEC-I imaging system.** The left image shows an image collected using the imaging configuration discussed in the text. The right image shows a cross-section extracted from the data in the image. The major noise sources apparent in the image are dominated by technical effects, such as the lines apparent in the data due to diffraction rings on the imaging laser light. These effects are largely compensated by the imaging processing algorithm, that compares all data images with reference images, but some residual noise remains. On the cross-sectional data, the observed signal-to-noise lies between  $\sim 10 - 50$  across the peak. This is significantly greater than the fundamental limits discussed throughout this chapter.

imaged cloud, *increasing* the number of pixels of information about the cold atoms or BEC. In practice, this was not an option for the ANU-BEC-I experiment, because the expansion times were severely limited by the diagonal imaging geometry. After more than 20 ms, the atom cloud had fallen far enough, under the influence of gravity, to move beyond the field of view of the imaging system.

The opposite resolution behavior is observed when increasing the dynamic range of the CCD camera, by, for instance, binning  $3 \times 3$  pixels across the CCD to increase the well depth by a factor of 9. This *decreases* the number of pixels of information collected from the cloud, in some sense degrading the quality of the imaging system. This is true, also, of true 16-bit CCD camera systems, exhibiting larger well depths, because these improvements are achieved by simply increasing the volume of the pixel material, with a corresponding increase in pixel size. A practical solution for both configurations, however, is to increase the magnification using a more advanced imaging system, which overcomes the above limitation as long as the image does not become larger than the physical size of the CCD.

Many of these considerations are overwhelmed by the practical realities of imaging in the ANU-BEC-I experiment. Figure 8.3 shows both processed image data, and a cross-sectional plot generated from the image data, collected from the ANU-BEC-I experiment. Clearly apparent in the data is technical noise, such as bands,

especially visible in the background, due to diffraction rings on the laser light.

The cross-sectional plot to the right illustrates the effect of such noise sources in the ANU-BEC-I imaging system. Specifically, the signal to noise ratio varies between less than 10 near the wings of the cloud, up to approximately 50 near the peak of the cloud. These values constitute a significant degradation of image quality below the 8-bit signal to noise ratio (256), discussed above. The ANU-BEC-II imaging system, although assembled from components sporting similar specifications, improved the performance with respect to the reduction of technical noise sources in the image data, thereby approaching the fundamental limits discussed throughout this chapter.

### 8.3 Improvements to the imaging system for ANU-BEC-II

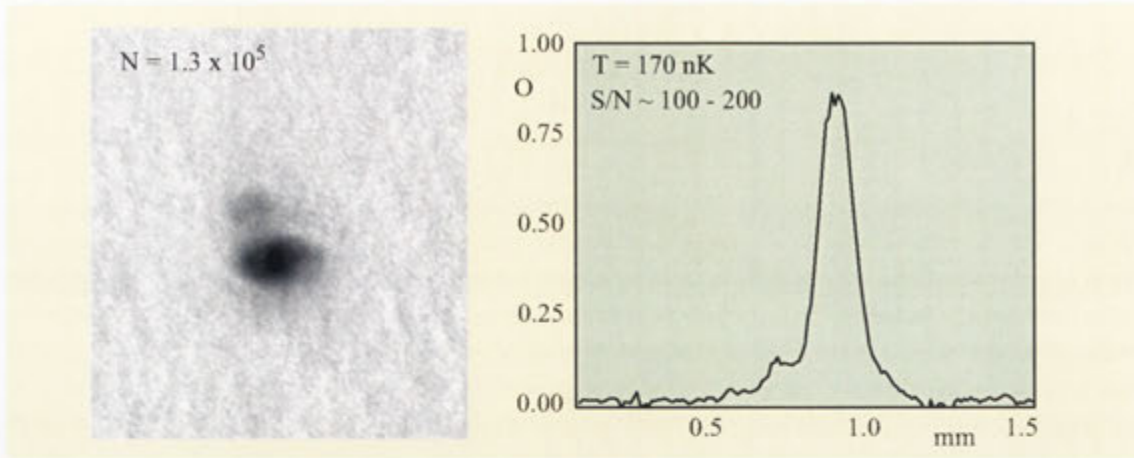
The ANU-BEC-II system received two key improvements to the optical system: NR redesigned the magnetic trap to provide better optical access, and a new CCD camera was purchased. The absorption imaging process remained fundamentally the same. However, the redesigned trap allowed easier alignment, much better placement of collection lenses, and imaging along trap axes to simplify analysis. It also allowed longer expansion times before the atom cloud fell out of the field of view of the imaging system.

The purchase of the new CCD camera involved some detailed consideration of CCD requirements, as discussed above. The CCD camera eventually purchased was a Photometrics SenSys, and represented a compromise between the understanding garnered by these considerations, and the confidence inspired by incrementally improving the previous imaging system to guarantee a minimum level of performance. In fact, it employed the same CCD chip, a Kodak KAF0402E, as the Princeton Instruments camera from the original BEC machine. It implemented a less stringent thermoelectric cooling scheme, simply temperature stabilizing the camera to 10 °C, which significantly reduced the cost. In addition, it provided a more modern form-factor and software interface.

The improved geometry of the imaging system allowed the diffraction-limited resolution to be improved, by closer positioning of a larger lens, to the vacuum chamber. In this configuration, a 50 mm diameter lens was placed 100 mm from the BEC, yielding 5  $\mu\text{m}$  resolution. This increased resolution was integrated with the CCD camera specifications using a magnification factor of 1.75, again matching the 9  $\mu\text{m}$  size of the pixels fairly well.

As the specifications of the CCD chip were identical, the analysis of the imaging process from the ANU-BEC-I machine applies just as well to the ANU-BEC-II imaging system. The key differences between the two systems is the use of greater than unity magnification in the ANU-BEC-II experiment and the ability to expand for longer times, both of which relax the tendency of the systems to saturate the well depth of the CCD pixels.

The major improvements associated with the upgrade to the ANU-BEC-II imaging system were consequent of a reduction in the amount of technical noise through-



**8.4: Data collected by the ANU-BEC-II imaging system.** The left image shows data collected using the imaging configuration on the ANU-BEC-II experiment. The plot on the right shows a cross-section extracted from this image data. The key improvement to the ANU-BEC-II imaging system was the reduction of technical noise sources, such as diffraction rings on the laser light, that limited the performance of the system as a whole. These improvements are manifested in the improved signal to noise ( $\sim 100$ ) apparent in the cross-sectional data.

out the system. Major progress towards this goal included minimization of diffraction rings on the imaging light generated by the imaging optics, the improved qualitative noise performance of the Photometrics SenSys CCD camera, and the significant effort expended upgrading the image-analysis algorithm to almost completely cancel common-mode noise between images.

The improvement to the imaging data is apparent in figure 8.4. Specifically, almost no evidence of distortion, due to diffraction or other technical effects, exist in the background areas of the image. This translates to a significantly improved signal to noise ratio on the cross-sectional data in the plot at the right of figure 8.4. The observed signal to noise ( $\sim 100 - 200$ ) approaches the fundamentally-limited signal to noise,  $\sim 256$ , discussed above. This provides great confidence that many of the technical noise sources throughout the ANU-BEC-II imaging system have been identified and controlled to the level necessary to achieve optimum imaging in this experiment.

## 8.4 Future improvements to the imaging system

The success of ANU-BEC-II has led to great confidence in the understanding of the requirements of the imaging system. The camera purchased for the ANU-BEC-III machine is (literally) an order of magnitude cheaper, and significantly smaller than either the ANU-BEC-I or ANU-BEC-II cameras. This improvement

was realized using the analysis performed above, yielding a significant relaxation in the required specifications, and the development of technology driven by consumer digital cameras.

The camera is a Micropix M1024, and employs a Sony ICX-204AK CCD sensor. This is a mass-produced sensor, with an array of  $1024 \times 768$  pixels. Although this pixel count far exceeds the requirements for a BEC imaging system, it is a commercially popular specification, and is therefore inordinately economic. In addition, the CCD features very small pixels,  $4.65 \mu\text{m}$  square. It does not implement any cooling or temperature control, and so is extremely small ( $7 \times 6 \times 4$  cm). It has a signal to noise of 58 dB, coupled with a 10-bit A/D convertor to yield approximately 9-bit information. It communicates over the PC-standard IEEE-1394 (firewire) bus, dispensing with custom electronic interfaces and further reducing the cost. It was rendered suitable for imaging at 780 nm by removing an infrared filter window in the imaging path.

This camera has not yet been implemented for imaging the ANU-BEC, so it is not possible to present field-proven results of its performance in such a system. However, the analysis presented in this chapter suggests that, coupled with a compatible imaging system, it should perform sufficiently well that the entire system is limited by fundamental noise sources. Several of the larger groups worldwide have reported success with similar units.

## 8.5 Future Plans

The above considerations apply specifically to absorption imaging of BECs, in current-generation machines. More advanced imaging systems, or bigger BECs, require their own analysis to determine the optimum configuration, but the same principles will apply.

One of the most anticipated advances in imaging systems is the ability to take true continuous, non-destructive images during the formation and manipulation of BECs. This has the potential to elucidate the dynamics of BEC experiments, and reduce the dependence on shot-to-shot stability in current machines.

The first non-destructive technique demonstrated on a BEC was phase contrast imaging by Ketterle et. al. [158]. They used the less destructive nature of phase contrast imaging, coupled with a frame-transfer architecture CCD camera, to take ten pictures of a cold cloud of atoms crossing the transition to BEC.

An advent that may improve the performance of such a system is the increased speed of CCD cameras. Taking typical trap frequencies as a characteristic timescale for BEC dynamics, cameras of the resolution and sensitivity derived above, that can capture images at  $\sim 1000$  frames per second, would provide an excellent new diagnostic. These cameras already exist, and await implementation [161]. What's more, some of these designs can transfer data out of the camera at this speed, opening the door to real-time dynamic feedback and control of BECs and atom lasers [11].

Even faster dynamics may be expected between atoms in the BEC, on timescales



determined by the natural linewidth of the optical transitions. These dynamics lie beyond the domain of CCD cameras at present, but can be captured using photodiodes. The increased dynamic response of the imaging system comes at the cost of 2D density information about the BEC, but does allow investigation of total number variations.

Finally, ideas for imaging systems that beat the quantum limit in either spatial resolution or sensitivity are in development. The ANU-BEC group has the advantage of excellent access to experts in generating, manipulating and detecting quantum states of light that may allow sub-quantum limited imaging in the future.

## 9 The control system

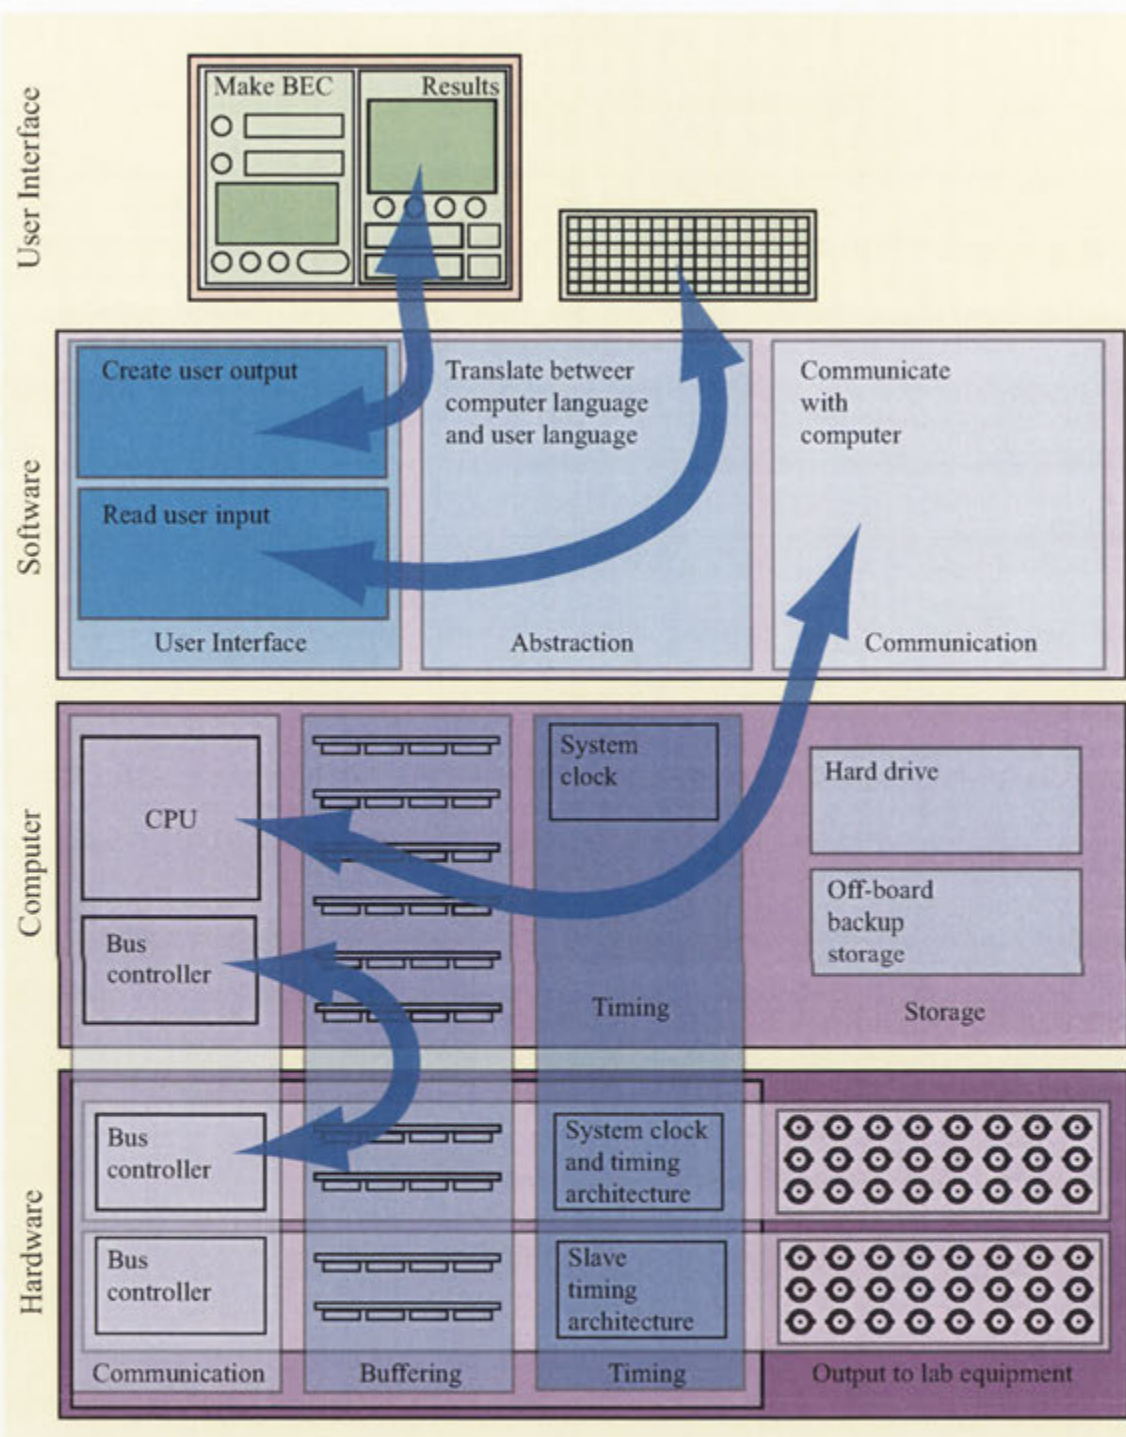
The computer control and associated electronics systems are an oft-neglected component in reports and theses summarizing BEC machine design. A well designed control system is one of the most important, and most enduring, legacies left for future machine operators, because essentially all user interaction with a running BEC machine is through the computer control user interface. The ability to quickly and easily alter pertinent aspects of the experiment, while maintaining an intuitive feel for each section of the the experimental run, palpably effects the speed and quality of the results that can be generated on the machine. A significant effort was expended designing the control system for the ANU-BEC, and its long term success justifies a substantial description.

### 9.0.1 The control system in BEC experiments

The detailed diagram in figure 9.1 illustrates the important parts of a computer control system. It is presented as the image of a specific system to aid the layout, but the descriptions of the components listed are quite general. The “hardware” is the apparatus that interfaces the computer control to the standard laboratory electronics that controls equipment such as lasers, modulators, power supplies, switches and shutters. The computer is the mediator that enables communication between the user and the hardware. It runs the software and has some direct connection to the hardware. The software interprets users’ commands and translates them into the format that the hardware requires for optimum operation. The user interface is the window through which the user enters commands and receives information about the experiment.

## 9.1 Hardware

The job of the “hardware” in figure 9.1 is to translate computer commands to a format that laboratory equipment can understand. This requires communicating in the correct “language”, and across the correct physical connections, to each individual piece of equipment in the BEC experiment. The menagerie of apparatus that comprises a typical BEC machine can require implementing many communication standards simultaneously. In addition, many processes, such as pulsing imaging light, require fast, accurate timing over time scales as short as microseconds. Combining the hardware to communicate at these rates, in many different languages, requires careful consideration of communications protocols, timing, and synchronization.



**9.1: Block diagram of the control software.** The four main components of the control system are illustrated: the hardware, the computer, the software and the user interface.

The specific method of interfacing laboratory equipment to a personal computer is usually determined by the connection methods supported on current laboratory equipment. However, if too many communication standards are implemented simultaneously, the performance of the experiment can suffer as a result of poor synchronization and timing between standards. Long term optimum operation of a complicated experiment like a machine to produce BEC requires, and deserves, consideration of these problems at the design and purchasing stage of the experiment.

Communication between hardware and laboratory equipment can be implemented in many ways. For instance, laboratory equipment increasingly supports some selection of common digital communication protocols, such as USB or ethernet, and can be connected, like any peripheral, to standard computer communication ports. The biggest problem with these standards is low control speed and lack of rigorous synchronization between equipment. At present, the best way to generate precise, high speed control of an entire BEC experiment is using hardware that directly controls equipment parameters via analog or digital voltage control ports. This is achieved using data acquisition (DAQ) expansion cards containing digital-to-analog converters (DACs) and digital line drivers, mounted in the PCI slot, or equivalent, of a personal computer.

The other important aspect of interfacing hardware to laboratory equipment is the physical connection. DAQ cards usually output control signals to a standard format high density connector (such as 68-pin SCSI-II) at the back of the PC. Most equipment in a physics laboratory expects signals delivered on co-axial cable fitted with BNC connectors. Signals are converted from high density computer connectors to a multitude of BNC connectors using a “break-out box”. If a computer system is being established for reliable long term operation of a complicated experiment, the break-out box needs to be extremely rugged and dependable.

One aspect of accurate timing is synchronization between various control hardware. This is very difficult to achieve, even at microsecond precision, unless all the DAQ cards run off a common clock. This vital feature is implemented by vendors selling high speed DAQ cards, such as National Instruments’ (NI) real time system integration (RTSI) bus. Even better is the PXI architecture for advanced DAQ systems. This extends the concept of the RTSI bus using rigorous hardware geometry standards to ensure precise synchronization of the communication of all signals to all DAQ cards, yielding jitter across various cards lower than 1 ns.

The feature of high speed DAQ cards guaranteeing that communication with the host computer will be seamless is hardware-side “buffered” data transfer. Buffering involves transmitting data to a DAQ card faster than is required, continuously replenishing the top of a buffer memory stack. Data being output from the card is read from the bottom of the stack. Even if communication between the computer and the DAQ card is interrupted, data continues to drain from the buffer stack. As long as the entire stack is not depleted before communication resumes, output from the DAQ card is smooth and continuous. Modern high-speed DAQ cards, such as those produced by National Instruments, use this feature to guarantee performance with normal operating systems [162].

Several methods of communication exist between the computer and DAQ systems. The most common is the PCI bus protocol for personal computers. In this system, DAQ cards are made on circuit boards that are fitted into a host computer, drawing power from its power supply, and communicating directly with the rest of the computer through the computer's own internal PCI bus. Other computer buses for interfacing with hardware include the now-redundant ISA bus, and the PXI bus explicitly designed to extend the PCI bus for data acquisition use. DAQ systems for use with standard high-speed digital communication protocols recently incorporated in personal computers, such as USB and ethernet, are becoming more common due to low cost and flexibility. The most important parameters of the communication method for high-speed DAQ systems are the bandwidth and support for advanced communication methods such as the direct memory access (DMA) channel on the PCI bus. However, the modern communication protocols should all be fast enough when used for *controlling* a BEC experiment with a well-design computer control system.

The best control of experimental apparatus by a computer is achieved using analog voltage control. This technique provides a very uniform, general control method for interfacing to various equipment, and can achieve fairly high data speeds. Commercial offerings include analog output cards containing eight 16-bit dynamic range (0.02 % accuracy) channels, simultaneously updated at up to 1 MHz. Weaknesses of direct analog voltage control include its reliance on relatively complicated and expensive interface cards, and the potential for degradation of the voltage signal during generation and transmission to effect the accuracy and precision expected from the controlled instrument. This second problem can make this type of system prone to difficulties arising from noisy cable runs or RF interference from other apparatus.

Digital control performs better in this regard. However, it increases the onus placed on laboratory equipment in terms of communicating with the computer. With multiple popular standards, the divergence of communication protocols can make integrating the entire control system difficult. The simplicity and ruggedness of digital communication does allow extremely high speeds for data transfer. Unfortunately, in traditional analog-based equipment, translation between the digital communication signal and the analog control signal at the piece of apparatus usually reduce the speed benefits of this technique to below that of analog control. The main benefit of a simple digital control method is for apparatus that performs a fundamentally digital function, such as shutters, switches or flipper mirror mounts that are either open or closed.

One digital communication protocol that is almost ubiquitous in laboratory equipment is IEEE-488, otherwise known as the general purpose information bus (GPIB). This protocol uses an ASCII text string programming language to enable general commands so that almost any scientific instrument can be controlled. It does require a dedicated controller, but this has become the standard among equipment vendors so that almost all modern equipment can be fitted with a GPIB controller. The main problem of such a general digital communication protocol is the overhead: all commands are converted to relatively long ASCII strings, which create a large

digital transfer, only to be retranslated to analog voltages or digital code at the receiving apparatus. Coupled with low speed transmission (due to the age of the GPIB standard) this overhead precludes high speed real time control.

## 9.2 Computer

The job of the computer is to mediate communication between the software and the hardware. Increasingly, with modern systems, it achieves this in an almost seamless manner by heavily abstracting the user from any specific interactions. This benefit can be amplified by carefully selecting compatible equipment at the planning stage. A prime example of this is complete data acquisition solutions, including hardware, computer and software, provided by vendors such as National Instruments.

The explicit choice of hardware will determine the protocol employed to communicate with the hardware. However, a good choice of computer system will so completely abstract the user from this process that they need hardly be aware what protocol is actually being used. This process is implemented by the concepts of “plug-and-play” hardware, “plug-and-play” operating systems, and elegant driver design by hardware vendors.

The only downside to the development (and improvement) of these concepts over the past five years, is the difficulty encountered in penetrating this abstraction if low-level changes need to be made. This generally requires changes at the lowest level of computer control: the BIOS. The extent of low-level control available depends on the choice of motherboard in the computer. Good motherboards for computer control systems allow all optional on-board equipment, such as sound cards, ethernet and even video cards, to be disabled, releasing I/O resources back to the computer. They also allow manual configuration of PCI parameters, in the case that there are equipment conflicts that the plug-and-play systems do not resolve. As plug-and-play matures, and these problems become less common, many of these controls are disappearing from BIOS settings on commonly available motherboards.

The most powerful solution, however, is one that provides full control of all aspects of the computer system. This is (approximately) the case with computers running real-time operating systems, such as NI Labview RT, or Realtime Linux. National Instruments offers computers designed explicitly for computer control systems, designed to work with Labview RT to give the user unprecedented control. With increased power comes increased responsibility, and a potentially more complicated programming structure. The benefit is performance: well designed real-time computers and operating systems currently represent the pinnacle of data acquisition performance.

Another important part of communicating with hardware is common to both sides of the conversation: buffering. The computer generally buffers information that is to be sent to the DAQ cards in system memory to ensure continuity of transmission if the CPU is needed for some other action. The benefit of this buffer over the hardware buffer is the scale. Computer memory is abundant and cheap, and

in extreme cases of communication breakdown, enough can be provided to buffer the control sequence of an entire experiment.

A much more elegant, and scaleable, solution is double circular buffering. This uses modest buffers on both the computer and the hardware. The “top” of the computer memory buffer is continuously replenished with the hardware-level code generated by the CPU from the software. It smoothes any CPU interruptions. The “top” of the hardware memory buffer is continuously replenished from the “bottom” of the computer memory buffer. It smoothes any interruptions in communication. The benefit of this scheme is that as long as the CPU can, on average, supply the computer memory buffer fast enough, the experiment can continue indefinitely, even infinitely. To do this with a non-circular computer memory buffer would require infinite memory.


The computer also provides the framework for recording information about the experiment. This information is stored on the hard drive of the computer. Even vaguely efficient programming should find modern hard drives sufficient for data storage. The biggest use of storage is image data collected from the experiment using CCD cameras. This is discussed further in chapter 8. Even by conservative estimates, a large modern hard drive should provide storage for  $10^5$  images, which should provide years of storage on the most active machine. This data should be backed up regularly.

The other side of the mediation process provided by the computer is communication with software. Again, this is a matter of abstraction. High level programming languages will abstract the user from the computer just as the computer and operating system abstract the hardware from the software. The only key point is ensuring good compatibility between the hardware, operating system and programming language.

## 9.3 Software

The main choices involved in software development are programming language and programming paradigm. Traditional programming languages have the benefit of familiarity for some users. High level graphical programming languages, such as Labview, produced by National Instruments, offer rapid pick-up and an intuitive interface for experimentalists, as illustrated in figure 9.2. The benefit of excellent driver support and hardware abstraction in Labview makes it a common choice for computer control systems.

Two programming paradigms of interest in BEC machine design are top-down and bottom-up programming. Top-down programming emphasizes the user interface (UI). The first step of coding is designing the UI to provide all the controls and displays needed by the user. The program is built down from the UI to provide the functionality required by the user. This is the paradigm that the Labview programming language is built around, and that National Instruments espouses when designing “virtual instruments”. Bottom-up programming begins with the code that is needed to interface with the hardware in the system. For simple programs,



**9.2: The main code window of the computer program.** Coding Labview programs involves wiring together “virtual instruments”, a model that experimentalist supposedly find intuitive. The inset shows a frame of the sub-routine that converts the short list of user-readable changes, generated by the interface shown in figure 9.3, to a large array of machine readable commands for every time step in the experiment.



Labview's excellent hardware abstraction makes this type of programming somewhat redundant. However, as an experiment gets more complicated, and low-level control of hardware is unavoidable, this technique gains merit.

An important job performed by the program written to control the machine is the preparation of the sum of all user commands into the format required by the hardware. The complexity of this process can be simplified when the program is written in a high level language with good organizational subroutine support. The extent to which this programming needs to be implemented depends on the quality and abstractedness of the user interface.

High speed analog and digital control cards are usually configured to receive a data point for each channel, at each time step during an experimental run. Even if the value on a given channel at a given time does not change, the card must receive the unchanged parameter. In an experiment that lasts 100 seconds, with 30 channels of control, operated at  $10\mu s$  resolution, a  $30 \times 10^7$  array must be gradually fed on to the control cards. On the other hand, the number of changes that must be implemented in during such a BEC experiment may total only 200 in entirety. It is inadvisable to generate and manipulate the entire output array directly from the user interface, so some subroutine must be executed at run time to generate the large output array from the much smaller array of changes.

This generally involves many steps of table-sorting and scaling from human-readable formatted values (ie. physical units) to hardware-readable formatted values (ie. 0 to 10V for analog, 0 or 1 for digital). Depending on the level of abstraction of the user interface, more advanced procedures can also be necessary, such as interpretation of strings, and interpretation of events defined by conditional and relational commands. A high level programming language is typically supplied with a large library of subroutines for carrying out relatively sophisticated tasks of this sort, and can greatly simplify the program.

