

**Progress towards a Portable Cold-Atom Pressure
Standard**

How I learned to stop worrying and love my Bill of Materials

by

Perrin A. Z. Waldock

B.Sc. of Physics, University of the Fraser Valley, 2019

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Master of Applied Science

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL
STUDIES

(Engineering Physics)

The University of British Columbia
(Vancouver)

October 2022

© Perrin A. Z. Waldock, 2022

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the thesis entitled:

Progress towards a Portable Cold-Atom Pressure Standard

submitted by **Perrin A. Z. Waldock** in partial fulfillment of the requirements for the degree of **Master of Applied Science in Engineering Physics**.

Examining Committee:

Kirk W Madison, Associate Professor, Physics and Astronomy, UBC
Supervisor

Valery Milner, Associate Professor, Physics and Astronomy, UBC
Supervisory Committee Member

Abstract

All existing pressure standards (classical and quantum-based) are large apparatuses—typically, the size of a van or larger. In this thesis, we report on the progress towards building a portable version of a cold-atom based pressure standard which can easily be transported by a van instead. Specifically, we discuss

- the flexible laser design that can produce 780 nm pump light anywhere in a 26 GHz range and repump light anywhere in a ± 6.8 GHz range within < 1 MHz with no realignment.
- the magnetic field coils that can produce gradients in excess of 400 G/cm without changing temperature (provided that > 14 mL/s of 10 C to 15 C water is available).
- the low-outgassing vacuum chamber that can supply rubidium for measurement while minimizing contamination in the chamber under test.
- preliminary results characterizing the performance of its rubidium source and magneto-optical trap.
- future plans for the apparatus.

This pressure standard leverages the linear relationship between the pressure of a vacuum chamber and loss rate of atoms from a magnetic trap in that chamber. Using the universality of loss rates with respect to trap depth, we can use this apparatus to measure the total collision cross section between the trapped rubidium and any species of collision partner. We can also take advantage of the flexible design of its laser system to quickly alternate measurements between ^{85}Rb and ^{87}Rb to

confirm the universal behaviour of both isotopes. We plan to take the apparatus to Physikalisch-Technische Bundesanstalt in the coming months to compare the cold-atom based standard against their orifice flow standard.

Lay Summary

Physicists need standards to translate units into something they can measure and calibrate their gauges with. Existing vacuum pressure standards are large, immovable devices that produce a known pressure of nitrogen or argon, but cannot calibrate pressure gauges for other background species. Worse, existing vacuum pressure gauges can lose calibration over time.

We can instead use the loss rate of atoms from a magnetic trap to measure pressure. Because atoms are immutable and do not drift with time, they can be used to create a pressure standard—one that works for all background species. However, existing atom-based pressure standards are still large and immovable.

This work describes our efforts to produce a portable version of a cold atom-based pressure standard that can be transported to other pressure standards to directly compare readings with. We have completed preliminary tests of the apparatus, and are working on its final assembly.

Preface

All the work presented in this thesis was conducted on at the Quantum Degenerate Gas (QDG) laboratory at University of British Columbia, Vancouver campus and at the MAKE+ lab at British Columbia Institute of Technology, Burnaby campus under the supervision of Professor Madison. The project to create a portable cold-atom vacuum pressure standard was started by Dr. James Booth under the Prometheus CFI grant¹, but I have designed or modified every part of the apparatus. Therefore, I am responsible for the work presented in this thesis unless it was specifically credited to someone else.

The overview of the theoretical background of the project in chapter 2 is based on three works. The guide to different traps (section 2.1) is heavily influenced by Harold Metcalf's Laser Cooling and Trapping of Neutral Atoms [30]. Fig 2.2 was taken from his book. The collision theory section (section 2.2) is primarily based on Chapters 2-5 of Mark Child's Molecular Collision Theory textbook [6], which I took fig. 2.3 from. This section also received heavy input and guidance from Dr. James Booth and Professor Madison. The discussion of the universality of quantum diffractive collisions (section 2.3 and section 2.4) and its implications are based on the PhD thesis of Dr. Pinrui Shen [43] and the paper that he, Dr. Booth, Professor Madison, and Professor Krems released [3]. table 2.1 was taken from their work in [3], and table 2.2 was compiled from their work in [45] and [46]. fig. 2.6, fig. 2.4, and fig. 2.5 were taken from their work in [3].

The initial vacuum chamber bakeout and assembly described in section 3.1 was conducted with Dr. Shen's assistance. The rebuilt design of the chamber (depicted

¹In a sense, it has become is closer to 'the apparatus of Theseus' than the 'apparatus of Prometheus'

in fig. 3.3) was completed in consultation with Professor Madison and Dr. Booth, and some of the assembly and leak testing of the rebuilt chamber was conducted with assistance from Dr. Booth.

Most of the work described in chapter 4 was conducted by the BCIT MAKE+ group, especially Matt Grieg. I wrote the coil simulation and optimization code. Matt, Fardin, and I assembled the coils I designed. Matt designed the coil controller with assistance from Fardin Barekat and Behnam Mohammad, and Matt and Fardin assembled the coil cooling hydraulics. They also donated many spare parts that were lying around the MAKE+ lab, including the motor that drove the pump and the water chiller, and assembled the ‘angry amp cart’ that currently houses the coil controller, hydraulics, and coils, and will eventually house the entire apparatus. Finally, Fardin provided assistance lifting the cart into the U-Haul van that I used to transport it to UBC for final testing and assembly. I designed the magnetic field coil mounts with some input from Mladen Bumbulovic of the UBC PHAS Technical Services, and Garrett Kryt of the BCIT MAKE+ group 3D printed them.

Dr. Booth originally designed the optics for the project (not discussed in chapter 5). Although the original design did not produce enough power and was prone to misalignment, the apparatus still uses the components he purchased, including the Vescent laser controllers and lock boxes. table 5.2 was adapted from a table in the Vescent D2-135 user manual fig. 5.5. Daniel A. Steck’s documents compiling the relevant theory and numbers for trapping rubidium-85 and rubidium-87 were of great value, and I used his work to calculate the frequencies used in the control code for the apparatus. Fig 5.1 is a modified version of figure 2 in his work in ref [48]. The Vescent wikis were also a useful resource for understanding how the saturated adsorption and offset locks work. The images showing the interiors of the Vescent D2-210 (fig. 5.3) and Vescent D2-250 (fig. 5.4), as well as the Vescent D2-135 lock box (fig. 5.5) were taken from their wiki. Also, fig. 5.18a, fig. 5.18b, and fig. 5.19 were taken from the ColdQuanta PICAS user manual [8].

The analog outputs, digital outputs, and direct digital synthesizers (section 6.1.2) used in this apparatus were designed by Dr. Todd Meyrath and fabricated by the BCIT MAKE+ group. The QDGBus FPGA controller (section 6.1.1) was designed by Wenjun Wu, Tristan Calderbank, Shi Jing (Koko) Yu, and James Liu. The shelf of electronics was assembled by the BCIT MAKE+ group, although I modified it

by removing excess amplifiers, installing the microwave electronics, and repairing/replacing broken components.

The magnetic trap measurement sequence described in section 7.2.2 and section 7.2.3 is heavily informed by the sequence used by Riley Stewart on the MAT apparatus in the QDG lab.

I would like to acknowledge CMC Microsystems and Canada's National Design Network (CNDN) for providing access to SolidWorks and Altium. I received financial support from the Patrick David Campbell Graduate Fellowship and the Arvind Gupta and Michelle Pereira Graduate Scholarship. I also acknowledge financial support from the Natural Sciences and Engineering Research Council of Canada (NSERC/CRSNG) and the Canadian Foundation for Innovation (CFI). This work was done at the Center for Research on Ultra-Cold Systems (CRUCS).

Table of Contents

Abstract	iii
Lay Summary	v
Preface	vi
Table of Contents	ix
List of Tables	xiv
List of Figures	xv
Acknowledgments	xxv
I Background	1
1 Introduction	2
1.1 Review of Existing UHV Pressure Standards	2
1.2 Motivation	5
1.3 Design Requirements	6
1.4 Thesis Overview	6
2 Background	8
2.1 Atom Traps	8
2.1.1 Magneto-Optical Traps	8
2.1.2 Magnetic Traps	12

2.2	Collision Theory	12
2.2.1	Classical Collisions	12
2.2.2	Quantum Collisions	15
2.2.3	Collisions with a Background Gas	16
2.3	Pressure Measurement using Cold Atoms	18
2.4	Quantum Diffractive Universality	19
2.5	Refinements	21
2.5.1	Other Loss Mechanisms	21
2.5.2	Trap Anisotropy and RF Knife	22
2.5.3	Energy Distribution Corrections	22
2.5.4	Limitations	23
2.6	Summary	25
II	Apparatus	26
3	Vacuum Chamber	27
3.1	Design	28
3.1.1	Conductance	29
3.1.2	Leak Checking	30
3.1.3	New Design	32
3.1.4	ColdQuanta Rubidium Source	34
3.2	Preparation	34
3.2.1	Leak Rate in Vacuum Chambers	34
3.2.2	Baking the Vacuum Chamber	36
3.3	Outgassing Performance	43
4	Magnetic Field Coils	47
4.1	Theory	47
4.1.1	Thermal Stability	48
4.1.2	Types of Coil Construction	49
4.2	Optimization of Design	49
4.3	Construction of Coils	51
4.3.1	Measured Coil Parameters	53

4.4	Coil Mounts	56
4.4.1	Design Goals	56
4.4.2	Design Overview	56
4.4.3	Performance	57
4.5	Magnetic Field Coil Driver	59
4.5.1	Theory of Operation	59
4.5.2	Water Cooling	60
4.5.3	User Controls	63
4.5.4	Error Conditions	69
4.5.5	Limitations	69
4.5.6	Performance	73
5	Optical Design	77
5.1	Locking Optics	77
5.1.1	Pumping Scheme	77
5.1.2	Specifications	77
5.1.3	Design	80
5.1.4	Performance	89
5.2	Vacuum Optics	100
5.2.1	2D MOT Optics	100
5.2.2	3D MOT Optics	100
6	Electronics	104
6.1	Pre-existing Electronics	104
6.1.1	QDGBus	104
6.1.2	QDGBus devices	105
6.2	Microwave Electronics	107
6.2.1	Microwave Control Board	107
6.2.2	Microwave Performance	114
6.3	Intensity Stabilization	115
6.3.1	Control algorithm	116
6.3.2	Performance	118
6.3.3	Conclusions	122

7	Software and Control	124
7.1	Software structure	124
7.1.1	AOM Throughput Optimization	125
7.2	Measurement Sequences	127
7.2.1	MOT Pressure Measurement	127
7.2.2	MAT Pressure Measurement	129
7.2.3	MT Pressure Measurement	131
8	Assembly	133
8.1	Carts	133
8.2	Light Shielding	135
8.2.1	Vacuum Optics	135
8.2.2	Locking Optics	136
8.3	Cart Walls	136
III	Results	139
9	Results	140
9.1	Analysis	141
9.1.1	MOT Analysis	141
9.2	Measurements	143
9.2.1	Optimal Push Beam Detuning	143
9.2.2	Dispenser Performance	145
9.2.3	Magnetic Trapping	153
10	Conclusions	155
	Bibliography	157
A	Magnetic Field Coil Modelling	164
A.1	Inductance	164
A.2	Fluid Flow	165
A.2.1	Flow Rates	165
A.2.2	Helical Corrections	166

A.3	Cooling Theory	167
A.3.1	Non-liquid Cooling	168
A.3.2	Liquid Cooling	169
A.4	Cooling Models	170
A.4.1	Hollow-core Coil Models	170

List of Tables

Table 2.1	The universal coefficients in eq. (2.23) calculated by fitting to the quantum scattering calculations. Each parenthetical number is the estimated error in each coefficient. Taken from ref [3]. . .	21
Table 2.2	Originally calculated collision cross sections (calculated using eq. (2.23)) and corrected collision cross sections (calculated using eq. (2.25)). The first number in parentheses is the statistical error, and the second is the systematic error. All cross sections are taken from [45] except for $_{87}\text{Rb-H}_2$, which is taken from [46].	23
Table 3.1	Rebuilt vacuum chamber components	33
Table 4.1	Simulated and measured parameters of the constructed coils. .	53
Table 4.2	Dynamic (while loaded) and Permanent (after loaded) displacement of the upper coil when a force was applied along its strong side (in line with the 1.5” support posts) and weak side (perpendicular to its 1.5” support posts). Displacement of the lower coil was significantly less as the mounts tended to torque about the support posts.	59
Table 4.3	Parts used in the coil cooling hydraulics.	61
Table 5.1	The primary components used in the locking and distribution optics.	82
Table 5.2	Offset Frequency Lock Ranges of the D2-135 [57]	88

List of Figures

Figure 1.1	A high-level schematic of a basic orifice flow standard.	3
Figure 2.1	A prototypical schematic of a MOT. The blue arrows represent the direction of the magnetic field generated by the magnetic field coils. The red cylinders represent the laser light, and the σ^+ and σ^- indicate the polarization of the incoming beams.	9
Figure 2.2	Energy splitting in a MOT due to the Zeeman effect (taken from ref [30], fig 11.4)	10
Figure 2.3	A diagram of a classical elastic scattering event as seen from the center of mass frame. The black curve represents the path the incoming particle takes, and the origin is the center of the inter-particle potential (taken from ref [6], fig 2.1)	13
Figure 2.4	Numerically computed velocity-averaged loss cross sections for He (red triangles), Ar (black circles), and Xe (blue squares) colliding with Li (dashed lines) and Rb (solid lines), taken from [3] [fig 2b]	20
Figure 2.5	Experimental data for trapped Rb and background He, Ar, Xe, H ₂ , N ₂ , and CO ₂ fitting to the same universal curve [3] [fig 3f]	20
Figure 2.6	Calculated cross sections as a function of total angular momentum for several different potential energy surfaces. Notice the initial oscillatory behaviour followed by the monotonic decrease. Taken from ref [3].	24
Figure 3.1	The original design of the vacuum chamber.	28

Figure 3.2	A set of MOT loss rates (Γ) taken using the original vacuum chamber (fig. 3.1). MOT loss rates provide a useful proxy for the pressure in the chamber. The orange dots were fitted using eq. (9.2) to extract the loss rate, while the blue dots were fitted using eq. (9.4) to extract the loss rate. The chamber has a turbo pump connected through the inline valve, and it can also be pumped by opening the NEG valve. At 17:22, the inline valve to the turbo pump was closed, causing the pressure to slowly rise. At 23:05, the valve to the NEG was closed, causing a sharp spike in the pressure. At 23:10, the valve to the turbo was re-opened, causing the pressure to fall to its original value.	31
Figure 3.3	The final revision of the vacuum chamber used in the apparatus	33
Figure 3.4	We baked this chamber with the source cell and science cell replaced with CF133 blanks. It was baked at 300 C for 55 days. The inline valve, NEG tube, zero-length reducer, and source cell adapter were re-used in the rebuilt vacuum chamber. . . .	38
Figure 3.5	Manually recorded temperatures of different thermocouples while conducting the first bake.	39
Figure 3.6	Pressure reported by the ion gauge while conducting the first bake.	39
Figure 3.7	The collection of spare parts baked with the intention of using some of them in an interconnect between the apparatus chamber and the chamber under test. The ion pump, leftmost CF275 angle valve, blanks, and uppermost elbow were used in the rebuilt chamber. We plan to use the 12" bellows for the interconnect. They were baked at ~ 400 C for 37 days. . . .	41
Figure 3.8	The temperatures in the chamber during the second bake. The thermocouples attached to the ion pump and 4 cross came loose and produced inaccurate readings for significant portions of the bake.	42

Figure 3.9	The pressure recorded by the ion gauge near the scroll pump during the second bake. We discovered that selecting a portion of text in a Windows shell pauses the command the shell is running. This inadvertently caused the script that was collecting the pressure data to freeze for much of the bake.	43
Figure 3.10	The temperature in the chamber during the third bake. While we were able to monitor the pressure on a dial, our pressure logging device encountered an issue and did not reproduce the readings on the dial.	44
Figure 3.11	The temperature of the cube, bellows connecting it to the turbo pump, ion pump, and turbo pump during the fourth bake. The temperature was increased slightly for the last week of the bake to increase the amount of removed hydrogen. Fluctuations in the temperature were due to fluctuations in the mains voltage.	44
Figure 3.12	Pressure recorded with the ion pump during the fourth bake. The slight spike in temperature on August 3rd was due to an increase in the temperature. An impending power shutdown in the building cut the bake short. We decided to not continue the bake because we expected that other vacuum parts baked at lower temperatures would outgas significantly more than the cube.	45
Figure 3.13	Recorded temperatures of the rebuilt vacuum chamber during the fifth bake. The five-cross' temperature was increased to 400 C to remove as much hydrogen as possible from it.	46
Figure 3.14	Recorded pressure in the rebuilt vacuum chamber during the fifth bake. The spike in pressure on September 26th was from activating the Rb dispenser.	46
Figure 4.1	Photographs from the coil assembly process.	54
Figure 4.2	Axial magnetic field at different axial distances from the shaved surface of each of the coils. Surprisingly, the less shaved coil produced a higher gradient. It's possible the more shaved coil has a short-circuit between some of its winds.	55

Figure 4.3	A collage of images of the coil mounts	58
Figure 4.4	A schematic of the hydraulics used to cool the magnetic field coils. Note that water flows counter-clockwise, following the direction of the arrows.	61
Figure 4.5	The vibration damping mounts the pump sits on.	62
Figure 4.6	A demonstration of the effects of the different stages of damping on the vibrations produced by the chiller.	64
Figure 4.7	A demonstration of the effects of the different stages of damping on the vibrations produced by the pump.	65
Figure 4.8	Acceleration on the floor of the room next to the cart with and without the pump and chiller turned on. Any vibrations transmitted to the floor are below the noise floor of the accelerometer we used.	66
Figure 4.9	Accelerations on the vacuum optics table and the cart when the cart is hit by a screwdriver.	67
Figure 4.10	An example image of the touchscreen controlling the coil driver.	68
Figure 4.11	The water filter with the red purge button.	70
Figure 4.12	The control panel for the coil controller and hydraulics. It houses the touchscreen, the pump controls, the BNC for remote control, and the high pressure gauge.	71
Figure 4.13	The underpressure valve with its control collar (the component with the white label).	72
Figure 4.14	The step response of the coil driver for different set currents. See fig. 4.16 for the long-term behaviour.	74
Figure 4.15	A view of the initial current ramps displayed in fig. 4.14. . . .	75
Figure 4.16	A long-term view of the current stability of the coil driver. The noise is likely primarily measurement noise, but the drift is due to flaws in the controller design. The settling time is longest at 65A because the controller is at the top of its range. It is also why the current never reaches the set current.	76
Figure 5.1	Hyperfine structure of Rubidium-87 D2 line (taken from ref [48], fig 2) and annotated with the pump and repump transitions.	78

Figure 5.2	Above is a diagram detailing outlining how the light is generated and transported to the 2D and 3D MOTs. Part numbers and abbreviations are in table 5.1	81
Figure 5.3	This illustrates what is inside of the Vescent D2-210 Saturated Adsorption Spectroscopy module (taken from [58]).	84
Figure 5.4	This illustrates what is inside of the Vescent D2-250 Heterodyne module (taken from [59]).	86
Figure 5.5	A block diagram for the offset lock (taken from [57])	87
Figure 5.6	Output frequency as measured from the D2-135 Beat Note Div 2 output vs the set frequency assuming the VCO frequency is linear across its input range. The calibration data used are stored here: https://qdg-code.phas.ubc.ca:2633/Perrin/PrometheusPython/src/branch/master/hardware/offsetdata	89
Figure 5.7	The spectrum of the reference laser. This was taken by beating the light from the reference laser against light from another locked laser in the lab, so the actual linewidth of the laser is ~ 3 MHz. The sidebands are caused by the 4 MHz dither used to create the lock signal. The sideband amplitude can be decreased at the expense of reducing the lock signal amplitude, and thus the lock stability.	90
Figure 5.8	The spectrum of the power laser. This was taken by beating the light from the reference laser against light from another locked laser in the lab, so the actual linewidth of the laser is < 3 MHz. The offset lock locks to the average frequency of the beat note, so it removes the sidebands that are present in the reference laser spectrum (seen in fig. 5.7). Technically, the fit should be a Voigt profile (the convolution of the gaussian and lorentz functions), but a gaussian fit suffices for extracting the FWHM.	91

Figure 5.9	The spectrum of the repump light (one of the sidebands generated by the EOM) beat against the same laser as fig. 5.8. The microwave frequency FWHM is ~ 35 kHz. Assuming the phase noise of the EOM is negligible, we expect spectrum of the repump and pump light to be nearly identical since the repump light is in the same field as the pump light but is frequency-shifted by the EOM. This spectrum confirms that the EOM does not add significant phase noise. Technically, the fit should be a Voigt profile, but a gaussian fit suffices for extracting the FWHM.	92
Figure 5.10	MOT loading curves without loading rate fluctuations.	93
Figure 5.11	Examples of fluctuating loading rates. The orange line is fit to the initial slope.	94
Figure 5.12	Fluctuations in the power in the two arms of the 2D MOT splitter prior to adding any polarization cleaning optics. Notice how the net power through the arms is roughly constant but the splitting is highly variable.	95
Figure 5.13	Fluctuations in the polarization in the white arm of the 2D MOT splitter prior to adding any polarization cleaning optics. Behaviour in the red arm was similar.	96
Figure 5.14	An image of the Thorlabs patch cables used on the apparatus, taken from https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=3345 . Light can be polarized along the fast axis (perpendicular to the connector key) or slow axis (parallel to the connector key). We used FC/APC (Fiber Connector/Angled Physical Contact) connectors to reduce back-reflections.	96
Figure 5.15	Polarization fluctuations in the white arm after adding a free-space GT in front of the 2D MOT splitter.	97
Figure 5.16	Power fluctuations between the red and white arms after adding a free-space GT in front of the 2D MOT splitter.	98
Figure 5.17	Polarization and power fluctuations after adding an inline polarizing fiber.	99
Figure 5.18	Coldquanta PICAS components (images taken from Ref [8]) .	101

Figure 5.19	Magnetic field generated from the permanent magnets on the PICAS (image taken from user manual)	101
Figure 5.20	Layout of the 3D MOT optics. The light shield is transparent in the model but is opaque in real-life.	102
Figure 5.21	Layout of the vacuum chamber, coils, and optics.	103
Figure 6.1	A schematic showing the wiring of all the control electronics connected to QDGBus	105
Figure 6.2	The DDS could only offer a limited tuning range before the microwave source railed out	108
Figure 6.3	Schematic for the input stage of the microwave controller. . .	109
Figure 6.4	Schematic for the FIFO stage of the microwave controller. . .	110
Figure 6.5	Schematic for the output stage of the microwave controller. . .	111
Figure 6.6	An example of the ADF5356's 'exciting' transfer function . .	112
Figure 6.7	The transfer function for the XM-A3B4-0404C-01 microwave voltage-variable attenuator.	112
Figure 6.8	Spectrum of the amplified microwave source signal before it is fed into the EOM.	115
Figure 6.9	Long timescale fluctuations in the unstabilized 3D MOT power. .	116
Figure 6.10	Step response of intensity feedback controller	119
Figure 6.11	Open loop performance compared to closed loop performance over short timescales.	120
Figure 6.12	Open loop performance compared to closed loop performance over long timescales.	121
Figure 6.13	Open loop performance compared to closed loop performance when a triangle wave disturbance was added.	123

Figure 7.1	A block diagram showing (at a high level) the basic components needed for the apparatus to operate. Methods have been renamed and helper methods have been removed for clarity. Other modules will be added in the future (such as a trigger for the DAQ, auto-locking by operating the reference laser servo and two laser controllers, and serial control of the TA and feedback controllers).	125
Figure 7.2	A grid of AOM throughput efficiency at different DDS amplitudes and frequencies	127
Figure 7.3	Timing of different amplitudes and frequencies in order to measure pressure using a MOT	128
Figure 7.4	Timing of different amplitudes and frequencies in order to measure pressure using Prometheus optics and electronics with a different experiment in the lab.	130
Figure 7.5	Timing of different amplitudes and frequencies in order to measure pressure using a magnetic trap.	132
Figure 8.1	A mockup of where all of the components in Prometheus can fit on the new cart. The lower two shelves house the coil controller and the hydraulics (not depicted). The acrylic and metal panels in the vacuum optics cover have been made transparent to show the internals - the panels are actually opaque.	134
Figure 8.2	Current layout of the locking optics. The fiber optics are spread through two levels - the lower level houses the pre-TA optics and the upper level houses the post-TA optics.	137
Figure 8.3	Proposed layout of the locking optics. Support structures have been suppressed and acrylic has been made transparent to allow a clearer view of the optics within.	138
Figure 9.1	Magnetic field vs axial position of the temporary coils when driven at 1 A	141

Figure 9.2	Load rate (in units of V/s) for different push beam parameters (push power is in μW and push detuning is in MHz from resonance with the pump transition). The MOT loads for 10 s, the pump light detuning is 12 MHz, and the load rate is in units of “normalized photodiode voltage” as discussed in eq. (9.6). . . .	143
Figure 9.3	Load rate (in units of V/s) for different push beam parameters (push power is in μW and push detuning is in MHz from resonance with the pump transition). The MOT loads for 10 s, the pump light detuning is 15 MHz, and the load rate is in units of “normalized photodiode voltage” as discussed in eq. (9.6). . . .	144
Figure 9.4	Load rate (in units of V/s) for different push beam parameters (push power is in μW and push detuning is in MHz from resonance with the pump transition). The MOT loads for 10 s, the pump light detuning is 18 MHz, and the load rate is in units of “normalized photodiode voltage” as discussed in eq. (9.6). . . .	145
Figure 9.5	Comparison of load rates (in units of V/s) across different dispenser currents with the push beam on compared to off. . . .	146
Figure 9.6	Comparison of loss rates (in units of Hz) at different dispenser currents. Re-analyzing the data revealed that fitting errors caused the apparent increase in the decay rate near the maximum load rate. Efforts to improve the loss rate fitting are ongoing. . . .	148
Figure 9.7	Comparison of well-fitted and poorly-fitted MOT decay curves. Points that are in red are ignored due to MOT saturation while points in blue were used to generate the fit. Curves such as the one on the right were manually rejected when re-analyzing the MOT decay curve data.	149
Figure 9.8	Loading rate (V/s) and fitted MOT loss rate (Hz) over time with the dispenser turned off and measurements being taken continuously. It is clear that after 10-20 minutes, the fitted loss rates cannot be trusted.	151
Figure 9.9	Loading rate (V/s) and MOT loss rate (Hz) over time with the dispenser turned off and measurements being taken every five minutes.	152

Figure 9.10	Comparison of loading amplitudes (as a proxy for loading rate) over time when measuring continuously as compared to measuring infrequently. It illustrates how taking frequent measurements does not significantly deplete the supply of free rubidium in the source cell.	152
Figure 9.11	A sample curve from the magnetic trap measurement	153
Figure 9.12	A preliminary decay curve generated from a series of magnetic trap measurements. The fitting code is not fully functional so the fit parameters were estimated by hand.	154

Acknowledgments

There are too many people to thank, and quite literally not enough time to thank them—this thesis was written in under six weeks while simultaneously completing as much reassembly of the apparatus as possible. Many sections, including this, did not get the treatment they deserve,

First of all, thank you to Megan, my ex-girlfriend (although she hates being called that. She prefers ‘wife’). You have provided me with encouragement and support (both emotional and financial) every step of the way. I love you, and I look forward to being able to spend more of my life with you (and Banjo) now that I don’t have a thesis hanging over my head.