## 9.4 The user interface

The user interface, shown in figure 9.3, is an aspect of the software running the experiment that is disproportionately important. Designing a control system to be intuitive is largely a matter of careful user interface design. This can be augmented using controls and displays that users are already familiar with: physicists respond well to controls that mimic typical hardware front panels. The commercial Labview programming language, produced by National Instruments, can do everything from design the user interface to access the lowest levels of hardware control. This graphical language is designed to achieve exactly the mix of powerful computer control and familiar, intuitive interface that makes setting up complicated control systems easy.



## 9.5 The ANU-BEC-I control system

The design of the ANU-BEC-I control system, like much of the experiment, depended heavily on equipment that was already in the laboratory. However, the control system probably bore the brunt of this situation, as it was the mediator through which these varied apparatus communicated. This entailed less freedom to optimize the design of the system, and much greater effort making it interface seamlessly, than may have otherwise been desired.

The ANU-BEC-I used “expansion cards” mounted in the ISA expansion slots of a personal computer. It used a single 6 channel, 12-bit analog output card, the National Instruments AT-AO-6. This was later upgraded to the 10 channel version of the card, the AT-AO-10, with minimal effort. It also used a digital output card, the NI DIO-32F. Finally, the initial design planned to use a NI GPIB interface card, but it was never implemented in the system because of the lack of a rigorous timing standard for this communication protocol.

Despite their age, the NI AT-AO-6/10 and DIO-32F cards did support the NI real-time system integration (RTSI) bus. Via this bus, the AT-AO-6/10 actually ran off the clock mounted on the DIO-32F card. This provided synchronization between the outputs of the cards of better than  $5 \mu\text{s}$ , which was easily sufficient for operating the BEC machine. The AT-AO-6/10 and DIO-32F cards also supported buffer memory stacks sufficient for preventing underrun errors in typical operation.

The ANU-BEC-I experiment used home made break-out boxes that employed a manufacturer-supplied screw-terminal connector attached via a complicated and fragile wiring harness to approximately 40 BNC connectors mounted in a custom-made enclosure. This worked well, but was a constant source of (usually) unfounded suspicion when something in the control system malfunctioned.

The computer that housed the control system was largely determined by the requirement of interfacing with the hardware cards. As they were legacy items, few modern (in 1999) computers were available that would accept them. The other requirement was that the computer have a good level of control over resource assignment at the BIOS level. This was important, because many systems were not rigorously plug-and-play compliant. Extra motherboard features, such as sound-cards, caused inconsistent operation of the computer system unless they were disabled in the BIOS. The computer was also fitted with 512 megabytes of RAM, after trouble with the low-level Labview subroutines prevented the use of double circular buffering (see below).

The programming language used to run the experiment was Labview 3. Like many of the other components of the control system, this was inherited and obsolete. This did not overly effect the main strengths of the program: quick up-take, easy to understand graphical programming style, and excellent user interface construction methods; but it did lead to difficulties in achieving optimum performance.

Labview did provide excellent abstraction from hardware. This quality was amplified because National Instruments hardware was used to provide DAQ capabilities. This feature allowed the same intuitive commands to be used to write values to different output channels independent of the actual model of hardware attached.

This meant that swapping the AT-AO-6 card for an AT-AO-10 card 18 months after the program was first implemented required almost no reprogramming. This abstraction, and the very high level programming structures provided by Labview, lead to an structurally simple and intuitive program.

All was not well, however. Great difficulty was encountered actually implementing the program, simply because the high level subroutines supplied by Labview did not perform as their documentation indicated. Writing a large, complicated program with a significant fraction of subroutines performing essentially randomly was a daunting and disheartening process. In the final stages it literally involved two people writing countless “test programs” for problematic subroutines and entering random parameters to understand how they should work.

The worst problem encountered was the failure of the continuous buffering subroutines. This issue was never resolved. Careful analysis of the problem showed that the subroutines performed properly, but forced all outputs on the analog output board to output a full range 10 ns pulse every 100 ms. This pulse was enough to trigger responses in some equipment, and although low-pass filtering reduced the problem, an unpredictable program was not deemed acceptable for long term use.

The fast but inferior solution to this problem was to create a massive buffer in computer memory that held the data for an entire run prior to execution. This ensured that the computer could send data to the card quickly enough, but it required enough computer memory to hold an experimental run. The vagaries of the process consumed a huge amount of resources, requiring the computer memory to be upgraded to 512 megabytes. This was a very large (and quite expensive) amount of memory for the time, but worse, limited the experiment to 1 ms precision.

The ANU-BEC-I control program favored the bottom-up coding paradigm, although much work was put into the user interface later in the process. This decision was based on the need for advanced control over the hardware required for this project. Also, as most “traditional” programming languages support bottom-up programming, it was a familiar, quickly implemented model. Improvements in both hardware and software in the last few years mean that many of these considerations are less relevant today. A well designed top-down program built in modern, robust Labview code, interacting with computers and DAQ cards many times faster than those used in ANU-BEC, should produce an excellent system.

The benefit of the bottom-up approach was that it used a minimum of manipulation of data between the user and the DAQ cards, as illustrated in figures 9.2 and 9.3. For instance, turning off the lower MOT shutter, attached to digital channel 12, at a time 57.012 seconds after the program began involved adding an “action” for “lower MOT shutter”, at time “57 seconds” and “12 milliseconds”. The user was only abstracted from the hardware by the fact that channel twelve was assigned the label “lower MOT shutter” in a configuration file. They could enter any an “action” for any channel at any millisecond during an experimental run. This made the entire program very flexible, as changing the purpose of channel 12 when the experiment was altered (making the lower MOT shutter redundant), was as simple as changing the label in the configuration file. The rest of the program consisted of “housekeeping” functions available to the user to simplify common tasks, such as

ramping voltages on analog channels.

The other benefit of this style of program was that there was essentially only one task to learn to use it to its full potential: the “add action” task (and its antithesis “delete action”). After mastering this one command any user could create every single permutation of control sequences by simply repeating it at will.

The problem with the control program was the absolute power and information it gave the user. Not only are most permutations of control sequences not useful, many would be harmful (such as running full current through uncooled magnetic coils). A machine to be used by a wide range of people with limited experience of the control program should probably have safety tests built in where appropriate (ideally to both software, and the experimental apparatus).

The provision of absolute information was problematic as well. Any fully experimental BEC machine does require some indication of the control sequences at the lowest possible level. This is what the ANU-BEC control program provided: a chronological list of every single change that occurred on every single channel throughout the experimental run. It did have some simple tools for improving the display of this information, but it was no means ideal. A better display of information about the machine configuration would make running the experiment more efficient.

## 9.6 Improvements to the control system during ANU-BEC-II

The ANU-BEC control program was a mitigated success. Testament to its quality and flexibility is the fact that it was one of very few parts of the ANU-BEC-I machine to remain in its original state during the ANU-BEC-II reincarnation of the machine. The program has proven extremely easy to use, in the superficial sense. However, building up complicated experiments did test users ability to develop an intuitive feeling for long lists of information.

The main limitation of the control system was the expected one: the failure of National Instruments hardware and software to perform as documented. This limited the precision of the system to 1 millisecond increments. This proved sufficient for producing BEC by augmenting the computer system with hardware pulse generators when more precise time increments were required. However, as the experiments performed with ANU-BEC-II have become more advanced this factor has begun to limit the usefulness of the program.

The biggest trade-off in the program turned out to be building the flexibility in at the user interface level. This was a result of the bottom-up design process, which created programming structures linked to the hardware systems that they interfaced with. This generated a flexible program that has never been significantly recoded since it was first implemented. This is lucky, because understanding the code lurking beneath the user interface requires a solid knowledge of the hardware systems it is interfaced with. A top-down design would trade UI-level flexibility for a whole-program organization that reflected the design of the experiment: something that experimentalists would probably find easier to understand intuitively.

## 9.7 Future improvements to the control system

The control system for the ANU-BEC-III BEC machine has been broadly designed and parts have been purchased. The system has not yet been implemented, so no real comparisons of performance can be made. The main difference between the ANU-BEC-I control system and the ANU-BEC-III system (other than the level of technology) is the level of integration across the entire system. As well as using National Instruments supplied hardware and software, the computer running the equipment is also made by National Instruments. Because the entire experiment was being designed from scratch, equipment that favored analog control could be specified. Hopefully this will not only lead to a simpler implementation process, but a system that performs two orders of magnitude faster.

The biggest difference with the hardware selected for this control system is that it uses the PXI protocol to communicate with the computer. This is an extension to the PCI protocol, pioneered by National Instruments specifically for data acquisition systems. It is implemented in a rugged chassis design that is extensible to allow many more DAQ cards to be used than in a standard PC. It provides a high performance bus dedicated to the distribution of timing and triggering signals, guaranteeing synchronization between DAQ card to 1 ns.

The cards selected to run the ANU-BEC-III experiment are essentially identical to the analog control cards used to run ANU-BEC-I. The main differences are the more modern PXI format, faster output (1 MHz) and higher resolution (16-bit). The system as a whole has almost twice as many output channels for control than ANU-BEC-I. The biggest difference is the inclusion of an analog input card for implementation of diagnostic recording and deterministic control. This is hoped to decrease the complexity of debugging small problems with the ANU-BEC-III machine on a day to day basis, increasing up-time and productivity. The new experiment also used home made break-out boxes, but designed a printed circuit board with rugged direct attachment of all connectors to the board.

The computer is potentially the biggest difference in the control system. Due to the trouble in generating reliable high speed communication between the computer and control cards in ANU-BEC-I, a National Instruments PXI computer was implemented for ANU-BEC-III. It also provides the option of running the Labview RT real time operating system for optimum operation. A second computer is used to actually host the user interface, but all time critical operations will be performed by the National Instruments computer.

The software for writing the control program is Labview 7. This is obviously many stages more advanced than Labview 3. The fundamental aspects of Labview 3, such as graphical programming, good abstraction from hardware and excellent user interface tools, have been maintained. The hope is that improvements in Labview and Windows, and the exclusive use of compatible components throughout the system, will prevent the problems encountered during ANU-BEC-I.

The ANU-BEC-III control program will be implemented with a top-down programming scheme. This is considered essential because the machine will be used primarily to produce results on advanced experiments utilizing BECs. Users will not

necessarily have a lot of experience on this particular BEC machine, and will need to pick up machine operation quickly through intuitive controls. This leaves them to implement more complicated experiments, increasing the efficiency of using the machine. Also, as the machine moves into a more rapid phase of operation, chances that significant changes will be required to some aspects of the setup to incorporate a new experiment will increase. A logical, intuitive program (underneath the user interface), will be very powerful in facilitating broad changes for even inexperienced Labview users.

## 10 ANU-BEC results

The first *published* results of BEC for groups new to the field are usually absorption images showing evidence of bimodal distribution in the atom cloud signifying some condensed component mixed with a mostly thermal cloud [20, 145]. Ideally, images showing the asymmetric expansion of the BEC may also be produced. These are the classic “signatures” of Bose-Einstein condensation.

However, the very first results that support the existence of a BEC on a new machine are usually, somewhat less dramatically, rapid increases in the observed optical depth, and derived calculations of the phase space density. ANU-BEC-I was no exception. The first BEC was created at ANU on 3 May 2001. Its existence was inferred from a distinct increase in optical depth after the evaporation run, and repeatably large calculated phase-space densities ( $\sim 30 - 100$ ).

With some fine-tuning, ANU-BEC-I went on to produce observations of two-component bimodal distribution. However, it was not until the ANU-BEC-II improvements were implemented on the experiment that truly asymmetric expansion was observed.

### 10.0.1 ANU-BEC results in context


Construction of ANU-BEC-I began mid-1999. BEC was observed in May 2001. This twenty-something months represents one of the most rapid realizations of BEC in the world. The feat is made even more impressive by the fact that, before this experiment, none of the original group members had any significant experience with laser cooling and trapping, or magnetic trapping and evaporation.

In addition, the machine produced was one of the most compact and elegant machines of its day. This was unarguably the result of judicious selection by ANU-BEC team members of the best parts of second generation machines of the more advanced groups around the world. However, making these cutting-edge implementations work together to form a BEC machine was a difficult task in itself.

## 10.1 First results of BEC

Figure 10.1 shows the one of the first confirmed BECs ever produced in Australia. The cloud consists of roughly  $5 \times 10^5$  atoms, about 70 % of which are condensed. The cloud (thermal + BEC component) has an average phase-space density of approximately 30. BECs were probably created on many previous experimental runs, but the magnitude of the calculated phase space density, and the repeatability with which it could be attained, proved that BEC was in fact being realized.



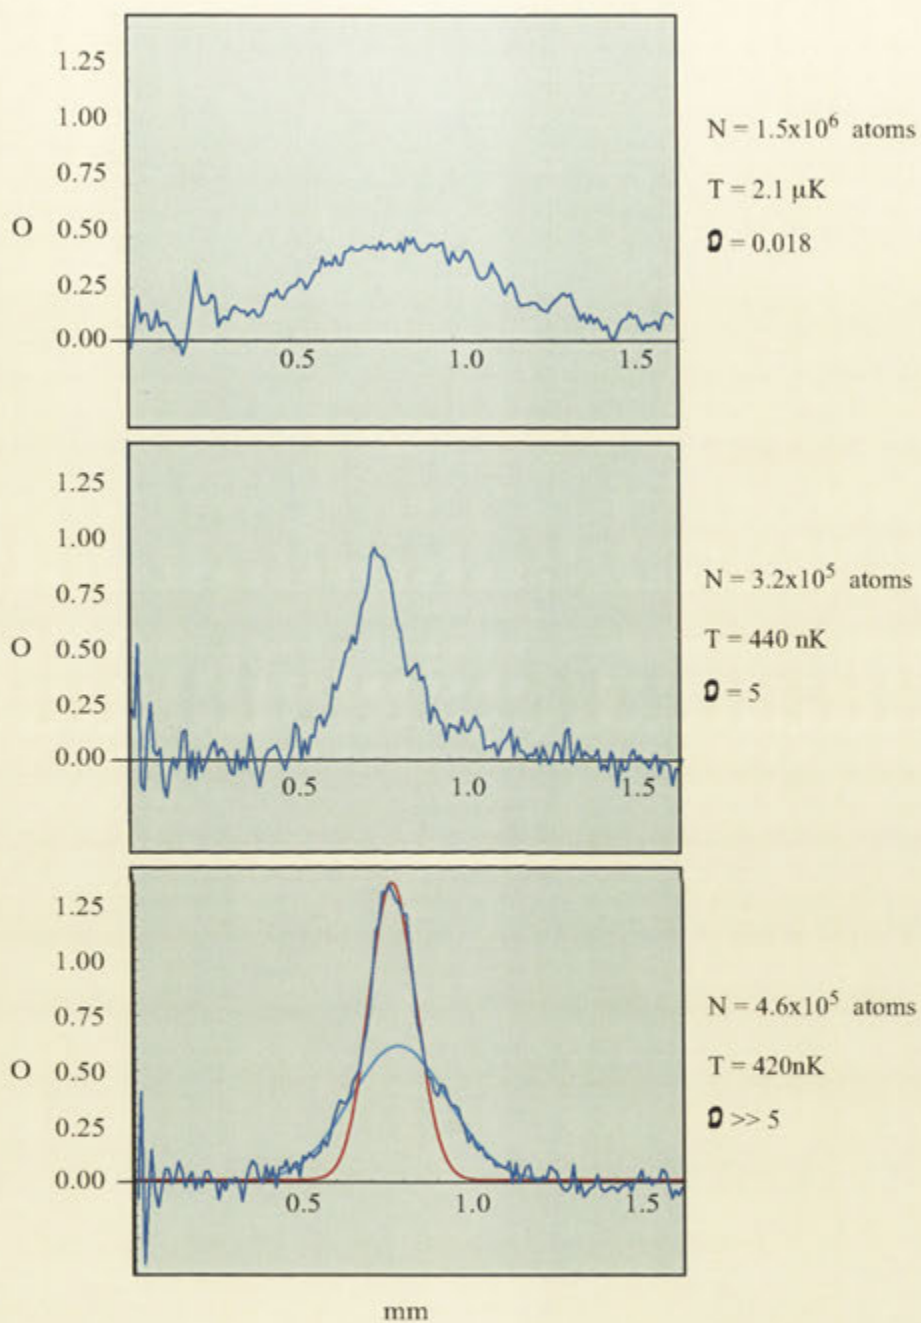


**10.1: Bose-Einstein condensation!** One of the first BECs created in the ANU-BEC lab, from the machine discussed throughout part I of this thesis. It contains approximately  $5 \times 10^5$  atoms, of which approximately 70 % are Bose condensed. It has a phase-space density of roughly 30.


The final progress towards producing BEC was actually purely equipment based. Excellent phase space densities just short of degeneracy were being obtained up to 3 months before BEC was finally realized. Known problems with two major pieces of equipment in the experiment, the main cooling laser and the magnetic trap power supply, were suspected to be the factors limiting increased phase space density.

The laser system had degraded, within specification, over the 15 months it had been used preparing the BEC experiment. The key degradation was an increase in the linewidth produced by the laser: from less than 1 MHz to over 20 MHz. This was an insidious problem, as most of the experiment still ran, just with poor performance. The laser system was sent back to Toptica for replacement of the amplifier diode, which fully corrected the problem.

The problem with the Powerbox SM3540 power supply was more fundamental: the noise specifications were on the border of what was considered acceptable for producing BEC. The SM3540 was a high current switch-mode power supply, and not only were its noise specifications marginal, the residual switching frequency signal concentrated this noise at specific frequencies. Given the sufficient operation of other aspects of the experiment, borne out by the high phase space densities already being achieved, the decision was made to significantly upgrade the power supply running the experiment. A very low noise F.u.G. Elektronik NLN 1400M-35



**10.2: Analysis of condensation.** Cross sections of ANU-BEC-I condensates showing the onset of condensation as a bimodal distribution. The bottom plot illustrates a cross-section of the cloud in figure 10.1.



**10.3: Logarithmic evaporation ramp.** Sequence of images showing the progression towards BEC using a radio-frequency ramp that decreases frequency logarithmically with time. By spending more time at lower frequencies, such an evaporation cycle eventually creates a BEC.

linear power supply was ordered.


Within two weeks of the upgraded supply and the repaired laser arriving in the laboratory, the definitive images of BEC at ANU were produced.

## 10.2 Fine-tuning evaporative cooling

Before these two components were replaced, however, much work was completed optimizing each aspect of the experiment to produce high numbers of cold atoms, and the efficient evaporation of these atoms to degeneracy. This was the “real” last stage of the experiment in terms of effort and importance.

The only aspect of this procedure reported in this thesis is the optimization of the evaporative cooling ramp. This work was published in *Journal of Optics B: Quantum and Semi-Classical Optics* [6]. The result is a comparison of evaporation using a single-slope linear ramp of the RF-knife frequency, to a single time-constant logarithmic ramp of the RF-knife frequency. The logarithmic ramp, shown in figure 10.3, produces BEC, while the linear ramp, shown in figure 10.4, does not. This dramatic indication of the sensitivity of the production of BEC to some aspects of the evaporation cycle.

The conditions describing efficient evaporative cooling have been considered in

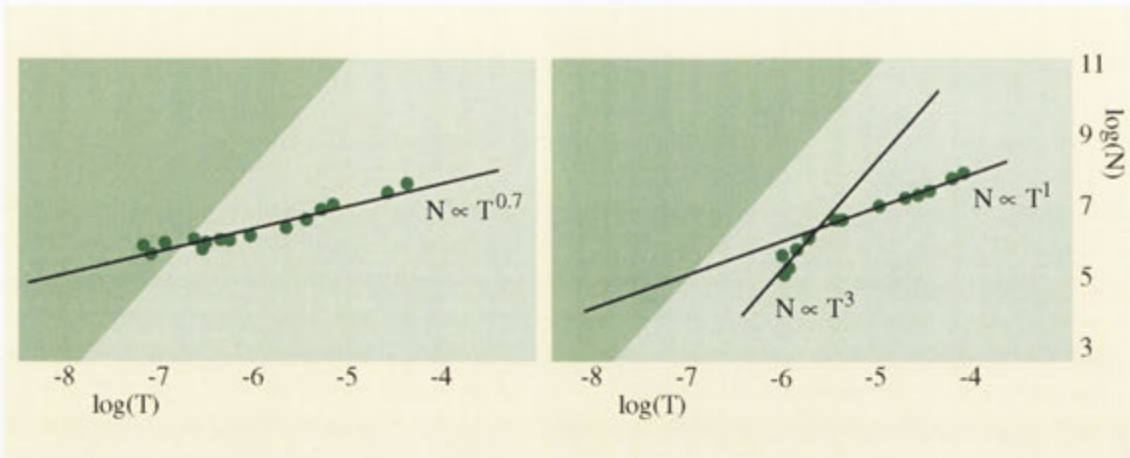


**10.4: Linear evaporation ramp.** Sequence of images showing the progression through phase space using a radio-frequency ramp that decreases frequency logarithmically with time. The ramp progresses through low frequencies too quickly, losing efficiency during the last stages of cooling, never producing a BEC.

chapter 7. Of most importance is the exploitation of changes in optical depth as a measure of elastic collision rate, and the consequent quantification of evaporation efficiency. The power of this analysis is illustrated in figure 10.5. The key result is the comparison between the increase in spatial density with decreasing temperature. The logarithmic ramp (left image), follows a path with a single exponent relation,  $N \propto T^{0.7}$ , all the way to BEC. In contrast, the linear ramp (right image), exhibits a distinct change in its progression through phase space, from a  $N \propto T^1$  path, to a  $N \propto T^3$  path, that prevents it from ever reaches the BEC transition.

The change of exponent in the path of the linear ramp through phase space corresponds to a sharp decrease in the elastic collision rate,  $\sigma_{2b} \propto N/T$ . This decrease in the efficiency of the linear evaporation ramp, coupled with other issues such as technical noise and low atom number, prevents the realization of BEC in this system.

The result is consequent of the relative importance imparted on sections of the evaporation cycle by the ramping process. Linear ramps weight each *absolute* frequency range equally, while logarithmic ramps weight each *relative* frequency range equally. In this case, the linear ramp progressed too quickly through the regions of low frequency, preventing adequate rethermalization and efficient evaporation. In



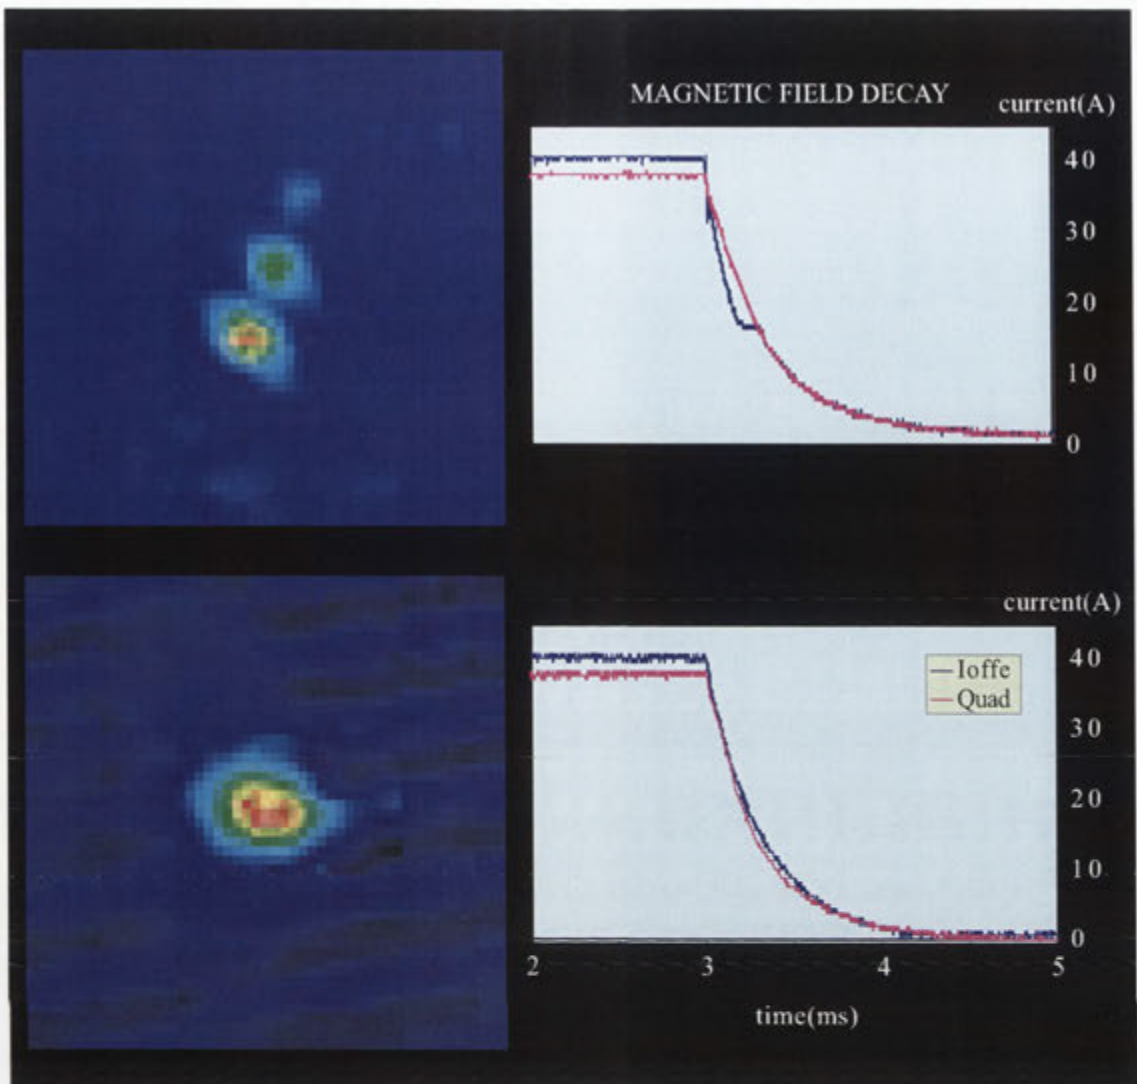
**10.5: Evaporative cooling comparison.** Phase-space diagrams of the data in figures 10.3 and 10.4. The plot points illustrate the path through phase space traced by clouds during a logarithmic (left plot) and linear (right plot) cooling ramp, respectively. The distinct change of gradient in the right hand plot reflects a loss of elastic collision rate as evaporation progresses through the linear ramp. This suppression of collisions proves fatal to the attempt to produce a Bose-Einstein condensate.

contrast, the logarithmic ramp ramps through the lowest frequencies most slowly, allowing rethermalization to occur at a similar level all the way through the evaporation cycle, as evidenced by the single exponent describing the process.