Thank you Kirk! You given me the opportunity to work on this project, and offered me free reign (possibly too much) to pursue side-projects. Some made it into this work, such as the microcontroller intensity stabilizers, but many others did not, such as the peltier-based water chiller. I appreciate your support through this extended master’s degree, and hope to provide you with a fully-functioning apparatus before the new year. I also will be eternally grateful for your non-academic advice and support when I encountered some challenges in my personal life throughout this degree.

Thank you Jim! Your good humour was always appreciated, and I am grateful for all of your help—for staying late to reassemble the vacuum chamber while protecting the rubidium source, reviewing this thesis at short notice, supplying the funds to purchase parts when Kirk’s money ran dry, and countless other things. I also appreciate that you splurged on the Vescent lockboxes, which are infinitely easier to lock using than the master lasers in the lab. I hope you and Patrick enjoy your retirement!

Thank-you to all of the people at the BCIT MAKE+ group who helped build the apparatus (including the original electronics, the coils, the coil power supply, the hydraulics, and the cart that everything sits on)! I am especially grateful for Matt Greig's work, but also the contributions of Fardin Barekat, Behnam Mohammad, and Garrett Kryt.

Thank you to everyone in Kirk's lab! I am especially grateful for the help of Pinrui Shen, who was very involved in the design and assembly of the apparatus before he graduated. I also am grateful for all of Denis Uhland's help with aligning my optics—I grew better at it as time went on, but I will never match his prowess. I look forward to (hopefully) working with him to set the apparatus up at PTB. I also want to thank Erik Frieling, Riley Stewart, Avinash Deshmukh, and Sudhang Varshney for their help at various points in this project.

There are a couple of other individuals in the ChemPhys basement I could not have done this without. Thank you Ian for your advice on vacuum systems and other mechanical engineering problems. I also appreciate hearing your strong opinions regarding urban planning and design. Thank you to Pavle for lending me equipment (including but not limited to a fiber splicer, a leak tester, a ion gauge and controller, KF valves, a scroll pump, and various other pieces of equipment), and refraining from physical violence when I broke the leak tester. I appreciate the patience of everyone in the basement with the apocalyptic mess I created while building this apparatus, and I promise to clean it up before I leave for Germany.

Thank you to the Physics machine shop crew (especially Nick, Mark, Dan, and Matt) for their advice and assistance with miscellaneous problems, from re-welding thermocouples to removing stuck gaskets from Conflat flanges. I also appreciate the advice of the Physics electronics lab regarding the PCBs developed for this project. A special thank-you to Janelle van Dongen, who has helped this project both by conducting some of the preliminary research as part of Kirk's group, and by helping me find miscellaneous equipment around the department as a member of it's technical services team.

Thank you for the support of all of the friends I made along the way, with a special thank-you to Ashley Warner for being a fantastic roommate for Megan and I.

And of course, thank you to my parents Tonya and Henry! Thank you for rais-

ing me to be curious and fostering an interest in science. Thank you for tolerating my intensely antisocial behaviour during exams, and for always trying to understand what my research is about. More so than everyone else, I wouldn't be here without you!

Finally, thank you to you the reader! I hope this makes sense and helps you in some way. If you have any questions² I will be happy to answer them—ask Dr. Madison how best to reach out to me!

²provided I haven't died from electrocution or in a canyoning accident

Part I

Background

Chapter 1

Introduction

1.1 Review of Existing UHV Pressure Standards

SI units are now defined in terms of physical constants, meaning that they are immutable quantities¹. The speed of light in a vacuum is fixed at 299792458 meters per second. The ground-state hyperfine transition frequency of caesium-133 is fixed at $9\,192\,631\,770\,\text{s}^{-1}$. Therefore, measuring ‘base’ units such as the second or meter is straightforward if technically challenging [21].

Measuring ‘derived’ units that are composed of some combination of base units can require more work. Metrologists build standards to link derived units back to the base units, which can be found by measuring the physical constants. These standards can be used to calibrate gauges that may be more responsive, cheaper, more portable, or in some other way more convenient to use.

Pressure is a derived unit measured in units of force per unit area, or mass per length per time squared. For roughly human-scale pressures (hundreds of pascals to hundreds of megapascals), metrologists use piston gauges. Piston gauges compare the measured pressure to a precisely calibrated weight in a precisely calibrated area ($P = mg/A$), moving the burden of calibration to easier-to-measure units. The pressure range of piston gauges can be extended with the use of a mercury manometer,

¹ Although some individuals have speculated about the possibility of time-variance in fundamental constants [41], the general consensus is that fundamental constants should be treated as constant unless measured otherwise.




Figure 1.1: A high-level schematic of a basic orifice flow standard.

but are still limited to measuring pressures above 1 Pa [26].

Lower pressures can be created using a static expansion standard. It uses a mercury manometer to measure the pressure in a small volume, then adiabatically increases the volume to create an inversely proportionally small pressure ($P_{\text{final}} = \frac{V_{\text{initial}}}{V_{\text{final}}} P_{\text{initial}}$). These can create pressures down to 10^{-4} Pa to 10^{-5} Pa (roughly 10^{-6} Torr to 10^{-7} Torr; 1 Torr = 133.322 Pa), although recent efforts have extended that range to 10^{-6} Pa [29]. This calibrated pressure can be used to calibrate a spinning rotor gauge, which measures pressure using the viscous drag on a magnetically-levitated spinning ball [15].

However, creating and measuring pressures in the ultra-high vacuum (UHV) (10^{-4} Pa to 10^{-8} Pa or 10^{-6} Torr to 10^{-10} Torr) is more challenging. Typically, an orifice flow standard (also known as a continuous expansion standard) is used

to create a known pressure in a known gas, which is used to calibrate a UHV gauge such as an ion gauge [37]. An orifice flow standard works by using a flowmeter to produce a calibrated flow of gas Q into a chamber (chamber 1), which has a calibrated medium-vacuum gauge (such as a spinning rotor gauge) and the gauge under calibration attached to it. The first chamber is attached to a second chamber via an orifice of known conductance C . The second chamber is pumped out by a UHV pump such as a turbomolecular pump. See fig. 1.1 for a schematic of the system.

Initially, a large flow Q_h is introduced to the first chamber so that the medium-vacuum gauge can measure P_1 . $Q = C(P_1 - P_2)$, so P_2 can be calculated. Provided the gas is not ‘sticky’ or reactive, the system is in steady-state, and the pressure is still within the molecular flow regime (below 1×10^{-3} Torr), the ratio of pressures $r = P_2/P_1 = (1 - \frac{Q_h}{CP_1})$ will always be constant [22]. The flow Q can be then reduced to Q_l , changing P_1 but preserving r . The new pressure in the first chamber can be calculated using $P_1 = \frac{Q_l}{C(1-r)}$, bringing it within the measurement range of the UHV gauge. This known pressure can be used to calibrate the UHV gauge [37].

One of the most commonly used gauges for the UHV regime is the ion gauge. Ion gauges measure pressure by ionizing the gas under measurement then detecting the new-created ions via the current they induce in a filament. However, ion gauges have several drawbacks:

- Ion gauges have a minimum measurable pressure (usually 2×10^{-11} Torr) due to x-rays produced during the measurement process [12].
- Ion gauges change the local pressure by ionizing the gas under measurement (removing it from the chamber) and by heating the local chamber (which increases the pressure both increasing the local outgassing rate and due to the ideal gas law [$P \propto T$]). [12]
- The calibration of the ion gauge can drift over time due to physical changes in the geometry of its mechanical components and/or chemical changes to the emission filament due to exposure to reactive gases [2, 18].
- They require calibration, which currently is only done with orifice flow systems. The orifice flow technique only works for inert, non-sticky gases—

only Argon and Nitrogen calibrations are routinely offered for ion gauges. Thus, the gauge readings for other species in a vacuum system are inaccurate [2].

In 2009, a different method of sensing UHV pressure was proposed: recording the loss rate of atoms from a trap due to collisions with the background gas. This loss rate is proportional to the pressure, and the proportionality constant hinges on the velocity-averaged collision cross section between the sensor atoms and the background particles [12]. These cross sections can be difficult to calculate *ab initio*, and measuring them by calibration against an orifice flow system only works for two species (N_2 and Ar), neither of which are the dominant background in UHV (typically H_2 or H_2O) [22].

In 2019, it was demonstrated that the velocity-averaged loss cross section is a universal function of the trap depth, allowing us to measure the velocity-average collision cross section directly for many background gas species [3]. Therefore, the cold atom approach offers several distinct advantages:

- Atoms are immutable objects with no manufacturing errors, so this technique requires no calibration.
- The measurement can be performed at room temperature. If a ‘sticky’ sensor atom is used, it minimizes the effect on the pressure of the chamber under test.
- This technique can be used for any type of background gas, not just the nonreactive gases that must be used in an orifice flow system.
- The universality provides a method of easily measuring the collision cross section between the sensor atoms and new background species.

1.2 Motivation

All published orifice flow standards and cold atom standards are large, immovable systems, so no direct comparison between them has been made. While an indirect comparison of this standard with an orifice flow system has been performed by using a calibrated ion gauge, it is possible the gauge lost calibration during the

shipment or installation [3]. We hope to remove any doubt of the validity of the technique by performing a direct comparison with an existing pressure standard. As such, our goal is to build a portable version of the cold atom standard, then transport it to different metrology groups around the world to perform a direct comparison with their standards.

A secondary goal is to test the validity and limits of the universality. Comparing with an orifice flow standard will allow us to validate our self-calibrated rubidium cross sections by comparing the pressure we measure using them with the pressure produced with an orifice flow standard. We do not yet know the limits of calculations, so proving that the universality approach is reliable could help guide the development of theoretical and computational collision models.

1.3 Design Requirements

Specifically, we need an apparatus that can

- produce a magneto-optical trap and magnetic trap (see section 2.1 for a description of such traps).
- measure loss rates of rubidium atoms from a magnetic trap.
- connect to other vacuum chambers to measure their pressure.
- contaminate other vacuum chambers as little as possible.
- be transportable without needing a lengthy reassembly stage at the destination.

Ideally, any calibration and realignment could be achieved within the time it would take to bake out the connection between the apparatus and the system under test (under a week).

1.4 Thesis Overview

This thesis discusses the progress in realizing a portable version of a cold-atom standard. Chapter 2 discusses the atomic traps, scattering theory, and the universality of quantum diffractive collisions that the apparatus will use to measure pressure

for a wide variety of background species. Chapters 3-7 discuss the technical details of how the apparatus was designed and constructed, as well as relevant calibration data for each of the components that compose the apparatus. All of the components are inter-connected, but can be roughly divided into five major subsystems, each of which have a chapter dedicated to them:

- The vacuum chamber that is connected to other chambers to measure their pressure.
- The magnetic field coils, their driver, and their cooling system.
- The lasers and optics that produce the necessary frequencies to control the sensor atoms,
- The electronics that control the light and magnetic field driver.
- The software that controls the electronics.

Chapter 8 discusses the remaining work that needs to be done—primarily assembling all of the subsystems together into a cohesive whole.

The apparatus has completed one series of full-system tests. It needed to use some temporary components that limited its performance, but offered us some valuable data characterizing how its atomic source performs and what parameters are optimal. These tests are discussed in Chapter 9. Finally, all of the above is summarized in Chapter 10.

Chapter 2

Background

This chapter gives some background on the quantum mechanics behind how the apparatus works. It starts with a description of atomic traps, then discusses some basic quantum collision theory, and ends with a description of pressure measurement using atomic traps, including a new form of universality that allows us to measure velocity-averaged collision cross sections.

2.1 Atom Traps

Two common ways of confining cold atoms are the magneto-optic trap (MOT) and the magnetic quadrupole trap (MT) [30]. MOTs can both cool (collect) and confine atoms, while all other traps (e.g. magnetic traps, electric field traps, optical dipole traps) are conservative and cannot cool their trapped atoms. The best a conservative trap can do is selectively retain its lowest-energy atoms, bringing the net temperature of the trapped ensemble down [27].

2.1.1 Magneto-Optical Traps

Magneto-optical traps are typically used to confine atoms and cool them down to millikelvin (mK) temperatures. To do this, MOTs use permanent magnets or magnetic field coils to create a magnetic field with a single extrema in the center, and three pairs of orthogonal, counter-propagating circularly polarized laser beams that coincide at the minimum of the magnetic field (see fig. 2.1). The light from




Figure 2.1: A prototypical schematic of a MOT. The blue arrows represent the direction of the magnetic field generated by the magnetic field coils. The red cylinders represent the laser light, and the σ^+ and σ^- indicate the polarization of the incoming beams.

the lasers is detuned by δ below some atomic transition frequency (a.k.a. ‘red-detuned’ light). The spatially-varying magnetic field induces a spatially-varying Zeeman splitting in the trapped atoms. The counterpropagating beams serve two purposes; they create a damping force that slows the atoms down, and work with the magnetic field to produce a spatial confining force [30]¹.

Suppose an atom is travelling through two equal-intensity counterpropagating beams of monochromatic light tuned below some atomic resonance (red-detuned). The atom will have a velocity component parallel to one beam and antiparallel to the other. The doppler effect will shift the antiparallel light closer to resonance and the parallel beam further from resonance, so the atom will preferentially absorb photons from the antiparallel beam. The spontaneous emission from the atom had no preferred direction, so the net effect of the momentum transfer from the photons

¹See sec. 11.4 of ref [30] for a more thorough description.




Figure 2.2: Energy splitting in a MOT due to the Zeeman effect (taken from ref [30], fig 11.4)

is a force decelerating the atom in the direction of the antiparallel beam.

If there are counterpropagating beams in three dimensions, then the atom will always preferentially experience a net force antiparallel to its direction of travel. The net effect is a damping force on the atom known as an ‘optical molasses’ (see sec. 7.2 of ref [30] for a more detailed explanation).

The optical molasses slows atoms down, but a position-dependent force is necessary to confine atoms in a specific location. The Zeeman effect causes the energy levels to separate according to eq. (2.1):

$$\Delta E = -\vec{\mu} \cdot \vec{B} = -\mu_B g_F m_F B \quad (\text{eq 2.19 in Ref [16]}) \quad (2.1)$$

where B is magnitude of the magnetic field, $\vec{\mu}$ is the magnetic field moment, g_f is the Lande g-factor of the atom, and m_F is the magnetic quantum number

If the magnetic field is generated from a pair of anti-helmholtz coils, then the magnetic field is described with eq. (4.3). Inserting this magnetic field into eq. (2.1) produces

$$\Delta E = \mu_B g_F m_F \left(\frac{dB}{dz} \Big|_{z=0} \right) \frac{\sqrt{x^2 + y^2 + 4z^2}}{2} \quad (2.2)$$

where $\left(\frac{dB}{dz} \Big|_{z=0} \right)$ is the magnitude of the magnetic field gradient in the z (axial) direction.

If the counterpropagating beams have opposite circular polarizations, then the σ_+ light will be able to drive the $\Delta F = +1, \Delta m_F = +1$ electric dipole transition, and the σ_- light drive the $\Delta F = +1, \Delta m_F = -1$ transition. This means that the atom will always ‘see’ the light source to the side of the magnetic field zero it is on as more resonant, so radiation pressure will always push it towards the magnetic field zero (see equation 11.4 of [30] for the specific force equations).

Magneto-optic traps are excellent tools, but have several problems:

- The recoil from atoms re-emitting light puts a lower bound on the temperatures achieved by atoms in a MOT. This bound is called the “Doppler Limit”, and is generally within an order of magnitude of $300 \mu\text{K}$ [30]². This is much too warm for some applications (such as creating a Bose-Einstein Condensate).
- The cycle of excitation and relaxation means that the atoms in the MOT are in a mix of electronic states, which makes it harder to extract results for a single state [17].
- The MOT trap depth is highly dependant on the trapping parameters, and is very sensitive to beam alignment. This makes the trap depth ‘unstable’—it can change day-to-day as the laser alignment changes [42, 43]. As the relationship between loss rate and collision rate is dependant on the trap depth, this variation of the trap depth with time makes measuring the decay rate from a MOT unsuitable for a pressure standard.

²The Doppler temperature of ^{87}Rb is $145.5 \mu\text{K}$ [48]

2.1.2 Magnetic Traps

A magnetic trap (MT) takes advantage of the Zeeman splitting (eq. (2.1)) to produce a trapping potential without using optical forces [30]. This apparatus uses a quadrupole magnetic trap (QMT), which produces the conservative potential shown in eq. (2.2).

The confining forces in a MT are not nearly as strong as those in a MOT, and they cannot cool their atoms, so only cold ensembles can be trapped within them. However, the depth of a MT is well-defined, which makes them an attractive choice for pressure metrology since the collision-induced loss rate of the trapped atoms will depend on the trap depth. Furthermore, if atoms are transferred to the $F=1$ ground state of rubidium-87, there is only one state they can be trapped in: the $|F=1, m_F=-1\rangle$ state. Unlike a MOT, a magnetic trap does not change the state of its trapped atoms, so this means that the trapped atoms remain in a definite state, dramatically simplifying our analysis of the experiment.

Our experiment initially confines atoms in a MOT, then transfers the atoms to a magnetic trap by turning extinguishing the light and increasing the magnetic field strength. This approach combines the trapping strength of a MOT with the well-defined depth of a magnetic trap.

2.2 Collision Theory

2.2.1 Classical Collisions

In a classical non-relativistic two-body elastic scattering event, one can replace the two colliding particles with a single reduced-mass particle passing through the inter-particle potential (described as a function of r , the inter-particle separation). One can extract original trajectories by viewing the origin as the center of mass of the two particles.

In this picture, the incoming reduced mass particle approaches on course to pass a distance b (the ‘impact parameter’) from a second particle, reaches a minimum separation a from the second particle, and is scattered away at angle Θ .




Figure 2.3: A diagram of a classical elastic scattering event as seen from the center of mass frame. The black curve represents the path the incoming particle takes, and the origin is the center of the inter-particle potential (taken from ref [6], fig 2.1)

Classically, that angle is

$$\Theta = \left| \pi - 2b \int_a^\infty \frac{dr}{r^2 \left(1 - \frac{b^2}{r^2} - \frac{V(r)}{E} \right)} \right| \quad (\text{eq 2.5 in ref [6]}) \quad (2.3)$$

where Θ (known as the deflection function) is the deflection angle relative to the initial velocity, r is the distance between the two particles, $V(r)$ is the potential between the two particles, and E is the energy of the incoming particle.

The magnitude of the deflection angle is called the ‘scattering angle’ ($\theta = |\Theta|$). One can find a by finding the smallest root of r in

$$1 - \frac{b^2}{r^2} - \frac{V(r)}{E} = 0 \quad (\text{eq 2.4 in ref [6]}) \quad (2.4)$$

In an atom trap, the atoms in the cold ensemble are colliding with the background gas particles, sampling a wide range of impact parameters and speeds for each measurement. It is impractical to try to calculate impact parameters for each

collision individually; instead, we take the ensemble view, calculating and measuring collision cross-sections. The differential cross section is flux of scattered atoms per unit solid angle, $d\Omega$, and the total elastic cross section is obtained by integrating across the entire solid angle:

$$\frac{d\sigma}{d\Omega} = \sum_b \left| \frac{b}{\sin(\theta) \frac{d\theta}{db}} \right| \quad (\text{eq 2.16 in ref [6]}) \quad (2.5)$$

$$\sigma = \int \frac{d\sigma}{d\Omega} d\Omega \quad (\text{eq 2.17 in ref [6]}) \quad (2.6)$$

where $\frac{d\sigma}{d\Omega}$ is the differential cross section, and σ is the total collision cross section. Note that eq. (2.3) must be inverted to find $b(\theta)$. The sum in eq. (2.5) is the sum over all b that correspond to a given θ .

The classical cross-section fails and becomes infinite when either $\frac{d\theta}{db}$ or θ is zero. Quantum mechanics fixes this problem and makes the cross section finite. The point when quantum mechanics becomes relevant can be estimated using the uncertainty principle, $\Delta x \Delta p \geq \frac{\hbar}{2}$ [6].

The position uncertainty is the impact parameter uncertainty (Δb), and the corresponding momentum uncertainty is the uncertainty in the transverse momentum p_{\perp} following the collision (Δp_{\perp}). Some algebra shows that

$$\Delta\theta = \frac{\Delta p_{\perp}}{p} \approx \frac{\hbar}{2\Delta b p} = \frac{\lambda}{\pi\Delta b}$$

where $\lambda = h/p$ is the de Broglie wavelength of the center of mass particle. So, for classical collision results to be valid, the scattering angle and impact parameter must satisfy

$$\theta \gg \Delta\theta \approx \theta_c = \frac{\lambda}{\pi\Delta b} \gg \frac{\lambda}{\pi b} \quad (2.7)$$

where θ_c is the critical angle below which quantum mechanical effects dominate.

Small-angle collisions where eq. (2.7) is not satisfied require a quantum mechanical description to determine the collision cross section [6]. These small-angle collisions contribute far more to the loss of atoms from shallow traps than the large-angle classical collisions, so understanding the quantum mechanical description is crucial for precise pressure measurements [3]. However, classical calculations can

still be used to understand the qualitative variation of the loss cross section of atoms at large trap depths.

2.2.2 Quantum Collisions

Provided that the potential between the trapped and background particle is spherically symmetric (such as the Coulomb, Van der Waals, and Lennard-Jones potentials), the center of mass Schroedinger equation can be rewritten as

$$(\nabla^2 + k^2 - U(r)) \psi(r) = 0 \quad (\text{eq 3.7 in ref [6]}) \quad (2.8)$$

where $k^2 = \frac{2\mu E}{\hbar^2}$, $U(r) = \frac{2\mu V(r)}{\hbar^2}$, and μ is the reduced mass of the collision partners.

At a point far from the center of the potential, solutions to the equation should resemble

$$\psi(r, \theta) \sim (e^{ikz} + f(k, \theta) \frac{e^{ikr}}{r}) \quad (\text{eq 2.55 in [16]}) \quad (2.9)$$

where the first term represents the incoming wavefunction, and the second term represents the spherical scattered wavefunction, with $f(k, \theta)$ as its amplitude.