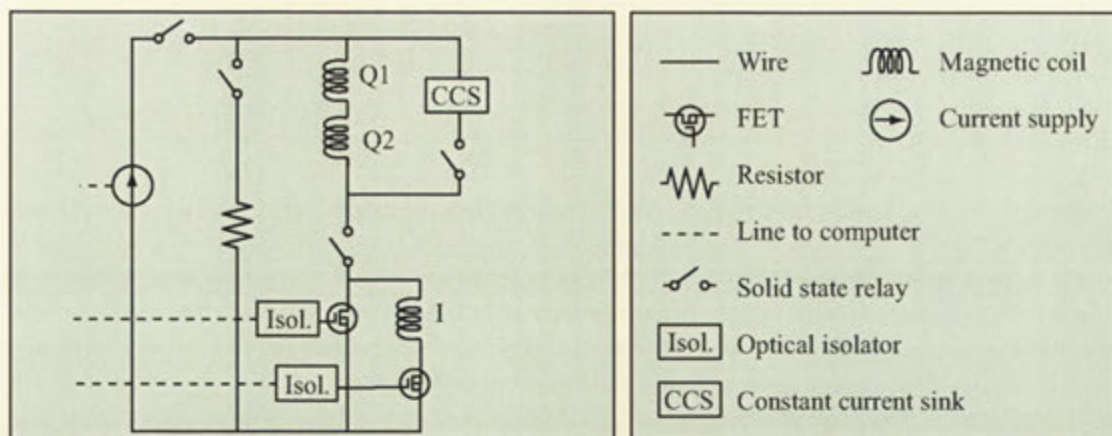
### 10.3 Fine-tuning magnetic trap switching

The first BECs created at ANU did not produce the picture-perfect images from other successful groups, that have since become icons for dilute gas BEC. The most obvious difference between the upper image in figure 10.6, and images of BECs off the improved ANU-BEC apparatus, such as figure 10.8, was the existence of multiple apparently-condensed clouds. This was a result of a problem during the switch-off process for the magnetic trap, which caused atoms to be spin-flipped into multiple  $m_F$  states. As the magnetic field decayed further, each  $m_F$  state experienced a different force due to the various magnetic gradients, physically separating the  $m_F$  components. In practice, this represented a significant loss of atoms from the major condensate component, that in the  $m_F = +2$  state, containing approximately  $1.3 \times 10^5$  atoms, into the  $m_F = +1$  state, containing  $7.3 \times 10^4$  atoms, and the  $m_F = 0$  state, containing  $4.2 \times 10^3$  atoms. Presumably, those atoms transferred into the  $F = 2$ ,  $m_F = -1$  and  $m_F = -2$  “anti-trapped” states have been dispersed by the time the expanded absorption images are collected.

This is actually a very common occurrence, and has been reported elsewhere, sometimes as a diagnostic “feature” of the experimental system [35]. The problem



**10.6: Switching off the magnetic trap** The upper figures show the data collected from a BEC released from a magnetic trap using a switch configuration as illustrated in figure 6.11 (a). The time constant in the plot of the switch off currents is approximately  $500 \mu\text{s}$ , but the dominant effect is the glitch during the switch off of the quadrupole coils. The resulting magnetic field decay contains a field zero swept through the BEC, populating  $m_F$  states other than the  $m_F = +2$  trapped state, which became physically separated due to magnetic field gradients during switch-off. The largest cloud ( $m_F = +2$ ) in the data image contains approximately  $1 \times 10^5$  atoms, while the smaller clouds ( $m_F = +1$  and  $m_F = 0$ ) contain approximately  $7 \times 10^4$  and  $4 \times 10^3$  atoms, respectively. The lower image shows a BEC released from a trap with all coils switched in series, using a circuit as shown in figure 10.7. The time constant of the field decay using this circuit is shorter, approximately  $250 \mu\text{s}$ , and, more importantly, the current glitch has been resolved. This created a markedly better switch-off, decreasing population of other  $m_F$  states, leaving approximately  $4 \times 10^5$  atoms in the  $m_F = +2$  state.




**10.7: The improved switching circuit.** The simultaneous decay of currents through the coils was forced by switching them off in series, with a single dissipative resistance to tune the switching time. This ensured that no zero-field region was swept through the atom cloud, yielding a significant improvement in the loss of atoms into other  $m_F$  states during switch off.

occurs when the fields from various magnetic coils in the trap decay at different rates, at some stage propagating a zero-field region through the condensate. In the ANU-BEC-I experiment this was a result of fine-tuning the switch off of the Ioffe coil independent of the quadrupole coils. The switch-off speeds were roughly matched by choosing switching resistances to complement the inductance of the Ioffe and quadrupole coils, as shown in figure 6.11 (a). However, a glitch in the switch off of the quadrupole coils, seen in the upper plot of figure 10.6, assumedly due to improper operation of a solid state relay, caused by the high reverse voltages created by switching the inductive coils quickly, affected the careful balance of switching rates.

The solution involved re-implementing the magnetic coil circuitry to switch all the coils in series through a single switching resistor, as shown in figure 10.7. This guaranteed the proportional decay of current in all coils simultaneously, maintaining the shape of the magnetic trap as it decayed. Implementing this solution produced the current decay seen in the lower plot of figure 10.6. This created a large single BEC, seen in the lower image of figure 10.6. The cloud contains  $3.9 \times 10^5$  atoms, increasing useful atom numbers for experiment with BECs.

More importantly, it showed the ability to sensitively manipulate the magnetic “polarization” of cold atom clouds and BECs. During experiments on ANU-BEC-II, NR perfected the manipulation of the zero-field scalpel during switch-off, allowing precise population of any number of Zeeman sublevels during switch off [2]. This skill was eventually put to use to create, manipulate and understand atom lasers.



**10.8: Optimized ANU-BEC-II BEC.** The “classic” illustration of BEC. Note the superior images collected as a result of increases in the size of BECs generated by the apparatus, and improvements to the imaging process itself. Many of these improvements developed from lessons learnt during the production of ANU-BEC-I, as outlined in part I of this thesis.

## 10.4 First results on ANU-BEC-II

The changes made to the experiment to create the ANU-BEC-II machine from the ANU-BEC-I machine have been well documented throughout this part of the thesis. These changes lead to significant improvements in the size, quality and, most importantly, reliability of BEC production, as borne out by figure 10.8. The figure shows the improvement in both the quality of images generated, and the techniques used to display them, that came with the ANU-BEC-II machine. These improvements increased the quality and reliability of the ANU-BEC-II machine so that it could be used to perform experiment *with* BECs, instead of just *producing* BECs.

The biggest changes affecting the production of images such as figure 10.8, were the implementation of continuous quantitative diagnostics on the collection and science MOTs, the re-implementation of the science MOT as a three-beam retroreflected MOT, and the implementation of a new magnetic trap to allow simpler imaging. The first two factors effected the day-to-day reliability, and hence the “average” level of optimization of the experiment. This lead to the production of bigger BECs. The third factor allowed these bigger BECs to be imaged optimally,



leading to the excellent results shown in figure 10.8.


## 10.5 The atom laser

The first experiment performed on the optimized ANU-BEC-II machine was the implementation of Australia's first "atom laser", shown in figure 10.9. The atom laser was created by carefully applying an RF-knife to a BEC to bleed atoms slowly from the trap. This requires precision and stability in both the knife frequency, and the bias field at the bottom of the magnetic trap: the condensate is typically only  $\sim 5$  kHz, or  $\sim 5$  mG, "wide". The work building on, and improving, the QUIC trap for ANU-BEC-II, was vital in producing this result: it would have not been possible in the magnetic trap used in ANU-BEC-I.

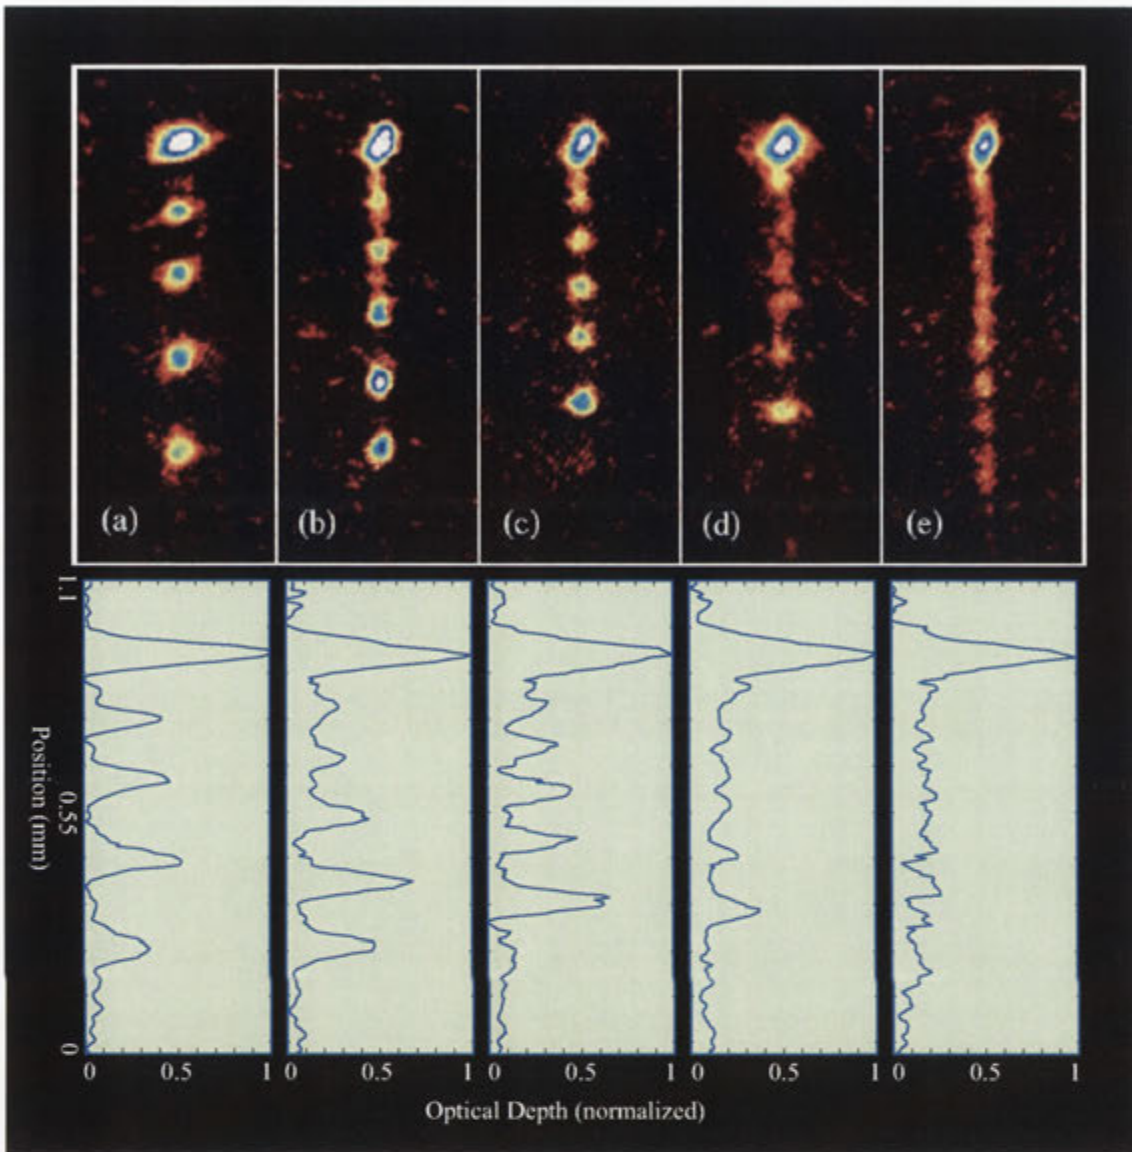
However, figure 10.9 represents the best of only a few "semi-continuous" atom lasers achieved in the ANU-BEC-II machine. The stability of the ANU-BEC-II QUIC trap still limits the precision available in applying the RF-knife, making continuous operation of an atom laser very difficult. Improving the stability of the trap to allow production of continuous atom lasers is one of the primary aims of the new ANU-BEC-III machine.

The solution for performing repeatable atom laser experiments on the ANU-BEC-II machine was to run the atom laser in pulsed mode. A very short pulse of RF can have a significant width in Fourier space, addressing the entire BEC, and relaxing the stringent stability requirements. By tuning the intensity of the RF signal, atom laser pulses of arbitrary intensity can be produced. Figure 10.10 shows results from such an experiment.

This work has produced some interesting physics results, published as a Rapid Communication in Physics Review A [13]. The spin-flip process coupling atoms out of the BEC is not an irreversible process: as subsequent pulses of RF are applied to "pulse" the atom laser, previously produced atom laser pulses are coupled into all  $m_F$  states, with amplitudes determined by their precise position in the magnetic field. The most striking example of this is the disappearance of the first (lowest) atom laser pulse in 10.10 (c). This effect has been modeled by Craig Savage, and its observation is explained in the publication mentioned above.



**10.9: The semi-continuous atom laser.** Detail of an optimized atom laser. The image contains approximately  $10^5$  atoms. The output beam was created using a weak RF field (0.8 mG), over 12 ms, 2 kHz above the bias field of the trap. The spatial axis are quantized in pixels, and the vertical axis is scaled to the optical depth. The ability to consistently produce high quality data that came with the development of ANU-BEC-II was a key development in producing these studies of atom lasers.



**10.10: The pulsed atom laser.** These images were part of a publication by the ANU-BEC group that provided experimental and theoretical confirmation of the non-Markovian nature of an atom laser. The period over which pulses were applied was kept constant, at 8 ms, for each data set. From left to right, four, five, six, seven and ten RF pulses are applied within this 8 ms window. Each output atom cloud is created by a short pulse of RF, with a significant width in Fourier space (200 kHz). Importantly, this is broad enough that each pulse addresses the entire BEC, minimizing the effect of trap bias instability on the repeatability of the data. The interesting result is that image (c) only exhibits five distinct pulses, although six pulses of RF have been applied. The first (bottom) pulse has been coupled back into other components, due to the reversible nature of the out-coupling process.

## Part II

# Degenerate Fermi gases

# 11 ANU-DFG

Following the demonstration of a BEC at ANU, the focus of this PhD turned to the development of Australia's first experiments towards creating a degenerate Fermi gas. The ANU-DFG aimed to cool Fermionic  $^{40}\text{K}$  atoms to degeneracy via sympathetic cooling using Bosonic  $^{87}\text{Rb}$  atoms. In contrast to the development of the ANU-BEC-I machine, this work was conducted without significant assistance from other people.

The results presented in this part of the thesis are almost entirely technical. It aims to give an overview of the system proposed for producing a dilute gas DFG, and then elucidate the *unique* progress made towards this goal during this thesis. All of the work presented in Part II, *Degenerate Fermi gases*, constituted a significant research investment towards this thesis.

## 11.0.1 DFGs in context

Degenerate Fermi gases were, for some time, the neglected half-brother of dilute gas BECs. This can be attributed to the difficulty of evaporatively cooling DFGs, and the lack of a dramatic, amplified and detectable phase transition at the onset of degeneracy. On the other hand, the suppressive nature of Fermi-Dirac statistics promises just as much interesting physics as the Bose amplification that drives BEC transitions. At the beginning of 2004, this is becoming increasingly apparent as the world of degenerate Fermi gases explodes into a diverse wave of new experiments contemplating DQG mixtures [107, 114, 115], molecular BECs [64–67], and, finally, Cooper pairing and BCS transitions [123–125].

Chapters 1 and 2 outlined the fundamental difficulties inherent in the pursuit of a degenerate Fermi gas. This led to the significant gap between the first observation of degeneracy in Bosons in 1995 and in Fermions in 1999, despite the fact that very similar techniques were used in both experiments. Even more telling, however, is that when development of ANU-DFG began at the start of 2001, still only one DFG existed worldwide. Not until mid-2001 was the second achievement of DFG reported, and at this point almost 40 BECs were in operation around the world.

Even at the start of 2004, however, degenerate Fermi gases are not yet “common” research. Although they have certainly piqued the interest of the scientific community, their comparative difficulty still keeps them out of reach of many groups worldwide. The groups that have achieved degeneracy in Fermi gases are generally large, well-established groups with access to extensive experience with BEC experiments. In contrast, ANU-DFG has been a low-key project that has invested three years gradually overcoming the significant obstacles facing a DFG experiment.

## 11.1 Overview of the DFG machine

The ANU-DFG machine is still being developed, and is yet to produce a degenerate gas of Fermions. The experiment described below is the planned process for producing a DFG. It is similar to the method used by Inguscio et. al., the only group in the world to have successfully produced a  $^{40}\text{K}$  DFG by sympathetic cooling with  $^{87}\text{Rb}$  [77]. The discussion of this “blueprint” is important for orienting the following chapters, explaining progress towards this goal.

The defining feature of the DFG machine is its implementation as a two-species experiment utilizing sympathetic cooling of the Fermionic atom sample via the Bosonic atom sample. This is necessary, as discussed in chapter 2, because the rethermalizing collisions vital to efficient evaporation are suppressed in very low temperature, spin-polarized Fermionic samples. The ANU-DFG machine will use  $^{87}\text{Rb}$  to sympathetically cool  $^{40}\text{K}$ . The  $^{87}\text{Rb}$  will be prepared and cooled essentially identically to the process described in part I of this thesis. The Fermionic atoms will be cooled due to interactions with the colder Bosonic gas after optically cooling, polarizing and trapping a gas of  $^{40}\text{K}$  in the same region as the  $^{87}\text{Rb}$  BEC.

The experimental configuration is very similar to that used in the  $^{87}\text{Rb}$  BEC experiment. Atoms are collected in a double MOT configuration, from a high pressure collection MOT to a UHV science MOT. The key difference is that the MOTs are made with the light for both  $^{87}\text{Rb}$  and  $^{40}\text{K}$ . Both species are released into the collection chamber from dispensers, and both are initially cooled and trapped, in the same volume, at the center of the dual-species collection MOT. This process is relatively efficient. Other experiments achieve MOTs of  $10^9$   $^{87}\text{Rb}$  atoms overlapped with  $10^8$   $^{40}\text{K}$  atoms. The atoms are transferred to the dual-species science MOT in the same way as the  $^{87}\text{Rb}$  BEC experiment.

From the science MOT, both species of atoms undergo polarization gradient cooling and optical molasses, by detuning and attenuating the MOT beams as described for the  $^{87}\text{Rb}$  BEC in part I of this thesis. Both species are pumped into weak-field seeking states:  $5^2S_{1/2}$   $F = 2$ ,  $m_F = +2$  for  $^{87}\text{Rb}$ , and  $4^2S_{1/2}$   $F = 9/2$ ,  $m_F = 9/2$  for  $^{40}\text{K}$ . A magnetic trap, practically identical to that used in the  $^{87}\text{Rb}$  BEC experiment, is switched on, and the ultracold, spin-polarized atoms are captured.

The  $^{87}\text{Rb}$  atoms are evaporated gradually, as in the BEC experiment. As they cool, they undergo collisions with other  $^{87}\text{Rb}$  atoms, as well as the  $^{40}\text{K}$  atoms. In this way, the relative energy distribution of both species remains fairly similar, and the evaporation of  $^{87}\text{Rb}$  also cools the  $^{40}\text{K}$ . A major benefit of this technique is that the  $^{40}\text{K}$  is not evaporated, leaving large numbers of atoms available after evaporation. This can be used to produce excellent signal to noise for sensitive experimentation. Alternatively, the  $^{40}\text{K}$  can be further evaporated (using mixed spin states), after all the  $^{87}\text{Rb}$  is lost, to reach even further into the degenerate regime.

## 11.2 Progress of the ANU-DFG experiment

The ANU-DFG experiment has seen significant progress, despite employing only one person and limited equipment. The development plan was based around demonstrating a prototype dual  $^{87}\text{Rb}$ ,  $^{40}\text{K}$  MOT, in an experiment independent from the ANU-BEC experiments. This necessitated the design and implementation of all the laser systems, the source of each atomic species, and preliminary vacuum, control and detection systems. Much of this complicated process was achieved, in a step-by-step manner, over a little less than three years.

The initial work focussed on developing a  $^{87}\text{Rb}$  MOT, as the technology and expertise already existed within the ANU-BEC group. This involved developing lasers in-house that were similar to those supplied commercially for the ANU-BEC-I experiment. This process involved an honors student, Jessica Trevena, and formed a significant part of this thesis. Although the designs implemented were very similar to those used in other lasers in the laboratory, they had never before been constructed, assembled and optimized by the ANU-BEC group. Understanding the subtleties of the design, and optimizing it for use in the new experiment, required a significant investment. This investment eventually paid off, however, when novel lasers for the  $^{40}\text{K}$  cooling and trapping were developed. The development of external cavity diode lasers is discussed in detail in chapter 12.

Along with the lasers, all the other components required for producing a  $^{87}\text{Rb}$  MOT were collected and implemented. This included vacuum systems, a  $^{87}\text{Rb}$  dispenser, optics and detectors. Several prototypes were constructed at various stages throughout the experiment, as they were refined and more equipment became available. Milestones in the implementation of the  $^{87}\text{Rb}$  MOT system included: the initial configuration using unlocked home-made external cavity diode lasers, plate beam splitters and a recycled glass cell; the improvement of this initial system with locked lasers and calibrated CCD camera fluorescence detection to allow the development of the self-locked magneto-optic trap and to test non-destructive imaging of BECs; the implementation and testing of a new vacuum system and square-section glass chamber identical to those to be used on the ANU-BEC-III machine; and the creation and analysis of the first  $^{85}\text{Rb}$  MOT in the ANU-BEC laboratory. These investigations took a very considerable amount of time, but most of them are not reported in detail in this thesis.

As discussed in chapter 2, one of the fundamental difficulties with creating a DFG is the lack of Fermionic atomic species available for consideration. Within the alkalis only two stable isotopes exist:  $^6\text{Li}$  and  $^{40}\text{K}$ . The primary *practical* difficulty that this engenders is a lack of choice when selecting the species with which to work. This means that there is little leeway for optimizing the experiment to the equipment available, either in the laboratory, or commercially. The two major difficulties experienced when choosing to work with  $^{40}\text{K}$  are sourcing lasers at the correct wavelength, and sourcing the  $^{40}\text{K}$  atoms themselves.

Both of these issues needed to be addressed before a suitable potassium MOT could be implemented. The source of  $^{40}\text{K}$  was particularly problematic. Standard abundance potassium consists of 93.26 %  $^{39}\text{K}$ , 6.73 %  $^{41}\text{K}$  and only 0.012 %  $^{40}\text{K}$ .

Making a MOT from a standard abundance source is difficult, because in producing sufficient  $^{40}\text{K}$ , so much  $^{39}\text{K}$  and  $^{41}\text{K}$  are created that the hot background gas in the collection chamber severely effects the operation of the MOT [163, 164]. The solution is to create a sample of potassium in which the fraction of  $^{40}\text{K}$  has been enriched. This has been achieved, and the method published, by DeMarco et. al., the first group to create a  $^{40}\text{K}$  DFG [70]. Their method involved purchasing KCl that had been enriched to a 5% isotope fraction of  $^{40}\text{K}$ , then mixing it with pure Ca as a fine powder, and holding the mixture in a small Nichrome boat. Placed inside the vacuum system, and attached to a pair of electrical feedthroughs, the enriched “dispenser” activates when several amps of current is passed through it, heating the Nichrome boat over 700 °C and activating a reduction-oxidation reaction between the Ca and the KCl. This system works in a manner similar to that in which the commercial Rb dispensers used on the ANU-BEC experiments operate. This path was followed during the ANU-DFG preparations. Unenriched dispensers were constructed and analyzed using a residual gas analyzer, but the results were inconclusive. The best test of performance involved using the dispenser to load a MOT, so other aspects of the ANU-DFG prototype were continued. As the experimental details to date are not significantly different from DeMarco’s publication, they are not discussed further in this part of the thesis.

The laser system required for cooling and trapping  $^{40}\text{K}$  induces significantly more consideration than that used for  $^{87}\text{Rb}$ . Laser diodes are commercially produced near the  $^{87}\text{Rb}$  780 nm cooling transition, making the use of diode laser systems, such as ECDLs, for  $^{87}\text{Rb}$  economical and efficient. Unfortunately, the potassium cooling transition at 767 nm is just far enough away from commercially available wavelengths to make the choice of laser system problematic. The major benefits of ECDLs, elegant simplicity and economy, can be lost if they become expensive or complicated in an effort to tune 780 nm diodes to 767 nm. More advanced Titanium:Sapphire or dye laser systems work at these wavelengths, but were not compatible with the ANU-DFG budget. The solution developed for the ANU-DFG harnessed the strengths of the commercially available 780 nm diodes, without relying on complicated or fragile components, in a manner that maintained the best features of the ECDL design. This resulted in the optimized design of a solid-state cooling system, insulation and enclosure for an ECDL, enabling extended temperature tuning from 783 nm to 767 nm. The detailed design process is presented in chapter 13.

The home-made ECDL lasers, for both  $^{87}\text{Rb}$  and  $^{40}\text{K}$ , achieved similar performance to the commercially produced lasers in the ANU-BEC laboratory. While these lasers could be used to make a single MOT, none of them produced sufficient power to drive an entire degenerate gas experiment. In the ANU-BEC experiments this issue was resolved using a commercial tapered amplifier diode (TA), injection locked to a low-power ECDL in a master-oscillator power-amplifier (MOPA) configuration. This is an elegant manner of producing approximately 500 mW of single mode laser light. Unfortunately, the TA chips are not extremely economical, so it was not possible to use three separate TAs for the various light needed throughout the dual-species degenerate gas experiment. The solution was to seed a single TA



chip with all four frequencies required throughout the experiment, as an extension of the technique reported by Salomon et. al. [165]. This provided moderate, tunable power (up to 400 mW total) in the cooling and repumping beams for both  $^{87}\text{Rb}$  and  $^{40}\text{K}$  simultaneously. This configuration was implemented, and used to run a  $^{87}\text{Rb}$  MOT while amplifying  $^{40}\text{K}$  light, in preparation for the operation of a dual-species MOT. It is reported briefly in chapter 12.

The other major result towards the production of a DFG machine in the ANU-BEC laboratories was the demonstration of a new frequency stabilization method for the lasers used to run a MOT. The self-locked MOT (SL-MOT) is reported in detail in chapter 14. It enabled the frequency of the lasers used to run the MOT to be stabilized via a feedback signal generated from a frequency measurement using the atoms in the MOT itself. The key development was a method of comparing the laser frequency with the atomic resonance without perturbing the MOT significantly. This was demonstrated as a proof-of-principle on the ANU-DFG  $^{87}\text{Rb}$  MOT. However, its main strength would be in frequency stabilizing the  $^{40}\text{K}$  lasers directly to the  $^{40}\text{K}$  transition. This is not possible in a standard abundance vapor cell, because of the small fraction of  $^{40}\text{K}$  available. Other groups working with  $^{40}\text{K}$  either have to lock to a more abundant isotope of potassium, then shift the frequency of their lasers using modulators, or use less accurate locking techniques.



## 12 External cavity diode lasers

Laser cooling and trapping is an extremely important aspect of all current dilute degenerate quantum gas systems. Magneto-optic traps, optical molasses and polarization gradient cooling combine to capture and cool atoms from room temperature samples to temperatures of the order of a millikelvin. This typically increases the phase space density of a sample of atoms by over 11 orders of magnitude, to approximately  $10^{-6}$ .

The lasers used in the ANU-BEC-I experiment were summarized in chapter 5. These consisted of a commercially produced, amplified diode laser system, and an external cavity diode laser system made by a group from the Research School of Physical Sciences and Engineering (RSPHYSSE) at ANU.

The DFG experiment had no such source of prepared lasers. Excepting a commercial tapered amplifier system configured at 780 nm, all diode lasers were designed, assembled and tested as part of this thesis. Developing the techniques to build standard external cavity diode lasers (ECDL) for interacting with rubidium involved some significant effort. Extending these techniques to produce ultracold ECDLs for potassium was an even larger investment. The first of these developments is detailed below, the second in chapter 13.


### 12.0.1 External cavity diode lasers in context

Diode lasers are ubiquitous in many areas of physics and engineering. Their production is driven by their use in the telecommunications industry, but their resulting low cost and availability, and their ease-of-use, has seen their implementation in many areas of physics, such as atom optics and high resolution spectroscopy [166].

Although diode lasers are cheap and readily available at many wavelengths across the visible and near infrared part of the spectrum, there are many intermediate wavelengths that are not produced commercially. Many of these intermediate regimes, although of no interest to big consumers such as the telecommunications industry, are of vital interest to researchers working in atom optics and spectroscopy.

Since diode lasers were first realized, techniques have been developed to change their operating characteristics from their nominal values to something that better suits a particular user. Their low cost, ease of use, and reliability, often makes performing such engineering feats preferable to pursuing other technologies. However, the risk of losing these benefits increases as the modifications to a diode laser system become more complicated.