The differential cross-section is simply the expectation value of the outgoing wave function into an infinitesimal solid angle $d\Omega$, and the total cross-section is the same as eq. (2.6).

$$\frac{d\sigma}{d\Omega} = |f(\theta)|^2 \quad (2.10)$$

A common technique for finding f is partial wave analysis, where ψ and f are expanded as a sum of Legendre polynomials [6]³:

$$\psi(r, \theta) = \sum_{l=0}^{\infty} \psi_l = \sum_{l=0}^{\infty} R_l(k, r) P_l(\cos(\theta)) \quad (2.11)$$

$$f(k, \theta) = \sum_{l=0}^{\infty} f_l(k) P_l(\cos(\theta)) \quad (2.12)$$

where P_l is a Legendre polynomial and R_l is the radial wave equation solution.

If the collisions are elastic, then each ‘partial wave’ ψ_l should contain an incoming and outgoing portion of equal amplitude due to conservation of kinetic

³ch. 3

energy before and after the collision. Thus, the only change in the wavefunction should be a relative phase between the incoming and outgoing portions. Solving for R_l and comparing to f gives

$$R_l(k, r) = B_l j_l(kr) + C_l n_l(kr) \quad (\text{eq 2.57 in ref [16]})$$

$$\delta_l(k) = \arctan\left(-\frac{C_l}{B_l}\right) \quad (\text{eq 2.58 in ref [16]}) \quad (2.13)$$

$$f(k, \theta) = \frac{1}{k} \sum_{l=0}^{\infty} (2l+1) e^{i\delta_l} \sin(\delta_l) P_l(\cos(\theta)) \quad (\text{eq 2.59 in ref [16]}) \quad (2.14)$$

where j_l are the spherical Bessel functions, and n_l are the Neumann functions. Each l corresponds to a “partial wave” contribution to the solution, with δ_l is its phase shift. These phase shifts can be solved numerically for specific interaction potentials using the procedure outlined in section 2.5.2 of ref [16].

2.2.3 Collisions with a Background Gas

In this work, we trap ultracold (<1 K) atoms inside of ultra-high vacuum (UHV) room-temperature container. We assume the background particles in the vacuum are in thermal equilibrium with their container, and follow a Maxwell-Boltzmann velocity distribution. Thus, on average, they should have significantly more kinetic energy than the trapped atoms [1].

When a background particle collides with a trapped atom, energy and momentum is transferred to the trapped atom. If the collision is elastic, the momentum, energy, and kinetic energies are conserved in the process. For a trapped atom with low initial momentum compared to the background particle

$$\Delta \vec{v}_t = \frac{\mu}{m_t} \Delta \vec{v}_r \quad (\text{eq. 1 in ref [1]}) \quad (2.15)$$

$$|\Delta \vec{v}_r|^2 = 2 |\vec{v}_r|^2 (1 - \cos(\theta)) \quad (2.16)$$

$$\Delta E_t = \frac{m_t}{2} \left(|\vec{v}_t + \Delta \vec{v}_t|^2 - |\vec{v}_t|^2 \right)$$

$$\Delta E_t \approx \frac{m_t}{2} |\Delta \vec{v}_t|^2 = \frac{\mu^2}{m_t} \vec{v}_r^2 (1 - \cos(\theta)) \quad (\text{eq. 1 in ref [17]}) \quad (2.17)$$

where Δx means the change in the parameter x after the collision, E_t is the kinetic

energy of the trapped atom, m_t is the mass of the trapped particle, $\vec{v}_t \approx 0$ is the initial velocity of the trapped atom, \vec{v}_r is the relative velocity between the trapped and background particles, and $\mu = \frac{m_t m_b}{m_t + m_b}$ is the reduced mass of the trapped and background particle.

The approximation in eq. (2.17) holds if the initial velocity of the background atom is much larger than the trapped atom. This will always be true unless the vacuum chamber is in a cryogenic chamber or outer space.

If the trapped atoms are in the bottom of a potential well of depth U , then if the energy imparted to a trapped atom (ΔE) is greater than U , then the atom is ejected from the trap. Using eq. (2.15) eq. (2.16), and eq. (2.17), one can see that when $\theta > \theta_{\min} = \arccos\left(1 - \frac{m_t U}{\mu^2 |\vec{v}_r|^2}\right)$, the trapped atom will be knocked out of the trap. If the scattering angle is less than the classical limit (θ_0), the trapped atom may stay inside of the trap, raising the average energy of the trapped atoms. In a ‘shallow’ trap, $\theta_0 < \theta_0$, so any heating of the trapped atoms is described by quantum mechanics [1].

Quantum Diffractive Energy

To apply the correct physics, it is important to know what trap depths are ‘shallow’ and therefore need to be treated quantum mechanically. In a diffractive collision, $\theta \ll 1$, so MacLaurin expanding eq. (2.16) gives $|\Delta \vec{v}_r| \approx |\vec{v}_r| \theta$ so eq. (2.15) becomes $|\Delta \vec{v}_t| \approx \frac{\mu |\vec{v}_r| \theta}{m_t}$. Quantum effects should occur when the impact parameter uncertainty is comparable to quantum collision radius of the trapped atom (r_t), so the important scattering angle is

$$\theta \approx \theta_c = \frac{\lambda}{\pi r_t} = \frac{2\hbar}{\mu v_r r_t}$$

Solving for v_t and using the modified eq. (2.15) gives $\Delta v_t \approx \frac{2\hbar}{m_t r_t}$. The quantum mechanical collision cross-section for a hard sphere is $\sigma = 2\pi r_t^2$, so $\Delta v_t \approx \frac{\sqrt{8\pi}\hbar}{\sqrt{\sigma} m_t}$ [1].

Finally, the important energy scale is when the change in kinetic energy of the trapped particle is equal to the trap depth. Using the above expression for change

in velocity gives

$$U_d = \frac{4\pi\hbar^2}{m_t\sigma_{\text{tot}}} \quad (\text{eq. 2 in ref [1]}) \quad (2.18)$$

as the ‘Quantum Diffractive Energy’—the energy scale below which quantum effects dominate [3].

2.3 Pressure Measurement using Cold Atoms

The total collision rate with an ensemble background particles of the same velocity is the product of the density of the particles, their velocity, and their collision cross-section with the trapped particle ($\Gamma_{\text{total}} = nv\sigma_{\text{total}}$).

To find the rate at which trapped atoms are ejected due to these collisions, the total collision cross section must be replaced with the loss cross section—the differential collision cross section integrated along all solid angles greater than θ_{min} .

$$\Gamma_{\text{loss}}|_{v_b} = n_b v_b \sigma_{\text{loss}} \approx n_b v_b \int_{\theta_{\text{min}}}^{\pi} \frac{d\sigma}{d\Omega} \sin(\theta) (2\pi) d\theta \quad (2.19)$$

where n_b is the density of background particles, v_b is their velocity, and σ_{loss} is the loss cross section.

The background particles follow a Maxwell-Boltzmann velocity distribution. Thus, eq. (2.19) must be averaged over the distribution to produce the actual loss rate of atoms from the trap:

$$\Gamma_{\text{loss}} = n_b \langle \sigma_{\text{loss}} v \rangle = n_b \left(\frac{2}{\pi} \right)^{1/2} \left(\frac{m_b}{k_b T} \right)^{3/2} \int_0^{\infty} \sigma_{\text{loss}}(v) v^2 e^{-\frac{m_b v^2}{2k_b T}} dv \quad (2.20)$$

There are three ways of determining the velocity-averaged loss cross section:

1. Calculate it: Calculate the cross section using a potential energy surface found using *ab initio* methods (from first principles) [14].
2. Measure it: Use a pressure standard to generate a known pressure, measure Γ_{loss} , then rearrange eq. (2.21) to find the cross section [3].
3. Fit for it using universality⁴: Record loss rates at different trap depths for a

⁴A prior draft of this thesis attempted to coin the word ‘Universalitate’ to describe this process but the word was rejected by a reviewer.

fixed but unknown background gas pressure, then use quantum diffractive universality (discussed later in section 2.4) to fit for the cross section and pressure simultaneously [3].

Regardless of the method used to find $\langle \sigma_{\text{loss}} v \rangle$, the procedure for measuring the pressure is the same. Particles in UHV are treated as an ideal gas, so $PV = N_b k_B T$. Rearranging it gives $P = n_b k_B T$ where n_b is the density of the background gas particles. Because $n_b = \frac{\Gamma_{\text{loss}}}{\langle \sigma_{\text{loss}} v \rangle}$,

$$P = \frac{\Gamma_{\text{loss}}}{\langle \sigma_{\text{loss}} v \rangle} k_B T \quad (2.21)$$

This method of measuring pressure is independent of the geometry of the detector and requires no calibration.

2.4 Quantum Diffractive Universality

Recent research shows that the loss cross section in shallow traps ($U \sim U_d$ or smaller) is universal function of the trap depth [3, 44, 45]:

$$\langle \sigma_{\text{loss}} v \rangle = \langle \sigma_{\text{total}} v \rangle \left(1 - p_{\text{QDU}} \left(\frac{U}{U_d} \right) \right) \quad (2.22)$$

where p_{QDU} is probability of a particle being retained in the trap after a collision, $\langle \sigma_{\text{total}} v \rangle$ is the total velocity-averaged collision cross section, and $U_d = \frac{4\pi\hbar^2 v_p}{m_t \langle \sigma_{\text{loss}} v \rangle}$.

The function can be approximated using a power series and fitted using measured data:

$$p_{\text{QDU}} \left(\frac{U}{U_d} \right) \approx \sum_j \beta_j \left(\frac{U}{U_d} \right)^j \quad (2.23)$$

The cross section for a new species can be measured by recording Γ_{loss} across a range of trap depths, then fitting eq. (2.23) for it.

There is no analytic proof for this universality, but there is numerical (fig. 2.4) and experimental (fig. 2.5) evidence for it.




Figure 2.4: Numerically computed velocity-averaged loss cross sections for He (red triangles), Ar (black circles), and Xe (blue squares) colliding with Li (dashed lines) and Rb (solid lines), taken from [3] [fig 2b]




Figure 2.5: Experimental data for trapped Rb and background He, Ar, Xe, H_2 , N_2 , and CO_2 fitting to the same universal curve [3] [fig 3f]

Term	β_j
1	0.673(7)
2	-0.477(3)
3	0.228(6)
4	-0.0703(42)
5	0.0123(14)
6	-0.0009(2)

Table 2.1: The universal coefficients in eq. (2.23) calculated by fitting to the quantum scattering calculations. Each parenthetical number is the estimated error in each coefficient. Taken from ref [3].

2.5 Refinements

2.5.1 Other Loss Mechanisms

Elastic collisions with background particles are not the only loss mechanism in magnetic traps. They will not compromise differential pressure measurements, where the relative increase in pressure due to an introduced background gas is measured by subtracting off the baseline from the measured value. However, they may affect direct pressure measurements.

Intra-trap collisions

Two-body collisions between trapped atoms can also eject trapped atoms. They can do this by elastically scattering one or more atoms out of the trap (evaporation) or by inelastically changing the state of one of the collision partners to an untrappable state. Three-body collisions can also induce losses. However, even if the density of the trapped atoms is $1 \times 10^7 \text{ cm}^{-3}$, the loss rate is $\sim 1 \times 10^{-5} \text{ Hz}$ – much smaller than the loss rates due to background collisions [43].

Majorana Losses

Quadrupole magnetic fields coils have a zero in the center. Atoms passing near the zero-field region can make a nonadiabatic spin-flip transition ($m_F \pm 1$), which can transfer atoms into an $m_F = 0$ state. This stops the atoms from interacting with

the magnetic field, so they fall out of the trap. These losses can be avoided by using a trap without a magnetic field at the center, such as an Ioffe-Pritchard trap [49]. Fortunately, the loss rate is constant for a constant magnetic field, so they can also be compensated for in differential measurements (along with the effects of outgassing in the measurement chamber) by pumping the chamber to its base pressure, recording the base loss rate, and subtracting it off of subsequent readings [43].

2.5.2 Trap Anisotropy and RF Knife

Most atom traps are created in rectangular glass cells, but the equipotential surfaces of a magnetic trap are oblate spheroids. As a result, atoms with some orbit paths may remain trapped, while other atoms with the same energy can collide with the cell wall and escape. To help address this issue, the energy distribution of the trapped atoms can be shaped using an ‘RF Knife’. A coil is placed near the glass cell and is driven at an RF frequency to produce an RF magnetic field. This creates an oblate spheroid of atoms whose $\Delta F = 0, \Delta m_F \pm 1$ magnetic dipole transition is driven. If they reach a non-negative $m_F = 0$ state, they cannot remain trapped and gravity pulls them out of the trap. Sweeping the RF frequency changes the size of the surface resonant with a specific energy level, removing atoms of energy $E \geq h\nu_{min}$ [3]. This creates an effective trap depth, which is useful when applying quantum diffractive universality to measure cross sections.

2.5.3 Energy Distribution Corrections

Not all atoms in the magnetic trap have zero energy—in fact, they need some energy or they will fall out [49]. A more rigorous calculation of p_{QDU} integrates over the energy distribution of the atoms in the trap:

$$\overline{\langle \sigma_{\text{loss}}(U)v \rangle} = \frac{\int_0^{E_{\text{max}}} \langle \sigma_{\text{loss}}(E_{\text{max}} - E)v \rangle \rho(E) dE}{\int_0^{E_{\text{max}}} \rho(E) dE} \quad (2.24)$$

$$= \langle \sigma_{\text{total}} v \rangle \left[1 - \beta_j \left(\frac{1}{U_d} \right)^j \langle (E - E_{\text{max}})^j \rangle \right] \quad (2.25)$$

Collision Partners	original σ_{tot} ($1 \times 10^{-15} \text{ m}^3 \text{ s}^{-1}$)	corrected σ_{tot} ($1 \times 10^{-15} \text{ m}^3 \text{ s}^{-1}$)
$^{87}\text{Rb-N}_2$	3.11(5)(2)	3.14(5)(2)
$^{87}\text{Rb-He}$	2.40(12)(8)	2.41(12)(8)
$^{87}\text{Rb-Ar}$	2.77(5)(2)	2.79(5)(2)
$^{87}\text{Rb-Xe}$	2.71(3)(3)	2.75(3)(3)
$^{87}\text{Rb-H}_2$	N/A	3.8(2)
$^{87}\text{Rb-CO}_2$	2.82(6)(2)	2.84(6)(2)

Table 2.2: Originally calculated collision cross sections (calculated using eq. (2.23)) and corrected collision cross sections (calculated using eq. (2.25)). The first number in parentheses is the statistical error, and the second is the systematic error. All cross sections are taken from [45] except for $^{87}\text{Rb-H}_2$, which is taken from [46].

where $\rho(E)$ is the energy distribution of the trapped atoms.

Using this method produces small ($\sim 1\%$) corrections to the total collision cross sections (see table 2.2). See ref [45] for a more explicit treatment of this error.

2.5.4 Limitations

The collision cross-section ($\sigma(v)$) as a function of velocity displays an oscillatory behaviour superimposed on a monotonically decreasing average cross-section (see fig. 2.6). These oscillations (known as Glory oscillations) occur due to the interference between the short-range and long-range inter-particle collision physics [6]. The quantum diffractive universality initially reported in ref [3] arises in part due to the effects of averaging $\sigma(v)$ over the Maxwell-Boltzmann velocity distribution of the background particles. This averages away the short-range oscillatory behaviour (which is different for different background species), leaving only the shared slowly-decreasing long-range behaviour responsible for the universal description of the collisions (eq. (2.22)). This universality can break down under several conditions:

- Extreme background temperatures: If the background particles are too cold ($\sim 1 \text{ K}$), then the velocity distribution is too narrow to smooth out the short-range physics. If the background particles are too hot ($\sim 1000 \text{ K}$), the highest velocities can reach a collision regime where the total cross section is domi-

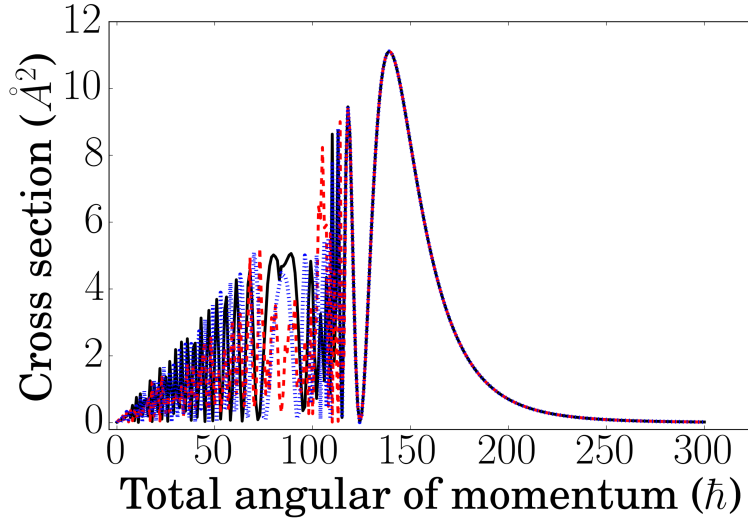


Figure 2.6: Calculated cross sections as a function of total angular momentum for several different potential energy surfaces. Notice the initial oscillatory behaviour followed by the monotonic decrease. Taken from ref [3].

nated by the short-range physics, and the universal approximation no longer applies [3]. This has not been experimentally verified yet, but is not a major limitation: ConFlat flanges (the standard method of connecting vacuum parts together at UHV) are rated from 77 to 723 K, so nearly all vacuum chambers operate within those temperature limits. [28]

- Low reduced mass: Low mass collision partners (several AMU or smaller) tend to have a small polarizability and C_6 Lennard-Jones coefficients, which govern their long-range behaviour. This makes their short-range interactions more prominent even at low collision velocities. Low-mass particles also have higher velocities, which increases the amount that they probe the short-range physics. These factors make the quantum diffractive universality description specific to C_6 interactions incomplete for these species [46]. This is a concern when using quantum diffractive universality to compute collision cross sections for low-mass collision partners, but fortunately it is less challenging to calculate cross sections between low-mass collision partners.

- Large trap depths: If the trap depth approaches U_d , then the sixth-order polynomial expansion in U/U_d will fail. More terms will be required to extend the universal description to larger U/U_d values.

2.6 Summary

Monitoring the loss rate of atoms in a magnetic trap provides an excellent and repeatable method of measuring pressure that works for all collision partners—not just inert gases. The portability of the apparatus will allow us to directly compare the efficacy of the technique against other state-of-the-art UHV pressure standards to promote international recognition of a new pressure standard realized with cold atoms. Furthermore, direct comparison will allow us to compare the collision cross sections calculated using quantum diffractive universality with those measured by comparing against a known pressure.

Part II

Apparatus

Chapter 3

Vacuum Chamber

In order to measure the pressure in a vacuum chamber, we need a vacuum chamber to contain the MOT and MT. Specifically, the vacuum chamber needs to include

- an optically transparent science section to create and image the trapped ensemble of rubidium atoms
- a rubidium source.
- a valve connecting the portable chamber to the chamber under test that can be closed during transportation.
- a pump to keep the vacuum chamber under vacuum while it is shipped between destinations.

We also imposed several design constraints:

- It should minimize the amount of Rb that leaks into the chamber under test.
- To minimize the disturbance to the system under test, it should
 - have a small interior volume.
 - outgas as little as possible.
 - have a good conductance to the chamber under test.




Figure 3.1: The original design of the vacuum chamber.

3.1 Design

The design of the vacuum chamber went through several configurations. First, we assembled the chamber shown in fig. 3.1.

We used a prebuilt Rubidium source from ColdQuanta with an integrated differential pumping tube, 2D MOT, and push beam (see section 3.1.4 for more details). Opposite the dispenser was the measurement cell (purchased from Precision Glassblowing) where the pressure measurements could be made.

To the right of the cross is a CF133 valve that could be used to valve off the non-evaporable getter (NEG)—a type of chemical pump that readily adsorbs certain common species seen in vacuum chambers (such as H_2 and N_2)¹. If connected directly to the chamber, NEG would disrupt pressure measurements as it reduces the observed pressure. However, a NEG is ideal for keeping the vacuum pressure down during transit as it can pump hydrogen (the most common background species) very efficiently, and does not need to be supplied power after its initial activation. Therefore, we planned to close the CF133 valve during measurements,

¹Its pumping rate was conductance-limited by the CF133 valve, but has a pumping rate of 312 L/s for hydrogen and 45 L/s for nitrogen

but open it when not conducting measurements to keep the base pressure down.

Opposite the CF133 pump cut-off valve is the CF275 inline valve used to connect the portable sensor to other vacuum chamber or pumps. It is critical that it remains sealed when in transit. First, it protects the walls of the chamber from atmospheric water - if left open, we would need to re-bake the entire chamber after connecting to a chamber under test, which would necessitate removing then reinstalling and realigning the optics and magnetic field coils. More importantly, it protects the rubidium dispenser from oxidization.

We used this chamber to collect the data discussed in chapter 9. During testing, we discovered several flaws with the chamber: its conductance to the measurement cell was poor, it had no way of pumping inert gases, and it leaked. This necessitated a redesign and rebuild of the vacuum chamber.

3.1.1 Conductance

Vacuum Conductance

Atoms in UHV are well within the molecular flow regime, where the mean free path of a particle is much larger than the size of the vacuum chamber (at 1×10^{-7} Torr [1×10^{-5} Pa] - the upper pressure limit of the apparatus - the mean free path of a particle is ~ 1 km as compared to the ~ 3 cm internal diameter of the vacuum chamber) [22]. This means the particles are virtually non-interacting.

The flow rate of particles through a pipe is

$$Q = C\Delta P \quad (3.1)$$

where C is the conductance of the pipe and ΔP is the difference in pressure between the ends of the pipe. If several pipes are joined together, the conductances add according to [56]

$$\frac{1}{C_{\text{tot}}} = \sum_i \frac{1}{C_i} \quad (3.2)$$

Ion pumps and NEG's are diffusion-limited - they can only pump an atom if it reaches them. Therefore, they are limited by the conductance between them and the region they are pumping. This means the effective pumping rate of a pump

with pumping rate S is [56]:

$$S_{\text{effective}} = \frac{SC_{\text{tot}}}{S + C_{\text{tot}}} \quad (3.3)$$

The conductance of a duct in a vacuum system depends on the rate that particles enter the duct, and the probability that they will pass through the duct [22]. The conductance of an aperture is

$$C_a = \frac{1}{4}Av_a \quad (3.4)$$

where A is the area of the aperture and $v_{av} = \sqrt{\frac{8k_B T}{\pi m}}$ is the average velocity of a thermal particle of mass m .

The conductance of a long tube of cross-sectional area A , length L , and circumference H is [22]

$$C = \frac{C_a}{1 + \frac{3LH}{16A}} \quad (3.5)$$

And if a tube has a bend in it, the conductance changes as if the length of the tube has increased [56]

$$L_{\text{effective}} = L_{\text{axial}} + 1.33 \frac{\theta}{\pi} D \quad (3.6)$$

where θ is the angle in radians and D is the diameter of the tube.

Conductance of Chamber

We calculated the conductance from the 3D MOT to the far side of the inline valve to be approximately 1.2 L s^{-1} , limited by the narrow diameter of the glass cell and the CF133 cross. The effective pumping rate of the 3D MOT by the NEG was estimated to be 0.9 L/s .

3.1.2 Leak Checking

Figure 3.2 shows that the pressure in the chamber rose steadily once only the NEG was pumping, which suggested that gases that the NEG could not pump easily (i.e. atmospheric gases) were present in the chamber - in other words, there was likely a leak.

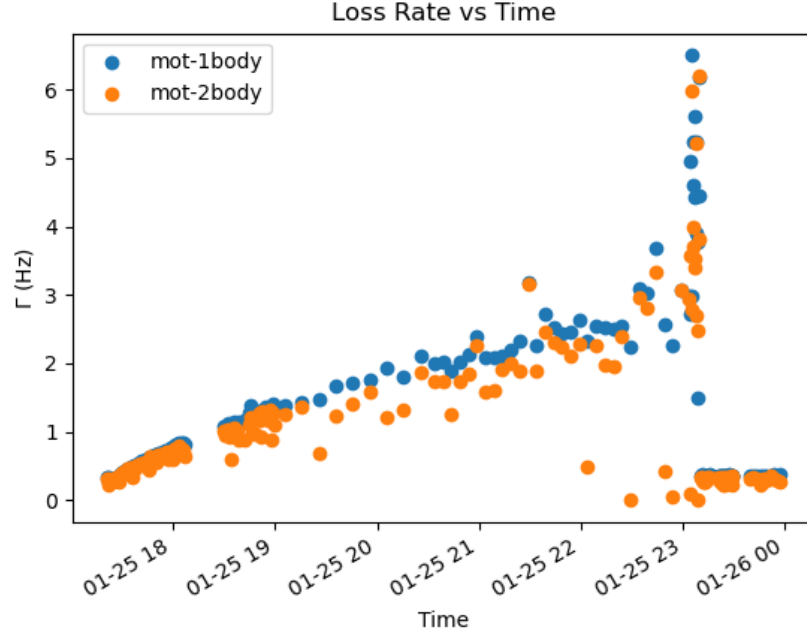


Figure 3.2: A set of MOT loss rates (Γ) taken using the original vacuum chamber (fig. 3.1). MOT loss rates provide a useful proxy for the pressure in the chamber. The orange dots were fitted using eq. (9.2) to extract the loss rate, while the blue dots were fitted using eq. (9.4) to extract the loss rate. The chamber has a turbo pump connected through the inline valve, and it can also be pumped by opening the NEG valve. At 17:22, the inline valve to the turbo pump was closed, causing the pressure to slowly rise. At 23:05, the valve to the NEG was closed, causing a sharp spike in the pressure. At 23:10, the valve to the turbo was re-opened, causing the pressure to fall to its original value.

We used an Agilent HLD MD30 Helium Leak Tester to look for the leak in the chamber. Typically, one tests for leaks using a leak tester by venting the chamber, attaching it to the leak tester, then spraying different sections of the chamber with Helium. If the leak tester detects helium on its inlet, it suggests that the sprayed section of the chamber has a leak that is letting the helium in.

However, we needed to avoid exposing the Rubidium source to atmosphere, and wanted to avoid attaching unbaked vacuum components directly to the chamber. We valved off the chamber, attached an angle valve to the exhaust line of the turbo pump between the turbo and the scroll pump, then baked out the connection between the chamber and the turbo pump using heater tape. Next, we valved off the turbo's exhaust, then replaced the scroll pump with the leak tester. We initiated the test sequence, which meant that the leak tester vented the roughing line but the turbo pump remained under vacuum. Finally, we slowly opened the roughing line valve, ensuring that the pressure on the leak tester did not rise more than 10x above its minimum value - if the pressure rose too high, it could damage the mass spectrometer filament.

When we leak-tested the chamber, we discovered a leak $\sim 3 \times 10^{-8}$ TorrLs⁻¹ between the measurement cell and the CF133 cross. This explained why the pressure in the chamber rose dramatically when the turbo pump was valved off. A leak that large would seriously compromise sensitive pressure measurements, so the vacuum chamber needed to be rebuilt.

3.1.3 New Design

To address those issues, we rebuilt the vacuum chamber (see fig. 3.3), correcting the following flaws:

- Prior to baking, the leak detector could not detect a leak (minimum detectable leak was 1×10^{-12} TorrL/s)
- It improved the estimated conductance to the 3D MOT to ~ 1.3 L/s and increased the estimated 3D MOT effective pumping rate to ~ 1.3 L/s.
- It included an ion pump to allow for pumping of inert gases.




Figure 3.3: The final revision of the vacuum chamber used in the apparatus

Part	Part No
3D MOT Cell	Custom 13mmx13mmx70mm cell
Source Cell	ColdQuanta Rubidium CASC
Cube	Lesker CU6-0275
Reducer	Lesker RF1000X275
Inline Valve	MDC 316001
Angle Valve	Norcal AMV-S04400
5-Cross	Lesker C5-0275
NEG Holder	Lesker FN-0275S
NEG	Edwards 300NP
Ion Pump	Gamma TiTan 25S
Blank	Lesker F0275X000N

Table 3.1: Rebuilt vacuum chamber components

3.1.4 ColdQuanta Rubidium Source

The apparatus uses a prebuilt Rb dispenser from ColdQuanta. The dispenser releases rubidium when it is heated by applying a current through the dispenser rod. Using a dispenser reduces the amount of rubidium contamination in the chamber under test because rubidium can be dispensed only when measurements are being made.

The source cell also has a differential pumping hole and a 2D+ MOT ². The differential pumping hole also reduces the level of rubidium contamination in the chamber under test. It consists of a wall at the end of the source cell with a small (1 mm diameter) hole in the center of it. Normally, rubidium has a very small chance of passing through the hole. However, activating the integrated 2D MOT reduces the atoms' radial velocity (relative to the axis of the differential pumping hole), which dramatically increases the probability that the atoms pass through the hole. To further encourage the atoms to pass through the hole, a 'push' beam of light resonant with the pump transition and parallel to the differential pumping hole can be turned on.

The push beam is not essential for loading the 3D MOT, but we found that it increased the loading rate by a factor of 10 (see section 9.2.2 for the details of this measurement).

3.2 Preparation

This section describes the procedure used to reduce the level of outgassing in the vacuum chamber used in the apparatus. It starts with some general theory about the pressure inside of a vacuum chamber before applying it to our apparatus.

3.2.1 Leak Rate in Vacuum Chambers

The pressure in a vacuum chamber

$$\frac{dP_0}{dt} = \frac{1}{V} \left(\sum_{i=0}^N C_i (P_0 - P_i) + R \right) \quad (3.7)$$

²A 2D+ MOT consists of a 2D MOT and a push beam

where P_0 is the pressure in the chamber, P_i is the pressure in chamber i with conductance to it C_i , and R is the leak rate into the vacuum chamber.

If a chamber is only connected to one other vacuum component and is being pumped out through it, its steady-state ('base') pressure is $P_0 = P_1 + \frac{R}{C_1}$. To decrease the base pressure of the system, one must decrease the leak rate, increase the conductance into the chamber, and/or decrease the pressure in the chamber it is being pumped into (get a better pump). Once under vacuum, the conductance and pumping rate are fixed, so the only way to reduce the base pressure is to change is the leak rate, R .

There are a number of contributions to the leak rate. One minor source is the atmosphere leaking through the seals in the vacuum chamber. However, unless the chamber has been improperly sealed, the dominant source is 'outgassing' - atoms that have stuck to the walls of the vacuum chamber and are slowly desorbing and bouncing through the chamber [22]. Provided the chamber was adequately cleaned prior to being pumped down, water will be the dominant species outgassing. In a typical metal vacuum chamber, the outgassing rate at room temperature is approximately [13].

$$q_{H_2O} \approx \frac{3e-9}{t} \quad (3.8)$$

where q_{H_2O} is the outgassing rate of water in millibar litres per second per centimetre squared, t is the time elapsed in hours since starting to pump down the vacuum chamber.

While metals outgas at a rate proportional to $\frac{1}{t}$, polymers tend to outgas at a rate proportional to $\frac{1}{\sqrt{t}}$. This means that pumping them down to a given pressure takes an order of magnitude longer. It is necessary to expose polymers to vacuum in a few applications (such as glass-to-metal seals), but their use in UHV chambers should be minimized as much as possible. [22].

The outgassing rates of water vapour and hydrogen gas increase with the temperature of the chamber. Therefore, to remove them rapidly, the entire chamber is heated or 'baked out'. Typically, heating the walls of the chamber to 150 C to 200 C for around several days removes almost all of the water from the walls of the chamber [22].

Once the water has been removed, hydrogen is the dominant background species.

It comes from the hydrogen that is dissolved into the steel in the vacuum chamber walls. It slowly diffuses towards the surface and desorbs, contaminating the chamber. The diffusivity of hydrogen is exponentially related to temperature—baking at 300 C for three months (theoretically) produces the same result as baking at 1000 C for an hour [4]. Consequently, CERN bakes their vacuum parts at 950C for two hours in a vacuum furnace [5]. Other groups have baked their vacuum chambers at 400C for several days to several weeks [14, 36]. Note that as the hydrogen concentration decreases, its desorbing rate becomes recombination-limited rather than diffusion-limited, further slowing the desorption rate and producing diminishing returns on the bakeout time [19].

3.2.2 Baking the Vacuum Chamber

Due to the vacuum chamber being rebuilt multiple times, and some parts being more sensitive to high temperatures than others, we conducted the high-temperature bakes in several stages.

First Bake

First, we baked the original chamber (minus the glass measurement and source cells) at a high temperature to remove hydrogen. Specifically, we baked the inline valve, coldquanta adapter plate, and zero-length reducer³ at 300 C for 55 days.

To conduct the bake, we first assembled the vacuum components and connected them to an ion gauge and a turbo pump backed by a scroll pump. We activated the scroll pump, waited 30 minutes, then activated the turbo pump. We waited another 30 minutes, then degassed and activated the ion gauge to monitor the pressure of the chamber during the bakeout.

Next, we started to build an enclosure around the chamber made in aluminum foil-wrapped ceramic bricks. The foil was primarily to reduce the amount of dust from the bricks. Once the enclosure wall was partially constructed, we placed three ceramic full trough infrared heaters (Salamander FTE-650-400-0-L6-Y-0) around the chamber, crimped them to fiberglass-insulated electrical cables, insulated the

³There were also some components baked at this time that were not used in the rebuilt vacuum chamber - a CF133 valve, two CF133 blanks, one CF275 blank, and a CF133 cross

crimp joint with glass tubes, then fed the fiberglass cables through the partially-constructed wall. We connected the electrical outlets to variacs, which were powered by Briskheat SDC120KC-A digital temperature controllers, which operated using bang-bang feedback control⁴. We placed the Briskheat controllers' thermocouples near the heaters to provide more accurate control of the temperatures.

We also attached separate thermocouples to areas that were more sensitive to high temperatures (the one attached to the CF133 valve was labelled 'small valve' and the one attached to the CF275 inline valve was labelled 'large valve') and areas that were close to heaters ('near side' and 'far side'). When installing the thermocouples, we first wrapped any exposed metal (such as the tip) in Kapton tape, then inserted the tip into the nearest leak-check slot so that we measured the temperature of the chamber, not the ambient air. Once a thermocouple was installed, we wrapped wire around the part of the vacuum chamber the thermocouple was measuring and twisted it to tighten the connection.

Because the turbo pump (turbo) could not be heated above 70 C, we used a bellows to connect the chamber to the turbo. The bellows was wrapped in aluminum foil (to help disperse the heat), heater tape (wrapped more densely close to the oven and less densely close to the turbo), and several more layers of aluminum foil. This reduced thermal gradients in the system and allowed more of the chamber to be heated, which reduced the surface area of relatively cold metal that gas can condense onto during the bake [12]. It is important to minimize the size of the cold spots because the condensed gas will outgas after the bake has concluded.

Once the thermocouples and heaters had been placed, we completed assembling the oven. We brought the oven up to target temperature by increasing the briskheat set temperatures and variac voltages, ensuring that the temperature never changed faster than 1C/min. Once the pressure in the chamber was no longer dropping, we slowly brought the temperature down and disassembled the oven, heaters, and thermocouples. See section 3.2.2 for a diagram of the vacuum chamber. See section 3.2.2 for a plot of the temperatures, and fig. 3.6 for a plot of the pressure during the bake.

The initial oscillations in pressure were due to the temperature of the chamber

⁴Bang-bang feedback control is the simplest form of feedback control. If the temperature was below the set point, they turned on. If it was above, they turned off




Figure 3.4: We baked this chamber with the source cell and science cell replaced with CF133 blanks. It was baked at 300 C for 55 days. The inline valve, NEG tube, zero-length reducer, and source cell adapter were re-used in the rebuilt vacuum chamber.

increasing in steps. We suspect that the pressure spike observed at 25 days and corresponding decrease at 35 days was due to one of the controllers' thermocouples shifting in position⁵, which would cause the controller to change its average output power. By day 45, the pressure had reached its steady-state value, limited by the compression ratio of the turbo pump. The ramp down in pressure at the 55 day point was due to the temperature of the chamber being decreased (due to the ideal gas law, $P_{\text{final}} = \frac{T_{\text{final}}}{T_{\text{initial}}} P_{\text{initial}}$). The slow increase in pressure following this may have been due to atmospheric hydrogen slowly backflowing through the turbo pump after the pressure temporarily fell below the pump's base pressure.

Second Bake

It is unlikely that the apparatus would be able to connect directly to the chamber under test due to geometric constraints. We wanted a selection of hydrogen-free conflat parts we could use to build an interconnect between the two chambers. We also wanted a low-outgassing ion pump and ion gauge available to reduce the

⁵The author was infected with COVID-19 during this period of time and was not paying close attention to the experiment.

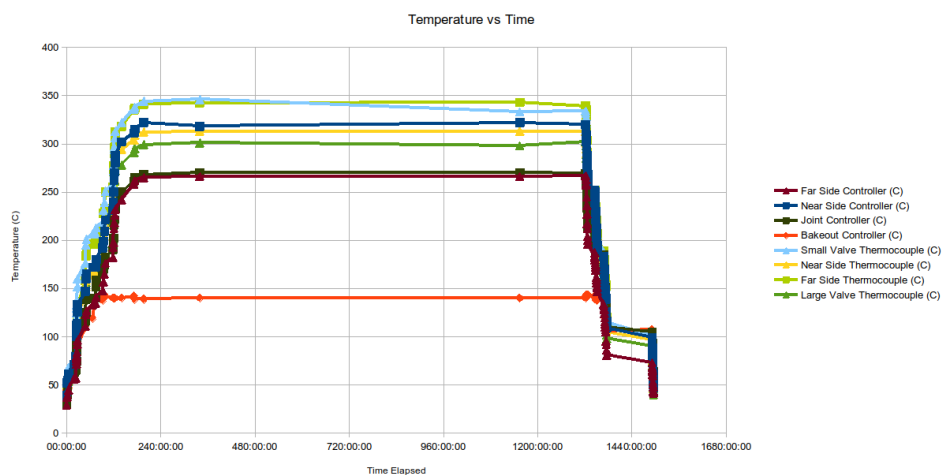


Figure 3.5: Manually recorded temperatures of different thermocouples while conducting the first bake.

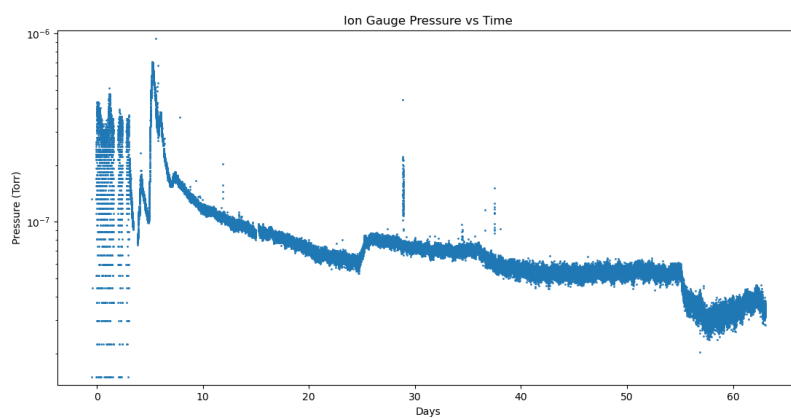


Figure 3.6: Pressure reported by the ion gauge while conducting the first bake.

system's base pressure and measure the pressure in the interconnect. As a result, we built a vacuum chamber out of spare parts and baked it at a high temperature to remove as much hydrogen as possible from those parts. Some of these parts (the ion pump, an elbow, both blanks, and a valve) were used later when we rebuilt the chamber.

We baked the ion pump (with its external magnets removed), angle valve, elbow, blanks, and a 12" bellows at 400 C.⁶ The procedure was very similar to the above except the bake was split into two ovens - one that went to 400 C heated by two heaters, and one that went to 300 C that was heated by a single heater. The chambers were held at temperature for 37 days. See fig. 3.7 for a diagram of the chamber that was baked. See fig. 3.8 and section 3.2.2 for the temperature and pressure during the bake.

Third Bake

Once we had baked the all-metal components in the original chamber at high temperature, we installed the glass measurement cell and coldquanta source and re-baked the chamber at a lower temperature to remove water (see fig. 3.10 for the temperature log). It was during this bake (on day 27) that we activated one of the rubidium dispensers and the NEG. After this bake, we recorded the data discussed in chapter 9.

Rubidium dispenser activation: Following the instructions in the source's user manual [7] we set the dispenser current to 3 A for 80 minutes (until the pressure in the chamber stabilized). We then set the current to 5.5 A for one minute, then reset it to 3 A until the pressure re-stabilized. We spiked the current, reset it, and waited for the pressure to settle two more times before declaring the source activated.

NEG activation: We activated the NEG by setting its current to 5.5 A for 70 minutes. We observed the pressure in the chamber rise from 2×10^{-8} Torr to 1×10^{-5} Torr in this time. This is because NEGs release their adsorbed hydro-

⁶We also baked another CF275 elbow, an ion gauge in a CF275 tube, a CF275 cross, two CF275 blanks, another CF275 valve, and a 6" CF275 bellows. We baked a CF275 tee, a CF275 valve, a CF275 elbow, a CF275 valve, and an SRS RGA-200 in a separate oven attached to the same chamber.




Figure 3.7: The collection of spare parts baked with the intention of using some of them in an interconnect between the apparatus chamber and the chamber under test. The ion pump, leftmost CF275 angle valve, blanks, and uppermost elbow were used in the rebuilt chamber. We plan to use the 12" bellows for the interconnect. They were baked at $\sim 400\text{ C}$ for 37 days.

gen when activated. The turbo pump we were using to pump out the system had a poor compression ratio for hydrogen (400), and the scroll pump we were using had a base pressure of 5×10^{-2} Torr, meaning that the pumps could not remove the released hydrogen from the chamber.

Fourth Bake

When we realized that we needed to rebuild the vacuum chamber, we decided to replace the CF133 cross with a CF275 cube to improve conductance. We purchased a cube from Lesker and baked the it at 420 C for 25 days to remove its hydrogen.

To conduct the bake, we connected a bellows that led to a turbo pump and ion gauge on one side of the cube. We blanked off the remaining five sides of the cube. Because only one small component was baked, we did not build a full oven.




Figure 3.8: The temperatures in the chamber during the second bake. The thermocouples attached to the ion pump and 4 cross came loose and produced inaccurate readings for significant portions of the bake.

Instead, we wrapped the cube in heater tape, then aluminum foil, then a second layer of heater tape, more aluminum foil, then enclosed the bundle with ceramic bricks. The bundle of heater tape was powered by a single variac, and another variac controlled the temperature of the bellows connecting to the turbo pump. See section 3.2.2 for the temperature of the chamber during the bake and fig. 3.12 for the pressure of the chamber.

Fifth Bake

We assembled the final version of the chamber (as depicted in fig. 3.3) in an oven and baked it out at 200 C using three IR heaters. Due to a last-minute change to the chamber design, we baked the five-cross connected to the other components but at a much higher temperature (400 C) by wrapping it in heater tape, then covering the heater tape in aluminum foil. This allowed us to remove hydrogen from it while protecting the more sensitive glass cells from high temperatures.

We re-activated the Rb dispenser in the ColdQuanta cell, which caused a sharp spike in the pressure as we ‘burnt off’ oxidized Rb from the outside of the dispenser.

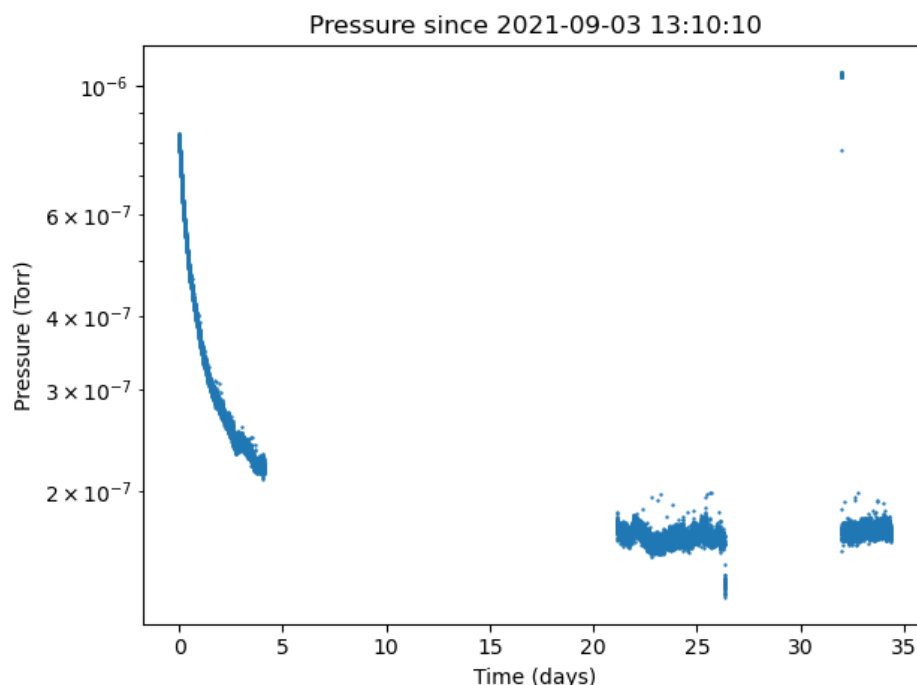


Figure 3.9: The pressure recorded by the ion gauge near the scroll pump during the second bake. We discovered that selecting a portion of text in a Windows shell pauses the command the shell is running. This inadvertently caused the script that was collecting the pressure data to freeze for much of the bake.

3.3 Outgassing Performance

We replaced the scroll pump with a leak tester and used the procedure described in section 3.1.2 to check for leaks in the rebuilt chamber. We could not detect any leaks larger than 1×10^{-8} TorrL/s. Unfortunately, a leak in the housing of the turbo pump prevented us from detecting leaks smaller than that; any helium we sprayed near the chamber would leak in through the turbo pump, ruining the measurement.

We do have some evidence to suggest the leak rate in the new chamber is minimal. Monitoring the ion pump current gives a measure of the pressure in the ion

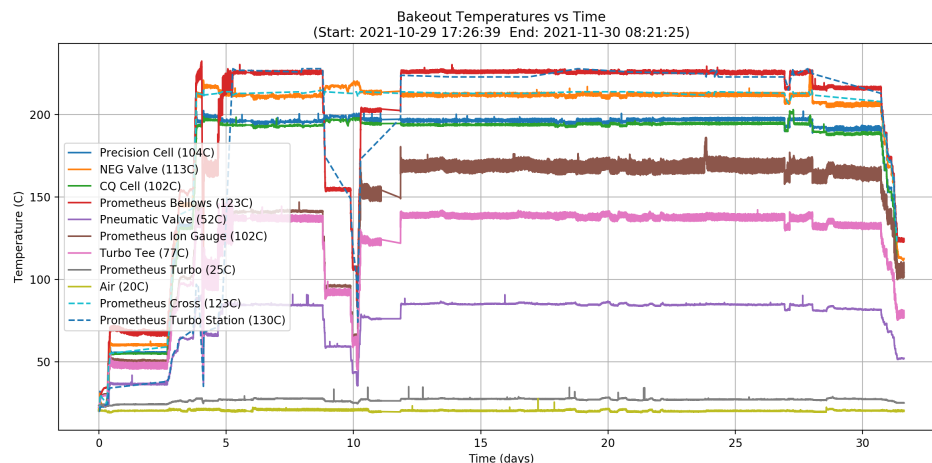


Figure 3.10: The temperature in the chamber during the third bake. While we were able to monitor the pressure on a dial, our pressure logging device encountered an issue and did not reproduce the readings on the dial.




Figure 3.11: The temperature of the cube, bellows connecting it to the turbo pump, ion pump, and turbo pump during the fourth bake. The temperature was increased slightly for the last week of the bake to increase the amount of removed hydrogen. Fluctuations in the temperature were due to fluctuations in the mains voltage.




Figure 3.12: Pressure recorded with the ion pump during the fourth bake. The slight spike in temperature on August 3rd was due to an increase in the temperature. An impending power shutdown in the building cut the bake short. We decided to not continue the bake because we expected that other vacuum parts baked at lower temperatures would outgas significantly more than the cube.

pump chamber. When the inline valve connecting the chamber to the turbo pump is opened, the pressure in the ion pump rises to 7×10^{-11} Torr, while the turbo pump ion gauge reads 3×10^{-9} Torr. When the inline valve is closed, the ion pump pressure quickly drops to the minimum value— 1×10^{-11} Torr. This suggests that the base pressure in the chamber is at least below 5×10^{-10} Torr. We look forward to being able to confirm the exact base pressure of the system with our atoms.




Figure 3.13: Recorded temperatures of the rebuilt vacuum chamber during the fifth bake. The five-cross' temperature was increased to 400 C to remove as much hydrogen as possible from it.




Figure 3.14: Recorded pressure in the rebuilt vacuum chamber during the fifth bake. The spike in pressure on September 26th was from activating the Rb dispenser.

Chapter 4

Magnetic Field Coils

4.1 Theory

There are many configurations of magnetic field coils that can trap cold atoms, such as the anti-helmholtz, ioffe-pritchard, and baseball configurations [38, 50]. We decided to build a quadrupole anti-helmholtz trap as they are the simplest to produce. Ideally, it consists of two axially concentric winds of wire with current i flowing in opposite directions (see fig. 2.1).

The on-axis magnetic field of a single coil loop is [20]

$$\vec{B}_z = \frac{\mu}{2} \frac{a^2 i}{((z + d/2)^2 + a^2)^{1.5}} \hat{z} \quad (4.1)$$

and the on-axis magnetic field of an ideal anti-helmholtz coil is

$$\vec{B}_z = \frac{\mu}{2} a^2 i \left(\frac{1}{((z + d/2)^2 + a^2)^{1.5}} - \frac{1}{((z - d/2)^2 + a^2)^{1.5}} \right) \hat{z} \quad (4.2)$$

where a is the radius of the wires, d is the spacing between them, μ is the magnetic constant, and z axial distance from their midpoint.