In atom optics one of the best examples of well engineered systems based around mass-produced laser diodes is the external cavity diode laser (ECDL) [143, 167]. It



**12.1: An external cavity diode laser.** The key components of an external cavity diode laser, including the diode, grating and temperature control. Not apparent is the small lens mounted in front of the diode to collimate output onto the grating. The arrows to the base plate and grating mount indicate the position of optics-grade bolts for adjustment of the cantilever mounts.

uses an elegantly designed external cavity around the gain medium provided by the laser diode to provide spectral definition and control over the lasing wavelength. It is capable of narrowing the spectral output of typical “single mode” laser diode to less than 1 MHz, while maintaining a tuning range of up to  $\pm 2$  nm. The resulting economical, simple and flexible laser has found extensive application in atom optics and laser spectroscopy experiments worldwide.

### 12.0.2 Summary of the DFG lasers

An external cavity diode laser, such as that shown in figure 12.1, consists of a typical single-mode laser diode, a high numerical aperture lens to collimate the output from the diode, and some device acting as a “half silvered mirror”, that effectively extends the length of the laser cavity by reflecting light back into the laser diode. In the Littrow configuration, an elegant example of which was detailed by Hansch et. al. in 1995, the “half-silvered mirror” is actually a diffraction grating [143].

The laser diode used in such a system typically exhibits free-running specifications of 1 - 100 mW output power, 100 MHz single-mode linewidth, and temperature tuning of  $0.25 \text{ nm} / ^\circ\text{C}$ . The semiconductor crystal is cleaved at the ends of the diode cavity, yielding  $\sim 30\%$  reflective mirrors. On standard high power diodes, such as those purchased for the ANU-DFG experiment, the rear facet is typically treated with a high-reflectance coating to maximize efficiency, while the output facet is treated with a low-reflectance coating to maximize output power.

The specifications required for laser cooling and trapping include  $\sim 100$  mW

power, sub-MHz linewidth, and a good stability against laboratory temperature fluctuations and other external disturbances. This is achieved by careful choice of the laser diode to be used, integration into a relatively long external cavity, and active temperature stabilization and careful design of the mechanics and electronics supporting the laser. In addition, various levels of precision tuning of the wavelength are necessary to provide the absolute frequency required by the atoms being cooled.

Figure 12.1 shows these features implemented: the external cavity is approximately 3 cm long, and reflects 30 % of the incident light back into the laser diode. This narrows the linewidth of a laser diode from 100 MHz to  $\sim 500\text{kHz}$ . The TEC indicated in the figure, and a thermistor, are configured in a control loop to stabilize the temperature of the laser diode and the external cavity to a precision of 5 mK. A heavy aluminum base, and precision driving electronics, effectively isolate the external cavity diode laser from environmental noise. Tuning can be implemented by very coarse, moderately coarse, or fine control of the grating position, by tuning the temperature of the diode laser system as a whole, or even by the drive current used to power the laser diode.

The extension of this system to provide extended temperature tuning of the laser is further discussed in chapter 13. This important development allowed the use of standard, mass-produced 780 nm diodes, to cool and trap potassium at 767 nm. Compared to figure 12.1, all that was added was a second stage of thermoelectric cooling, a carefully designed and implemented system of insulation, and a sealed enclosure. Temperature tuning a standard diode focussed all necessary technical developments on simple, cheap and easily replaceable parts of the laser, yielding an economical and rugged design.


## 12.1 Fundamental considerations

Figure 12.2 shows a typical semiconductor laser diode as used in an external cavity diode laser. Its key specifications are its gain curve, its physical dimensions, and the reflectivity at its output facets. Like any laser, its precise operating point and characteristics are determined by competition between gain and loss in the system. However, diode lasers occupy a fairly unique position in the parameter space available to laser systems.

### 12.1.1 Basic properties of laser diodes

Diode lasers have extremely short, low reflectivity cavities, and broad, extremely high gain. The first two characteristics ensure laser diode cavities have high free-spectral range and low finesse, yielding large tuning ranges. A laser is able to be constructed within such high-loss tunable cavities because of the extremely high gain.

Light is produced in diode lasers by transitions between direct gap energy bands in a heavily doped p-n semiconductor junction. The spectral width of the gain medium is broad compared to other types of lasers because transitions occur between



**12.2: A laser diode.** The basic construction of a laser diode. The important parameters for consideration in an external cavity diode laser include the length of the cavity, and details of the antireflection coatings of the front facet.

energy bands of the semiconductor rather than discrete energy levels of atoms or molecules. The full width half maximum of the gain curve of diodes used in these experiments was approximately 10 nm. The gain can be made extremely high because of the high density of carriers accessible in solid semiconductors.

The cavity of a diode laser is typically provided by cleaving the output facets of the active region along a crystal axis. The large difference in refractive index between the semiconductor crystal and air leads to a significant reflectivity, typically 30 %. Additionally, the facet at each end of the cavity may receive an interference coating to enhance or decrease its reflectivity. These experiments used single mode diode lasers, with typical cavity lengths of hundreds of micrometers. This yielded a free spectral range (FSR), or frequency step between mode-hops, of hundreds of gigahertz, or tenths of a nanometer.

The breadth of the gain curve, and the significant but insufficient free spectral range of the diode cavity, may suggest that typical laser diodes would support multiple mode operation. However, transitions between energy bands in semiconductors occur over very short timescales. The subsequent homogenous broadening in frequency space is therefore quite significant. As it is the dominant broadening mechanism in semiconductor lasers, carefully constructed diode lasers exhibit significant mode competition to enhance the spectral purity of the light produced. While multi-mode diodes are common, the lasers used in these experiments were carefully designed to be single-mode devices.

The linewidth exhibited by diode lasers is due to a convoluted combination of factors. Diode laser cavities have high FSR ( $\sim 500$  GHz) and very low finesse ( $\sim 5$ ), leading to high cavity linewidths ( $\sim 100$  GHz). However, in a manner similar to inter-mode competition, diode lasers exhibit very significant narrowing of the

individual cavity modes. This may be expected to decrease the observed linewidth significantly below that due to the finesse of the cavity, even to the fundamental Schawlow-Townes limit ( $\sim 1$  MHz). However, the small extent of the laser cavity generates another effect: the instantaneous carrier density throughout the gain medium can vary significantly during operation, changing the refractive index in the gain medium, and the effective length of the cavity. This manifests itself as fast frequency instability: an increased linewidth. In total, these processes yield a typical linewidth of  $\sim 100$  MHz for the lasers used in these experiments.

The wavelength of the actual mode produced by a laser diode depends on the coincidence of particular cavity modes across the gain curve of the semiconductor medium. Both the cavity specifications and the exact form of the gain curve are sensitive to the temperature of the p-n junction. Temperature changes cause thermal expansion of the semiconductor, microscopically changing the length of the laser cavity and the modes supported by the cavity. Temperature variation also changes the density of carriers, which modifies the consequent position of the gain curve in frequency space, as well as perturbing the refractive index of the semiconductor medium, and the effective length of the cavity. The magnitude of these effects is of the order of  $0.05$  nm /  $^{\circ}\text{C}$  and  $0.2$  nm /  $^{\circ}\text{C}$  respectively. A typical temperature tuning curve is shown, in the next chapter, in figure 13.8.

Not visible in the discrete data of figure 13.8, is the characteristic tuning pattern of regions of small, smooth wavelength tuning, separated by large, discontinuous frequency jumps, or mode-hops. The large jumps are due to the movement of the gain curve within the supported cavity modes, causing the dominant mode to hop from one cavity mode to the next. A few of these hops occur for each degree of heating, as the peak of the gain curve shifts in frequency one free spectral range of the cavity. Within each mode hop, the length change of the cavity provides a slower, continuous tuning. Optimizing the tuning of a free-running diode laser exhibiting these large mod-hops can be very difficult. The external cavity diode laser design, especially the careful linking of the diode current and the grating position, discussed in section 12.1.4, below, helps overcome these problems.

In many ways, these temperature effects are a double edged sword: while they allow some tuning of the operation wavelength of the diode laser, they can lead to instability in the operation wavelength with environmental temperature changes. In an operating experiment, the laser frequency is typically actively stabilized via a spectroscopy-based frequency measurement and feedback to the laser current or grating position. However, the locking circuit is only able to stabilize the laser within a certain range of operation, typically 1 GHz. For the locking circuit to maintain the frequency of a laser within 1 GHz, or 0.002 nm, of transition, the tuning coefficients above imply a temperature stability of  $\sim 10$  mK. Diode laser systems for atom optics incorporate the diode package into a temperature-controlled loop that allows both tuning and stabilization of the diode temperature and operation wavelength.

### 12.1.2 The external cavity

Atom optics, especially laser cooling and trapping experiments, rely on manipulating atoms via their naturally broadened hyperfine transitions. This requires lasers that produce light with a position and frequency spread well defined relative to the natural linewidth of these naturally broadened transitions. Alkali atoms typically support natural linewidths of approximately 5 MHz. Standard single mode laser diodes, producing linewidths of 100 MHz, are not sufficiently precise to manipulate atoms within these individual hyperfine levels.

The linewidth of the lasers can be reduced by creating an *external cavity* around the laser diode. The external cavity has several effects. By increasing the length of the cavity relative to the length of the gain medium, the influence of the carrier density fluctuations on the effective cavity length can be almost eliminated. This removes the broadening limiting the linewidth of the free-running laser diode. The external cavity can also provide an extremely powerful way of tuning the wavelength of the diode laser by directly changing the physical length of the cavity. [143, 166]

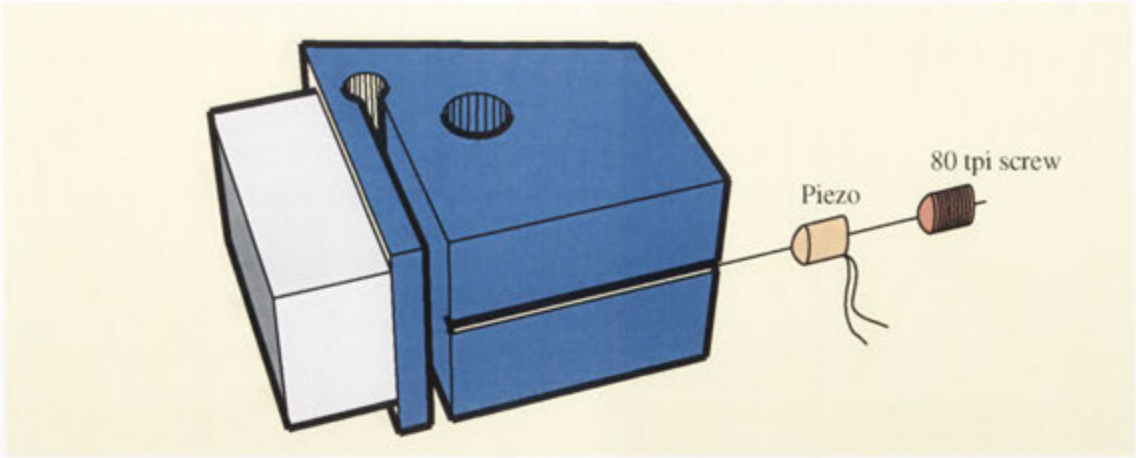
The linewidth of the ANU-DFG external cavity diode lasers was quantified, to the limit of precision of a beat experiment, to be better than 1 MHz. Other reports have indicated that, with care, the linewidth of such lasers can be minimized even closer to the fundamental limit. The residual frequency noise has been attributed to physical vibration of the external cavity components. The ANU-DFG ECDLs used a design based on that reported by Hansch et. al., based around solid cantilever joints, to minimize acoustic vibration noise. The actual linewidth may be even better than that measured. However, as the atomic transitions that need to be addressed are naturally broadened to  $\sim 5$  MHz, the performance is easily sufficient for these experiments.

The resonant frequency of a cavity could be tuned if the position of one of the cavity reflectors was able to be adjusted precisely. Although the ECDL design used for the ANU-DFG lasers did not involve anti-reflection coating the output facet of the diode, the operation of the laser as a single, large cavity, was well approximated. This allowed the tunability of the ECDLs to be significantly improved by mounting the grating on a cantilever mount. The cantilever was loaded by a screw via a piezoelectric transducer, as shown in figure 12.3. This was vitally important for enabling access to all wavelengths across the tuning range of the laser, between the mode-hops observed in the free-running tuning curve. Considerations based on this design are detailed below.

### 12.1.3 The grating

The back facet of the laser diode and a diffraction grating form the external cavity used in the ANU-DFG ECDLs. The grating acts as a partially reflective mirror. It is configured to reflect the first order diffraction beam directly back into the laser diode, while the zeroth order beam is reflected out of the cavity as if off a plane mirror. By careful choice of the grating parameters, the modal operation of the laser, the tunability of the external cavity, and the efficiency of the output can





**12.3: Detail of the grating mount.** A key component of both the grating mount and the ECDL base plate is the use of cantilever mounts, in AlBr, loaded by optics grade 80 tpi screws. These mounts are inherently stable, and the AlBr achieves sufficient elasticity whilst maintaining conductivity and thermal stability. The exploded view illustrates the use of piezoelectric transducer between the cantilever and the optics grade screw, allowing fine tuning and electronic scanning of the laser wavelength.


all be engineered to optimize the ECDL design for a specific use. The specifications of the grating include the line density, the blaze angle, and the surface finish, as illustrated in figure 12.4.

The external cavity has a much smaller FSR than the diode cavity due to its significantly greater length. However, the use of a grating as a cavity reflector provides a secondary method of discriminating the operation frequency of the laser. A key property of a grating is its ability to disperse an incident light beam based on its wavelength. The *resolving power*,  $R$ , is a dimensionless quantity expressing the minimum fractional wavelength that may be resolved by a planar grating,  $\Delta\lambda$ . The efficiency with which it achieves this is indicated in equation 12.1.

$$R = \frac{\Delta\lambda}{\lambda} = \frac{2nd}{\lambda} \tag{12.1}$$

$d$  represents the distance between adjacent rulings on the grating, while  $n$  quantifies the number of rulings illuminated by the incident beam. The high numerical aperture - short focal length lenses used in these ECDLs create a 5 mm collimated beam at the grating. The gratings used typically exhibit rulings at 1200 lines / mm. Approximately 6000 lines are illuminated, and a wavelength discrimination of just over  $1 \text{ \AA}$  results. This corresponds to a frequency resolution of less than 100 GHz.

As this resolution spatially separates adjacent cavity modes sufficiently that only one mode can be efficiently reflected back into the laser diode at any one time, only that one mode will lase. More than this, if the grating is gently moved to change its angle to the intra-laser beam, the resonance condition changes, and the



**12.4: Detail of a reflection grating.** The important parameters for use in an external cavity diode laser system are the rule density (lines per mm), the blaze angle, and the surface finish. The rule density determines the frequency discrimination abilities of the grating within the external cavity. The blaze design wavelength is chosen *away* from the operation wavelength to reduce diffraction into the first order mode (back into the diode) to 30 %. The surface finish effects the quality of the grating, and the overall efficiency and output power of the laser.

frequency of the laser is tuned, without the risk of mode-hopping to an adjacent cavity mode. This increase in the mode-hop-free scan range and easy tunability is the fundamental benefit of implementing the output cavity reflector as a diffraction grating.

The other major decision associated with choosing a grating for an ECDL is the blaze angle. The blaze angle affects how much light is reflected into the first order relative to the zeroth order. In practice, it determines both the tunability of the ECDL, and the output power available with a given diode. These parameters are inversely related, and the experimenter must choose the compromise between them that best suits their needs. Increasing the power into the first order, by choosing a blaze wavelength near the wavelength of operation, increases the influence of the external cavity over the diode cavity, and hence the extent and uniformity with which the laser can be tuned. Decreasing the power into the first order reduces the tunability, but increases the fraction of laser light that is emitted from the ECDL.

The diodes used in the ANU-DFG lasers were nominally 785 nm units, wavelength selected at 782 - 783 nm. Their integration into an external cavity was performed, predominantly, to take advantage of the ensuing linewidth reduction. The two to three nanometers of tunability required to bring them on to resonance with 780 nm was achievable with a combination of temperature tuning, current selection, and slight external cavity tuning. In addition, the entire manifold of laser cooling excited states for both  $^{85}\text{Rb}$  and  $^{87}\text{Rb}$  extends only 10 GHz, so only moder-

ate mode-hop free tuning range was required. Strong tuning of the external cavity was not a priority for these lasers, so low first order reflectance guaranteed that the lasers produced significant output power. Using a paper by Loewen et. al., the chosen parameters were selected, and several 1200 lines/mm, 450 nm blaze gratings were purchased [168]. These evinced approximately 30 % light reflected back into the first order, yielding a usable compromise between tunability and output power for the ANU-DFG ECDLs.

The construction of the potassium lasers represented the opposite extent of parameter space. Extensive tunability of the lasers would have eased the process of temperature tuning the lasers over 15 nm. However, as is often the case, these lasers simply used the gratings available in the laboratory, those purchased for the rubidium lasers.

#### 12.1.4 Optimizing tunability

There are several methods available for tuning the operation frequency of an external cavity diode such as those constructed as part of this thesis. Bulk mount temperature tuning provides a broad, but somewhat slow and imprecise means of selecting the operation wavelength. The current can tune the laser wavelength by changing the position of the gain curve in frequency space. Changing the length of the cavity, or the angle at which the grating reflects light back to the cavity, both constitute a well-controllable means of tuning the laser over a relatively broad regime. The first of these effects tunes the modes of the external cavity. The second changes the precise cavity mode that lases, due to the wavelength-selecting nature of the grating.

Each of these techniques experiences regions of continuous tuning, separated by discontinuous mode-hops, in the frequency domain. This occurs when the gradual changes of one or more characteristics of the laser cumulate to suddenly change the lasing mode favored by the balance of mode competition. Temperature tuning of free-running diodes results in middling regions of continuous tuning, perhaps several GHz, separated by mode hops of the free spectral range (FSR) of the diode, or larger [143]. Current tuning achieves much lower mode-hop free scan ranges, and suffers distinct intensity variation during tuning. Changing the length of cavity allows tuning ranges of up to half the FSR of the external cavity, typically several GHz, separated by mode hops to adjacent longitudinal modes. Carefully changing the angle of the grating simultaneous to the length adjustment performs somewhat more admirably, tuning up to half the FSR of the laser diode cavity, typically tens of GHz. This can be implemented at the design stage by extremely careful choice of pivot point in the mount that secures the grating to the ECDL base-plate.

A final method, the one employed in the ANU-DFG lasers, involves tuning the grating angle and position simultaneously with the laser drive current [169]. This technique is referred to as “feed-forward” in commercial ECDLs produced by Toptica [170]. It is able to compensate for imprecise choice of pivot point in the grating mount, yielding mode-hop free scan ranges of tens of GHz.

The configuration was implemented completely empirically, exploiting the in-

timate knowledge of the mode-hop behavior of the ANU-DFG ECDLs garnered by months of manual tuning. The piezo-drive voltage change correlated with the mode-hop free scan range of the grating was measured, and the corresponding current change required to reverse the mode hop was observed. The output from a simple circuit monitoring the current change at the laser, was filtered and amplified appropriately before being fed to the high voltage amplifiers driving the piezo. This resulted in scan ranges significantly larger than 7 GHz, the frequency extension of the  $^{87}\text{Rb}$  D2 manifold.

## 12.2 Practical considerations

The implementation of an external cavity diode laser is not a simple project. There are many aspects to consider if the the simplicity and ruggedness of commercial laser diodes is to be harnessed and improved without compromising these benefits in the process. An excellent practical solution, providing a tunable external cavity around a commercial laser diode, was published by Hansch et. al. in 1995. It used a diffraction grating as a partially reflective mirror, mounted on a aluminium bronze (AlBr) cantilever base to provide stable adjustment and good thermal properties. The diffraction grating could be turned about a pivot, changing the length of the external cavity to tune the wavelength of the laser.

### 12.2.1 Acoustic stability

The position of each part of the laser is vitally important for determining the operation wavelength of the laser. Building adjustment into the mounts allows the laser to be tuned, but it also introduces a source of instability. The static design of the laser is shown in figure 12.1. It employs a variety of materials to achieve the optimum blend of properties throughout. The primary requirement of the diode mount is excellent thermal conductivity between the diode and the ECDL base plate. The diode is push-fit mounted to a copper block, which is in turn bolted to the ECDL baseplate via a thin layer of thermally conductive paste. The baseplate itself is made of aluminum bronze (AlBr), which represents a compromise between conductivity, and machinability and flexibility for the cantilever mount that it supports. It presses a thermoelectric cooler down on to a large aluminum base, via bolts between the baseplate and base.

The grating mount is a particularly important aspect of the laser design. It is shown in figure 12.3. It must secure one end of the laser cavity, while allowing both coarse and fine adjustments of the position of the grating. It is made from a single-piece cantilever AlBr mount, bolted directly to the AlBr base plate, to provide good stability against acoustic vibrations. It employs an 80 tpi optics grade screw pushing a low voltage piezo crystal to support and adjust the cantilever. Positioning before tightening the mount bolt provides very coarse adjustment, the adjustment screw provides medium coarse adjustment, and the piezo provides fine adjustment of the grating position.

## 12.2.2 Temperature stability

As part of the laser, changes in length of the external cavity couple directly to changes in operation wavelength. This puts strict limits on both the acoustic and thermal stability of the construction methods used. Generally, the entire external cavity must be temperature controlled. This is significantly more difficult than temperature controlling an individual laser diode due to the large thermal loads and time constants presented by the macroscopic system.


A simple analysis of the temperature sensitivity of an ECDL elucidates the level of control required. As noted above, there is a fundamental requirement for  $\sim 10$  mK temperature stability due to the temperature tuning of the diode itself. This is a fairly general requirement, and there are many commercially available temperature control solutions for laser diode packages. These typically provide better than 10 mK stability using a feedback system, consisting of a thermistor, a small thermoelectric cooler and electronics to filter and scale the locking signal with time constants and gains appropriate to laser diode packages.

However, in adding an external cavity, another temperature stability requirement has entered the design. Thermal expansion of typical metals provides a significant length variation over a 3 cm long external cavity. This must be minimized through both careful choice of construction materials and active temperature control. The first of these is achieved through the use of aluminum bronze as the material from which the ECDL baseplate is constructed. Aluminum bronze has a thermal expansion coefficient of  $17 \times 10^{-6} K^{-1}$ , and is a compromise between economy, stability, and flexibility for the cantilever mount that it supports. Relating the 1 GHz frequency requirement to the length stability of the cavity yields a maximum length change of 100 nm in 3 cm. This is easily achieved with a thermal stability of 100 mK. While this appears a less stringent criterion than the diode temperature control, it is somewhat problematic because there are fewer commercially available precision temperature controllers that are designed for use with a system as large as the ECDL baseplate. Laser temperature controllers cannot generally be adjusted to provide suitable parameters for controlling a system with a thermal time constant as long as minutes. Additionally, if only one temperature controller is to be used for both the baseplate and the diode, the whole system must achieve a very precise level of stability.

## 12.3 The ANU-DFG external cavity diode lasers

The first implementation of an external cavity diode laser in the ANU-BEC group was performed as part of this thesis with assistance from an honors student with the group, Jessica Trevena. The basic design of the laser system was inherited from the Research School of Physical Sciences and Engineering at ANU. They, in turn, based their design on the original by Hansch et. al. [143].

The important specifications of the external cavity diode lasers are shown in figure 12.5. The length of the cavity is approximately 3 cm. The grating is an Edmund Industrial Optics 1200 lines / mm aluminum coated, ruled grating, blazed for use at



**12.5: External cavity diode laser specifications.** The supplier, specifications and part numbers of all important components used in the ANU ECDLs.

450 nm. The base plate and grating mount are constructed from aluminum bronze for good thermal conductivity, and good elasticity to support the cantilever mounts. The diode block is constructed from copper for excellent thermal conductivity. The diodes used were predominantly Sanyo DL7140-201S, 80 mW nom. 785 nm, single mode, wavelength selected to approximately 782 nm by Photonic Industries. Other diodes tested included Hitachi and Sharp models. These other models performed better than the Sanyo diodes, but produced lower output powers. The output of the diode was collimated using a 5mm focal length, 0.55 numerical aperture geltech lens provided by Thorlabs. The driving electronics was provided in a 19" rack mount case by the RSPHYSSE. The current driver was based on a design developed by Jan Hall at JILA [171]. The temperature controller was a LD'T-5100 OEM board provided by ILX Lightwave.

These lasers worked very well. They exhibited linewidths lower than 1 MHz, quantified via a beat-linewidth measurement. They provided a mode-hop free scan range of the order of 50 MHz using a current scan, 500 MHz to 2 GHz using a grating voltage scan, and at least 10 GHz using a combined current-grating scan. They demonstrated good day to day wavelength stability similar to the commercial lasers in the laboratory. Different lasers produced between 2 mW and 20 mW output power in single mode, although individual lasers exhibited good power stability over time. This variation was considered a result of well-known single mode performance issues with the Sanyo diodes [69]. The low efficiencies achievable in these systems was largely a result of using economical, but low quality, aluminum gratings from Edmund Scientific.

## 12.4 Considerations for improvement

As with many aspects of experiment physics, the process of designing the ECDLs used in the ANU-DFG was a combination of self-instruction and imitation of other groups' designs. As these investments have matured, it has become apparent that many aspects of the designs around which the ECDLs were originally based were not optimal.

There are several key performance issues with the ECDLs used in the ANU-DFG experiment. The first is lack of power. This is a result of two factors: very low efficiency of the grating system used, and poor performance of the diodes. Another problem is the sensitivity of the system to acoustic shock and environmental temperature changes. This results in the lasers unlocking too frequently. These problems are not absolute: the ECDL design used already manages them to some degree. However, other designs do perform better in these areas.

The amount of light reflected off the grating, back into the laser diode, determines the trade-off set between two aspects of operation: the ability of the grating to "pull" the operation wavelength of the diode, and the overall efficiency of the ECDL. The gratings used in the current ECDLs are cheap and readily available. However, they produce a fairly low efficiency because they reflect a lot of light back into the laser diode. Although this configuration yields excellent tunability, wavelength-selected 785 nm diodes generally operate close enough to 780 nm that this is not a large issue. On the other hand, the ultracold ECDLs developed in chapter 13 rely on this strong tunability to eke out the final "pull" of the external cavity. More careful selection of the gratings for their chosen purpose would lead to higher output power from the current ECDL systems.

The thermal design of the current ECDLs works, but it is not optimized. The process of creating the ultracold ECDLs elucidated many points of thermal design that were not carefully considered in the ECDL designs that were imitated for ANU-DFG. Each aspect of the thermal design: insulation, thermoelectric cooling, and enclosure, should be improved to yield better long term locking performance.

The benefit of well-implemented insulation was observed in the ultracold ECDLs, which exhibited excellent unlocked frequency stability over fairly long time scales ( $\sim 1$  hour). This characteristic implies that, when locked, the laser requires less active correction via the feedback circuitry. This in turn implies longer times between saturation of the integrated locking signal, the typical cause of long-term loss of lock.

The TEC should be selected with an understanding of the important aspects of the system which it is designed to control. This includes the expected thermal loads, plus the characteristics of the temperature controllers used to drive the TEC module. The original TEC selection was based on geometric factors only. Chapter 13 contains a detailed and general discussion on TEC selection.

The other factor of implementing TEC cooling is the expulsion of heat from the hot surface of the TEC. The current design "heat lags" the hot face of the TEC to the massive aluminum base. This provides a good thermal reservoir, but very little heatsinking. Eventually, the thermal capacity of the aluminum base is exceeded,

and the entire system starts to heat up, affecting the stability of the laser. The massive aluminum base is good for stabilizing the system from thermal transients, but should be complemented with a heatsink system.