In the real-world, the size of the coils may be comparable to the inter-coil spacing, so the spacing between the winds needs to be taken into account by summing the contributions from many pairs of radii a_i spaced d_i apart.

The off-axis axial and radial components of the magnetic field can be found

using elliptical integral functions [24] , but the magnitude of the magnetic field is approximately¹

$$B(x,y,z) \approx \frac{dB}{dz} \frac{\sqrt{x^2 + y^2 + (2z)^2}}{2} \quad (4.3)$$

Therefore, to maximize the trap depth of magnetic field coils, the axial magnetic field gradient must be maximized.

4.1.1 Thermal Stability

Because pressure is proportional to temperature, it is critical that the temperature of the atoms in the glass cell - and thus the glass cell - remain at the same temperature as the chamber under test. Generating a large magnetic field gradient (more than a two hundred Gauss per cm) requires a lot of power (more than fifty watts), so it is important to consider how the coils will be cooled.

Air cooling occurs primarily due to convection. The rate that heat is lost from a plate at temperature T of area A into a fluid of ambient temperature T_a follows Newton's law of cooling (eq. (4.4)) [35]:

$$P = hA(T - T_a) \quad (4.4)$$

where $h \sim 10 \frac{W}{cm^2 K}$ [35].

One can use the surface area of the coils for A to give a crude estimate of its ambient cooling capacity. One can increase the cooling capacity by

- increasing the area A by adding a heatsink. However, this introduces extra bulk to the coils which will likely get in the way of the optics.
- increasing h by adding a fan to blow extra air across the plate. However, this can also introduce stray magnetic fields, vibrations, and potentially extra dust

There are also two other problems with air-cooling:

- The vacuum chamber and optics will be enclosed in a box, so the air would need to be pumped in and out

¹For our coils, this approximation is within 5% of the true value for any point within the glass cell.

- Most importantly, air cooling is proportional to the temperature difference between the air and the object. Keeping an object close to the temperature of the air minimizes the cooling capacity of the air.

Because of these issues, we decided to build water-cooled coils.

4.1.2 Types of Coil Construction

Water-cooled coils can be constructed in two main ways: They can either be made from a solid conductor then enclosed in a watertight container that water flows through, or they can be made from metal tubing that the water flows through.

Several different geometries of such coils made from a solid conductor exist, including using

- magnet wire (solid-core wire with a thin enamel coating on the exterior) wrapped around a spindle.
- conductive foil topped with an insulator (such as Kapton tape) and wrapped in a helix (a ‘jelly roll’ coil).
- spiral-cut plates stacked on top of each other (also known as a Bitter coil) [47].

Two main geometries of hollow-core coils can also be constructed. They can be wrapped from a single long tube, or created from a set of smaller coils that are hydraulically in parallel but electrically in series. Although the latter design offers significant advantages for cooling (creating a coil from four 2x8 coils connected axially offers 16x less flow resistance than a monolithic 8x8 coil), we struggled to manufacture them such that they didn’t develop leaks or internal short-circuits. Consequently, we built a water-cooled, hollow-core coil. To compensate for the high flow resistance of its narrow tubes, we used a high-pressure pump.

4.2 Optimization of Design

We needed a coil with

- An inductance of less than 500 μH (due to constraints in the coil power supply design).

- A resistance of less than 100 m Ω (due to constraints in the coil power supply design).
- Able to produce the largest magnetic field gradient possible with a supply current of up to 60 A (due to constraints in the coil power supply design).
- A power draw at maximum current of less than 200 W (due to our power budget and cooling capacity of the chiller)
- The cooling capacity to maintain its temperature within 3 C of the water temperature fed to the coils if supplied water at 100 psi (to maintain a worst-case 1% error due to heating of the background atoms).
- An outer diameter of less than 90 mm and an axial length of less than 50 mm
- An inner diameter large enough to easily shine a 12 mm diameter beam through without clipping, even if the coils are misaligned. We decided this meant the inner diameter must be at least 15 mm.

It also must be made from inexpensive commercially available copper tubing.

We created software that brute-force modelled every coil configuration that met the above specifications and compared axial gradients. The model with the largest promised gradient was the 8 axial turn, 8 radial turn coil made from 0.125 in OD 0.65 in ID copper tubing. It would have 78 mm OD, 16 mm ID, and would be 31 mm long. When driven at maximum power (200 W), it should produce a gradient of 386 G/cm when spaced 22 cm apart, consuming 59 A. It should heat up an average of 1.9 K and heat its water 3.6 K, which flows through both coils at 13.3 mL/s. Finally, it should have a combined series resistance of 57 m Ω and inductance of 139 μ H.

The software also can model air-cooled coils, coils that are hydraulically in parallel but electrically in series, enclosed foil ‘jelly roll’ coils, and coils that are cooled by water flowing through a heatsink at their base. Details of the models used are discussed in appendix A.

4.3 Construction of Coils

Once we decided on an 8x8 design of 1/8" tubing, we needed to construct the coils. Winding magnetic field coils can be challenging. Copper work-hardens, so the copper tubing needs to be reshaped as little as possible during construction. The coils need to be kept taut so they do not unwind, but not so taut they push lower winds out of position. Finally, the copper tubing must be insulated with a material that is thin but tough enough to avoid damage when the coils winds are tightened.

We experimented with a variety of insulating materials (electrical tape, kapton tape, plasti-dip insulation, heatshrink, and spray-on magnet wire varnish) before deciding on physically separating the winds with thin sheets of plastic film.

We purchased 50' lengths of 1/8" OD .065" ID soft copper tubing from McMaster (5174K1) to create the coils. We threaded a 3/8" bolt through a 5/8" steel tube and machined 1.5" thick plastic end caps with 5/8" ID holes to act as end plates for compressing the winds. One of the end plates had a 5/16" hole drilled through it to allow the innermost winding to escape without kinking (see fig. 4.1a and fig. 4.1b). We wound the copper tubing eight times around the spindle, tightened it, then loosened it slightly to make room for the insulation.

We laser-cut 3.88 mil plastic film² circles of similar inner and outer diameter to the current copper tube layer, and cut them in one spot so they formed a single-turn helix. We then inserted them between the copper windings by forcing the copper tubing apart with a screwdriver and feeding the film into the gap. This insulated adjacent windings on each layer from each other.

On the last turn, we inserted a longer sheet the 3.88mil polyamide film in a 'C' shape around the copper to form an insulating layer between it and the next layer of windings. We then wrapped the copper in a layer of Kapton tape (to hold the windings in place), then a layer of 6.8mil plastic film³ (to provide a flat surface for the next set of winds to build upon), then a layer of thin yellow electrical tape⁴ to hold everything in place.

²3M polyester antifog film 9962 70-0000-6490-0

³3M polyester antifog film 9962 70-0000-6443-9

⁴3M #1350 YL1 ZC 5159-03

Once each coil had been wound, we vacuum-potted⁵ the coils with thermally and electrically insulating epoxy⁶. This epoxy holds the coil windings in place and reduces the amount of heat that can get transferred from the coils to the ambient air (specifically, the air near the measurement cell). This process left an uneven surface and several millimeters of material between the copper tubing and the surface of the coil.

To first order, the magnetic field gradient is inversely proportional to the inter-coil spacing. As we expected the coils to be mounted approximately 20 mm apart, this significantly reduced the maximum magnetic field gradient that could be achieved. To reduce the effect of this ‘dead space’, we used a belt sander to grind down the surface of the coil opposite the entry and exit tubing (see fig. 4.1c to see the sanded surface of one of the coils). There was a slight mistake made when sanding one of the coils - a small section of the copper tubing was also sanded away. To reduce the risk of the water pressure causing the weakened section to fail, we ensured that the high-pressure water entered the coil on the far side of the weakened tubing. This coil is the ‘more shaved’ coil in fig. 4.2.

We needed a low-resistance method of connecting the power supply to the coils. We clamped two 1/8”x1/2”x1” copper tabs together, and drilled two holes through them: a 1/4” hole normal to the two tabs so a bolt can clamp the tabs together, and a 1/8” hole along the mating surface so the tabs could clamp around the tubing. We polished the mating surfaces of the tabs and copper tubing, tinned them with solder, then used a blowtorch to solder the surfaces together (see fig. 4.1c).

Finally, we needed to be able to connect the coils to their mounts. We created an alignment tool out of a 5/8” solid aluminum rod, a 5/8” ID and 1” OD washer, and thick 3” metal plate with a 5/8” hole. We slid the rod into the hole and the washer onto the rod, then the whole assembly onto the coil mounting plate. We then slid the coil down the rod until it was approximately 1/2” from the mounting plate, then tightened the plastic grub screws in the mounting plates until they touched the coil. This ensured that the hole through the center of the coils and the hole in center of the mounting plates were concentric. The alignment tool also ensured that the coil

⁵Vacuum potting is a process where epoxy is mixed and dispensed under vacuum to prevent air bubbles or other voids from forming.

⁶MG chemicals 832TC thermally conductive epoxy

Parameter	Simulated	Measured
Resistance ⁸	57.3 mΩ	78 mΩ
Inductance ⁹	139 μH	155 μH
Combined Flow Rate ¹⁰	13.3 mL/s	14 mL/s
Magnetic Field Gradient	6.53 G/cm/A	6.4 G/cm/A
Water Temperature Increase ¹¹	2.1 K	1.5 K

Table 4.1: Simulated and measured parameters of the constructed coils.

was normal to the mounting plate. We then clamped the coil into place and built a mold of hot glue around the inside and outside gap of the coil, leaving a small hole on the outside to inject epoxy into. The mold was filled with epoxy seeded with glass microbeads⁷ to increase its strength. Once the epoxy had dried, we removed the hot glue, leaving only the epoxy (see fig. 4.1d). Finally, we clamped the electrical connection tabs to the holders, and attached nipples to the ends of the copper tubing to connect the high pressure water to.

We thank Matt Grieg and Fardin Barekat from the Make+ group at BCIT for their help with the assembly of the coils.

4.3.1 Measured Coil Parameters

Although not perfect, the simulated values provided a good starting point to make design decisions. The inductance and resistance were close enough to the simulated values for the coil controller to meet its performance specifications, and the flow rate was very accurate. Unfortunately, the coil transferred 30% less heat to the water than expected, which means that the the coils will be 30% warmer than predicted.

⁷fibertek glass micro bubbles

⁸The simulated resistance does not include the resistance of the 2 m cables connecting to the coils nor the contact resistance to the coils, which is why the measured resistance is significantly higher.

⁹We attempted to measure the inductance of the coils using impulse steps of 5 A and 60 A, which produced readings of 102 μH and 208 μH respectively.

¹⁰The flow rate was measured at a pump pressure of 150 psi.

¹¹The water temperature increase is the increase in temperature of the water as it passes through the coils. It should be approximately constant for different inflow temperatures. The increase was measured when driving 45 A through the coils—the current corresponding to the maximum steady-state power that the water cooling system could accommodate.




(a) The jig used to wind the coils



(b) The plastic insulation between winds.
Notice the 'canoe' on the inter-layer wind.



(c) Soldering the copper tabs to the copper tubing



(d) The finished product - a coil epoxied to its holder

Figure 4.1: Photographs from the coil assembly process.

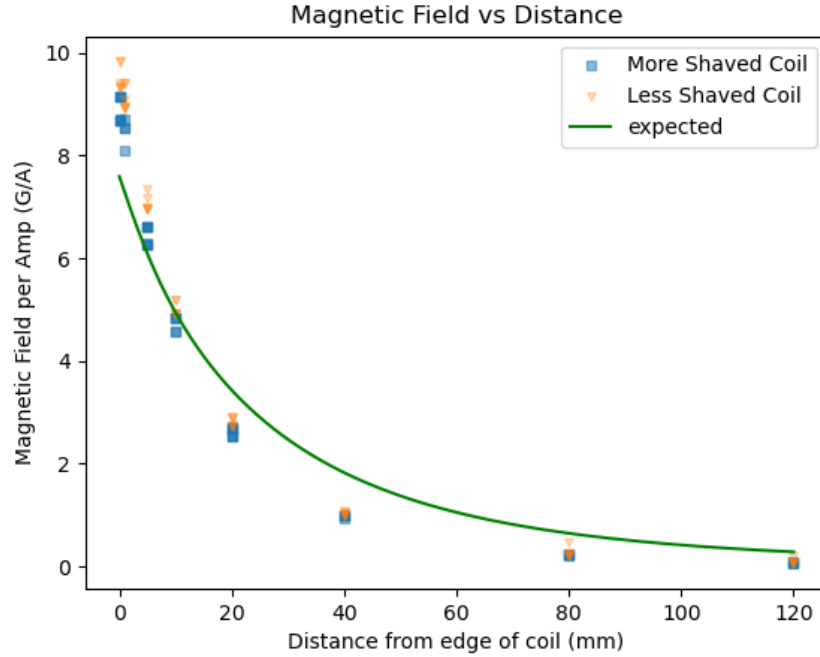


Figure 4.2: Axial magnetic field at different axial distances from the shaved surface of each of the coils. Surprisingly, the less shaved coil produced a higher gradient. It's possible the more shaved coil has a short-circuit between some of its winds.

Different measurements of the magnetic field produced different results. The magnetometer we used produced different values for different gain settings, so we took the average of all measured values. Adding the field of the two coils as measured in fig. 4.2 produces an average slope of 6.0 G/cm/A , but measuring the field directly between the coils produced a slope of 6.8 G/cm/A . It is likely that these discrepancies are due to poor calibration of the magnetometer. A more accurate value for the gradient can be measured using the atoms once the apparatus is fully reassembled.

4.4 Coil Mounts

Once we had decided on a coil design, we needed a method of mounting the coils in place.

4.4.1 Design Goals

We identified several design goals:

- No closed conducting loops that are not perpendicular to the magnetic field: This prevents eddy currents from forming, which could introduce aberrations in the magnetic field
- Able to withstand 10 N of force without deforming more than 1 mm or moving more than 100 μm : This ensures that the coils remain aligned even after being bumped or adjusted with a screwdriver.
- Able to withstand 500 N of force vertically without snapping: This is to ensure the glass cell is partially protected by the coils and mounts.
- Able to be adjusted up to 5 mm in three dimensions with 100 μm resolution: To improve the SNR, it is important to transfer as many atoms from the MOT to MT as possible. Being able to move the coils in three dimensions allows us to bring the magnetic field center to the center of the three optical beams rather than needing to keep the beams aligned while trying to find the magnetic field center. It is not critical that the magnetic field movement is *repeatable* because the exact intersection of the laser light is not precisely known. It is more important to be able to finely adjust the coil positions than it is to know how much the coil position was adjusted by.

4.4.2 Design Overview

We decided build the mounts out of 3D-printed acrylonitrile butadiene styrene (ABS) because it is non-conductive, relatively stiff, and easy 3d-print. Each mount comes in three parts: a plate that holds the coils (the ‘coil plate’, fig. 4.3a), a plate that connects to the outside world (the ‘base plate’, fig. 4.3e), and a translation

stage with dovetail slots (the ‘translation plate’, fig. 4.3b). The base plate and the coil holder plate attach to the translation plate using dovetail tabs (fig. 4.3c).


We can translate each coil independently in the horizontal plane using screws that push against the dovetail tabs. To move them, one must loosen the screws on one side of the translation plate, then tighten the screws on the other. The top base plate is mounted to the bottom base plate using four 1/2"x7" posts, which should leave ~0.8 in (~20 mm) between the coil faces. This means that the axial distance between the coils is fixed, but the magnetic field center can be adjusted in the vertical direction by adjusting the height of the two posts the bottom base plate is connected to.

We thank Garrett Kryt from the BCIT MAKE+ group for printing the coil holders.


4.4.3 Performance

We measured the vertical force specifications by placing weights on top of the coil mounts. We measured the horizontal force by pushing a scale against the screws and converting the mass reading to newtons.


- The coil mounts could withstand 50kg being placed on top of them without snapping.
- We found we could easily adjust the coil position in increments of 100 μm by turning the screws, and could move along the full 10 mm range in each direction.
- We needed to add a third adjustable support post to the base plate, but once it was installed, the coil mounts dynamically moved 1 mm or less when 10 N was applied and shifted 100 μm or less afterwards. This is more flexible than is ideal—future mounts will be printed out of a stiffer plastic—but still should be adequate for our purposes.




(a) The plate that the coils are epoxied to. The inner threaded holes are for the grub screws for aligning the coil prior to epoxing it, and the outer threaded holes are for mounting the RF knife or other accessories.




(b) This plate is what the dovetail tabs slide along. One set of slots is 2.75" apart to interface with the coil plate, and one is 2" apart to interface with the base plate. The threaded holes are 1" deep so that the tab threads will give before the translation plate threads.




(c) These dovetail tabs are what connect the base and coil plates to the translation plate. The hexagonal cutout in the bottom houses a metal nut that the screws screw into. We printed the tabs at a range of widths from equal to the width of the dovetail slot in the translation plate to .02" narrower. We found that .007" narrower generally provided close sliding fit.



(d) These tabs screw into the translation plates so that the push screws have something threaded to push against. They are separate from the translation plate to protect against stripping - because everything is plastic, it is possible to over-torque the push screws and strip the threads. These tabs are easy to replace in-situ, while the replacing the translation plate is more expensive and complicated.



(e) A single mounted coil



(f) Both coils mounted together.

Figure 4.3: A collage of images of the coil mounts

Direction	Force (N)	Dynamic Displacement (mm)	Permanent Displacement (mm)
Weak Side	10	1	0.01
Weak Side	50	4	0.1
Strong Side	10	0.1	0.01
Strong Side	50	0.5	0.1

Table 4.2: Dynamic (while loaded) and Permanent (after loaded) displacement of the upper coil when a force was applied along its strong side (in line with the 1.5” support posts) and weak side (perpendicular to its 1.5” support posts). Displacement of the lower coil was significantly less as the mounts tended to torque about the support posts.

4.5 Magnetic Field Coil Driver

To power the coil, we wanted a coil power supply that could supply current to a $\sim 60\text{ m}\Omega$, $\sim 140\text{ }\mu\text{H}$ coil and

- drive the output current from zero to within 1% (preferably .1%) of its full-scale value (0 A to 60 A) in under 1 ms.
- maintain the set current within 1% (preferably .1%) of its set value.
- programs its set current via a -10 V to 10 V BNC input.

Matt Grieg from the MAKE+ group at BCIT (with help from Fardin Barekat and Behnam Mohammad) built a custom power supply that partially.

4.5.1 Theory of Operation

The coil driver operates on a 50 kHz control loop on a microcontroller. The control loop operates in two separate modes depending on how far the output current is from the setpoint.

In normal operation, the control loop measures the input signal voltage, as well as the voltage across and the current through the coils. It controls the duty cycle of a 14-bit 200 kHz PWM signal alternating between -5.75 V and $+5.75\text{ V}$, which is fed through a two-stage LC filter to produce V_{out} , which is calculated using

$$V_{out} = V_{target}(t) + k_p e(t) + k_i \int e(t) dt \quad (4.5)$$

where $I_{set} = 6.5 \text{ A}/VV_{in}$, $V_{target} = RI_{set}$, and $e(t) = V_{target} - V_{measured}$.

However, if the difference between the set current and the measured current is greater than 1.5 A, then a high-voltage rail discharges directly into the coil, forcing the current to rise at approximately $\frac{dI}{dt} = \frac{V_{rail}}{L_{coil}}$. The control loop monitors the current through the coil and disconnects the high-voltage rail when it predicts the current will reach the target value. The high-voltage rails are kept stable at $\pm 10 \text{ V}$ with a 16 mF bank of capacitors.

The coil controller does not directly control the coil current. Instead, it monitors the resistance of the coil and calculates the necessary set voltage to produce the desired current. The resistance of the coil is calculated as follows:

1. Collect 100 samples of the current and voltage. Each sample is taken at 50 kHz, so it takes 2 ms to collect 100 samples.
2. If none of the samples were taken while the high-voltage rail was engaged, divide the mean voltage by the mean current and push it onto a 64-word FIFO.
3. Take the mean value of the FIFO.

4.5.2 Water Cooling

Matt and Fardin also assembled a water-cooling system to ensure that the coils remained cool while being driven at high current. See fig. 4.4 for a schematic of the water cooling system.

The water chiller's pump can only provide ~ 1.7 psi of pressure, so it chills a 2.4 L reservoir of water¹² that the high-pressure pump draws from.

The high-pressure pump is a rotary vane pump driven by an electric motor at 30V (1/3 of its maximum voltage of 90V) and produces ~ 150 psi of pressure. The pump produces some vibration in the water which the accumulator¹³ smooths out. The high-pressure gauge is used to monitor the pressure going into the coils and can be used to set the underpressure release switch, which turns off the pump if

¹²The reservoir was created by gluing PVC plates to a 6" tube of PVC.

¹³The accumulator is just an intentional air void in the system, which forms a pneumatic spring for damping out vibration.

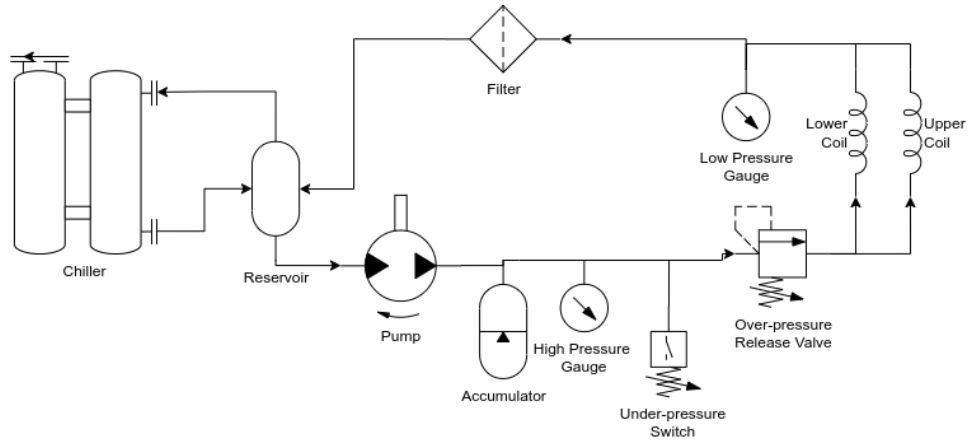


Figure 4.4: A schematic of the hydraulics used to cool the magnetic field coils. Note that water flows counter-clockwise, following the direction of the arrows.

Part	Part No
Pump	PB 101x Brass Rotary Vane Pump
Pump Motor	Baldor GPP2530
Chiller	Haake Circulating Bath D1-G
Reservoir	N/A (homemade)
Accumulator	N/A (homemade)
High Pressure Gauge	McMaster 4089k64
Pressure Switch	P88G-5-C15TB
Over-pressure Release Valve	Internal to pump, set to 300 PSI
Low Pressure Gauge	Unknown
Filter	McMaster 4422k4
Chiller Hose	Kurivama 3/8" ID 275 PSI
High-Pressure Hose	Flex-Loc 300 Non-Conductive 3/8" ID IC-114/03

Table 4.3: Parts used in the coil cooling hydraulics.




Figure 4.5: The vibration damping mounts the pump sits on.

the pressure drops below a certain threshold (currently set to ~ 120 psi). The low-pressure gauge is used to monitor the pressure going into the filter. The filter is in place because we discovered a small residue in the chiller after running the pump for several days - possibly from the pump. All components in the high-pressure loop are connected with high-pressure hose rated for pressures up to 300 psi.

Vibration Damping

The pump and water chiller produce some noise and vibration that needs to be dampened for the lasers to remain locked. The heavy side of the pump is mounted on a spring with rubber washers on either end. The light end is held in place with a rubber vibration damping sandwich mount ((McMaster 93945K117)) with the screw isolated using a rubber washer. See fig. 4.5 for an image of the mounts.

The chiller sits on top of four vibration-damping rubber mounts (McMaster

60525K25) to reduce the vibration from its pumps. All of the hydraulics are mounted on a plastic sheet that is bolted to the frame of the cart using rubber vibration damping sandwich mounts (McMaster 93945K118) to further reduce transmitted vibrations. The vacuum optics and locking optics were mounted on a third set of rubber vibration damping sandwich mounts (McMaster 93945K117) to add a final layer of protection from vibrations. The end result was a reduction in amplitude of the vibrations by an order of magnitude (see fig. 4.6 and fig. 4.7). No significant vibrations were transmitted to the floor (fig. 4.8), so the chamber under test does not need vibration isolation. It is also more likely that the chamber under test will introduce more vibration than this apparatus introduces as orifice flow standards are typically pumped out by a turbo pump backed by a scroll pump.

Anecdotally, we discovered that a person could not feel the vibrations from the pump or the chiller when they placed their hand on the vacuum optics table. In fact, we found that we could hit the cart with a screwdriver and even the vibration from that was difficult to feel on the vacuum optics tables. After installing the vibration damping, hitting the rack with a screwdriver also did not unlock the lasers—a marked improvement from prior performance, where merely tapping the rack could unlock the lasers.

4.5.3 User Controls

Coil Driver Touchscreen

The operating mode of the coil controller is set via a touchscreen (see fig. 4.10 for an image of the touchscreen). The screen also provides useful information about the operation of the coil:

- The output current is the current that the board measures it is producing. We found that this value was slightly inaccurate (discussed in section 4.5.6) due to a potential difference across the ground plane of the board.
- The output voltage is voltage between the board's two output terminals. Again, we found this value was slightly inaccurate due to the ground plane potential difference issue.

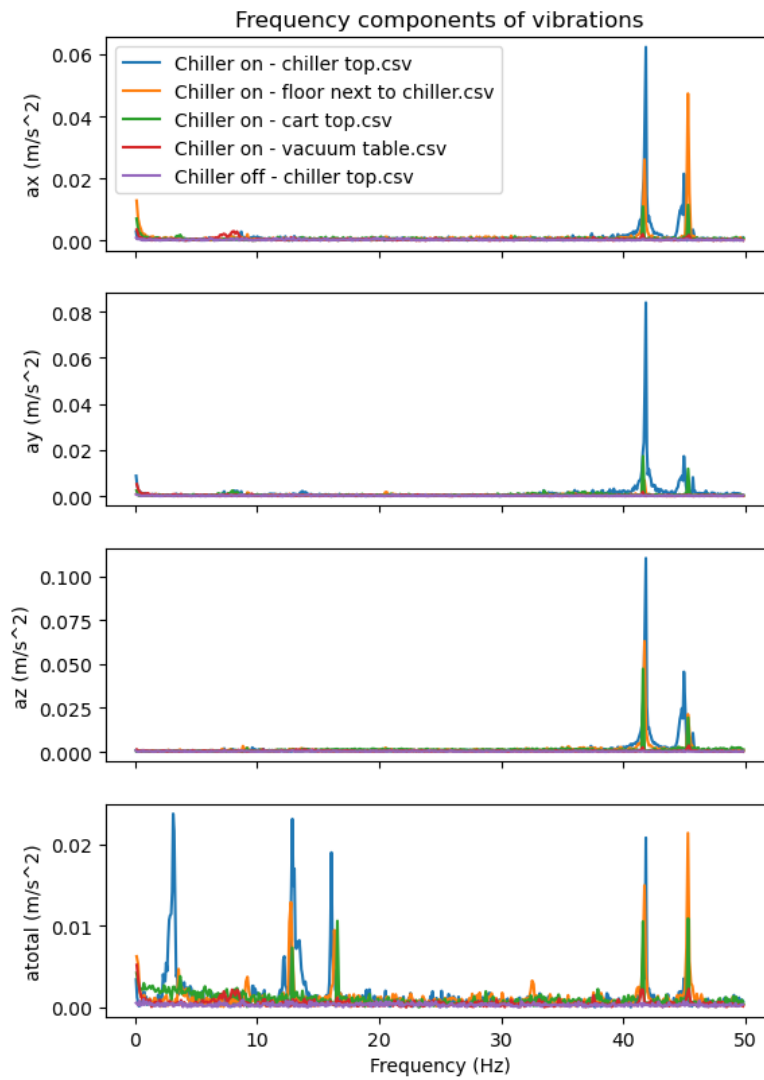


Figure 4.6: A demonstration of the effects of the different stages of damping on the vibrations produced by the chiller.

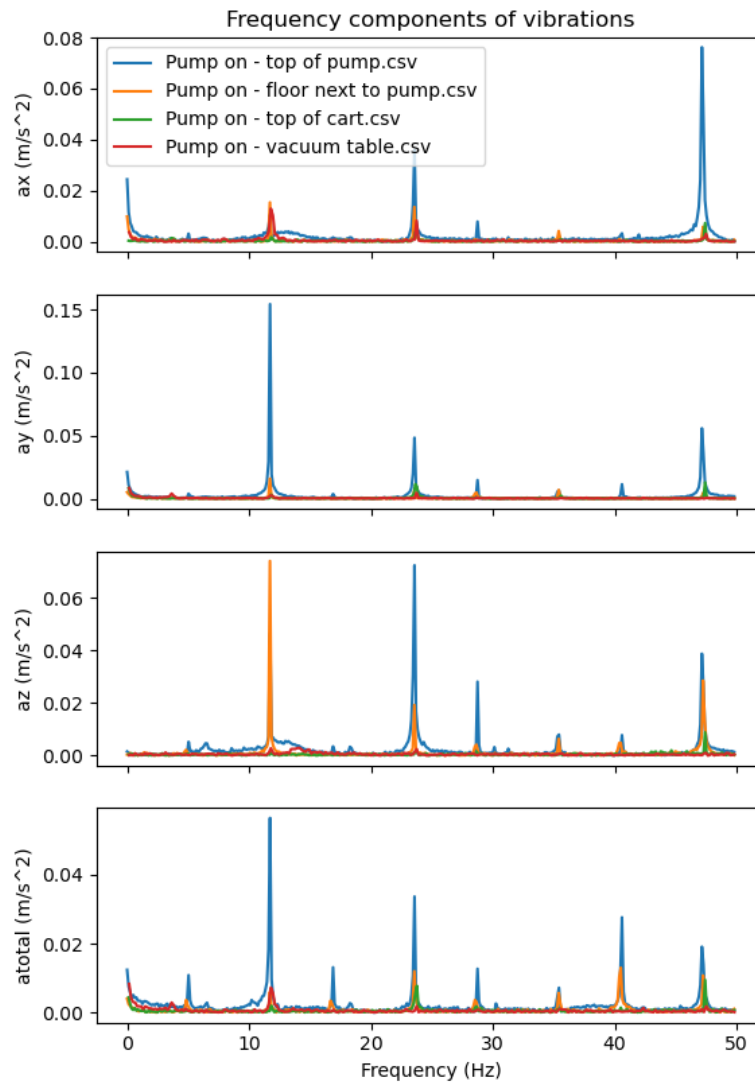


Figure 4.7: A demonstration of the effects of the different stages of damping on the vibrations produced by the pump.




Figure 4.8: Acceleration on the floor of the room next to the cart with and without the pump and chiller turned on. Any vibrations transmitted to the floor are below the noise floor of the accelerometer we used.

- The left radio buttons select what controls the output current. ‘Off’ turns off the output, ‘Remote’ allows control via the BNC input, and ‘Local’ generates a hardcoded test pattern.¹⁴
- There are three displayed temperatures. ‘PCB’ refers to the temperature of the display controller PCB. Coil1 and Coil2 measure the temperature from two K-type thermocouples connected to two thermocouple terminals on the display controller PCB. They currently are terminated with 1” long thermocouple cables, but can be replaced with longer thermocouples for monitoring the coil temperature.

¹⁴The hardcoded test pattern alternates between producing 5 A for 505 ms, 0 A for 5 ms, 65 A for 1000 ms, then 0 A for 105 ms. It then repeats the cycle for negative currents, then repeats the entire sequence infinitely.




Figure 4.9: Accelerations on the vacuum optics table and the cart when the cart is hit by a screwdriver.

- The displayed resistance is the resistance that is described in section 4.5.1. It is updated every 128 ms. The displayed inductance is calculated once at startup.
- The High bus voltages are the high-voltage rails used to drive fast current transitions. They can be user-adjusted using potentiometers R32 and R34, but it is not advised.
- The user can enable or disable the high-voltage kick using the ‘HV enable’ button
- ‘Test Load’ is not currently implemented as of 2022-09-13. It allows the user to re-test the coil’s inductance and resistance
- The counter next to HV enable is not currently implemented. In a future




Figure 4.10: An example image of the touchscreen controlling the coil driver.

board revision, the goal is to allow the user to control the voltage of the rails directly from the touchscreen.

- ‘Spare’ is not currently implemented.

Non-Touchscreen Coil Driver Controls

- The BNC on the control panel (fig. 4.12) allows the user to control the output current when the board is set to ‘Remote’ mode. The output current should be 6.5A per volt fed in, to a limit of ± 10 V.
- The green and red buttons control the water cooling pump. Green turns it on, and red turns it off. If the green button is released before the pressure switch threshold is reached (~ 120 psi), then the pump will turn off.
- The collar on the pressure switch behind the control panel can be turned to adjust the pressure switch threshold (see fig. 4.13)

- The red button on the water filter (see fig. 4.11) can be depressed to purge the filter of air. It works by opening a small hole on the top of the filter housing. The pump pressurizes the filter, so the more-dense water displaces the air out through hole. Releasing the button closes the hole, preventing pressurized water from escaping once the air has been purged.

4.5.4 Error Conditions

- On startup, the coil driver checks to make sure the resistance and the inductance of the coil is within bounds ($75\text{ m}\Omega$ to $95\text{ m}\Omega$, $100\text{ }\mu\text{H}$ to $225\text{ }\mu\text{H}$). If one or both of the values are out of bounds, the controller reverts to an error state which requires a power cycle to clear.
- If at any point the calculated resistance of the coil falls out of bounds, the controller will revert into an error state which requires a power cycle to clear
- If two successive high-voltage ramps are applied in the same direction within 20 ms, the second ramp is ignored until 20 ms has elapsed since the first ramp.

4.5.5 Limitations

High-voltage rails: The high-voltage rails are supplied via voltage regulators. The regulators can supply up to 1.5 A to the rails (depending on thermal effects). Because voltage regulators cannot sink current, each rail also has a $100\text{ }\Omega$ drain resistor between it and ground, which means they can supply energy at a rate of $>14\text{ J/s}$ and drain energy at a rate of 1 J/s . In the worst-case scenario, jumping from 0 A to 65 A requires putting $\sim 0.36\text{ J}$ of energy into the magnetic field of the coils, which means it should take $\sim 26\text{ ms}$ to recharge the rail for the next jump (ignoring resistive effects). However, jumping 65 A to 0 A takes 14 times longer to dissipate the energy from the coil, so it takes $\sim 360\text{ ms}$ for the rail to return to its nominal voltage. This may limit how quickly successive measurement can be made using the system.

As the above scenario is the worst-case scenario - lower-current jumps involve less energy, and crossing past zero allows some of the energy that was stored in



Figure 4.11: The water filter with the red purge button.




Figure 4.12: The control panel for the coil controller and hydraulics. It houses the touchscreen, the pump controls, the BNC for remote control, and the high pressure gauge.

the rail to be drawn again - but it limits the rate of successive high-speed jumps the controller can execute. The controller has no lockout for the voltage of the rails moving too far away from the nominal value, so the user should ensure the rail voltages don't move outside of 9 V to 11 V by monitoring the voltages shown on the display and adjusting the measurement duty cycle if the power rail voltages drift outside of the desired range.

If the low rate of power dissipation on the rails is limiting the desired behaviour, smaller drain resistors can be connected in parallel to the existing resistors. It is not recommended that the net drain resistance fall below 25 Ω as it will put a lot of strain on the voltage regulators.

Coolant Temperature: While the cooler datasheet claims it can dissipate 200 W of heat when the air is 20 C, we found the cooling system can only dissipate ~ 125 W from the coils while keeping the water temperature below the air temperature of 20 C. This inefficiency is likely due to losses of heat through the tubing and inefficiencies in the reservoir. At 40 A, the coils produce around 125 W of heat, so the coils can be run from 0 c to 40 continuously without adverse heating effects. However, at 65 A, the coils produce almost 260 W of heat. If currents from 40 A to 65 A are expected to be used for extended periods of time, the user should monitor the water temperature to make sure the coils are not heating to unacceptable levels. If the coil temperature becomes too high, the user should increase the delay between successive measurements to allow time for the water to cool down. If that is not possible, one can use workarounds such as introducing




Figure 4.13: The underpressure valve with its control collar (the component with the white label).

open-loop water cooling, or slowly replacing water with ice cubes inside of the chiller, but replacing the chiller with a higher capacity model such as the H50-500 is recommended.

The high-current duty cycle limitations are not expected to dramatically limit the performance of the apparatus as each measurement cycle includes times of low current (when loading the MOT) or no current (when recording the background light level). Current experiments in our lab use magnetic field gradients of 250 G/cm or less, which means this controller will not have to exceed 40 A to replicate their results. The thermal limitation of the cooler may only limit performance when making measurements at very low pressures ($<1 \times 10^{-10}$ Torr) using very deep magnetic traps.

Water Filter: An unknown source—likely the high-pressure pump—slowly produces debris in the water. To prevent the debris from clogging the coils, we installed a water filter. The water filter has a pressure gauge in front of it; replace the filter if the low pressure gauge reading rises above ~ 10 psi.

Displayed Inductance: The inductance value displayed on the touch screen is not always accurate - the cause is not known as of 2022-09-13. Because the calculated inductance is not used for anything other than validating the coils are attached to the driver, the issue was not addressed in the current prototype.

Only Input Step Functions: The controller was designed to respond to step inputs. If the input signal changes, it should vary quickly (such as a ramp from 0 V to 10 V that occurs within 1 μ s. If given a continuous input function (such as a sinusoidal or triangle wave signal of frequency 1×10^5 Hz to 10 Hz), it may cause a resonance in the control loop, causing erratic or damaging behaviour.

Disabling the high-voltage rails can improve the controller's resilience to continuous input functions, but will diminish its response time to large step inputs, increasing its settle time by an order of magnitude.

4.5.6 Performance

Fig. 4.15 shows that the controller can ramp from 0 A to 64 A or back in under 1 ms—important for quickly switching from a MOT to a deep magnetic trap and back. This meets our maximum current and rise time requirements. However,




Figure 4.14: The step response of the coil driver for different set currents.
See fig. 4.16 for the long-term behaviour.

fig. 4.14 shows that that there are some flaws to the output:

- It takes roughly 5 ms for the current to stabilize. This is because the PWM output filter takes too long to reach its target value. Future board revisions will have a higher knee frequency for the filter which will increase the response time at the cost of a higher ripple frequency.
- The initial ‘settle value’ that the controller reaches after ~ 5 ms can be up to 2% off of the set current, and it may vary slightly (by up to 2% of the set current) over the next 500 ms. This is because there is a small voltage across the ground plane of the board, causing inaccurate voltage and current readings. This means that the controller initially controls the coil to a slightly wrong voltage, causing the initial offset. Then, as the new resistance values propagate through the FIFO, the controller changes the output to a new set




Figure 4.15: A view of the initial current ramps displayed in fig. 4.14.

point over the next few hundred milliseconds (see fig. 4.16). Future board revisions will include differential measurements to eliminate problems due to voltages across the ground plane.

The output current instability is not optimal, but should not critically compromise performance of the apparatus. As the trap depth is ultimately set with an RF knife, we decided to accept the flaws with the coil controller so that we could assemble the apparatus more quickly. A revised board with differential sensing, a faster output filter, and an improved microcontroller is currently being designed.




Figure 4.16: A long-term view of the current stability of the coil driver. The noise is likely primarily measurement noise, but the drift is due to flaws in the controller design. The settling time is longest at 65A because the controller is at the top of its range. It is also why the current never reaches the set current.

Chapter 5

Optical Design

5.1 Locking Optics

5.1.1 Pumping Scheme

We plan to trap rubidium-87 atoms using the D_2 line hyperfine transitions (see fig. 5.1). The majority of the light needs to be resonant with the $F=2$ to $F'=3$ transition (the ‘pump’ light), but we also need some light resonant with the $F=1$ to $F'=2$ transition to ‘repump’ atoms that fall into the $F=1$ state back into the $F=2$ state

5.1.2 Specifications

Therefore, in order to trap ^{87}Rb atoms, we needed to generate the following frequencies and amplitudes of light:

- For the 3D MOT, we wanted one fiber-coupled beam with
 - $>40\text{ mW}$ of ^{87}Rb $F=2-3'$ ($384.228\,115\,2\text{ THz}$) pump light whose
 - * frequency can be detuned from 0 MHz to -30 MHz .
 - * frequency can be set within 1 MHz .
 - * linewidth is less than 6 MHz (the atomic transition linewidth).

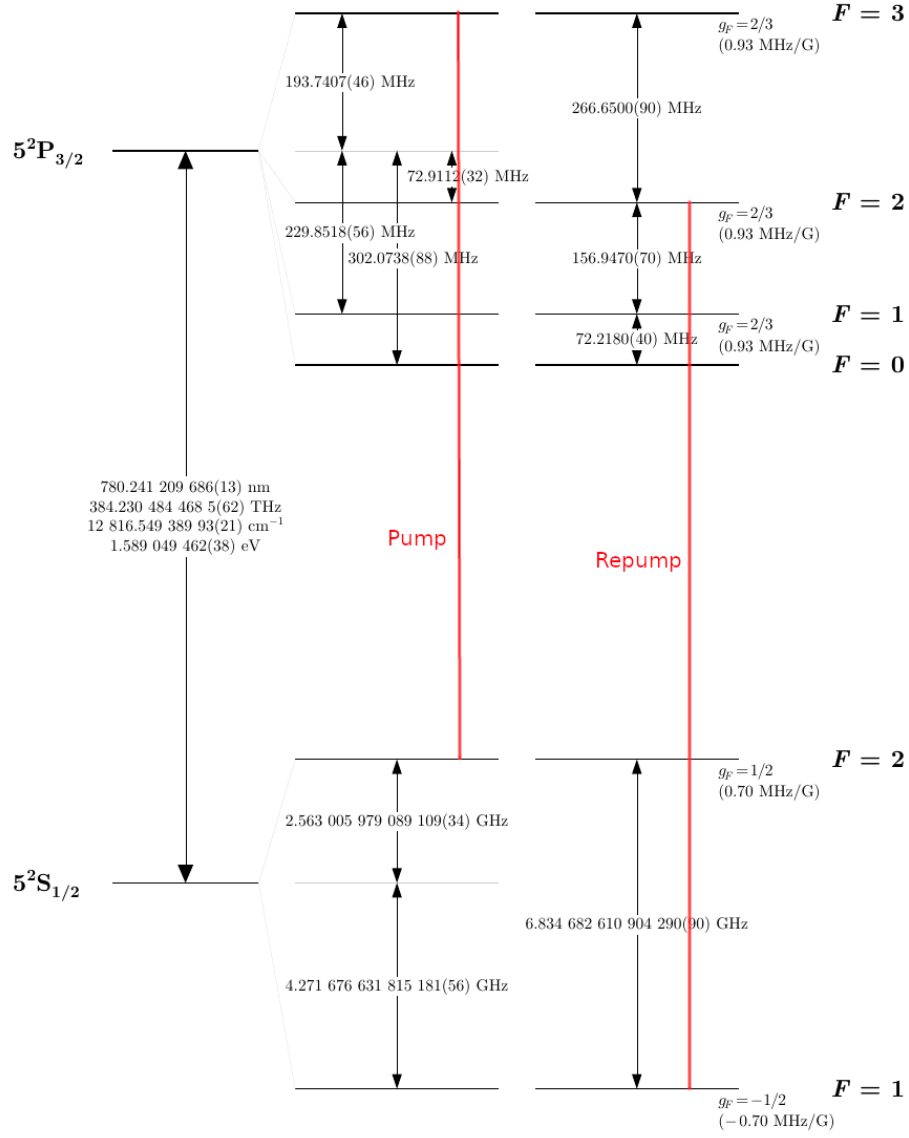


Figure 5.1: Hyperfine structure of Rubidium-87 D2 line (taken from ref [48], fig 2) and annotated with the pump and repump transitions.

- * power can be attenuated to less than $10\ \mu\text{W}$ within $10\ \mu\text{s}$ and can be attenuated to less than $100\ \text{nW}$ within $10\ \text{ms}$.
- $>1\ \text{mW}$ of ^{87}Rb $F=1-2'$ ($384.234\,683\,2\ \text{THz}$) repump light whose
 - * frequency can be detuned from $0\ \text{MHz}$ to $-30\ \text{MHz}$.
 - * frequency can be set within $1\ \text{MHz}$.
 - * linewidth is less than $6\ \text{MHz}$ (the atomic transition linewidth).
 - * power can be attenuated to less than $100\ \text{fW}$ within $10\ \text{ms}$.
- For the 2D MOT, we wanted two fiber-coupled beams with polarization parallel to the key of the fiber, each with
 - $>20\ \text{mW}$ of ^{87}Rb $F=2-3'$ ($384.228\,115\,2\ \text{THz}$) pump light that can be detuned from $0\ \text{MHz}$ to $-30\ \text{MHz}$ and whose frequency can be set within $1\ \text{MHz}$.
 - $>0.4\ \text{mW}$ of ^{87}Rb $F=1-2'$ ($384.234\,683\,2\ \text{THz}$) repump light whose frequency can be set within $1\ \text{MHz}$.
 - The ability to effectively attenuate the repump light to at least $<100\ \text{nW}$.
- For the push beam, we wanted one fiber-coupled beam with
 - $>2\ \text{mW}$ of ^{87}Rb $F=2-3'$ ($384.228\,115\,2\ \text{THz}$) pump light that can be detuned from $10\ \text{MHz}$ to $-10\ \text{MHz}$ and accurate within $1\ \text{MHz}$.
 - The ability to attenuate the light to less than $10\ \mu\text{W}$ within $10\ \mu\text{s}$ and less than $100\ \text{nW}$ within $10\ \text{ms}$.

The pump detuning range was specified so that the atoms can be trapped in a MOT (requiring a pump detuning of $-12\ \text{MHz}$ to $-18\ \text{MHz}$), and cooled (requiring a pump detuning of at least $-30\ \text{MHz}$) before being transferred to a magnetic trap. Maintaining a resolution of at least $1\ \text{MHz}$ gives us flexibility in what detuning is set, allowing us to find the optimal MOT settings for capture of and imaging the rubidium atoms.

The steady-state scattering rate of cold ^{87}Rb atoms is given by eq. (9.5). Atoms can fall out of a magnetic trap by if they leave the $|F = 1, m_F = -1\rangle$ state. The longest lifetimes we can expect to see in a magnetic trap are in the order of $10\ \text{mHz}$

[43], so to maintain an error of less than 1%, we need the ultimate scattering rate of all light reaching the atoms to be less than $100 \mu\text{Hz}$. We use eq. (9.5) to estimate the intensities of light to achieve this. The specifications for the 2D MOT are looser because the light has no direct path to the 3D MOT.

We based the power requirements on the input powers suggested in the ColdQuanta 2D MOT user manual [8], although we removed the repump light from the push beam. This is because we planned to use fiber-coupled components to shutter the push beam light. Fiber-coupled shutters and AOMs never completely block the light passing through them, and the push beam light shines towards the 3D MOT. Any repump light shining into the MOT would be disastrous—it only takes 1 pW of on-resonance repump light to bring the scattering rate above 1 mHz, which could disrupt low-pressure measurements. This decision reduced our maximum attainable loading rate, but section 9.2.2 illustrates how we could still achieve adequate loading rates without any push beam.

5.1.3 Design

Both lasers are powered by Vescent D2-105 Laser Controllers. One ‘reference’ laser (a Vescent D2-100-DBR standard model) is locked to the 85Rb D2 $F=3-3'/4'$ crossover transition using the Vescent D2-210 saturated absorption spectroscopy cell and a Vescent D2-125 laser servo. Its light is combined with the light from a second ‘power’ laser (a Vescent D2-100-DBR high-power model) using a Vescent D2-250 heterodyne module. A Vescent D2-135 offset-locks the light to 90 MHz above the desired pump transition - the 87Rb D2 $F=2-3'$ hyperfine transition. If the pump light should be on-resonance, then the power laser should be 976.2 MHz below the reference laser (90 MHz above the desired output frequency).

After splitting off some light for the heterodyne, the power laser’s light goes through a quarter-wave plate (QWP), half-wave plate (HWP), then a Glan-Thompson linear polarizer (GT). This is to clean up the polarization before it goes into the fiber; despite using polarization-maintaining fibers, we noticed that if light is launched into the fiber off-axis, the EOM’s insertion loss skyrockets.