The original enclosures were designed to prevent air currents destabilizing the laser cavity, but they are ill-considered in terms of their thermal properties. The main problem is the confinement of both the cooled ECDL base plate and the massive aluminum base, acting as a thermal reservoir, inside the enclosure. These two parts of the laser should be separated, and, ideally, insulated to isolate the cooled system from the environment. An optimum configuration towards this end is this use of a “cold plate” as described in chapter 13.

Very little consideration has been given to advanced methods of optimizing the design of the ECDLs to maximize acoustic stability. The current lasers rely on the massive aluminum base, and the inherently stable “cantilever” mounts used to create the cavity, to prevent shocks and vibrations being transmitted through the optics table to the external cavity. Rob Scholten’s group at the University of Melbourne have reported excellent performance from their ECDLs by implementing a layer of Sorbathene under a massive aluminum base. [172]

#### 12.4.1 Final thoughts

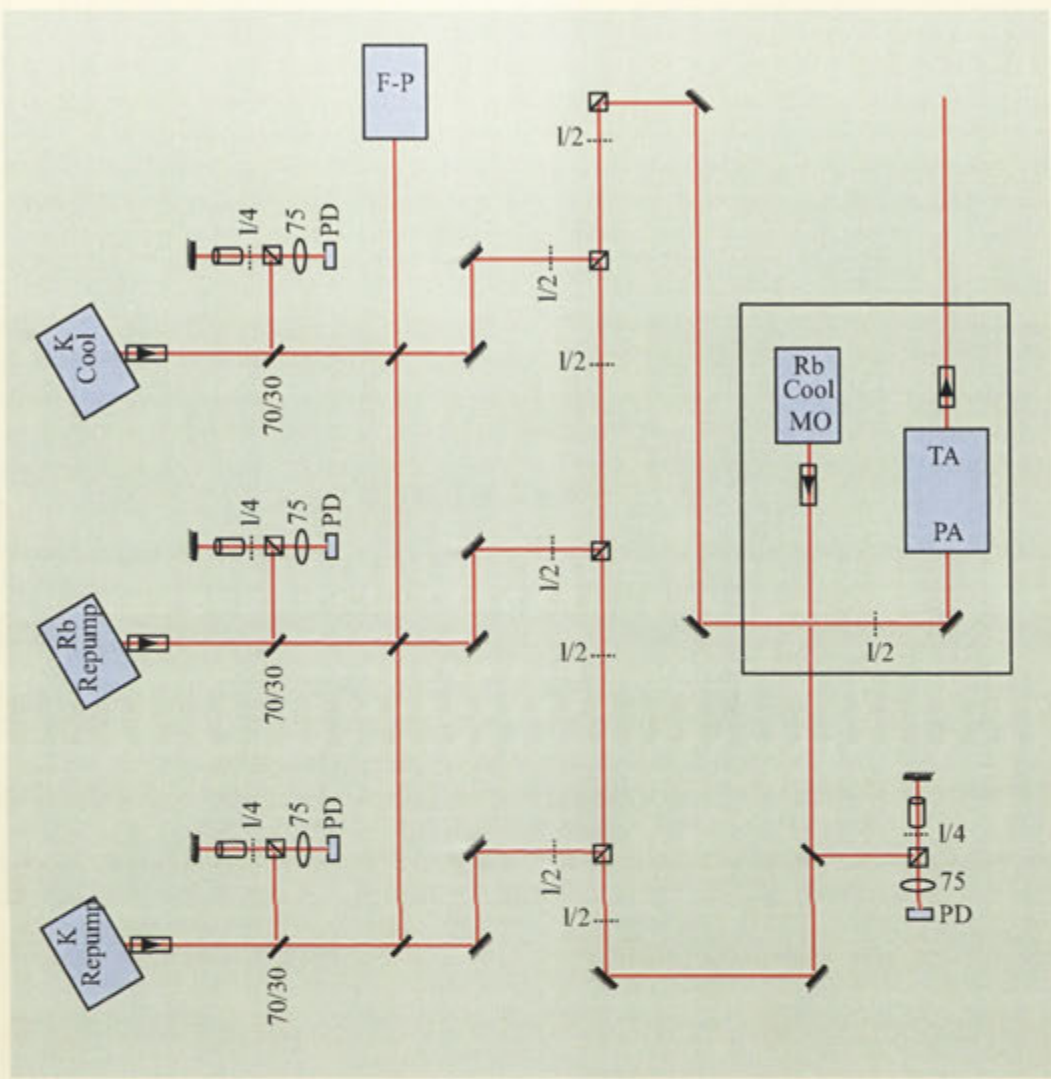
Since the first in-house ECDL was developed within the ANU-BEC group, many ECDLs have been constructed, implemented and tested as part of this thesis. Small changes have evolved through the ECDL designs. This chapter has summarized the basic principles behind the design of ECDLs in general, and considered the specific parts of the design used within the ANU-BEC laboratory. There are many aspects of assembling, aligning and maintaining these lasers that are too nebulous to report as part of this thesis. However, it was the investment in gaining these tidbits of knowledge that eventually allowed the improvement of the common ECDL design to create the extended temperature tunable lasers reported in chapter 13.

### 12.5 Seeding a tapered amplifier with multiple wavelengths

One method of overcoming the low powers produced by the ECDLs described above, is using an optical amplifier after the master laser. The optical amplifier simply amplifies the power of the laser beam created by the ECDL, without significantly perturbing the linewidth, spatial distribution, or noise characteristics of the light produced by the master laser.

An excellent optical amplifier for atom optics experiments is provided by tapered amplifier (TA) diodes [173–175]. These are simply large area diodes, antireflection coated on both the front and back facet, with a tapered gain medium. The double-ended nature of the diode allows simple injection locking in a master oscillator - power amplifier (MOPA) configuration: the light is shone into the back of the amplifier chip. The tapered gain medium allows power to build up through the diode with relatively constant power density, preventing damage to the chip itself. The





/ Mirror    / Beamsplitter    □ PBS cube    ▢ Isolator    ○ Lens    - - - Waveplate

**12.6: Injecting a tapered amplifier with multiple wavelengths.** The optical setup for the simultaneous injection of multiple wavelengths into a commercial tapered amplifier. The configuration is straightforward, if somewhat detailed and time consuming to implement. The gain curve of the tapered amplifier is very wide ( $\sim 15$  nm), allowing simultaneous amplification of both 780 nm light, for  $^{87}\text{Rb}$ , and 767 nm light for  $^{40}\text{K}$ .

high quality anti-reflection coatings ensure that the device operates predominantly via a single-pass of injected light: any laser amplification due to residual etaloning within the amplifier chip is severely suppressed due to the much larger injected light signal.

Commercial systems are available in this MOPA configuration, using a single ECDL master oscillator to seed a single tapered amplifier [170]. In fact, new designs dispense with the master oscillator completely, carefully constructing an external cavity directly about the tapered amplifier itself to provide the necessary frequency discrimination [176].

The use of tapered amplifiers as sources for atom optics and degenerate quantum gas experiments is well established [177]. The ANU-BEC laboratories contain three commercial tapered amplifier units: two in a MOPA configuration, one as a standalone TA in an external cavity. They have produced BEC, as outlined in part I of this thesis, as well as being used, over the long term, for a wide variety of atom optics experiments.

The difficulty inherent in the design of the laser systems of the ANU-DFG was due simply to the scale: up to three separate tapered amplifier systems would have been required for this experiment alone. Both the cooling light for  $^{87}\text{Rb}$  and  $^{40}\text{K}$  required amplification beyond that provided by the ECDLs. For  $^{40}\text{K}$ , the repumping light also required amplification, due to the high rate of pumping into the “dark” hyperfine state during optical cooling, because of the relatively small hyperfine splitting in the upper ( $4^2P_{3/2}$ ) manifold.

Although tapered amplifiers are economic relative to other sources of tunable, high power laser light, the purchase of three tapered amplifier systems was beyond the resources of ANU-DFG. The solution, following the work of Salomon et. al., was to inject multiple wavelengths, from multiple master oscillator ECDLs, into a single tapered amplifier [165]. In this publication, Salomon et. al. injected two wavelengths, separated by up to 12 GHz, into their tapered amplifier, carefully investigating the effect of non-linearities and power sharing due to operation of the tapered amplifier.

As a demonstration of the suitability of this type of system for the configuration required for the ANU-DFG, a tapered amplifier was used to simultaneously amplify three wavelengths. Moreover, although two of those wavelengths were within 7 GHz of each other ( $^{87}\text{Rb}$  cooling and repump light), the third was over 13 nm distant ( $^{40}\text{K}$  cooling light). This wavelength spread represented almost the outer limit of possible amplification combinations, realized by temperature tuning the TA gain curve such that the outer wings overlapped both the  $^{87}\text{Rb}$  and  $^{40}\text{K}$  transitions.

The experimental configuration is shown in figure 12.6. The various master oscillator lights are mixed with varying ratios using half-wave plates ( $\lambda/2$ ) and polarizing beam splitter cubes (PBS). Perfect co-alignment is assured by the long path length between the master oscillators and the tapered amplifier. One counterintuitive behavior of the circuit is the unavoidable loss of a significant fraction of the light at each PBS. This occurs because, although the PBS serves as a useful spatial mixer for two beams of light, the beam that is transmitted maintains an P-polarization, while the beam reflected maintains an S-polarization. Another  $\lambda/2$  plate and PBS

are used to subsequently mix the two polarizations at some tunable angle between the S and P polarization to select the relative power of each component (45 °yields 50 / 50 split).

The performance of this circuit appeared to be completely adequate for the purposes of the ANU-DFG. A  $^{87}\text{Rb}$  MOT was created using amplified light, consisting of all three components. This proved that sufficient light for both the  $^{87}\text{Rb}$  and  $^{40}\text{K}$  MOTs could be created, simultaneously, using a single tapered amplifier. The long term implementation would have aimed to include a second  $^{40}\text{K}$  master oscillator component, for the repump wavelength required by the  $^{40}\text{K}$  MOT.



# 13 Extended temperature tuning of external cavity diode lasers


Chapter 12 espoused the benefits of external cavity diode lasers. ECDLs achieve an enviable synergy between the benefits of mass-produced semiconductor lasers and the flexibility of home-built apparatus. They are capable of narrowing the spectral output of typical “single mode” laser diode to less than 1 MHz, while maintaining a tuning range of the order of  $\pm 2$  nm. With the addition of temperature tuning, the operating wavelength of an ECDL can be pulled even further. This chapter details the development of a system allowing extended temperature tuning, well beyond that normally implemented in ECDL designs. This resulted in tuning of over 15 nm, from a commercially produced 785 nm laser diode, wavelength selected near 782 nm, to the potassium transition below 767 nm.

## 13.0.1 Ultracold lasers in context

At wavelengths 10 nm or more from a commercially produced diode an inordinate amount of equipment, expertise or expense is required to harness these useful tools. At some significant cost diodes can be custom made or they can be wavelength selected from the very small number of commercially produced diodes that lase more than 10 nm from their nominal design wavelength. Mass produced diodes can be anti-reflection coated in small batches and built into an external cavity configuration to force them to lase  $\pm 5$  nm from their free-running wavelength [166].

Alternatively, the temperature of the diodes can be reduced, changing the position of the gain curve and microscopically shrinking the laser cavity, reducing the operating wavelength of the diode. Typical diode lasers have a temperature tuning coefficient of 0.25 nm/K [166]. Chapter 12 detailed the manner in which the diode is temperature controlled using a thermoelectric cooler in a feedback configuration to prevent unwanted temperature tuning of the diode. This technique is often coupled with some small amount of constant cooling, or heating, to “pull” the diode wavelength several nanometers around its design wavelength.

To pull the wavelength of a typical laser diode more than 15 nm, its temperature must be reduced approximately 60 °C. This sort of cooling has been achieved in the past by combinations of closed-cycle refrigeration, liquid nitrogen cooling, or multistage thermoelectric cooling and vacuum or nitrogen purged hermetically sealed chambers [178, 179]. These techniques are also used in other fields, such as cooled CCD arrays and cryogenic photodetectors and amplifier electronics. These techniques are relatively complicated, expensive and potentially unreliable, defeating many of the benefits of using cheap, rugged laser diodes.



**13.1: An ultracold diode laser system.** The key components of an ultracold system are the insulation, the cooling and the sealed laser enclosure.

### 13.0.2 Ultracold lasers in the DFG experiment

Figure 13.1 shows an “ultracold” diode laser system, although the basic components are the same in all efficient cooling devices: the cooling system, the insulation and the sealed enclosure. Each of these components has been implemented in various ways in other ultracold devices. However, the discussion below shows that, by judicious design and understanding of the system, it is possible to construct a simple, cheap, ultracold diode laser using solid-state components requiring low maintenance. The method of implementing each component of the ultracold system is outlined below. Quantitative background on each of these components is provided, allowing the extension of the design to more general systems. The final section of the chapter shows how these principles were successfully applied to produce an ultracold external cavity diode laser that operates more than 15 nm below its room temperature wavelength.

## 13.1 Fundamental and practical considerations

Thermoelectric coolers are heat pumps that move heat from their cold face to their hot face. If the hot face is attached to a heatsink held near room temperature, the cold face can fall up to 70 °K below the ambient temperature, although more typical experimental values are of the order of 40 °K . In order to reduce the

temperature of a laser diode by more than 60 °K , as well as removing heat generated in the diode, it is necessary to use multiple stages of TEC cooling.

In many ways, TECs are the ideal device for cooling and temperature stabilizing a laser diode. They are small, cheap, cause no vibration and are reliable and maintenance free. Unfortunately they are also very inefficient, creating many times as much heat as they can transport away from their cooled surface. Removing this large amount of heat from the system requires very careful design of the entire TEC system, including selection of specific TECs, design and use of “cold plates” to allow sufficient insulation between stages, and selection of the heatsink. These factors are modeled roughly and considered further in sections 13.1.1 and 13.1.2.

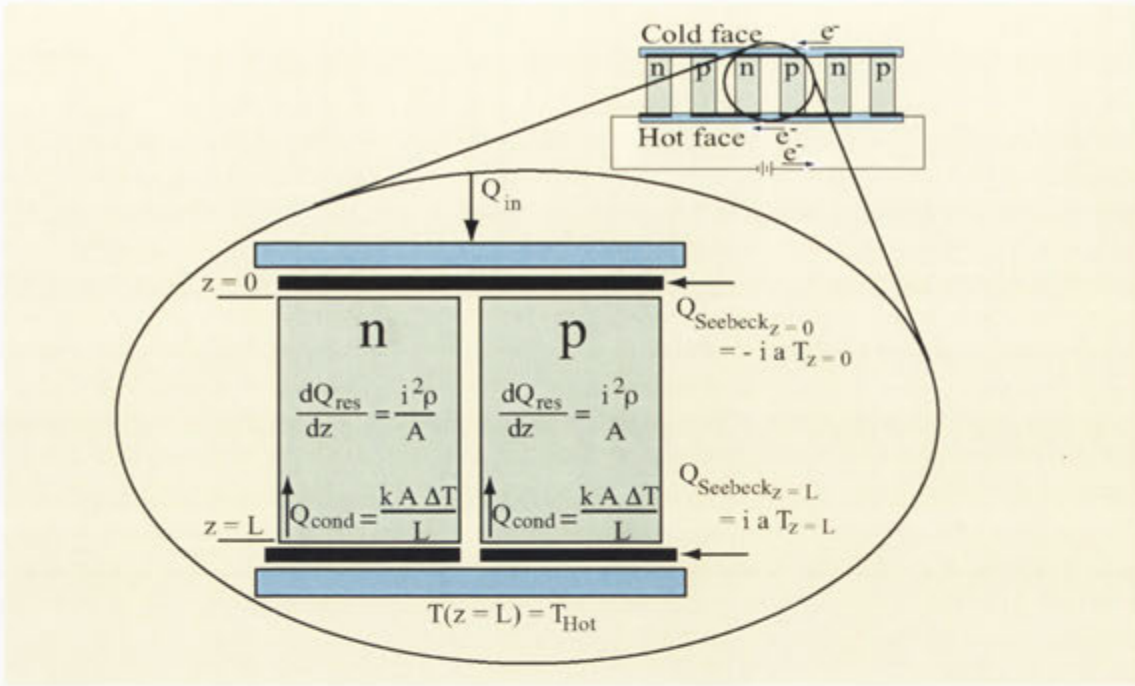
Any unnecessary coupling of heat in the system is bad, but unnecessary coupling at the coldest stage is the most wasteful because it must then be pumped out of the system, and magnified many times, by the inefficient TECs. The laser diodes used in the ANU laser cooling and trapping experiments typically dissipate less than 200 mW as an active heat load, but the passive heat flow from the surrounding laboratory, due mostly to convection, can act as a much larger heat load. Insulation is necessary to reduce these passive heat loads to reasonable values. Some ultracold systems have been built in vacuum chambers to minimize passive convective heat transfer, but section 13.1.3 illustrates how simple design and careful construction can yield a straight-forward polystyrene insulation scheme that is sufficient for the ANU-DFG ultracold system.

As an ultracold system is cooled below room temperature, condensation will form on the coldest parts of the system, which as the temperature drops becomes ice. In more complicated ultracold systems this too has been overcome by enclosing the laser in a vacuum chamber. Section 13.1.4 demonstrates that for this sort of laser, a simple IP65 neoprene-sealed cast aluminum jiffy box is a far cheaper and simpler solution.

### 13.1.1 Thermoelectric coolers

TECs are very inefficient, and the design of a multiple TEC stack is critical in ensuring any sort of useful operation of a given system. Unfortunately, manufacturer’s handbooks and web sites seem biased towards engineering issues such as efficiency of single TECs, or specific extrema of performance, such as maximum temperature differential with zero heat load [180, 181]. Published works on multi-TEC stacks are understandably general in nature, and often interested in engineering figures of merit such as cost of construction and running cost. However, by generating a simple model, it is possible to perform “custom” optimizations on specific systems, such as a multi-stage TEC stack for ultracold ECDLs, or for choosing a TEC under non-standard limiting conditions, e.g. for use with a specific temperature controller.

Very little information seems to be available on modeling and choosing TECs for a given application, especially for multi-TEC stack configurations. All TEC manufacturers use semiconductors with similar performance at room temperature. Any given TEC has well-defined optimum voltage and current for any given configuration. The only options available to the user buying TECs off-the-shelf are the



**13.2: Heat flows in a TEC.** The pertinent heat flows in a single junction of a TEC include those due to resistive, conductive and thermoelectric effects. The thermal equations noted below apply in the middle region of the TEC - the localized Seebeck heating at each face is applied purely as a boundary condition to the standard equations.

“shape factor”,  $G = \text{Area} / \text{Length}$  of individual thermocouples (0.01 - 1 cm), and  $n$ , the number of thermocouples in the TEC (20 - 200). However, even with stacks of only two TECs, the optimization of  $n$  and  $G$  for each TEC becomes non intuitive and must be modeled. Selection criteria for these parameters are derived below.

### 13.1.2 A Thermal Model

Figure 13.2 shows a thermal model for a single thermocouple system. The thermoelectric effect is quantified by the Seebeck coefficients of the semiconductors,  $a$ , which for commercially used materials is of the order of  $2 \times 10^{-4} \text{VK}^{-1}$ . Other heating effects inherent in the system are resistive heating inside the semiconductors, quantified by the resistivity,  $\rho$ , and heat flow through the semiconductors due to the temperature gradient created, quantified by the conductivity,  $\kappa$ . Equation 13.1 shows the 1D thermal model with boundary conditions.



$$\begin{aligned}
& -\kappa GL \frac{\partial^2 T}{\partial z^2} = \frac{i^2 \rho}{GL} \\
B.C.1 : & -\kappa GL \frac{\partial T}{\partial z} \Big|_{z=0} = Q \Big|_{z=0} = Q_{in} - iaT_c \\
B.C.2 : & T \Big|_{z=L} = T_H
\end{aligned} \tag{13.1}$$

In the equations above,  $L$  is the length of the semiconductor junction,  $i$  is the current flowing through the junction, and  $Q_{in}$  is the heat load at the cold face of the thermoelectric cooler. It is straight forward to solve this model and find the temperature at the cold side of the TEC in terms of the TEC parameters, the heat load on the system,  $Q_{in}$ , and the hot side temperature,  $T_H$ . Note that equation 13.2 extends the model above to  $2n$  junctions, as commercial TECs are typically constructed from 20 - 200 semiconductor junctions.


$$T_C = T \Big|_{z=0} = \frac{T_H + \frac{i^2 \rho}{2\kappa G^2} + \frac{Q_{in}}{2n\kappa G}}{1 + \frac{ia}{\kappa G}} \tag{13.2}$$

Very quickly, it is possible to make some general statements about the design of TEC-based systems by investigating equation 13.2 in various limiting cases. The heat that can be pumped from a system attached to the TEC cold face is  $Q_{in}$ . The heat that a TEC expels from its hot face is  $Q_{out}$ .  $T_C$  for a single, insulated ( $Q_{in} = 0$ ) TEC attached to a perfect heatsink ( $T_H$  fixed) is minimized with a current of  $i_{max}$ . Although this may seem an unlikely design criterion, it represents an upper bound on the optimum operating current in real systems and is widely used as a TEC specification by manufacturers.

$$\begin{aligned}
Q_{in} &= 2n(iaT_C - \frac{i^2 \rho}{2G} - \kappa G \Delta T) \\
Q_{out} &= Q_{in} + \frac{2ni^2 \rho}{G} + 2nia\Delta T \\
i_{max} &= \frac{G\kappa}{a} \left( \sqrt{1 + \frac{2a^2}{\rho\kappa} T_H} - 1 \right)
\end{aligned} \tag{13.3}$$

These equations contain information useful in developing an intuitive understanding of how systems using TECs should be designed.  $Q_{in}$  is maximized when  $\Delta T = 0$ .  $Q_{out}$  is simply  $Q_{in}$  plus the resistive heating inside the semiconductor legs of the TEC, plus the difference between the heat absorbed by the Seebeck effect at the cold face and the *heat liberated at the hot face*. The heat liberated at the hot face does not enter equation 13.2 above because it assumes a perfect heatsink that can maintain  $T_H$  independent of heat load, but the more advanced models discussed below do require consideration of this heat load.

Ultracold systems push the limits on the lowest temperatures that can be achieved using simple TEC stacks and polystyrene insulation. One measure of

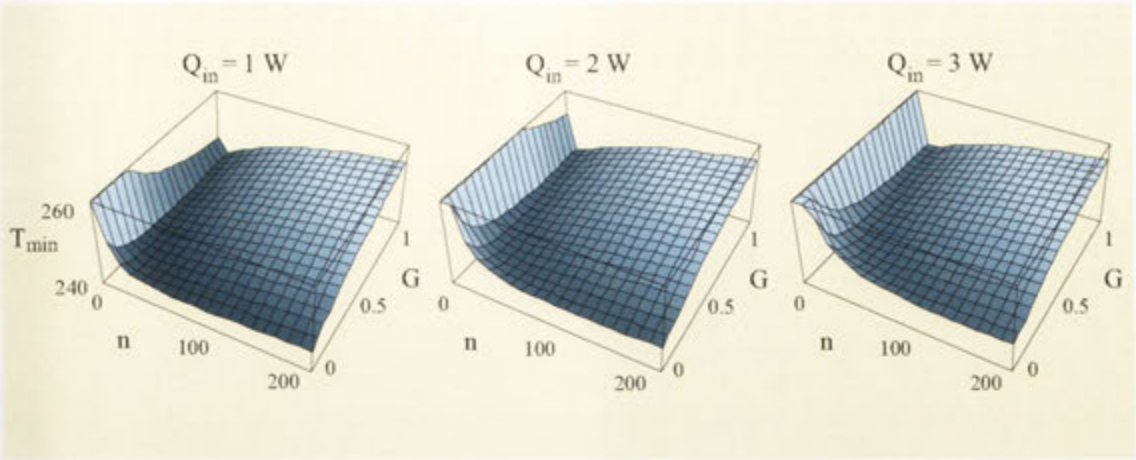


**13.3: Performance of a TEC lagged to a reservoir.**  $T_{min}$  versus  $n$  and  $G$  for a series of system configurations with fixed  $T_{Hot} = 300$  °K .

how well a given design performs is the lowest achievable temperature of the system. Equation 13.2 yields this quantity directly for given choosable TEC parameters ( $n$  and  $G$ ), given system parameters ( $Q_{in}$  and  $T_{min}$ ), and a given operating condition ( $i$ ). Equations 13.3 yield other important information such as the optimum operating current and the amount of heat output from the system.

Figure 13.3 shows a series of plots of  $T_{min}$  at  $i_{opt}$  (similar to  $i_{max}$ , but rederived to account for  $Q_{in}$ ) versus  $n$  and  $G$ , at various  $Q_{in}$  typical of the ANU-DFG system. The key result to note, is that except for the very smallest TECs (a few, small thermocouples = low  $n$ , low  $G$ ), most selections of TEC will yield the theoretical best  $T_{min}$ , approximately 70 °K below room temperature. This indicates that the expected loads of  $\sim 1W$  of heating (investigated more fully in section 13.1.3) are low compared to the design specification of most commercially available TECs. However, further consideration shows that simply selecting the largest TEC available, to ensure that it can maintain  $\Delta T \sim 70$  °K, is not an intelligent course. There is more to the calculation, as shown below.

An important example of the necessity of deriving the TEC equations above, is the extension of equation 13.2 to include a real heatsink. As it stands, equation 13.2 assumes that  $T_H$  is “set” by attaching the TEC to a reservoir, in practice some sort of heat sink. Heatsinks come in many shapes and sizes. Typically, one thinks of a finned block of black aluminium, which loses heat to its surrounding environment “efficiently” due to predominantly convective losses from its large surface area. An efficient heatsink is one that can be only slightly hotter than the surrounding environment while losing a lot of heat due to this temperature differential. This property can be quantified as “thermal resistance”,  $\Theta[^\circ\text{K}/\text{W}]$ , which yields a number describing how much hotter the heatsink must get to dissipate an extra watt of power. A reservoir, or ideal heatsink, has  $\Theta = 0^\circ\text{K}/\text{W}$  so that it can maintain “room” temperature independent of the power it must dissipate. A normal, passive



**13.4: Performance of a TEC lagged to a heatsink.**  $T_{min}$  versus  $n$  and  $G$  for a series of system configurations using a real heatsink with thermal resistance  $\Theta = 0.3^\circ\text{K/W}$  with a room temperature of 300 K.


aluminium heatsink, as described above, has  $\Theta \sim 1^\circ\text{K/W}$ , at best. A standard forced air aluminium heatsink can have as low as  $\Theta \sim 0.3^\circ\text{K/W}$  - more advanced designs can perform even better.

These considerations are actually vitally important, as a well-designed multi-stage TEC system rapidly ends up limited by the heatsinking arrangement implemented. If the heatsink temperature rises due to the heat being dissipated, it pulls  $T_H$  up with it. Since a TEC, at best, maintains a temperature differential,  $\Delta T$ , across its faces,  $T_C$  rises correspondingly. Hence, the thermal resistance of the heatsink directly determines how cold the system can get. Increasing the size and power of a TEC in a heatsink limited system increases  $T_{min}$ .

Including a heatsink in the TEC model is easy using a mathematical software package, such as Mathematica. Instead of  $T_H$  being fixed, it becomes dependent on  $Q_{out}$ , that is,  $T_H = T_{room} + \Theta Q_{out}$  for a heatsink with thermal resistance  $\Theta$ .

Figure 13.4 shows plots similar to figure 13.3, but includes the effect of using a real heatsink. Obviously this effect is very important in selecting the TEC for an ultracold application. Even though figure 13.3 indicated that any “large” TEC would reach the theoretical lowest possible temperature, figure 13.4 shows that the extra heat load generated by a large TEC ends up raising the temperature of the heatsink so much that it degrades the final performance of the system as a whole. There is a clear optimum TEC for the system, albeit quite weakly dependent on the shape factor,  $G$ , and number of the thermocouples,  $n$ . This relative insensitivity to  $n$  can permit the impedance of a TEC (drive voltage/current) to be matched to the available TEC power supplies.