Roughly 32 mW of the power laser’s light is coupled into a polarization-maintaining fiber where it is split into two separate fibers. 25% of the light goes through a fiber-

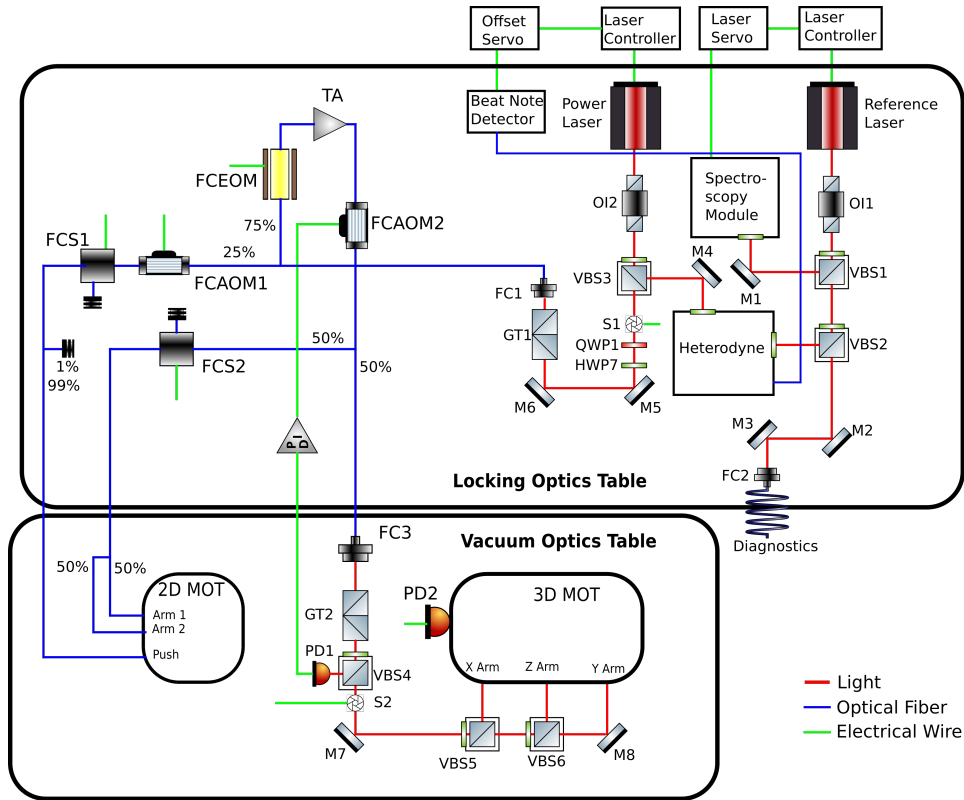


Figure 5.2: Above is a diagram detailing outlining how the light is generated and transported to the 2D and 3D MOTs. Part numbers and abbreviations are in table 5.1

coupled acousto-optic modulator (AOM) and a fiber-coupled shutter to be used as a push beam. The AOM shifts the frequency of the light by -90 MHz and can attenuate it by 40 dB to 50 dB within 60 ns. The shutter takes 3 ms to respond, but adds an additional 50 dB of attenuation.

75% of the light goes into the electro-optic modulator (which has a maximum optical input power of 25 mW) to be amplified by the 1W tapered amplifier. Once amplified, the light is coupled into a fiber-coupled AOM (which also drops the frequency of the light by -90 MHz and provides up to 40 dB to 50 dB of attenuation), then split equally into two paths. One path goes through a fiber-coupled shutter, a polarizing fiber, then through a 50:50 splitter to produce the two arms of the 2D

Part	Part No
Pump Laser	Vescent D2-100-DBR High Power
Reference Laser	Vescent D2-100-DBR Standard
Spectroscopy Module	Vescent D2-210
Heterodyne	Vescent D2-250
Beat Note Detector	Vescent D2-260
Offset Servo	Vescent D2-135
Laser Servo	Vescent D2-125
Laser Controller	Vescent D2-105
Optical Isolator (OI1)	IO-D-780-VLP
Optical Isolator (OI2)	Vescent D2-130
Variable Beamsplitter (VBS)	Thorlabs VA5-780
Glan-Thompson Polarizer (GT)	Thorlabs GTH10M-B
Fiber Coupler 1 (FC1)	Thorlabs CFC2-B
Fiber Coupler 2 (FC2)	Vescent D2-120
Fiber Coupler 3 (FC3)	C40APC-B
Fiber-coupled AOM (FCAOM)	Brimrose TEF-90-10-780-2FP-HP/PM
Fiber-coupled EOM (FCEOM)	iXblue NIR-MPX800-LN-10
Fiber-coupled Shutter (FCS)	Photonwares FFSW-123740323
Diaphragm Shutter 1 (S1)	Thorlabs SHB05
Diaphragm Shutter 2 (S2)	Thorlabs SHB1
Taper Amplifier (TA)	Toptica BoosTA 1W
Photodiode (PD)	Thorlabs PDA100A2
50:50 splitter	Thorlabs PN780R5A1
75:25 splitter	Thorlabs PN780R3A1
99:1 splitter	Thorlabs PN780R1A1

Table 5.1: The primary components used in the locking and distribution optics.

MOT. The other path is coupled into a free space 8 mm diameter beam to produce the 3D MOT. The 3D MOT light goes through a GT to clean up its polarization, then some light is picked off by a variable beamsplitter to use for feedback control of the light intensity. The rest of the light passes through a shutter (for additional isolation), then through two more variable beamsplitters to produce the three beams of light that are retroreflected for the 3D MOT.

This scheme offers several advantages:

- There is an excess of power available at every stage that requires fiber-coupling, so it is easy to re-align the optics to an acceptable level.
- Many of the difficult-to-align optics (such as AOMs) are fiber-coupled, eliminating the need for alignment.
- It is possible to set the pump frequency anywhere in a 26 GHz range depending on the lock point, and the repump frequency can be anywhere from 0.05 GHz to 6.8 GHz away from the pump frequency. Any pump frequency within a ± 0.25 GHz to ± 10 GHz range can be reached within a few microseconds, and any repump frequency can be reached within a few milliseconds. This means that the lasers are able to trap either Rubidium-87 or Rubidium-85, and can switch between trapping them in a few milliseconds (provided the reference laser is locked to a known repump transition).

However, this scheme has one major weakness: Polarization drifts can accumulate through the fibers, causing variations in splitting ratio of the fiber splitters. This is discussed in further detail in section 5.1.4.

Reference Laser Lock

This module allows the user to perform doppler-free saturated adsorption spectroscopy [40]. Naively, one might hope to lock a laser to a certain hyperfine transition by shining the laser through a Rb vapour cell and using a photodiode to see how much light passes through. The laser could be scanned in frequency, and the amount of light transmitted through the cell would dip when a transition is encountered. Feedback control electronics could keep the laser centered on the bottom of the correct ‘dip’.

Unfortunately, at room temperature, Rubidium atoms move at hundreds of meters per second. When the atoms encounter photons of a given frequency in the lab frame, they see a doppler-shifted frequency. Because the atoms are thermal, their velocities follow the Maxwell-Boltzmann distribution, meaning that light even 1 GHz away from a transition frequency shone into a cloud of room-temperature Rb atoms will interact with at least a few atoms. This makes the observed transition width ‘broadened’ to $\Delta f_{FWHM} = \sqrt{\frac{8kT \ln 2}{mc^2}} f_0$ which is ~ 500 MHz for room-




Figure 5.3: This illustrates what is inside of the Vescent D2-210 Saturated Adsorption Spectroscopy module (taken from [58]).

temperature Rb atoms [32]. The hyperfine transitions have a natural linewidth of ~ 6 MHz and are separated by ~ 100 MHz, so they are swamped by the doppler broadening [48].

One cannot use the scheme as designed above, but it can be modified to see through the doppler broadening by adding a more powerful ‘pump’ beam of the same frequency as the ‘probe’ beam that propagates antiparallel to the probe beam. When the pump light interacts with a given velocity class of atoms ¹, it excites them. If the probe light is off-resonance from the excited atoms, then it will be attenuated as it will be absorbed by a different velocity class of atoms. However, if the probe light is on-resonance with the atoms excited by the pump beam, then

¹A velocity class of atoms is the ensemble of atoms that share a velocity component along the propagation direction of the light; all atoms where \vec{v}_z is the same number.

it will see an increase in its transmittivity because there are fewer atoms that can absorb the light. This means that there will be small spikes in the optical spectrum when the pump light excites the velocity class of atoms that the probe light is on-resonance with. There are two cases that cause this:

- **Resonance Peaks:** In this case, the pump and probe beams are on resonance with the same transition (in the $v_z = 0$ velocity class). This produces a spike in the optical spectrum of the same frequency as the excited transition.
- **Crossover Peaks:** In this case, the pump and probe beams interact with two different transitions in two different velocity classes of equal magnitude and opposite direction (e.g. in the $+v_z$ class, pump beam excites the $F=2-2'$ transition while the probe beam is on-resonance with the $F=2-3'$ transition, and in the $-v_z$ class the pump beam excites the $F=2-3'$ transition while the probe beam is on-resonance with the $F=2-2'$ transition). In this case, the spike in the optical spectrum is the average of the two transition frequencies.

We use a Vescent D2-210 to generate the signal for the laser lock. The light enters the module via a half-wave plate that feeds into a PBS so the user can adjust the power feeding into the optics. One arm passes through the PBS and out a hole in the back (used for alignment). The other arm goes to the rest of the optics, where it passes through a cleanup PBS, then into a 50-50 non-polarizing beamsplitter. One arm goes to a reference photodiode, while the other arm passes through a fixed half-wave plate and PBS to split into a pump beam and a probe beam. The probe beam continues straight through the PBS, bounces off of a mirror, passes through the vapour cell, passes through the PBS, and hits the signal photodiode. The pump beam bounces off of a mirror and PBS so that it propagates the opposite to the probe beam through the vapour cell.

The spectroscopy cell produces a local extrema at the lock point. However, analog feedback control electronics work best at locking to a zero-crossing of a slope, so there needs to be a way of locking to the frequency derivative of the signal. The Vescent D2-125 accomplishes this by modulating the output signal with a 4 MHz dither signal and demodulating it on the input (similar to how an FM radio works). This effectively takes the frequency derivative of the system at the




Figure 5.4: This illustrates what is inside of the Vescent D2-250 Heterodyne module (taken from [59]).

cost of imposing 4 MHz sidebands on the output. Because the spectroscopy signal has more than one peak, the zero-crossing of the slope in the derivative signal may not occur at the precise location of the peak. Therefore, the controller allows the user to add a DC offset to the error signal to make the zero crossing of the error signal occur at the extrema of the spectroscopy signal. From there, a PI^2D loop filter (standard PID plus an additional integrator to boost gain at low frequencies) keeps the laser locked to the zero crossing.

Offset Lock

The D2-250 combines the light from the two different lasers and fiber-couples them into the beat note detector. Both inputs have half wave plates to set the amplitude of the light from that laser. We found that it is more important to keep the beams coupled into the fiber at precisely the same angle than it is to optimize the coupling efficiency - if the beams are misaligned, the beat note signal is significantly

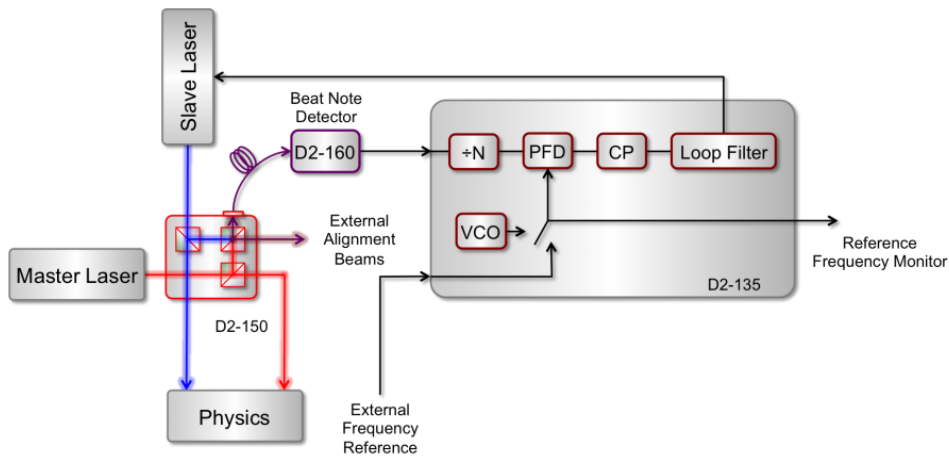


Figure 5.5: A block diagram for the offset lock (taken from [57])

noisier.²

The D2-260 is a 9 GHz beat-note detector. It digitizes the input signal by switching the state of its output every rising and falling edge. The D2-135 divides this digital signal by N (a multiple of 2) before using a phase-frequency detector (PFD), charge pump, and loop filter to lock the lasers $N \times f_{PFD}$ apart.

The D2-135 has a number of inputs that we use to control the offset lock:

- 9-pin connector:
 - Front Panel Active: Pull low to enable computer control (and disable front panel)
 - Gain sign: Set high to lock the controlled laser's frequency *above* the reference laser's frequency and low to lock it *below*.
 - VCO External: Set high to use an external VCO rather than the internal one
 - VCO High Range: If using the internal reference, effectively multiplies N by 2.

²This issue is due to the multimode fiber that Vescent uses to connect the heterodyne to the beat note detector. This exact issue is described here: https://www.vescent.com/manuals/doku.php?id=d2:260.beat_note_detector. The issue has been slightly mitigated by first coupling the light into a PM fiber before butt-coupling it to the Vescent multimode fiber.

N	Internal VCO State	Min Frequency (MHz)	Max Frequency (MHz)
8	Low	385	850
8	High	770	1700
16	Low	770	1700
16	High	1540	3400
32	Low	1540	3400
32	High	3080	6800
64	Low	3080	6800
64	High	6160	10000

Table 5.2: Offset Frequency Lock Ranges of the D2-135 [57]

- N1 and N2: $N = 8 \times 2^{2N_2+N_1}$
- Lock On/Off: Pull low to disable lock.
- VCO Freq. Adjust: This input sets the output frequency of the internal VCO (it is summed with the manual adjustment potentiometer).

Initially, we assumed that a -10 V to 10 V input would map linearly to the full range given in the D2-135 user manual (table 5.2).

However, we found that the VCO frequency was not precisely linear with the input voltage (see fig. 5.6). To compensate for the nonlinearity, we built a spline fit of the D2-135 transfer function and inverted it so that the correct voltage to get the desired output frequency was known. Currently, the control software re-fits the calibration data and creates the inverted model when an instance of the Prometheus class (the object that controls the experiment) is created. Future versions of the software will load the model from file unless a refit flag is specified to improve startup times.

There is also the option to use an external reference rather than the internal VCO. This configuration eliminates the temperature-dependant drift in the frequency generated by the VCO. We have experimented with using a DDS to generate the reference frequency, and have successfully trapped atoms using this configuration. However, it also seems to suffer from some slight nonlinearity issues that we have not fully explored at the time of writing this document.




Figure 5.6: Output frequency as measured from the D2-135 Beat Note Div 2 output vs the set frequency assuming the VCO frequency is linear across its input range. The calibration data used are stored here: <https://qdg-code.phas.ubc.ca:2633/Perrin/PrometheusPython/src/branch/master/hardware/offsetdata>.

5.1.4 Performance

Reference Laser Lock

The peak-locking circuit in the saturated adsorption lock produces large sidebands spaced 4 MHz apart. If the dither amplitude is increased enough, they can overwhelm the locking signal. This is one of the reasons we chose to use only the light from the offset-locked laser — it removes these sidebands.

The reference laser lock is very sensitive to vibration—dropping a 15/64" allen key onto the locking optics table from the height of six inches can unlock the lasers. To address this issue, we mounted the optical table on four rubber vibration




Figure 5.7: The spectrum of the reference laser. This was taken by beating the light from the reference laser against light from another locked laser in the lab, so the actual linewidth of the laser is ~ 3 MHz. The sidebands are caused by the 4 MHz dither used to create the lock signal. The sideband amplitude can be decreased at the expense of reducing the lock signal amplitude, and thus the lock stability.

damping sandwich mounts from McMaster (P/N 93945K117). These dramatically decoupled vibrations on the mounting rack from the optical table. We found we could hit the mounting rack with a screwdriver and the lasers would remain locked (see section 4.5.2).

Power Laser Lock

The power laser has a linewidth better than ~ 3 MHz—better than specified. The exact linewidth is unknown as we did not have a low-linewidth laser to compare against. The lock remained insensitive to vibration - the only thing that caused it




Figure 5.8: The spectrum of the power laser. This was taken by beating the light from the reference laser against light from another locked laser in the lab, so the actual linewidth of the laser is <3 MHz. The offset lock locks to the average frequency of the beat note, so it removes the sidebands that are present in the reference laser spectrum (seen in fig. 5.7). Technically, the fit should be a Voigt profile (the convolution of the gaussian and lorentz functions), but a gaussian fit suffices for extracting the FWHM.

to unlock was setting the gain too high or removing its lock signal.

The repump light beat note had a FWHM of 3.2 MHz. This suggests that the phase noise of the EOM is negligible.





Figure 5.9: The spectrum of the repump light (one of the sidebands generated by the EOM) beat against the same laser as fig. 5.8. The microwave frequency FWHM is ~ 35 kHz. Assuming the phase noise of the EOM is negligible, we expect spectrum of the repump and pump light to be nearly identical since the repump light is in the same field as the pump light but is frequency-shifted by the EOM. This spectrum confirms that the EOM does not add significant phase noise. Technically, the fit should be a Voigt profile, but a gaussian fit suffices for extracting the FWHM.


Power and Polarization Fluctuations

We observed fluctuations in the loading rate of the 3D MOT some of the time but not all of the time (see fig. 5.10 for expected behaviour and fig. 5.11 for the fluctuations).

These had a similar timescale to fluctuations in the power ratio between the two arms of the 2D MOT (fig. 5.12), as well as fluctuations in the polarization of the light in each of the two arms (fig. 5.13). Note that the white and red arms in the




(a) An example loading curve for the MOT where the MOT has not saturated.




(b) An example saturated loading curve (where the MOT is filled to its maximum size, producing a fluorescence signal that is not linear with the atom number).

Figure 5.10: MOT loading curves without loading rate fluctuations.



(a)



(b)

Figure 5.11: Examples of fluctuating loading rates. The orange line is fit to the initial slope.




Figure 5.12: Fluctuations in the power in the two arms of the 2D MOT splitter prior to adding any polarization cleaning optics. Notice how the net power through the arms is roughly constant but the splitting is highly variable.

figure refer to the two output arms of the Thorlabs PN780R5A1 50:50 fiber splitter. The white arm connects to the original fiber while the red arm is the tap.

Fluctuations in the power ratio affect the loading rate by reducing the overall number of atoms in the 2D MOT. It also deforms the 2D MOT so atoms are less likely to pass through the differential pumping hole. Polarization drifts also affect the power in each arm. Inside of the 2D MOT optics, each arm of light passes through a quarter wave plate. Light that is not polarized perpendicular to the key of the fiber (fig. 5.14) is not transformed into circularly polarized light by the quarter wave plate, reducing the effective power in that arm of the 2D MOT.

We believed that both issues were caused due to polarization fluctuations in the fiber. FC/APC PM fiber connectors only have a polarization extinction ratio (PER) of 20 dB [51] so we expect that the light passing through multiple FC/APC fiber connectors compounded the polarization fluctuations. Furthermore, the fiber-coupled AOMs are rated for a PER of 12 dB to 16 dB, which further exacerbates the polarization fluctuations.

To test the theory, we coupled the light from the 2D MOT switch into free




Figure 5.13: Fluctuations in the polarization in the white arm of the 2D MOT splitter prior to adding any polarization cleaning optics. Behaviour in the red arm was similar.




Figure 5.14: An image of the Thorlabs patch cables used on the apparatus, taken from <https://www.thorlabs.com/newgrouppage9.cfm?objectgroup.id=3345>. Light can be polarized along the fast axis (perpendicular to the connector key) or slow axis (parallel to the connector key). We used FC/APC (Fiber Connector/Angled Physical Contact) connectors to reduce back-reflections.




Figure 5.15: Polarization fluctuations in the white arm after adding a free-space GT in front of the 2D MOT splitter.

space, through a GT polarizer, then into the 1:2 splitter for the 2D MOT. This reduced the polarization fluctuations by two orders of magnitude (see fig. 5.15 and fig. 5.16). However, this solution was not optimal. It was difficult to couple the light efficiently into the fiber splitter. We also wanted to minimize the number of free-space optics to minimize the necessary realignment after shipping the apparatus.

To address the issue more permanently, we took two measures: We added GT1 before the light is coupled into the fiber, and added an in-line polarizing fiber from OZ Optics (FOP-11-11-780-5/125-P-P-40-3A3A-3-1-ER=25) directly before the 2D MOT splitter. It acts like a fiber-coupled GT - it eliminates stray polarizations from the fiber leaving only light polarized perpendicular to the key with a rated PER of 25dB. We found that it kept the relative power fluctuations within 10% and the polarization fluctuations within 1% (see fig. 5.17). We believe the residual fluctuations may be caused the imperfect polarization extinction seen in FC/APC connectors.


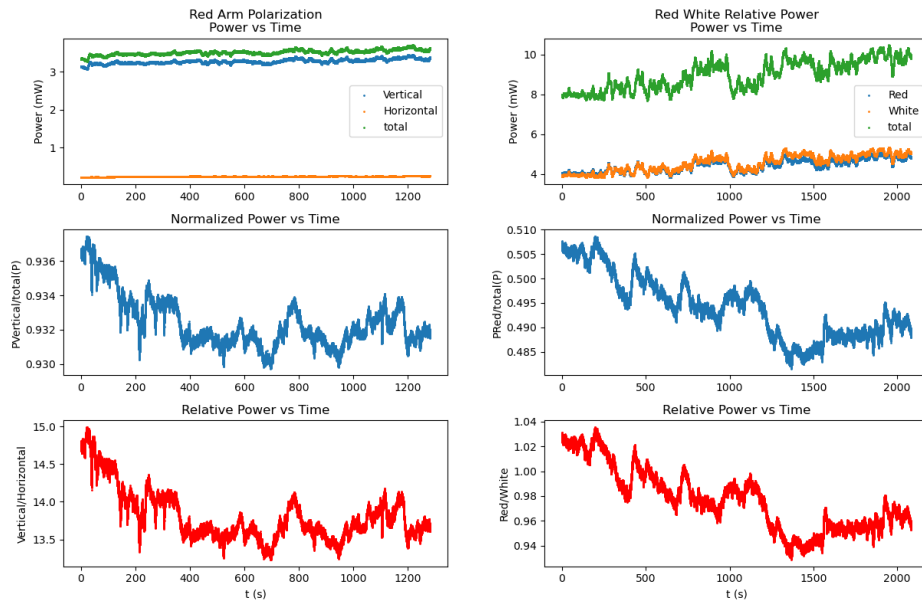


Figure 5.16: Power fluctuations between the red and white arms after adding a free-space GT in front of the 2D MOT splitter.



(a) Polarization fluctuations in the red arm after adding an in-line polarizing fiber.

(b) Power fluctuations between the red and white arms after adding an in-line polarizing fiber.

Figure 5.17: Polarization and power fluctuations after adding a inline polarizing fiber.

5.2 Vacuum Optics

Once light of the correct frequency has been generated, it needs to be directed to the MOTs.

5.2.1 2D MOT Optics

We purchased a prebuilt 3D printed optical assembly from ColdQuanta (part of the Photonically-Integrated Cold Atom Source [PICAS], see fig. 5.18). It takes three FC/APC fiber-optic inputs: one for each of the two 2D MOT arms, and one for the push beam. Each perpendicular arm of the 2D MOT light passes through a QWP then is reflected off of a series of beamsplitters to produce a large, rectangular beam of light through the cell. It is then retro-reflected through a QWP to produce the counterpropagating beam. The push light is collimated in a large ~ 1 cm beam that shines through the cell. This reduces the intensity of the push light, but means it requires no alignment.

The optical assembly clamps onto permanent magnet rods that produce the magnetic field for the 2D MOT (fig. 5.19). The rods can be translated along a 2D plane parallel to the axis of the source cell so that the center of the 2D MOT can be aligned with the differential pumping hole. We checked the alignment by centering a camera on the end of the source cell and moving the magnets until the differential pumping hole was blocked by the 2D MOT.

When the push beam shone directly through the center of our 3D MOT, it shifted the position of the MOT over by ~ 2 mm. This caused our MOT fluorescence signal to change artificially due to the 3D MOT no longer being at the focus of our imaging optics. To solve this issue, we shimmed the 2D magnet assembly using washers until the push beam shone along the top of the 3D MOT cell.

5.2.2 3D MOT Optics

The 3D MOT light is collimated using a Thorlabs C40APC-B into an 8.6 mm beam. It passes through a GT to remove any stray polarization drift that could affect the intensity stabilization, then through a HWP and PBS to split off some light (~ 1 mW) for intensity stabilization. It then passes through a shutter that provides additional optical isolation during magnetic trap holds. All of the above are covered by a




Figure 5.18: ColdQuanta PICAS components (images taken from Ref [8])




Figure 5.19: Magnetic field generated from the permanent magnets on the PICAS (image taken from user manual)

3D-printed cover that absorbs reflected light and prevents it from reaching the 3D MOT.

After the shutter, the light reflects off of a mirror (due to space constraints, the light must be reflected), then passes through two HWP-PBS cubes to split the light into three beams of equal amplitude. Each horizontal beam is reflected off of Polaris MA-45 45-degree mounting adapters vertically through a 1" hole in the horizontal optics table, then reflected horizontally by a second 45 degree mount




Figure 5.20: Layout of the 3D MOT optics. The light shield is transparent in the model but is opaque in real-life.

before passing through a QWP to produce circularly polarized light. The light passes through the 3D MOT cell before being retro-reflected through a QWP. The vertical light is reflected off of a 45 degree mount and passes through the center of the coils, and its retro-reflection mirror and QWP are mounted to the optical table.