Figure 13.5 shows how TECs can be run in series in a “multistage” configuration to provide greater temperature drops than a single unit. Their inefficiency puts severe limitations on how many stages should be used, because each extra stage



**13.5: Multiple TECs stacked in series.** Multiple TECs stacked in series to provide total temperature drops of greater than  $70\text{ }^\circ\text{K}$ . Note the use of cold plates to allow extra thickness of insulation between TEC stages (insulation not shown).

must remove both the heat pumped, and more significantly, the heat generated by the previous TEC. Of course, the final stage of the system, the heatsink, must dissipate all this heat efficiently to yield the lowest possible temperature at the cold point of the system. Any unnecessary inefficiency in the choice of the number of stages, or the type of TEC to use in each stage of the system, will directly impact the ultimate achievable temperature.

The theory presented above can be extended to multiple stages of TEC cooling. Equation 13.2 is set up for each stage. Assuming the TECs are connected by negligible thermal resistance, and there is suitable insulation from external heat loads, all the heat from the output of a given TEC stage must be transmitted to the input of the next stage. Similarly, the output of a given TEC stage should be at the same temperature as the input of the next TEC stage. The equations are straight-forward but can be quite detailed and are best dealt with in a mathematical software package.

Optimization of a multi-stage system is somewhat difficult due to the size of the parameter space available to probe. However, in many realistic experimental situations, not all parameters will be truly free; it may be necessary to operate with a given TEC power supply that will limit both the drive current available and the number of thermocouples in a given TEC. In practice, using some combination of the general design techniques for individual TECs elucidated above, integrating the limitations imposed by a specific experiment, and then optimizing over the remaining available parameters computationally, yields a workable technique for selecting TECs in multi-stage configurations.

### 13.1.3 Insulation

Section 13.1.2 showed that a compromise existed between the ultimate cold face temperature and the heat pumped through the cold face of a given TEC, that TECs are very inefficient, producing much more heat than they pump, and that most well-designed multi-stage TEC systems were limited by the heat dissipated at the heatsink of the system. Clearly, the lowest ultimate operating temperature of a TEC system will be very strongly affected by the amount of heat being pumped at the coldest face in the system.

The laser diodes used in the ANU laser cooling and trapping experiments typically dissipate less than 200 mW as an active heat load, which makes them ideal for use in a thermoelectric cooler system. However, the passive heat flow from the surrounding laboratory to the significantly colder laser diode, predominantly due to convection heating via air, can act as a much larger heat load. Preventing heat transfer by convection is part of the reason ultracold systems are sometimes built in vacuum chambers.

The passive heat loads in a real system can be estimated quantitatively from the very simple formulae in equation 13.4 for heat flow due to radiation, conduction and convection [180].


$$\begin{aligned} Q_{conv} &= hA\Delta T \\ Q_{cond} &= \frac{\kappa A\Delta T}{\Delta z} \\ Q_{rad} &= \epsilon\sigma A(T_{hot}^4 - T_{cold}^4) \end{aligned} \tag{13.4}$$

Typical values for the constants are  $h = 10\text{Wm}^{-2}\text{K}^{-1}$  (rough value for still air),  $\kappa = 386\text{Wm}^{-1}\text{K}^{-1}$  (copper),  $\epsilon = 1$  (worst case),  $\sigma = 5.67 \times 10^{-8}\text{Wm}^{-2}\text{K}^{-4}$ . The value of  $A$  is the exposed surface area of the cold device,  $\Delta z$  is the thickness or length of the conducting medium.

For a 2 cm x 2 cm x 2 cm copper block exposed on five faces, cooled 60 °C below a room temperature of 20 °C, connected to six 30 AWG copper wires at room temperature, the loads are approximately 1.2 W, 0.001 W, and 0.5 W respectively. Obviously, the convective and radiative loads are a severe limitation on the operation of the thermoelectric cooler system.

These loads can be greatly reduced by the judicious use of insulation around the cooled components. Although this concept is straightforward, its importance cannot be overstated. Proper design and implementation of the the insulation around an ultracold laser is critical, and small changes can change the ultimate minimum temperature of the system by more than 20 °K. However, it is vital to note, that by a combination of following some general design principles, calculated below, and using careful but simple construction techniques, it is possible to quickly, ruggedly, and extremely economically transform a standard ECDL into an ultracold laser lasing 15 nm below its design wavelength.

The amount of insulation required can be estimated using the formulae above. By cladding the ultracold diode laser in polystyrene insulation (white, closed cell



**13.6: Typical use of a cold plate** The use of a cold plate to increase the distance between the hot and cold surfaces either side of a TEC is a well known and vital part of TEC system design. In an ultracold ECDL the straightforward implementation of this simple idea is critical to the optimum performance of the cooling system.

foam, conductivity  $0.04\text{Wm}^{-1}\text{K}^{-1}$ ) the convective and radiative loads calculated above are essentially removed, replaced by a conductive heat load through the insulation. The amount of heat passing through a slab of polystyrene of area  $10\text{cm}^2$ ,  $t\text{cm}$  thick, due to a heat gradient of  $60\text{ }^\circ\text{K}$ , can be calculated to be approximately  $0.5/t\text{W}$ . This shows that simply cladding the coldest part of the diode laser in a couple of centimeters of polystyrene insulation will reduce the passive loads on the system to the order of magnitude of the active load. More accurate modeling of the heat flows in the system is simply not required, due to the experimental realities of constructing such a system. However, it is important to note the effect of several difficulties faced when constructing such a system.

Typical TECs are approximately 3 mm thick, and can hold a temperature differential of up to  $40\text{ }^\circ\text{K}$  in a realistic experimental setup. If  $10\text{ cm}^2$  hot and cold surfaces attached to such a TEC are insulated with 3 mm of polystyrene, the conduction between the surfaces can be calculated from the equations above to be approximately 0.5 W. Clearly, this is another significant load on the cold face of the system, but again, careful design of the insulation is sufficient to minimize the problem. "Cold plates", such as those shown in figure 13.6, are copper or aluminum spacers, the same size as the face of the TEC, that allow more insulation to be added between the cold and hot faces of the system. Two 1 cm cold plates in the above example, allowing 23 mm of insulation, reduces the load on the system to less than 70 mW, which is much less than the active load of the diode itself.

As desirable as it might be, it is usually impossible to encase an ultracold system completely in insulation. An ultracold diode laser needs electrical wires attached to the diode itself, and a hole from the diode, through the insulation, to let light

out of the system. Electrical wires for passing low-power signals into the system are normally not problematic - by using very fine gauge wires the conduction heat load into the system can be kept small (e.g. 0.01 W for six 30 AWG wires, as calculated above), and even very fine gauge wires have low resistive heating - 10 cm of 30 AWG copper wire has a resistance of the order of  $0.05\Omega$ , so could pass an amp of current with negligible heating. Holes in the insulation connecting the cold face of the system with the outside world are a serious problem, however. In practice this can end up a significant limitation on the performance of the cooling system, so holes should be kept to an absolute minimum, both in size and number.

### 13.1.4 Sealed Enclosure


As an ultracold system is cooled below room temperature, condensation will form on the coldest parts of the system, which, as the temperature drops, becomes ice. Although this rarely permanently damages a hermetically sealed laser diode package, it can impede the operation of an external cavity laser diode system if there is a large build up of ice that scatters light from the main beam path.

Water will tend to condense on the first part of the laser enclosure that drops below the dew point of the system as the laser is cooled. In the case of the ANU-DFG ultracold lasers, this is not the laser diode itself, but the copper block in which the diode laser is mounted, and any exposed parts of the TEC. However, on an average  $20^{\circ}\text{C}$  day with a relative humidity of 50 percent, the room air will contain about  $8.7\text{gm}^{-3}$  of water. If a 10 cm by 10 cm by 10 cm enclosure is sealed under these conditions, it will contain approximately  $10^{-2}\text{g}$  of water. As ice, this will occupy a volume of less than  $5\text{mm}^3$ . By ensuring that the first point in the enclosure to pass the dew point is a relatively large area away from the direct path of the laser light - in this case, the 10 mm x 10 mm top of laser diode block, it is possible to guarantee that condensation will not significantly effect the operation of the laser.

This fact affects the design of the ultracold laser in one important way. The enclosure used around the laser must be relatively air-tight. The pressure of water vapor inside the enclosure drops when it condenses as the laser is cooled, and unless the enclosure is sealed, water vapor will be pumped from the surrounding room air into the chamber. This will cause a large amount of ice to form inside the laser system, eventually impeding its operation. In practice, a simple die-cast aluminum jiffy box, sealed to IP65, gives many months of hassle-free operation.

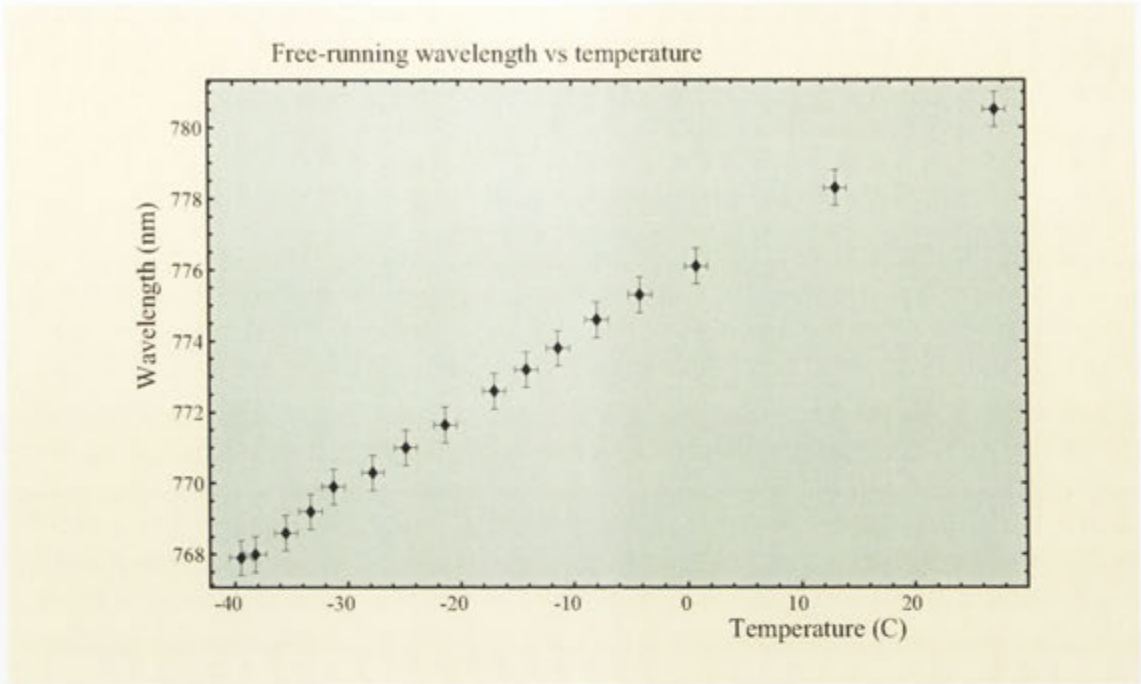
## 13.2 A Solid-State, Economical Ultra-Cold Diode Laser

Figure 13.7 shows a diagram of the ultracold external cavity diode laser system constructed using the principles outlined throughout this chapter. It uses a \$ 50 Sanyo DL-7140-201S laser diode (80 mW), wavelength selected at 782 nm, in the standard ECDL configuration discussed in chapter 12. The cooling system is comprised of two thermoelectric coolers (TECs) in a multi-stage configuration, utilizing



**13.7: The ultracold external cavity diode laser system.** The key components of an ultracold diode laser system are the insulation, the multi-stage thermoelectric cooling and the hermetically sealed laser enclosure.





**13.8: Temperature dependance of the free-running wavelength.** The temperature dependance of the free-running wavelength produced by a nominally 782 nm diode (at room temperature), as it is cooled to  $-40^{\circ}\text{C}$ .

cold plates and a pair of fan-forced heatsinks to provide robust and quiet cooling of the ECDL. The insulation is constructed from 6 pieces of polystyrene, carefully cut to provide snug-fitting insulation between all components and the enclosure. The enclosure is a simple die-cast aluminum “jiffy” box, sealed with a neoprene gasket to IP65.

The exploded diagram in figure 13.7 shows all the parts that are mounted inside the enclosure. Components are inserted and coarsely aligned at room-temperature before the top-most insulation is inserted and the enclosure sealed. Once the laser is cooled, fine adjustments of both the horizontal and vertical cantilever mounts can be made via optics-grade 80 tpi screws using the insulated hex keys mounted in the enclosure. The grating cantilever mount can be adjusted even more precisely by applying a voltage to a piezoelectric transducer inside the mount. This amount of adjustment is sufficient to retune the laser from its room temperature configuration to allow ultracold operation.

Figure 13.8 shows the dependance of the diode laser operating wavelength on temperature. The wavelength was measured using a fiber-coupled ANDO AQ-6315A optical spectrum analyzer. The resistance of a thermistor mounted in the copper diode cooling block was measured and converted to a temperature using the manufacturer’s specifications. Note that the figure shows data collected from a free-running laser diode, as opposed to one in an external cavity configuration. Figure


13.10 shows a saturated absorption spectrum [167] of the D2 transition of potassium at 766.7 nm, collected using a cooled diode laser in an external cavity configuration. The data was collected on an uncalibrated photodiode as the laser wavelength was scanned by tuning the voltage across a piezoelectric transducer attached to the grating of the external cavity. This data shows stable, narrow linewidth operation of a diode laser at 766.7 nm, the wavelength required for cooling and trapping experiments on potassium.

The key design specification of the cooling system was the ability to cool to  $-40^{\circ}\text{C}$  from a room temperature of  $20^{\circ}\text{C}$ . There were several limitations on the design peculiar to the ANU-DFG experiment. It was necessary to use an existing, small capacity temperature controller built into the laser controller, for the cold TEC. The two hex screws in the external cavity diode laser mount needed to be accessible for alignment purposes, and, obviously, it was important to get the laser light out of the ultracold system. Also, as these enhancements were being retrofitted to an existing ECDL, the physical size of the TECs had to be consistent with the rest of the laser. The design was based around two TECs, a small “cold” TEC that cooled the laser diode block, and a much larger “hot” TEC that cooled the ECDL base plate.

### 13.2.1 Cooling

A general model for a two stage TEC system attached to a real heatsink was constructed, as outlined in section 13.1.2, yielding an equation for  $T_C$  in terms of the system parameters ( $Q_{in}$  and  $\Theta_{heatsink}$ ), the choosable TEC parameters ( $n_{Hot}$ ,  $n_{Cold}$ ,  $G_{Hot}$  and  $G_{Cold}$ ), and the operating points ( $i_{Hot}$  and  $i_{Cold}$ ). The heat load on the cold face of the system (with the insulation outlined below) was estimated at somewhere around 0.5 - 1 W, and the thermal resistance of the implemented heatsinks was specified as approximately 0.4K/W. The available temperature controller supplied 2 A at 2 V, and this severely limited the choice of the cold TEC and its operating point. In addition, the power supply for driving the hot TEC was relatively modest, supplying 4 A at up to 20 V. The model that was left to optimize involved the selection of the hot TEC ( $n_{Hot}$  and  $G_{Hot}$ ) within the accessible current ranges. This parameter space is quite easy to optimize, and figure 13.9 shows a plot of the results of the model.

The TEC chosen was close to the optimum value of  $G$  (within the choices available), although the relative insensitivity to  $n$  enabled the drive impedance to be tuned to the power supply. The final cooling system design obeys the general principles expected for a multi-stage TEC system, within the specific requirements of this experiment. The “hotter” TEC ( $n_{Hot} = 67$ ,  $G_{Hot} = 0.118$ , 40 mm square, running at 60 W) has significantly larger capacity to transport the heat generated at the laser diode as well as the heat generated internally by the “colder” TEC ( $n_{Cold} = 31$ ,  $G_{Cold} = 0.08$ , 15 mm square, running at 4 W). The upper TEC is controlled by the laser controller in a feedback loop governed by a thermistor mounted in the laser diode copper block, but the lower TEC free-runs at constant current. The heat sinks are the largest useful models available:  $0.4^{\circ}\text{C}/\text{W}$  “tube” heatsinks



**13.9: Minimum temperature versus  $n_{Hot}$  and  $G_{Hot}$ .** Minimum temperature versus  $n_{Hot}$  and  $G_{Hot}$  for our real ultracold diode laser system, including all its limitations. There is a clear minimum at  $G \sim 0.1$ .

cooled by forced convection provided by four 80 mm 12 V fans. They typically sit only 5 - 10 °C above ambient room temperature. The cold face of the system can reach as low as -45°C.

### 13.2.2 Insulation

The construction of the insulation followed the straightforward concepts outlined in section 13.1.3, but the specific implementation and quality of the assembly had dramatic effects on the performance of the ultracold ECDL. The enclosure was chosen to provide a reasonable amount (several centimeters) of insulation around the laser components. Multiple pieces of insulation were cut, and often re-cut, to provide a snug fit between the enclosure and various components. The consequent piecing together of the 3-dimensional foam-and-laser jigsaw puzzle was a “once-only” event. Dismantling the laser inevitably required re-cutting some foam pieces to ensure a snug fit. However, careful construction paid off; a well constructed ultracold laser would reach temperatures up to 15 °K below a poorly constructed specimen.

An ultracold ECDL may produce light 15 nm below its design wavelength, but unless you can get that light out of the system it is of little use. The ANU-DFG ultracold ECDL had the smallest possible hole in the insulation that would still allow

coarse wavelength tuning using the grating, approximately 8 mm in diameter and 50 mm long. Even slightly larger holes (12 mm diameter), significantly increased the ultimate minimum temperature at the laser diode ( 5 K).


The ability to quickly, easily, and, most of all, accurately and repeatably cut polystyrene, significantly affected the quality and performance of the finished ultracold laser. A foam cutter board was constructed by passing a few amps of current through a 22 AWG nichrome wire perpendicular to a board similar to a draughtsman's table that enabled quick, square, accurate cuts that fit together tightly to produce a well insulated system.

ECDLs require some tweaking, typically through hex key driven adjustment screws mounted in the baseplate and grating mount. Normal steel hex keys would provide a large heat path into the cooled laser, so an insulating hex key was developed using two halves of a single key separated by a length of 8 mm diameter perspex rod. It passed through the die-cast aluminum enclosure via standard IP65 cable glands. These keys were simple to construct and insulated the laser so effectively that they could be left attached to the system permanently without affecting the cooling performance. In this configuration the end of the hex key attached to the system is also cooled. This is vital, as the thermal shock induced by touching a room temperature hex key to the laser induced so much frequency drift that it became impossible to tune the laser effectively.

Electrical signals were passed into the laser enclosure through standard multi-core cables inserted via a standard IP65 cable gland. Once inside the enclosure, the wires were soldered to 30 AWG enameled copper wire that passes through the insulation to the cold face of the system. In addition, the cables that must be attached to the ultracold copper block that houses the diode can be lagged to the less cold ECDL baseplate to ensure that no unnecessary heat is transferred into the ultracold region.

### 13.2.3 Enclosure

The size of the enclosure was roughly selected based on readily available models and the rough size of the laser and insulation. Most electronics suppliers carry a line of enclosures fitted with neoprene gaskets to provide dust-and-moisture sealing to IP65. They also supply standard "cable glands", plastic feed-throughs with a highly compressible o-ring that clamps down on the cable passing through it when tightened. Two of these cable-glands were used to pass electric wires into the enclosure; another two were used to allow the perspex allen keys to pass into the enclosure, and even be turned, without disrupting the sealing. The laser light was passed out of the system by gluing an anti-reflection coated window over a hole drilled in the enclosure and sealing it with silicone adhesive. Similarly, the thermoelectric coolers were attached to the heatsinks via a cold plate glued with silicone into a precisely cut hole in the base of the enclosure.



**13.10: Saturated absorption spectrum of potassium at 766.7 nm.**

The diode used to produce this spectrum ran at 782 nm at a room temperature of 20°C, but has been cooled, in an external cavity configuration, to below -40°C to reduce its wavelength. The spectrum was collected using a vapor cell containing standard abundance potassium vapor.

### 13.3 Summary

These experiments demonstrated a simple, cheap, and rugged design for pulling the wavelength of a laser diode in an external cavity diode laser 15 nm below its nominal design wavelength. Because the design uses only carefully designed polystyrene insulation, a sealed die-cast aluminum box and one more TEC than a traditional ECDL, it maintains the inherent strength of the original ECDL design: a simple, rugged home-built structure that harnesses the cheap, easily integrated nature of mass produced telecommunications laser diodes to provide the sensitive optoelectronics of a laser. This makes it superior to complicated systems based on liquid nitrogen cooling or vacuum sealed cooled laser heads for many uses in atom optics and other physics laboratories. A nominally 782 nm laser diode was cooled to below -40°C, and used to perform saturated absorption spectroscopy on the potassium D2 transition at 766.7 nm, as shown in figure 13.10. These lasers represent an economic, elegant solution for laser spectroscopy or degenerate quantum gas experiments working with potassium.



## 14 The self locked MOT

Potassium 40 is a rare, radioactive isotope, accounting for less than 0.05 % of a natural potassium sample. DeMarco et. al. illustrated the construction of a dispenser of potassium in the chloride salt form, enabling production of  $^{40}\text{K}$  enriched to the 5 % isotopic abundance level [70]. This was a key step in producing a DFG of potassium: it increased MOT capture numbers by four orders of magnitude, to approximately  $10^8$ .

### 14.0.1 The SL-MOT in context


MOTs cool and trap a sample of gas to temperatures useful for atom optics experiments. However, a separate atomic sample is generally required for stabilization of the laser frequency. Normally this is achieved by performing saturated absorption spectroscopy on a room temperature vapor cell containing the species being cooled. This process creates naturally broadened locking signals absolutely stabilized to the pertinent atomic transitions.

It is not easy to construct a vapor cell of enriched  $^{40}\text{K}$ . Vapor cells are typically constructed using macroscopic amounts of the required species, and operated at near-room-temperature vapor pressure. As a result, the two other DFG experiments working with  $^{40}\text{K}$  do not lock directly to a  $^{40}\text{K}$  transition. Instead, they lock to transitions in the more abundant isotopes: Jin et. al. use DAVLL to lock at a tunable offset to the room temperature Doppler broadened distribution [69]; while Inguscio et. al. use saturated absorption spectroscopy to lock directly to naturally broadened transitions, then shift the frequency of their light to the  $^{40}\text{K}$  transitions using a complicated circuit of AOMs [164].

Both of these techniques have their strengths and their weaknesses. Jin's method achieves economy and tunability at the price of questionable absolute stability. Inguscio's method is absolutely stable and extensively used for other species, but requires large expense to subsequently tune the lasers back to the  $^{40}\text{K}$  transitions with AOMs.

### 14.0.2 The SL-MOT in brief

As part of the work towards ANU-DFG, a new method of frequency stabilization for laser cooling was developed: the self-locked magneto-optic trap (SL-MOT). It is a novel technique to lock the frequency of the cooling laser of a magneto-optic trap using a signal generated from the ultracold atoms in the MOT itself. This technique utilizes the large number of very cold, slow-moving atoms in the MOT, to generate a signal free of Doppler-broadening. A system analogous to Pound-Drever-



**14.1: Saturated absorption spectroscopy.** Saturation of atomic transitions using a pump-probe beam configuration. The figure on the left shows the configuration used in the ANU-BEC laboratories. The figure on the right shows a more traditional layout, in which the pump and probe beams are clearly distinguished.

Hall locking allows a shot-noise limited error signal to be generated on resonance without significantly perturbing the MOT.

The key feature of the SL-MOT is that the laser cooling experiment no longer requires a separate locking circuit: the magneto-optic trap provides the atom sample used to generate the locking signal. This may be important for miniaturization of MOTs, or for systems trapping rare isotopes, such as ANU-DFG, where an extra atom sample may be hard to come by. The ultracold MOT atoms are an ideal source for generating a naturally broadened locking signal. As a proof-of principle test, the frequency of the main trapping laser of a  $^{87}\text{Rb}$  MOT was successfully stabilized, using this technique, to generate a self-locked MOT.

## 14.1 Fundamental considerations

To create a MOT, the trapping laser light must have a linewidth that is at least as narrow as the natural linewidth of the atomic transition used to cool the gas. Lasers used to trap alkali metals, such as the ECDL lasers described in chapters 12 and 13, typically have short-term linewidths of 1 MHz or less. However, the laser frequency must be stabilized against longer-term drifts by locking it to a naturally broadened atomic transition.

The standard technique for frequency stabilizing the MOT laser, saturated absorption spectroscopy (SAS), derives a naturally broadened spectrum by optically selecting only the low velocity atoms from a room-temperature Doppler-broadened sample [144]. It does this using one of the pump-probe configurations shown in figure 14.1. Pump and probe counter-propagate and operate at the same frequency,



which may be scanned across a resonance to produce a spectral signal similar to that shown in figure 13.10. Exactly on resonance, the pump beam saturates the transition for those atoms with almost no velocity along the axis of the beam. Very little light is absorbed from the weak probe beam, as it also addresses the zero axial velocity atoms. A little below resonance, the pump beam is absorbed by atoms moving towards it with an appropriate velocity to Doppler-shift the laser on to resonance. The counter-propagating probe beam addresses atoms with the same speed, but the opposite direction of motion. The transition measured by the probe beam is no longer saturated, and light is absorbed from the probe beam. In this way, the Doppler-broadened absorption peak is resolved, with the naturally broadened transitions appearing as “anti-absorption-peaks”.

This process effectively isolates the atoms with zero, or very low, velocity along the beam direction, in order to measure the natural transition. The fraction of atoms contributing to the signal can be estimated as the ratio of the natural linewidth to the Doppler-broadened linewidth. For the alkali metals, the Doppler linewidth ( $\sim 1$  GHz) is two orders of magnitude larger than the natural linewidth ( $\sim 10$  MHz) of the optical transitions, and less than 1% of the room-temperature population of the vapor cell contributes to the SAS signal. In contrast, a MOT provides an ideal spectroscopic sample for locking a laser to a naturally broadened atomic transition, because the entire population contributes to a SL-MOT locking signal.

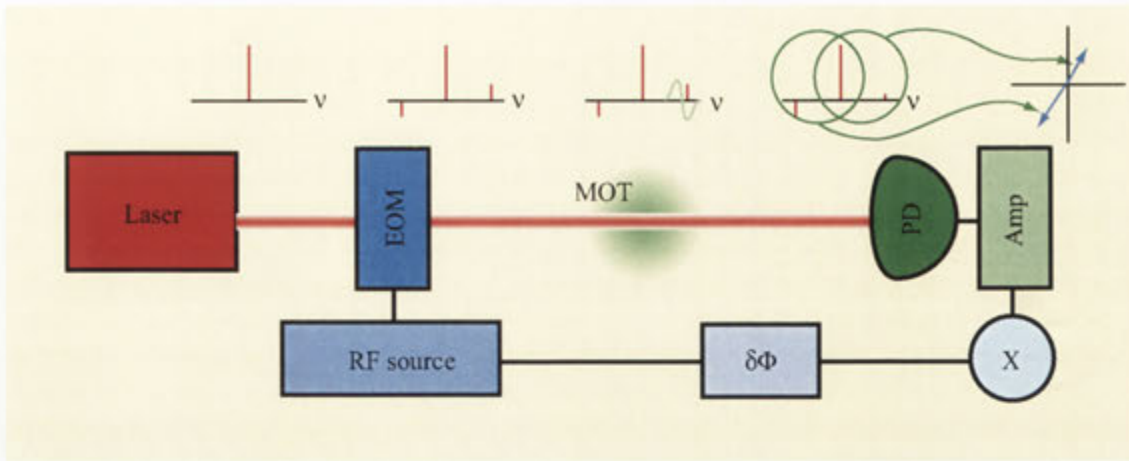
In addition, the MOT is an excellent “filter” for rare isotopes in the presence of their more dominant counterparts. Not only does the MOT collect together naturally-broadened atoms, it collects only the specific isotope for which it is configured. This helps isolate the signal from large, nearby signals due to other isotopes.

## 14.2 Practical considerations

The largest practical difficulty in the development of a self-locked MOT, is implementing a non-destructive measurement of the laser frequency. To create an error signal, the measurement technique must yield *phase sensitive* information about the laser frequency relative to the atomic transition. In addition, any on-resonant light will distort the MOT, decreasing its performance. The solution is to use an interferometer to separate a high-power local oscillator from a low-power resonant probe beam that passes through the atom cloud.

In the SL-MOT a particularly elegant system was employed: frequency modulation spectroscopy (FMS) [182]. FMS is a single-beam frequency-space interferometer, in which both the high power local oscillator and the low-power probe pass through the atom cloud, as shown in figure 14.2, making it impervious to acoustic noise and it easy to align. The high power beam does not perturb the atoms because it is separated in frequency space from the atomic resonance: it is off-resonance. The destruction of the MOT by the measurement system can be minimized by reducing the power in the resonant probe beam.

Frequency modulation spectroscopy (FMS) measures the rapid phase change of a beam near an atomic resonance similar to the way in which Pound-Drever-



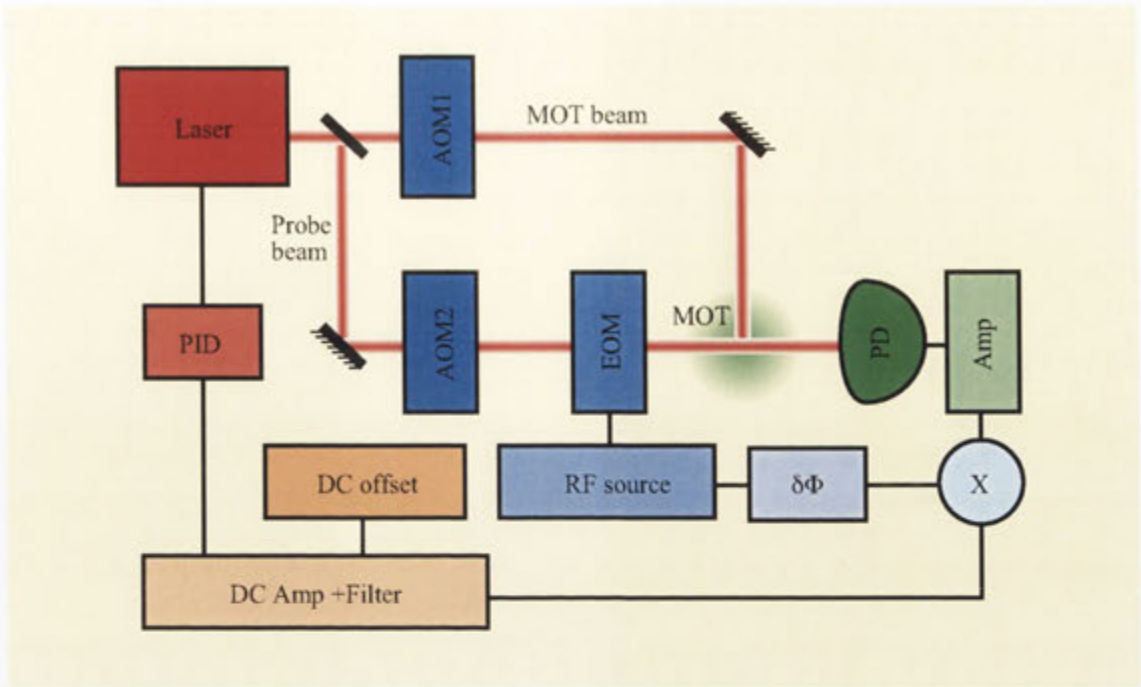
**14.2: Frequency modulation spectroscopy** The interaction of carrier, sidebands and atomic resonance to create a phase-sensitive error signal using frequency modulation spectroscopy.

Hall (PDH) [183] locking generates a signal from a cavity resonance. It uses the interference of a near-resonant carrier with two out-of-phase sidebands significantly outside the linewidth of that resonance, as shown in figure 14.2. The two off-resonant frequency components are negligibly affected by the resonance and act as a local oscillator, each generating a beat signal with the near-resonant component. Exactly on resonance, the phase of the near-resonant component is unaffected by the sample and the beat signals produced with the out-of-phase off-resonant components cancel to produce a null signal. As the near-resonant component moves off resonance, it is phase shifted closer to one of the off-resonant components, generating a non-zero beat, which can be demodulated to produce an error signal. Using FMS, the power in the resonant component can be reduced to make the perturbation to the MOT caused by on-resonant light acceptably small, whilst power in the off-resonant components can be increased to generate shot noise-limited signals from arbitrarily dilute vapors [184].

### 14.3 The self-locked magneto-optic trap

The SL-MOT was demonstrated as a proof-of-principle experiment using  $^{87}\text{Rb}$  atoms. This system was extremely well known and provided ample means for quantifying the performance of the SL-MOT. The experimental set up is shown in figure 14.3.

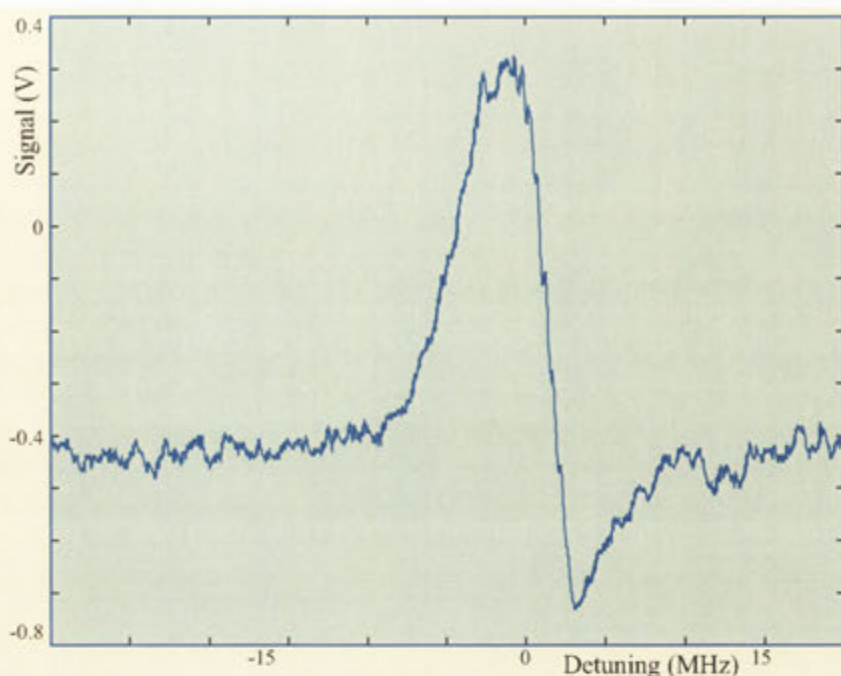
A probe beam was split from the main laser and shifted 14 MHz to the blue of the trapping beams using an acoustic-optic modulator (AOM). The probe was frequency modulated with an electro-optic modulator (EOM) giving sidebands at 850 MHz, well outside the linewidth of the naturally broadened transition. The



**14.3: The experimental set up.** An AOM and an EOM shift the probe beam 14 MHz to the blue of the MOT beam and give it 850 MHz out-of-phase sidebands. The RF beat of the carrier with the sidebands is detected on a high frequency photodiode (PD) and amplified (Amp) before being mixed to DC using a phase-adjustable ( $\delta\Phi$ ) reference from the RF source. The DC signal is further amplified, filtered, offset and processed by a PID controller before being fed back to the SL-MOT laser.

frequency of the trapping laser was manually adjusted until the MOT appeared, bringing the probe beam on to resonance and creating the FMS error signal. The probe passed through the MOT and was detected on a high frequency photodiode. The beat signal generated at the modulation frequency was demodulated, filtered and fed back via a proportional-integral-differential (PID) controller to the grating of an external cavity diode laser. Figure 14.4 shows the locking signal as the probe frequency is scanned across the cooling transition. Figures 14.5 (a) and (b) show lock acquisition plots for the SAS and SL-MOT techniques. In both systems, the low-frequency drifts are suppressed when the laser is locked.

This proof-of-principle experiment did not mimic the full system proposed for locking  $^{40}\text{K}$  in every way. The repump laser was locked using saturated absorption spectroscopy to guarantee its stability during tests on the cooling laser. The light for the MOT beams was shifted 14 MHz to the red of the atomic transition using one AOM, while the FMS beam was shifted on to resonance using a second AOM. An optimized set up would generate the probe beam and the repump beam from the main laser using direct modulation of the laser diode drive current, synchronizing



**14.4: The demodulated FMS signal.** Shows the demodulated FMS signal as the probe frequency is scanned across resonance. This signal is indicative of the locking signal generated by an SL-MOT system before it is fed to the PID controller.


them to provide a single-knob frequency control and locking technique.

### 14.3.1 A $^{40}\text{K}$ SL-MOT

The SL-MOT has not yet been implemented using  $^{40}\text{K}$ . An SL-MOT would lock the laser directly to the trapping transition of  $^{40}\text{K}$  by generating a shot-noise limited locking signal from the entire population of the MOT, significantly reducing the cost and complexity of this type of experiment.

The SL-MOT, by definition, stabilizes the laser frequency only while the MOT exists. In current degenerate gas experiments, the collection MOT is turned off during other parts of the experimental run. Utilizing a continuously operated 2D MOT as the collection MOT would provide a method of continuously stabilizing the laser frequency during other parts of the experiment. This configuration would require excellent optical shuttering, both outside and inside the vacuum system, during the magnetic trapping and evaporation stages of the experiment.

Another possible system would automate the lock re-capture process, allowing the lasers to unlock during parts of the experiment that the collection MOT was not running, then automatically relock when required. This proposal becomes more applicable as the SL-MOT is implemented with all laser frequencies generated from



**14.5: The lock acquisition plots.** Figs. (a) and (b) show lock acquisition plots for frequency locking the trapping laser using SAS and an SL-MOT respectively. Both plots show significant suppression of low-frequency drifts after 20 s, when the error signal is switched to feed back to the laser.

a single source. When the frequency of all laser light is controlled by a single “knob”, or voltage, the automated relocking algorithm becomes straightforward to implement.

### 14.3.2 Other plans

An SL-MOT can be made more compact than a standard MOT by directly modulating the diode current so that the MOT beams contain the FMS sidebands and a locking signal can be generated without requiring a separate vapor cell and optics. This may be important for experiments such as high precision space-based atomic clocks, which will allow millimeter scale GPS-type positioning and tests of the theory of relativity [185]. The size and weight benefits of an SL-MOT ideally complement the compact diode laser systems already built for these experiments.



## 15 Thinking back and looking forward

The investigations outlined herein lead to the production of the first Bose-Einstein condensate in Australia [6]. The massive efforts expended designing, making and assembling all manner of apparatus and experiments provided the foundation of experience that drove the development toward a degenerate Fermi gas machine at ANU. The very essence of this thesis is manifested in equipment constructed, and lessons learnt, that encompass every aspect of degenerate quantum gas production.

The major contribution of this work to the ANU-BEC-I machine was undoubtedly the design, modeling, and construction of the magnetic trap. The embryonic decision to develop the relatively untested QUIC trap for the ANU-BEC, initiated with the research reported in this thesis. The consequent development lead to a large investment in modeling and design before the trap was even constructed. Every aspect of the trap received extended consideration, from the minutiae of the coil design, and the practical implementation of water cooling, to material selection for the mounts, and the high-performance design of the driving electronics. The result was an excellent trap, that facilitated the first observation of Bose-Einstein condensation in Australia. Each successive iteration of the ANU-BEC machine has built on the understanding and experience developed through the research outlined in this thesis, to improve and upgrade the QUIC magnetic trap, and further the ability of the ANU-BEC team to perform experiments with BECs.

A most visible legacy of the work reported in this thesis is the computer control system that runs the experiment. Over the four and half years since its implementation, it has operated flawlessly. This result is especially impressive when the structural changes to the ANU-BEC machine over this period are considered. Even as it was originally conceived, the design of the control system fought to reconcile the vintage of the available, legacy control equipment, with the more modern BEC apparatus of a younger electronic generation. Two generalized solutions were eventually implemented: the considered restriction of communication protocols; and the careful, brute force application of modern computing power to overcome specific limitations of old, slow I/O equipment and unreliable software. These unique solutions, for unique problems, have since enabled solid communication between the user and the equipment, through every experiment, and every result, so far garnered from the machine.

The development of the ANU-BEC-I vacuum system was a group effort, to which all four members contributed greatly. The fruit of this collaboration is apparent for all to see, as the vacuum system sits staid and solid, after five years, at the center of the ANU-BEC apparatus. The construction, bake-out, and leak-testing took months to implement. The authoritative report of the general principles of

vacuum system design supplied in this thesis, belie the practical inconsistencies and vagaries of fortune that can plague any ultra-high vacuum system. Only by careful, sustained application, and re-application, of these principles, did the ANU-BEC vacuum system reach the performance level required for BEC production.

The implementation of CCD cameras in the imaging system of the ANU-BEC-I received little consideration during the early days of the ANU-BEC group. Much more thought was given to the effective implementation of an optimized imaging system with the equipment already available in the laboratory. However, as experiments have progressed, and as new machines have been constructed, much consideration has been invested in the topic of detection of cold atoms using CCDs, primarily as part of this thesis. The conclusions so generated have governed, to a greater or lesser degree, the purchase of the last two CCD cameras by the ANU-BEC group. These await full implementation, but early results indicate that they represent a more intelligent application of limited resources to the problems of imaging BECs.

The lasers used in the ANU-BEC-I experiment were provided commercially, or by collaborators with extensive experience in diode laser system design and construction. The ANU-DFG experiment reveled in no such luxury. As any who have tried, know, coincidental operation of commercial diode laser systems in the laboratory is a far cry from the design and implementation of such systems from scratch. The developments reported in this thesis perhaps understate the gradual process of gathering knowledge and experience in these matters, as tens of systems were designed and constructed over a period of years. The results speak for themselves. The ANU-BEC group now has the capability to assemble, align and tune onto resonance, a usable external cavity diode laser, from a stock of prepared parts, within half a day.

This knowledge lead to an appreciation of the strengths of the original ECDL design: the synergy of delicate, but mass-produced, laser diodes with rugged, simple mechanical and optical components; reaping a more precise and usable system from a small, but considered, investment. It was this understanding that lead to the development of the ultracold ECDLs designed to produce light for the potassium MOT in the ANU-DFG experiment. They, too, harnessed the economy and simplicity of mass-produced commercial laser diodes, then extended the operation of the standard ECDL design, using only simple, cheap parts away from the critical photonic components of the system. Their successful demonstration constituted a major step towards the production of a degenerate Fermi gas at ANU.

The self-locked magneto-optic trap was developed early during the work on the ANU-DFG. It is an elegant experimental technique, derived from studies of non-destructive imaging techniques for BECs. It finds application to the world of degenerate Fermi gases due to its ability to simplify a traditional laser cooling setup, through a solid understanding of sensitive modulation-based detection techniques, and how they may be elegantly employed using a minimum of components. This ability becomes especially important when cooling and trapping a rare isotope, such as  $^{40}\text{K}$ , where multiple samples of the gas may not be available, and equipment minimization becomes an utmost priority.

But to the future. The lessons learnt, and reported, throughout this thesis, pro-



vide a solid foundation from which to develop experiments using BECs and DFGs. A primary focus of the ANU-BEC group is the generation of non-destructive measurements of BECs, using novel interferometers, cavities and squeezed states of light [11]. Coupled with the simultaneous development of an improved atom laser apparatus, these techniques will further the study of atoms lasers in the ANU-BEC laboratories [14]. The ANU degenerate Fermi gas machine promises further such success. As the field of degenerate Fermi gases raises interest with the demonstration of mixed DQG systems [107], molecular BECs [64], and BCS-type transitions [123], we look forward to the time when we realize control of room temperature superconductors [186], superfluids [187] and understandings of atomic structure and matter itself [89].



# Bibliography

- [1] J. E. Lye. *Dynamic Non-destructive Detection of Bose-Einstein Condensates and Atom Lasers*. PhD thesis, Australian National University, 2003.
- [2] N. P. Robins. *Bose-Einstein condensation and the atom laser*. PhD thesis, Australian National University, 2003.
- [3] C. S. Fletcher, J. D. Close, U. Kallmann, and J. E. Lye. ANU-BEC. Carver Presentation Series 2001, 2001.
- [4] J. E. Lye, C. S. Fletcher, U. Kallmann, and J. D. Close. Bec at anu. Spain San Feliu de Guixols BEC Series poster presentation, 2001.
- [5] J. D. Close, J. E. Lye, C. S. Fletcher, and U. Kallmann. The ANU BEC. ACOLS Brisbane 01 oral presentation, 2001.
- [6] J. E. Lye, C.S. Fletcher, U. Kallman, H. A. Bachor, and J. D. Close. Images of evaporation to bose einstein condensation. *J. Opt. B, Quantum and Semiclass. Opt.*, 4:57, 2002.
- [7] C. S. Fletcher and J. D. Close. The self-locked magneto-optic trap. AIP Congress Sydney 02 oral presentation, 2002.
- [8] C. S. Fletcher, J. E. Lye, N. P. Robbins, and J. D. Close. A self locked magneto-optic trap. *Opt. Comm.*, 212:85, 2002.
- [9] C. S. Fletcher, J. E. Lye, N. P. Robins, and J. D. Close. Dynamic imaging of Bose-Einstein condensates and the self-locked magneto-optic trap. BECQI Coloundra 03 poster presentation, 2003.
- [10] A.K. Morrison, C. M. Breme, M.C. East, C.S. Fletcher, N.P. Robins, M.B. Gray, and J.D. Close. Atomic tilt locking. ICOLS Cairns 03 poster presentation, 2003.
- [11] J. E. Lye, J. J. Hope, and J. D. Close. Nondestructive dynamic detectors for Bose-Einstein condensates. *Phys. Rev. A*, 67:043609, 2003.
- [12] C. S. Fletcher and J. D. Close. Extended temperature tuning of an external cavity diode laser. *Appl. Phys. B*, 78:305, 2004.
- [13] N.P. Robins, J. E. Lye, C. S. Fletcher, S. A. Haine, J. Dugue, C. Breme, J. J. Hope, and J. D. Close. Dynamical effects of backcoupling on an atom laser. In P. Hannaford, A. Sidorov, H. Bachor, and K. Baldwin, editors, *Laser Spectroscopy, Proc. XVI Int. Conf.* World Scientific, 2004.

- [14] N. P. Robins, C. M. Savage, J. J. Hope, J. E. Lye, C. S. Fletcher, S. A. Haine, and J. D. Close. Classical noise and flux: the limits of multi-state atom lasers. *Phys. Rev. A. (in press)*, 2004.
- [15] A. Einstein. Quantentheorie des einatomigen idealen Gases. Zweite Abhandlung. *Sitzungber. Preuss. Akad. Wiss.*, 1925:3, January 1925.
- [16] Charles E. Hecht. The possible superfluid behaviour of hydrogen atom gases and liquids. *Physica*, 25(10):1159, October 1959.
- [17] Harald F. Hess. Evaporative cooling of magnetically trapped and compressed spin-polarized hydrogen. *Phys. Rev. B*, 34(5):3476, September 1986.
- [18] Steven Chu, L. Hollberg, J. E. Bjorkholm, Alex Cable, and A. Ashkin. Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure. *Phys. Rev. Lett.*, 55(1):48, July 1985.
- [19] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science*, 269(0):198, July 1995.
- [20] K. B. Davis, M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.*, 75(22):3969, November 1995.
- [21] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet. Evidence of Bose-Einstein condensation in an atomic gas with attractive interactions. *Phys. Rev. Lett.*, 75(9):1687, August 1995. *ibid.* **79**, 1170 (1997).
- [22] Tino Weber, Jens Herbig, Michael Mark, Hanns-Christoph Nägerl, and Rudolf Grimm. Bose-Einstein condensation of cesium. *Science*, 299:232, January 2003.
- [23] S. L. Cornish, N. R. Claussen, J. L. Roberts, E. A. Cornell, and C. E. Wieman. Stable  $^{85}\text{Rb}$  Bose-Einstein condensates with widely tunable interactions. *Phys. Rev. Lett.*, 85(9):1795, August 2000.
- [24] G. Modugno, G. Ferrari, G. Roati, R. J. Brecha, A. Simoni, and M. Inguscio. Bose-Einstein condensation of potassium atoms by sympathetic cooling. *Science*, 294(5545):1320, November 2001.
- [25] A. Robert, O. Sirjean, A. Browaeys, J. Poupard, S. Nowak, D. Boiron, C. I. Westbrook, and A. Aspect. A Bose-Einstein condensate of metastable atoms. *Science*, 292:461, April 2001.
- [26] F. Pereira Dos Santos, J. Lonard, Junmin Wang, C. J. Barrelet, F. Perales, E. Rasel, C. S. Unnikrishnan, M. Leduc, and C. Cohen-Tannoudji. Bose-einstein condensation of metastable helium. *Phys. Rev. Lett.*, 86:3459, 2001.

- [27] Dale G. Fried, Thomas C. Killian, Lorenz Willmann, David Landhuis, Stephen C. Moss, Daniel Kleppner, and Thomas J. Greytak. Bose-Einstein condensation of atomic hydrogen. *Phys. Rev. Lett.*, 81(18):3811, November 1998.
- [28] Y. Takahashi, K. Maki, Y. Takasu, K. Komori, T. Takano, K. Honda, A. Yamaguchi, Y. Kato, M. Mizoguchi, M. Kumakura, and T. Yabuzaki. Bose-Einstein condensation of ytterbium atoms. In P. Hannaford, A. Sidorov, H. Bachor, and K. Baldwin, editors, *Laser Spectroscopy, Proc. XVI Int. Conf.* World Scientific, 2004.
- [29] M. R. Andrews, C. G. Townsend, H.-J. Miesner, D. S. Durfee, D. M. Kurn, and W. Ketterle. Observation of interference between two Bose-Einstein condensates. *Science*, 275(0):637, January 1997.
- [30] D. S. Hall, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Measurements of relative phase in two-component Bose-Einstein condensates. *Phys. Rev. Lett.*, 81(8):1543, August 1998.
- [31] E. W. Hagley, L. Deng, M. Kozuma, M. Trippenbach, Y. B. Band, M. Edwards, M. R. Doery, P. S. Julienne, K. Helmerson, S. L. Rolston, and W. D. Phillips. Measurement of the coherence of a Bose-Einstein condensate. *Phys. Rev. Lett.*, 83(16):3112, October 1999.
- [32] L. Deng, E. W. Hagley, J. Wen, M. Trippenbach, Y. Band, P. S. Julienne, J. E. Simsarian, K. Helmerson, S. L. Rolston, and W. D. Phillips. Four-wave mixing with matter waves. *Nature*, 398(6724):218, March 1999.
- [33] M.-O. Mewes, M. R. Andrews, D. M. Kurn, D. S. Durfee, C. G. Townsend, and W. Ketterle. Output coupler for Bose-Einstein condensed atoms. *Phys. Rev. Lett.*, 78(4):582, January 1997.
- [34] E. W. Hagley, L. Deng, M. Kozuma, J. Wen, K. Helmerson, S. L. Rolston, and W. D. Phillips. A well-collimated quasi-continuous atom laser. *Science*, 283:1706, March 1999.
- [35] Immanuel Bloch, Theodor W. Hänsch, and Tilman Esslinger. Atom laser with a cw output coupler. *Phys. Rev. Lett.*, 82(15):3008, April 1999.
- [36] A. P. Chikkatur, Y. Shin, A. E. Leanhardt, D. Kielpinski, E. Tsikata, T. L. Gustavson, D. E. Pritchard, and W. Ketterle. A continuous source of Bose-Einstein condensed atoms. *Science*, 296:2193, June 2002.
- [37] D. M. Stamper-Kurn, M. R. Andrews, A. P. Chikkatur, S. Inouye, H.-J. Miesner, J. Stenger, and W. Ketterle. Optical confinement of a Bose-Einstein condensate. *Phys. Rev. Lett.*, 80(10):2027, March 1998.

- [38] J. Stenger, S. Inouye, A. P. Chikkatur, D. M. Stamper-Kurn, D. E. Pritchard, and W. Ketterle. Bragg spectroscopy of a Bose-Einstein condensate. *Phys. Rev. Lett.*, 82(23):4569, June 1999.
- [39] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard. Vortex formation in a stirred Bose-Einstein condensate. *Phys. Rev. Lett.*, 84(5):806, January 2000.
- [40] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard. Vortices in a stirred Bose-Einstein condensate. *J. Mod. Opt.*, 47(14/15):2715, 2000.
- [41] J. R. Abo-Shaeer, C. Raman, J. M. Vogels, and W. Ketterle. Observation of vortex lattices in Bose-Einstein condensates. *Science*, 292(5516):476, April 2001.
- [42] B. P. Anderson and M. A. Kasevich. Macroscopic quantum interference from atomic tunnel arrays. *Science*, 282:1686, November 1998.
- [43] Markus Greiner, Olaf Mandel, Tilman Esslinger, Theodor W. Hänsch, and Immanuel Bloch. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. *Nature*, 415:39, January 2002.
- [44] C. Orzel, A. K. Tuchman, M. L. Fenselau, M. Yasuda, and M. A. Kasevich. Squeezed states in a Bose-Einstein condensate. *Science*, 291(5512):2386, 2001.
- [45] D. S. Petrov, G. V. Shlyapnikov, and J. T. M. Walraven. Regimes of quantum degeneracy in trapped 1D gases. *Phys. Rev. Lett.*, 85:3745, 2000.
- [46] V. Dunjko, V. Lorent, and M. Olshanii. Bosons in cigar-shaped traps: Thomas-Fermi regime, Tonks-Girardeau regime, and in between. *Phys. Rev. Lett.*, 86:5413, 2001.
- [47] D. Blume. Fermionization of a Bosonic gas under highly elongated confinement: A diffusion quantum Monte Carlo study. *Phys. Rev. A*, 66:053613, 2002.
- [48] W. D. Phillips, J. H. Huckans, B. Laburthe Tolra, K. M. O'Hara, J. V. Porto, S. L. Rolston, and M. Anderlini. Experimental study of a bose gas in one dimension. In P. Hannaford, A. Sidorov, H. Bachor, and K. Baldwin, editors, *Laser Spectroscopy, Proc. XVI Int. Conf.* World Scientific, 2004.
- [49] W. Hänsel, P. Hommelhoff, T. W. Hänsch, and J. Reichel. Bose-Einstein condensation on a microelectronic chip. *Nature*, 413(6855):498, October 2001.
- [50] H. Ott, J. Fortagh, G. Schlotterbeck, A. Grossmann, and C. Zimmermann. Bose-Einstein condensation in a surface microtrap. *Phys. Rev. Lett.*, 87:230401, 2001.

- [51] Ph. Treutlein, P. Hommelhoff, T. W. Hansch, and J. Reichel. Coherent atomic states in microtraps. In P. Hannaford, A. Sidorov, H. Bachor, and K. Baldwin, editors, *Laser Spectroscopy, Proc. XVI Int. Conf.* World Scientific, 2004.
- [52] J. Fortagh, H. Ott, S. Kraft, A. Gunther, C. Truck, and C. Zimmermann. On-chip laboratory for Bose-Einstein condensates. In P. Hannaford, A. Sidorov, H. Bachor, and K. Baldwin, editors, *Laser Spectroscopy, Proc. XVI Int. Conf.* World Scientific, 2004.
- [53] L. Feenstra, K. Brugger, R. Folman, S. Groth, A. Kasper, P. Kruger, X. Luo, S. Schneider, S. Wildermuth, and J. Schmeidmayer. Atom optics with microtraps and atom chips; assembling tools for quantum information processing. In P. Hannaford, A. Sidorov, H. Bachor, and K. Baldwin, editors, *Laser Spectroscopy, Proc. XVI Int. Conf.* World Scientific, 2004.
- [54] Roahn Wynar, R. S. Freeland, D. J. Han, C. Ryu, and D. J. Heinzen. Molecules in a Bose-Einstein condensate. *Science*, 287:1016, February 2000.
- [55] Elizabeth A. Donley, Neil R. Claussen, Sarah T. Thompson, and Carl E. Wieman. Atom-molecule coherence in a Bose-Einstein condensate. *Nature*, 417:529, May 2002.
- [56] Stephan Durr, Thomas Volz, Andreas Marte, and Gerhard Rempe. Observation of molecules produced from a Bose-Einstein condensate. *cond-mat*, page 0307440, 2003.
- [57] Jens Herbig, Tobias Kraemer, Michael Mark, Tino Weber, Cheng Chin, Hanns-Christoph Nagerl, and Rudolf Grimm. Preparation of a pure molecular quantum gas. *Science*, 301:1510, 2003.
- [58] S. Jochim, M. Bartenstein, A. Altmeyer, G. Hendl, C. Chin, J. Hecker Denschlag, and R. Grimm. Pure gas of optically trapped molecules created from Fermionic atoms. *Phys. Rev. Lett.*, 91:240402, 2003.
- [59] Kevin E. Strecker, Guthrie B. Partridge, and Randall G. Hulet. Conversion of an atomic Fermi gas to a long-lived molecular Bose gas. *Phys. Rev. Lett.*, 91:080406–1, 2003.
- [60] J. Cubizolles, T. Bourdel, S. J. J. M. F. Kokkelmans, G. V. Shlyapnikov, and C. Salomon. Production of long-lived ultracold Li<sub>2</sub> molecules from a Fermi gas. *cond-mat*, page 0308018, 2003.
- [61] Cindy A. Regal, Markus Greiner, and Deborah S. Jin. Lifetime of molecule-atom mixtures near a Feshbach resonance in 40K. *cond-mat*, page 0308606, 2003.
- [62] S. Jochim, M. Bartenstein, A. Altmeyer, G. Hendl, S. Riedl, C. Chin, J. Hecker Denschlag, and R. Grimm. Bose-Einstein condensation of molecules. *Science*, 302:2101, 2003.

- [63] K. Xu, T. Mukaiyama, J.R. Abo-Shaeer, J.K. Chin, D. Miller, and W. Ketterle. Formation of quantum-degenerate sodium molecules. *cond-mat*, page 0310027, 2003.
- [64] S. Jochim, M. Bartenstein, A. Altmeyer, G. Hendl, C. Chin, J. Hecker Denschlag, and R. Grimm. Pure gas of optically trapped molecules created from Fermionic atoms. *Phys. Rev. Lett.*, 91:240402, 2003.
- [65] Markus Greiner, Cindy A. Regal, and Deborah S. Jin. Emergence of a molecular Bose-Einstein condensate from a Fermi gas. *Nature*, 426:537, 2003.
- [66] M. Bartenstein, A. Altmeyer, S. Riedl, S. Jochim, C. Chin, J. Hecker Denschlag, and R. Grimm. Crossover from a molecular Bose-Einstein condensate to a degenerate Fermi gas. *cond-mat*, page 0401109, 2004.
- [67] M.W. Zwierlein, C.A. Stan, C.H. Schunck, S.M.F. Raupach, S. Gupta, Z. Hadzibabic, and W. Ketterle. Observation of Bose-Einstein condensation of molecules. *Phys. Rev. Lett.*, 91:250401, 2003.
- [68] B. DeMarco and D. S. Jin. Onset of Fermi degeneracy in a trapped atomic gas. *Science*, 285(5434):1703, September 1999.
- [69] B. DeMarco. *Quantum behavior of an atomic Fermi gas*. PhD thesis, University of Colorado, Boulder, Colorado, 2001.
- [70] B. DeMarco, H. Rohner, and D. Jin. An enriched  $^{40}\text{K}$  source for Fermionic atom studies. *Rev. Sci. Inst.*, 70:1967, 1999.
- [71] Andrew G. Truscott, Kevin E. Strecker, William I. McAlexander, Guthrie B. Partridge, and Randall G. Hulet. Observation of Fermi pressure in a gas of trapped atoms. *Science*, 291(5513):2570, March 2001.
- [72] D. J. Larson, J. C. Bergquist, J. J. Bollinger, W. M. Itano, and D. J. Wineland. *Phys. Rev. Lett.*, 57:70, 1986.
- [73] C. J. Myatt, E. A. Burt, R. W. Ghrist, E. A. Cornell, and C. E. Wieman. Production of two overlapping Bose-Einstein condensates by sympathetic cooling. *Phys. Rev. Lett.*, 78(4):586, January 1997.
- [74] Z. Hadzibabic, S. Gupta, C. A. Stan, C. H. Schunck, M. W. Zwierlein, K. Dieckmann, and W. Ketterle. Fiftyfold improvement in the number of quantum degenerate Fermionic atoms. *Phys. Rev. Lett.*, 91:160401, 2003.
- [75] K. M. O'Hara, S. L. Hemmer, M. E. Gehm, S. R. Granade, and J. E. Thomas. Observation of a strongly interacting degenerate Fermi gas of atoms. *Science*, 298:2179, December 2002.
- [76] S. R. Granade, M. E. Gehm, K. M. O'Hara, and J. E. Thomas. All-optical production of a degenerate Fermi gas. *Phys. Rev. Lett.*, 88(12):120405, March 2002.



- [77] G. Roati, F. Riboli, G. Modugno, and M. Inguscio. Fermi-Bose quantum degenerate  $^{40}\text{K}$ - $^{87}\text{Rb}$  mixture with attractive interaction. *Phys. Rev. Lett.*, 89(15):150403, October 2002.
- [78] Z. Hadzibabic, C. A. Stan, K. Dieckmann, S. Gupta, M. W. Zwierlein, A. Grititz, and W. Ketterle. Two-species mixture of quantum degenerate Bose and Fermi gases. *Phys. Rev. Lett.*, 88:160401, 2002.
- [79] J. Goldwin, S. B. Papp, B. DeMarco, and D. S. Jin. A two-species magneto-optical trap with  $40\text{K}$  and  $87\text{Rb}$ . *Phys. Rev. A*, 65:021402, 2002.
- [80] M. G. Moore and P. Meystre. Atomic four-wave mixing: Fermions versus Bosons. *Phys. Rev. Lett.*, 86(19):4199, May 2001.
- [81] C. P. Search, H. Pu, W. Zhang, and P. Meystre. Diffraction of a superfluid Fermi gas by an atomic grating. *Phys. Rev. Lett.*, 88:110401, 2002.
- [82] Sierk Pötting, Marcus Cramer, Weiping Zhang, and Pierre Meystre. Raman coupler for a trapped two-component quantum-degenerate Fermi gas. *Phys. Rev. A*, 65:063620, June 2002.
- [83] Wolfgang Ketterle and Shin Inouye. Does matter wave amplification work for fermions? *Phys. Rev. Lett.*, 86(19):4203, May 2001.
- [84] S.-K. Yip and T.-L. Ho. Zero sound modes of dilute fermi gases with arbitrary spin. *Phys. Rev. A*, 59:4653, 1999.
- [85] K. Helmerson, M. Xiao, and D. Pritchard. In *International Quantum Electronics Conference 1990 book of abstracts*, page No. QTHH4. IEEE, New York, 1990.
- [86] B. DeMarco and D. S. Jin. Exploring a quantum degenerate gas of fermionic atoms. *Phys. Rev. A*, 58(6):R4267, December 1998.
- [87] T. Busch, J. R. Anglin, J. Cirac, and P. Zoller. *Europhys. Lett.*, 44:1, 1998.
- [88] J. Ruostekoski and J. Javanainen. Optical linewidth of a low density Fermi-Dirac gas. *Phys. Rev. Lett.*, 82:4741, 1999.
- [89] B. I. Schneider and H. Wallis. *Phys. Rev. A*, 57:1253, 1998.
- [90] G. M. Bruun and K. Burnett. Interacting Fermi gas in a harmonic trap. *Phys. Rev. A*, 58(3):2427, September 1998.
- [91] M. Amoruso, I. Meccoli, A. Minguzzi, and M. P. Tosi. Collective excitations of a degenerate Fermi vapour in a magnetic trap. *Euro. Phys. J. D*, 7:441, 1999.
- [92] M. Amoruso et. al. Density profiles and collective excitations of a trapped two-component Fermi vapour. *Euro. Phys. J. D*, 8:361, 2000.

- [93] Georg M. Bruun and Charles W. Clark. Hydrodynamic excitations of trapped Fermi gases. *Phys. Rev. Lett.*, 83(26):5415, December 1999.
- [94] Georg M. Bruun and Charles W. Clark. Ideal gases in time-dependent traps. *Phys. Rev. A*, 61:061601, 2000.
- [95] L. Vichi and S. Stringari. Collective oscillations of an interacting trapped Fermi gas. *Phys. Rev. A*, 60(6):4734, December 1999.
- [96] Anna Minguzzi and Mario P. Tosi. Scissors mode in a superfluid Fermi gas. *Phys. Rev. A*, 63:023609, January 2001.
- [97] M. Houbiers, R. Ferwerda, H. T. C. Stoof, W. I. McAlexander, C. A. Sackett, and R. G. Hulet. Superfluid state of atomic  ${}^6\text{Li}$  in a magnetic trap. *Phys. Rev. A*, 56(6):4864, December 1997.
- [98] M. J. Holland, B. DeMarco, and D. S. Jin. Evaporative cooling of a two-component degenerate Fermi gas. *Phys. Rev. A*, 61:053610, April 2000.
- [99] Subhadeep Gupta, Zoran Hadzibabic, James R. Anglin, and Wolfgang Ketterle. Collisions in zero temperature fermi gases. *cond-mat*, page 0307088, 2003.
- [100] M. E. Gehm, S. L. Hemmer, K. M. O'Hara, and J. E. Thomas. Unitarity-limited elastic collision rate in a harmonically trapped Fermi gas. *Phys. Rev. A*, 68:011603, 2003.
- [101] B. DeMarco, J. L. Bohn, J. P. Burke, Jr., M. Holland, and D. S. Jin. Measurement of  $p$ -wave threshold law using evaporatively cooled fermionic atoms. *Phys. Rev. Lett.*, 82(21):4208, May 1999.
- [102] D. Jin, B. DeMarco, and S. Papp. Exploring a quantum degenerate fermi gas. In *ICAP Proceedings*, 2000.
- [103] B. DeMarco, S. B. Papp, and D. S. Jin. Pauli blocking of collisions in a quantum degenerate atomic Fermi gas. *Phys. Rev. Lett.*, 86(24):5409, June 2001.
- [104] S.D. Gensemer and D.S. Jin. Transition from collisionless to hydrodynamic behavior in an ultracold Fermi gas. *Phys. Rev. Lett.*, 87:173201, 2001.
- [105] B. DeMarco and D. S. Jin. Spin excitations in a Fermi gas of atoms. *Phys. Rev. Lett.*, 88:040405, 2002.
- [106] M. E. Gehm, S. L. Hemmer, S. R. Granade, K. M. O'Hara, and J. E. Thomas. Mechanical stability of a strongly interacting Fermi gas of atoms. *Phys. Rev. A*, 68:011401, 2003.

- [107] Giovanni Modugno, Giacomo Roati, Francesco Riboli, Francesca Ferlaino, Robert J. Brecha, and Massimo Inguscio. Collapse of a degenerate Fermi gas. *Science*, 297(5590):2240, September 2002.
- [108] Francesca Ferlaino, R. Brecha, Peter Hannaford, Francesco Riboli, Giacomo Roati, Giovanni Modugno, and Massimo Inguscio. Dipolar oscillations in a quantum degenerate Fermi-Bose atomic mixture. *J. Opt. B: Quantum Semi-class. Opt.*, 5:S3, 2003.
- [109] T. Bourdel, J. Cubizolles, L. Khaykovich, K. M. F. Magalhães, S. J. J. M. F. Kokkelmans, G. V. Shlyapnikov, and C. Salomon. Measurement of the interaction energy near a Feshbach resonance in a  $^6\text{Li}$  Fermi gas. *Phys. Rev. Lett.*, 91(2):020402, July 2003.
- [110] S. Gupta, Z. Hadzibabic, M. W. Zwierlein, C. A. Stan, K. Dieckmann, C. H. Schunck, E. G. M. van Kempen, B. J. Verhaar, and W. Ketterle. Radio-frequency spectroscopy of ultracold fermions. *Science*, 300:1723, June 2003.
- [111] T.-L. Ho and V. Shenoy. The hydrodynamic equations of superfluid mixtures in magnetic traps. *J. Low Temp. Phys.*, 111:937, 1999.
- [112] C. P. Search, H. Pu, W. Zhang, and P. Meystre. Quasiparticle spectrum and dynamical stability of an atomic Bose-Einstein condensate coupled to a degenerate Fermi gas. *Phys. Rev. A*, 65:063615, June 2002.
- [113] M. Lewenstein, L. Santos, M. A. Baranov, and H. Fehrmann. Atomic Bose-Fermi mixtures in an optical lattice. *Phys. Rev. Lett.*, 92:050401, 2004.
- [114] Michele Modugno, Francesca Ferlaino, Francesco Riboli, Giacomo Roati, Giovanni Modugno, and Massimo Inguscio. Mean-field analysis of the stability of a K-Rb Fermi-Bose mixture. *Phys. Rev. A*, 68:043626, 2003.
- [115] C. Fort, G. Modugno, F. S. Cataliotti, J. Catani, E. De Mirandes, L. Fallani, F. Ferlaino, M. Modugno, H. Ott, G. Roati, and M. Inguscio. Quantum degenerate Bosons and Fermions in a 1D optical lattice. In P. Hannaford, A. Sidorov, H. Bachor, and K. Baldwin, editors, *Laser Spectroscopy, Proc. XVI Int. Conf.* World Scientific, 2004.
- [116] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Theory of superconductivity. *Phys. Rev.*, 108:1175, 1957.
- [117] M. Mackie, E. Timmermans, R. Cote, and J. Javanainen. Driving superfluidity with photoassociation. *Opt. Exp.*, 8:118, 2000.
- [118] D. van Oosten, P. van der Straten, and H. T. C. Stoof. Quantum phases in an optical lattice. *Phys. Rev. A*, 63:053601, April 2001.
- [119] Andrea Simoni, Francesca Ferlaino, Giacomo Roati, Giovanni Modugno, and Massimo Inguscio. Magnetic control of the interaction in ultracold K-Rb mixtures. *Phys. Rev. Lett.*, 90:163202, 2003.

- [120] K. Dieckmann, C. A. Stan, S. Gupta, Z. Hadzibabic, C. H. Schunck, and W. Ketterle. Decay of an ultracold fermionic lithium gas near a Feshbach resonance. *Phys. Rev. Lett.*, 89(20):203201, November 2002.
- [121] K. M. O'Hara, S. L. Hemmer, S. R. Granade, M. E. Gehm, J. E. Thomas, V. Venturi, E. Tiesinga, and C. J. Williams. Measurement of the zero crossing in a Feshbach resonance of fermionic  $^6\text{Li}$ . *Phys. Rev. A*, 66:041401(R), October 2002.
- [122] D. SPetrov, C. Salomon, and G. VShlyapnikov. Weakly bound dimers of Fermionic atoms. *cond-mat*, page 0309010, 2003.
- [123] C. A. Regal, M. Greiner, and D. S. Jin. Observation of resonance condensation of Fermionic atom pairs. *Phys. Rev. Lett.*, 92:040403, 2004.
- [124] M.W. Zwierlein, C.A. Stan, C.H. Schunck, S.M.F. Raupach, A.J. Kerman, and W. Ketterle. Condensation of pairs of Fermionic atoms near a Feshbach resonance. *Phys. Rev. Lett.*, 2004.
- [125] T. Bourdel, L. Khaykovich, J. Cubizolles, J.Zhang, F.Chevy, M.Teichmann, L.Tarruel, S. J. J. M. F. Kokkelmans, and C. Salomon. Experimental study of the BEC-BCS crossover region in Lithium 6. *cond-mat*, page 0403091, 2004.
- [126] D.E. Pritchard, K. Helmerson, and A.G. Martin. In S. Haroche, J.C. Gay, and G. Grynberg, editors, *Atomic Physics 11*, page 179. World Scientific, Singapore, 1989.
- [127] T. W. Hijmans, O. J. Luiten, I. D. Setija, and J. T. M. Walraven. Optical cooling of atomic hydrogen in a magnetic trap. *J. Opt. Soc. Am. B*, 6(11): 2235, November 1989.
- [128] W. Ketterle et al. Dark spontaneous-force optical trap. In *OSA Annual Meeting, Toronto, Canada, October 3-8, 1993*.
- [129] J. Sakurai. *Modern Quantum Mechanics*. Addison-Wesley, New York, 1994.
- [130] C. J. Myatt, N. R. Newbury, R. W. Ghrist, S. Loutzenhiser, and C. E. Wieman. Multiply loaded magneto-optical trap. *Opt. Lett.*, 21(4):290, February 1996.
- [131] Tilman Esslinger, Immanuel Bloch, and Theodor W. Hänsch. Bose-Einstein condensation in a quadrupole-Ioffe-configuration trap. *Phys. Rev. A*, 58(4): R2664, October 1998.
- [132] K. B. Davis, M.-O. Mewes, M. A. Joffe, M. R. Andrews, and W. Ketterle. Evaporative cooling of sodium atoms. *Phys. Rev. Lett.*, 75:2909, 1995.
- [133] O. J. Luiten, M. W. Reynolds, and J. T. M. Walraven. Kinetic theory of the evaporative cooling of a trapped gas. *Phys. Rev. A*, 53(1):381, January 1996.

- [134] E. A. Cornell, J. R. Ensher, and C. E. Wieman. Experiments in dilute atomic Bose-Einstein condensation. In M. Inguscio, S. Stringari, and C. E. Wieman, editors, *Proceedings of the International School of Physics - Enrico Fermi*, page 15. IOS Press, 1999.
- [135] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn. Making, probing and understanding Bose-Einstein condensates. In M. Inguscio, S. Stringari, and C. E. Wieman, editors, *Proceedings of the International School of Physics - Enrico Fermi*, page 67. IOS Press, 1999.
- [136] C. J. Myatt. *Bose-Einstein Condensation Experiments in a Dilute Vapor of Rubidium*. PhD thesis, University of Colorado, Boulder, 1997.
- [137] Wolfgang Ketterle and N. J. van Druten. Evaporative cooling of trapped atoms. *Adv. At. Mol. Opt. Phys.*, 37(0):181, 1996.
- [138] Elstein-Werk M. Steinmetz GmbH and Co. KG Spezialfabrik keramischer Infrarotstrahler, 6D- 37142 Northeim, Postfach 1227, D- 37154 Northeim, Stettiner Strae 14 , <http://www.elstein-werk.de>.
- [139] K. Dieckmann. *Bose-Einstein condensation with high atom number in a deep magnetic trap*. PhD thesis, University of Amsterdam, 2001.
- [140] H. Metcalf. *Laser Cooling and Trapping*. Springer, 2002.
- [141] R. Folman, P. Krüger, J. Schmeimayer, J. Denschlag, and C. Henkel. Microscopic atom optics: from wires to an atom chip. *Adv. At. Mol. Opt. Phys.*, volume=.
- [142] S. Bali, K. M. OHara, M. E. Gehm, S. R. Granade, and J. E. Thomas. Quantum-diffractive background gas collisions in atom-trap heating and loss. *Phys. Rev. A.*, 60:R29, 1999.
- [143] L. Ricci, M. Weidemüller, T. Esslinger, A. Hemmerich, C. Zimmermann, V. Vuletic, W. König, and T. W. Hänsch. A compact grating-stabilized diode laser system for atomic physics. *Opt. Comm.*, 117:541, 1995.
- [144] K. B. MacAdam, A. Steinbach, and C. Wieman. A narrow band tunable diode laser system with grating feedback, and a saturated absorption spectrometer for Cs and Rb. *Am. J. Phys.*, 60(12):1098, December 1992.
- [145] Wolfgang Petrich, Michael H. Anderson, Jason R. Ensher, and Eric A. Cornell. Stable, tightly confining magnetic trap for evaporative cooling of neutral atoms. *Phys. Rev. Lett.*, 74(17):3352, April 1995.
- [146] M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. M. Kurn, D. S. Durfee, and W. Ketterle. Bose-Einstein condensation in a tightly confining dc magnetic trap. *Phys. Rev. Lett.*, 77(3):416, July 1996.

- [147] Wolfgang Ketterle and David E. Pritchard. Trapping and focusing ground state atoms with static fields. *Appl. Phys. B*, 54:403, 1992.
- [148] W. Wing. *Prog. Quant. Elect.*, 8:181, 1985.
- [149] David E. Pritchard. Cooling neutral atoms in a magnetic trap for precision spectroscopy. *Phys. Rev. Lett.*, 51(15):1336, October 1983.
- [150] V. S. Bagnato, G. P. Lafyatis, A. G. Martin, E. L. Raab, R. N. Ahmad-Bitar, and D. E. Pritchard. Continuous stopping and trapping of neutral atoms. *Phys. Rev. Lett.*, 58(21):2194, May 1987.
- [151] Y.V. Gott, M.S. Ioffe, and V.G. Tel'kovskii. *Nuc. Fus. Supp.*, 3:1045, 1962.
- [152] Biotsavart. Ripplon Software Inc., 4794 Fernglen Drive, Burnaby BC, Canada V5G 3V7, <http://www.ripplon.com>.
- [153] Mathematica. Wolfram Research, Inc., 100 Trade Center Drive, Champaign, IL 61820-7237 USA, <http://www.wolfram.com>.
- [154] P. Hannaford and R. J. McLean. *Spectrochim. Acta B*, 54:2183, 1999.
- [155] J.E. Lye, B.D. Cuthbertson, H.-A. Bachor, and J.D. Close. Phase modulation spectroscopy: a non-destructive probe of Bose-Einstein condensates. *J. Opt. B: Quant. Semiclass. Opt.*, 1:402, 1999.
- [156] D. M. Stamper-Kurn. *Peeking and poking at a new quantum fluid: Studies of gaseous Bose-Einstein condensates in magnetic and optical traps*. PhD thesis, Massachusetts Institute of Technology, 2000.
- [157] O. Hadar and G. D. Boreman. Oversampling requirements for pixelated-imager systems. *Optical Engineering*, 38:782-785, 1999.
- [158] M. R. Andrews, M.-O. Mewes, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Direct, nondestructive observation of a Bose condensate. *Science*, 273(0):84, July 1996.
- [159] V S. Letokhov and V G. Minogin. Laser radiation pressure on free atoms. *Phys. Rep.*, 73:1, 1981.
- [160] M.A. Joffe, W. Ketterle, A. Martin, and D.E. Pritchard. Transverse cooling and deflection of an atomic beam inside a zeeman slower. *J. Opt. Soc. Am. B*, 10:2257, 1993.
- [161] DALSA Waterloo, 605 McMurray Road, Waterloo, Ontario, Canada N2V 2E9, <http://www.dalsa.com>.
- [162] National Instruments Corporation, 11500 N Mopac Expwy, Austin, TX 78759-3504 USA, <http://www.ni.com>.

- [163] R. S. Williamson III. *Magneto-optical trapping of potassium isotopes*. PhD thesis, University of Wisconsin - Madison, 1997.
- [164] F. S. Cataliotti, E. A. Cornell, C. Fort, M. Inguscio, F. Marin, M. Prevedelli, L. Ricci, and G. M. Tino. Magneto-optical trapping of Fermionic potassium atoms. *Phys. Rev. A*, 57(2):1136, February 1998.
- [165] G. Ferrari, M.-O. Mewes, F. Schreck, and C. Salomon. High-power multiple-frequency narrow-linewidth laser source based on a semiconductor tapered amplifier. *Opt. Lett.*, 24:151, 1999.
- [166] C. E. Wieman and C. Hollberg. Using diode lasers for atomic physics. *Rev. Sci. Instr.*, 62:1, 1991.
- [167] K. B. MacAdam, A. Steinbach, and C. Wieman. *Am. J. Phys.*, 60:1098, 1992.
- [168] E.G. Loewen, M. Nevire, and D. Maystre. Grating efficiency theory as it applies to blazed and holographic gratings. *Appl. Opt.*, 16:2711, 1977.
- [169] H. Leinen, D. Glaessner, H. Metcalf, R. Wynands, D. Haubrich, and D. Meschede. GaN blue diode laser: a spectroscopists view. *Appl. Phys. B*, 70:567, 2000.
- [170] TOPTICA Photonics AG, Fraunhoferstr.14, D-82152 Martinsried/Munich, Germany, <http://www.toptica.com>.
- [171] K. G. Libbrecht and J. L. Hall. A low-noise high-speed diode laser current controller. *Rev. Sci. Instrum.*, 64:2133, 1993.
- [172] L. D. Turner, K. P. Weber, C. J. Hawthorn, and R. E. Scholten. Frequency noise characterization of narrow linewidth diode lasers. *Opt. Communic.*, 201:391397, 2002.
- [173] D. Mehuys, D. F. Welch, and L. Goldberg. 2.0 W CW, diffraction-limited tapered amplifier with diode injection. *Electron. Lett.*, 28:1944, 1992.
- [174] F. Z. Cruz, M. Rauner, J. H. Marquart, L. Hollberg, and J. C. Bergquist. In J. C. Bergquist, editor, *Proceedings of the Fifth Symposium on Frequency Standards and Metrology*, page 511. World Scientific, Singapore, 1995.
- [175] A. C. Wilson, J. C. Sharpe, C. R. McKenzie, P. J. Manson, and D. M. Warrington. *Appl. Opt.*, 37:4871, 1998.
- [176] Sacher Lasertechnik GmbH, Hannah-ArendtStrasse 3-7, D-35037 Marburg, Germany, <http://www.sacher-laser.com>.
- [177] S. Stry, L. Hildebrandt, J. Sacher, C. Buggle, M. Kemmann, and W. von Klitzing. Compact tunable diode laser with diffraction limited 1w for atom cooling and trapping. In *Photonics West*, 2004.

- [178] D. Weidmann and D. Courtois. *Infrared Phys. Technol.*, 41:361, 2000.
- [179] B. D. Carter and M. C. B. Ashley. *Proc. ASA*, 9:158, 1991.
- [180] Device performance formulae, melcor thermal solutions. Melcor Corporation, 1040 Spruce Street, Trenton, NJ 08648 USA, <http://www.melcor.com>.
- [181] Mathematical modeling of thermoelectric cooling modules, technical reference for ferrotec thermoelectric cooling modules. Ferrotec (USA), 40 Simon Street, Nashua, NH 03060-3075 USA, <http://www.ferrotec.com>.
- [182] G. C. Bjorklund, M. D. Levenson, W. Lenth, and C. Ortiz. *Appl. Phys. B*, 32:145, 1983.
- [183] R. W. P. Drever, J. L. Hall, and F. V. Kowalski. *Appl. Phys. B*, 31:97, 1983.
- [184] M. D. Levenson, W. E. Moerner, and D. E. Horne. *Opt. Lett.*, 8:108, 1983.
- [185] Ph. Laurent, P. Lemonde, E. Simon, G. Santarelli, A. Clairon, N. Dimarcq, P. Petit, C. Audoin, and C. Salomon. *EPJ D*, 3:201, 1998.
- [186] H. Heiselberg, C. J. Pethick, H. Smith, and L. Viverit. Influence of induced interactions on the superfluid transition in dilute Fermi gases. *Phys. Rev. Lett.*, 85(12):2418, September 2000.
- [187] M. Holland, S. J. J. M. F. Kokkelmans, M. L. Chiofalo, and R. Walser. Resonance superfluidity in a quantum degenerate Fermi gas. *Phys. Rev. Lett.*, 87(12):120406, September 2001.