Imaging

We use a 1" tube-mounted system to house all of the imaging optics in front of the PDA100A2 amplified photodetector. We use a 75 mm convex lens located 75 mm from the MOT (perpendicular to the glass cell) to collimate the light. A 50 mm lens focuses the light onto the photodetector. An iris is located directly in front of the photodetector blocks the reflections from the horizontal beams of light while allowing the light from the MOT through. The optimal position was determined by maximizing the ratio of light detected with the MOT on compared to the MOT being off.




Figure 5.21: Layout of the vacuum chamber, coils, and optics.

Chapter 6

Electronics

6.1 Pre-existing Electronics

Several devices had been assembled by Dr. James Booth and the MAKE+ group at BCIT prior to the author taking over this project.

6.1.1 QDGBus

Our lab uses a custom communication protocol to control our experiments called ‘QDGBus’. All devices are connected to a unidirectional 25-pin 1MHz parallel bus that is controlled by an FPGA. There are 16 data pins, 8 address pins, and 1 clock pin on the bus. It works as follows:

1. A python script compiles the commands into a series 64-bit bytecode instructions (8 bits for address, 16 bits of data, 8 bits of flags , and 32 bits specifying the time that the instruction should be sent.
2. The bytecode is sent over ethernet to a linux daemon on a discrete processor on the FPGA board
3. When the ‘start’ command is sent to the daemon, it starts to feed the bytecode instructions to the FPGA via two ping-pong FIFO buffers. Meanwhile, the FPGA starts a 1MHz counter.

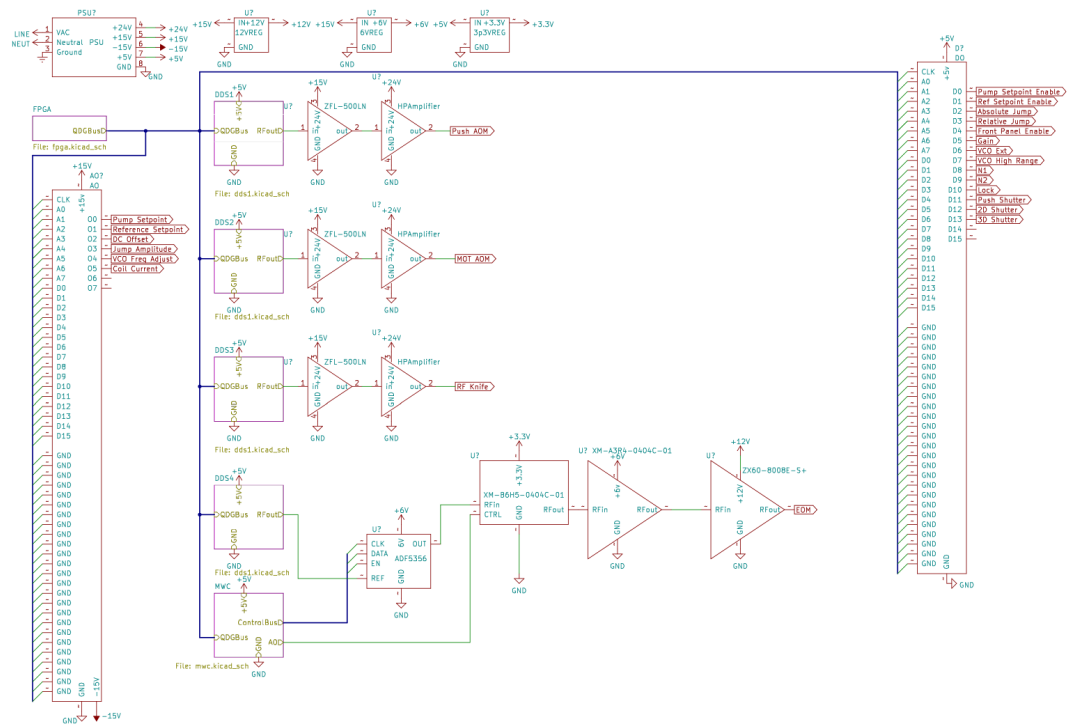


Figure 6.1: A schematic showing the wiring of all the control electronics connected to QDGBus

4. When the counter reaches the same value as the next timestamp in the instruction buffer, the FPGA sends the command over the parallel bus
5. If the address bits match the address of the device on the bus, the device stores the data bits and executes the corresponding command

6.1.2 QDGBus devices

Three types of devices that were compatible with QDGBus had been built by the BCIT MAKE+ lab under the Prometheus CFI project. They were based off of a set

of designs by Todd Meyrath for the Raizen group at University of Texas [31].

Digital Outputs

Digital output boards can set 16 BNC outputs that latch to either TTL high or TTL low, and could source over 100 mA of current. They have an address that could be set using an 8-position DIP switch. Every new QDGBus command sent to a Digital Output board assigns each output one bit of the data value, so the state of every pin on a device can change simultaneously.

The apparatus uses one digital output board, which is used to control the lasers and the shutters.

Analog Outputs

Analog output boards can set up to 8 BNC channels to a value between ± 10 V. Each channel can source up to 250 mA if properly cooled. Each output has its own address, so the address of the analog output board is set via a 5-pin DIP switch. The output is set by sending a 16-bit signed integer, scaled such that the minimum and maximum values produce the minimum and maximum voltage.

This apparatus uses one analog output board, which is used to control the lasers and the coil current.

Direct Digital Synthesizers

The direct digital synthesis (DDS) devices produce frequencies between DC and 135 MHz. They use all 16 bits of the data and two address bits to program the registers in an AD9852, which produces an output frequency that the DDS board filters to get the set frequency. The address is set with a 6-pin DIP switch because the DDS occupies four addresses in the address space.

The boards we use have a slight design flaw - the data values sent to the chip are latched but the address values are not. Occasionally, noise on the address lines can result in effective timing delays, which can cause the DDS to fall into an invalid state, requiring a reset.

Four DDSs are used on the apparatus - two to control AOMs, one to control the RF knife, and one to provide a 122.88 MHz reference signal for the microwave

source.

6.2 Microwave Electronics

To produce the repump light, we needed to power the EOM with a ~ 1 W signal of ~ 6.568 GHz. To be able to change the repump amplitude, we needed to be able to attenuate that signal. Finally, we needed to be able to control that signal using the QDGBus protocol.

The optimal solution would have been to produce a custom microwave board, but it can be challenging to produce electronics that work well at that frequency. Instead, we used an evaluation board (EV-ADF5356SD1Z) for a microwave frequency source (ADF5356) that could produce frequencies from 50 MHz to 6800 MHz. Unfortunately, the board could only be programmed via a GUI application that only ran on Windows. This meant that it would take seconds for the microwave frequency to change, rather than the milliseconds needed to effectively control the experiment.

We addressed this limitation two ways. Following the instructions in the evaluation board's user manual, we removed R12 to disconnect the internal reference from the power rails, and R27 to decouple the differential reference inputs [10]. We terminated Reference A with a 50 Ohm terminating cap and connected a DDS to Reference B. This allowed us to change the output frequency by 1.5% in either direction within $10\ \mu\text{s}$ (the programming time of the DDS). When operating at 6.57 GHz, this allowed us to tune the frequency of by up to 100 MHz in either direction (fig. 6.2).

6.2.1 Microwave Control Board

The DDS tuning approach only allowed us a small range of programmable frequencies, and still required the board to be initially programmed via the GUI application on startup. To solve the problem, we built a custom board that interfaced with QDGBus and hijacked the evaluation board's communication lines via the test points. It uses some digital logic to interface with the address, a bank of FIFOs to buffer the data, and a microcontroller to process the commands and program the ADF5356.




Figure 6.2: The DDS could only offer a limited tuning range before the microwave source railed out

Input Address Comparison

First, we needed to receive commands from QDGBus and power the board. We used an 8-pack DIP switch to set the address of the device and compare it to the input address. If it matched, it passed the signal on to the next stage.

Data Buffering

The Teensy 3.2 takes roughly $400\ \mu\text{s}$ to service an interrupt service request (ISR), and $40\ \text{ns}$ to read a set of digital inputs. Reading 16 pins required reading from three separate digital input registers. This meant that it could not reliably pull all of the data from the data lines during the $500\ \text{ns}$ read window.

To solve this problem and to reduce the complexity of the control code, the incoming data is fed into a bank of 16-deep, 4-bit-word FIFO buffers (it was easier to use four buffers than to get a 16-bit word FIFO due to component shortages).

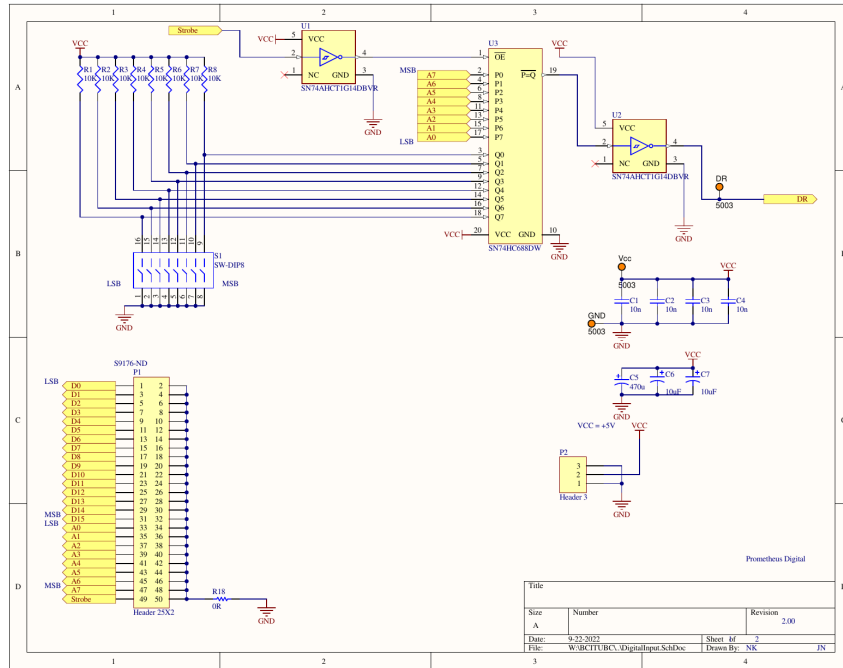


Figure 6.3: Schematic for the input stage of the microwave controller.

The microcontroller can pull data from the buffer at its leisure, and does not need an ISR to store incoming instructions while past instructions are being executed.

Microcontroller and Output

The microcontroller receives the incoming data, and programs the microwave chip. It uses a voltage divider to bring the 3.3V logic of the microcontroller down to the 1.8V logic of the ADF5356. We soldered wires to the eval board and to the PCB, then used a three-pin crimp connector to connect them (CLK to TP3, DATA to TP4, and EN to TP5). Ground was not connected; Digital and Analog ground on the ADF5356 are shared, so the two boards already share a common digital ground from their power supplies.

The microcontroller's 12-bit analog output is sent to a BNC, which is connected to a microwave voltage-variable attenuator (VVA). This is for two reasons: It allows for easy tunability of the repump power, and it allows for compensation




Figure 6.6: An example of the ADF5356’s ‘exciting’ transfer function




Figure 6.7: The transfer function for the XM-A3B4-0404C-01 microwave voltage-variable attenuator.

$\text{INT} \in \text{Integers} : 75 \leq \text{INT} \leq 65535,$
 $\text{FRAC1} \in \text{Integers} : 0 \leq \text{FRAC1} \leq 16777215,$
 $\text{FRAC2} \in \text{Integers} : 0 \leq \text{FRAC2} \leq 268435455,$
 $\text{MOD1} = 16777216,$
 and $\text{MOD2} \in \text{Integers} : 2 \leq \text{MOD2} \leq 268435455$

The microcontroller calculates their values prior to programming the appropriate registers. See the ADF5356 datasheet for further information [9].

Control Sequence

The microcontroller could be controlled one of two ways: First, it could be sent commands via a 115200 baud serial connection. This interface allowed every register on the ADF5356 to be programmed. It also accepted high-level commands such as ‘set frequency’ or ‘set output amplitude’. This interface was designed for debugging the control code and testing to make sure that the QDGBus interface was working correctly.

Second, it could receive high-level commands via QDGBus. It received a 16-bit integer (n_{in}) linearly mapping to the desired output frequency:

$$f_{\text{set}} = \frac{6.8 \text{ GHz}}{2^{16} - 1} n_{\text{in}} \quad (6.1)$$

However, as the minimum output frequency of the ADF5356 was 50 MHz, it left numbers 0-127 of the 16-bit input integer without a purpose. In other words, if the nine most significant bits were set to zero, the seven least significant bits could be used to send further configuration information to the microcontroller. If bit 7 was set, it allowed the lower 6 bits to control the output amplitude (after linearizing the output power using the lookup table). The remaining bits controlled the output power of the ADF5356, whether or not the output was enabled, whether the ADF5356 should be powered on or off, and whether the microcontroller should reset its state.

6.2.2 Microwave Performance

Response Time

- It takes roughly 8 ms from the start of programming for the new frequency for it to turn back on and another 2 ms for it to settle.
- The analog output fed into the voltage-variable microwave attenuator allowed the output amplitude to be set anywhere in a 30 dB range within tens of microseconds.
- Enabling or disabling the outputs takes several milliseconds, and reduces the amplitude of its output by ~ 30 dB. Turning off the ADF5356 by programming its “power off” register takes a similar amount of time and completely disables the output signal. Turning on the ADF5356 by programming its “power off” register returns the signal to its previous frequency.

Precision and Linewidth

Although the ADF5356 has in principle a resolution of Hz, we measured the FWHM of its amplified output signal as ~ 40 kHz. Furthermore, it imposes sidebands on the pump light, which has a linewidth in the order of 1 MHz. As a result, most of the precision with which the output frequency can be set is not useful.

When programming the frequency via QDGBus, the frequency resolution is $\frac{6.8\text{GHz}}{2^{16}-1} = 104\text{kHz}$. As this is still less than the linewidth of the pump light, it is an acceptable tradeoff for the simplified communication protocol.




Figure 6.8: Spectrum of the amplified microwave source signal before it is fed into the EOM.

6.3 Intensity Stabilization

We use the fluorescence from trapped rubidium atoms in the MOT as a proxy for the number of atoms in the MOT. The fluorescence is proportional to the atom number (provided the MOT is not too full), but is also dependant on the intensity of the light in the MOT. Thus, it is critical to have a constant intensity of light entering the 3D MOT so that the atom fluorescence linearly corresponds to the atom number. We found that the amount of light entering the 3D MOT fluctuated by $\sim 15\%$ over timescales of tens to hundreds of seconds (see fig. 6.9). These fluctuations directly affect our measurements, so we needed to stabilize them.

We could monitor the amplitude of the light by directing some of the light reflected from VBS4 (see fig. 5.2) and to amplified photodetector PD1 (a Thorlabs PDA100A2). We decided the best way to stabilize it would be to control the voltage




Figure 6.9: Long timescale fluctuations in the unstabilized 3D MOT power.

of a voltage-controlled RF attenuator on the output of the MOT AOM DDS. We purchased an XM-B6H4-0404C-01 to use as a VVA. As other groups have had success using microcontrollers as low-cost feedback controllers [11, 23, 34], we purchased a Teensy 3.2 to implement a modified PID feedback control algorithm.

6.3.1 Control algorithm

Nominally, the microcontroller runs its control loop at 50 kHz, taking 4 averages of a 12-bit ADC read of the photodetector voltage and controlling a 12-bit output voltage. Its output is connected directly to the VVA, while the photodiode's output is divided from a 5V maximum down to a 3.3V maximum via a voltage divider to protect the microcontroller's ADC input. A digital input allows it to switch between two setpoints which can be programmed (along with the PID constants, read averaging, set point bounds, and output bounds) via a 9600 baud serial connection.

In normal operation, the microcontroller's output is nominally governed by a discretized version of eq. (6.2).

$$y(t) = K_P e(t) + K_I \int e(t) dt + K_D \frac{de(t)}{dt} + F(s(t))e(t) = s(t) - x(t) \quad (6.2)$$

where $y(t)$ is the microcontroller's output, $x(t)$ is the microcontroller's input, $s(t)$ is the set voltage, K_P , K_I , and K_D are the standard PID feedback constants, and $F(s(t))$ is the feedforward function.

The PID control is implemented using the “FastPID” arduino library. It uses 32-bit bitshifted integers to perform the operations, which are significantly faster than the floating-point numbers used in the default library. Its output is added to the output generated by the feedforward function.

The feedforward function is stored in a lookup table. At startup, it loads a lookup table from EEPROM. The user can also request the microcontroller regenerate the lookup table via the serial connection. It generates it by ramping through 256 output voltages twice and recording the average input voltage at each output. It then uses linear interpolation to determine what output voltage should correspond to what input voltage for all 2^{12} possible setpoints, unless the setpoint is out of range. In that case, it sets the output to the corresponding maximum/minimum output value. It stores each value in a lookup table such that the necessary output voltage is at the index of the desired input voltage.

If the setpoint is set to an out-of-range value, the output is set to the corresponding maximum/minimum value, and the PID loop is disabled. As soon as the setpoint returns to an in-bounds value (either by reprogramming it serially or via the digital input changing state), the feedback loop runs once using feedforward control, then operates normally. This minimizes the response time for the controller as it does not wait to collect the feedback input or to run the PID calculation.

A digital input was selected to control the output state for several reasons:

- Our current experimental sequences do not require on-the-fly adjustment of the light intensity beyond ‘on’ and ‘off’, so an analog input is not necessary.
- The feedback loop can run more quickly because only one analog measurement is required. The feedback control calculation takes $\sim 8 \mu s$, and each

analog measurement takes $\sim 2 \mu\text{s}$. We found that at least four averages were needed for good operation, so reducing one set of reads sped up the control loop by a factor of a third.

- The feedback loop removes a source of noise by having the setpoint remain constant.

The control code also includes preprocessor directives to enable or disable

- feedforward control.
- negative control (used if a positive change in the output decreases the feedback signal).
- analog control of the setpoint.
- the serial port (to reduce noise and increase performance).

should we discover that we need to change the behaviour of the feedback controller in future.

6.3.2 Performance


The data in this section were recorded with $K_I = 10000$ and $K_P = .01$. Two separate outputs were logged (both the feedback light split off VBS4 [reference] and the transmitted light to the rest of the MOT [throughput]) using PDA102 amplified photodetectors.

Step Response


The feedback controller responded to an input signal within $10 \mu\text{s}$ (not shown). The input settled to its steady-state value within $240 \mu\text{s}$ (see fig. 6.10)—better if its calibration data were more recent. It could switch off the light within $2 \mu\text{s}$. As the shutters take 10 ms to open/close, this performance is more than adequate.

Steady-State Performance

Over short timescales, the feedback controller increased the amount of noise by a factor of two: the feedback controller output had a standard deviation of 0.30% of




(a) Step response of intensity feedback controller




(b) Zoomed-in step response of intensity feedback controller

Figure 6.10: Step response of intensity feedback controller




(a) Open-loop performance




(b) Closed-loop performance

Figure 6.11: Open loop performance compared to closed loop performance over short timescales.



(a) Open-loop performance



(b) Closed-loop performance

Figure 6.12: Open loop performance compared to closed loop performance over long timescales.

the fullscale value while the open loop control had a standard deviation 0.16% the fullscale value (see fig. 6.11). However, over long timescales, the trend reversed; the feedback controller had a standard deviation ratio of .29% while open loop control had a ratio of .77%, with excursions of up to 2.5%.

This behaviour is expected - the analog input has some noise, which introduces some noise to the output of the feedback controller. However, each data point in a measurement involves recording tens of thousands of points that can be easily averaged away. Therefore, the feedback controller's response to low-frequency (<1 kHz, large-amplitude (greater than 1% of the full-scale amplitude) disturbances is much more important than the amount high-frequency noise it generates.

Disturbance Rejection

To test the microcontroller's ability to reject large disturbances, we introduced a 520 Hz triangle wave to the laser current. The intensity controller completely suppressed it (see fig. 6.13).

The triangle wave caused a standard deviation of 0.76% of the fullscale amplitude of the signal, with excursions of over 3% of the fullscale amplitude. The intensity controller reduced the standard deviation ratio to 0.33% - very similar to its disturbance-free performance.

6.3.3 Conclusions

The intensity controller provided more than adequate stability, disturbance rejection, and response time for the purposes of this experiment. However, its performance could be further increased by

- Using a Teensy 3.6 microcontroller with double the processing power and 16 bit ADCs and DACs.
- Reducing expense by using a Minicircuit ZX73-2500-S+ (which costs \$50 instead of the \$330 XM-B6H4-0404C-01).
- Building a PCB with an active filter to reduce the DAC noise and increase output stability.
- Pipelining the ADC reads so they are performed in parallel with the feedback calculations.
- Automating the PID tuning to further increase output stability [33].



(a) Open-loop control demonstrating the added triangle wave disturbance.



(b) Closed-loop control of the disturbance.

Figure 6.13: Open loop performance compared to closed loop performance when a triangle wave disturbance was added.

Chapter 7

Software and Control

Building hardware is useless if it cannot be controlled. This chapter describes the software used to control the hardware. The software used to control the experiment is stored here: <https://qdg-code.phas.ubc.ca:2633/Perrin/PrometheusPython>, and the underlying driver that controls the QDGBus is stored here: <https://qdg-code.phas.ubc.ca:2633/QDG/QDGBusV1>.

7.1 Software structure

Different parts of the apparatus need to work together to produce a cohesive result. For example, detuning the pump frequency requires interfacing with an analog output or a DDS and seven digital outputs. Keeping the repump frequency fixed while detuning the pump frequency requires all of the above in addition to controlling a microwave source and its DDS. Additionally, many components (such as the AOMs and Offset Lock) require feedforward compensation. Including all of this information in a single file would clutter it and be challenging to keep track of. Consequently, the control of the apparatus is divided into several layers. At the base layer, there are the device drivers for the raw digital outputs, analog outputs, DDSs, and microwave generators. The next layer models the physical devices those outputs control (coil drivers, AOMs, EOMs, shutters, laser controllers, offset lock box, RF Knife). Finally, a top layer ‘Prometheus’ object calculates the necessary outputs for the physical outputs to produce the desired result.

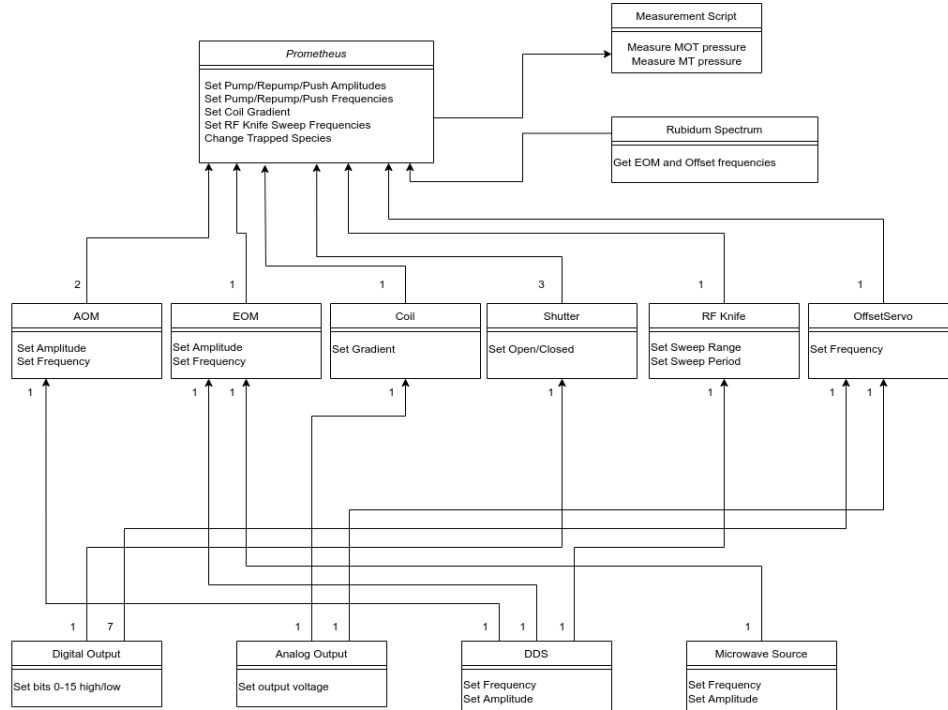


Figure 7.1: A block diagram showing (at a high level) the basic components needed for the apparatus to operate. Methods have been renamed and helper methods have been removed for clarity. Other modules will be added in the future (such as a trigger for the DAQ, auto-locking by operating the reference laser servo and two laser controllers, and serial control of the TA and feedback controllers).

Specific measurement sequences control the Prometheus object without needing to interface with the devices underneath, simplifying the control process and reducing the risk of error in script development. They also control the USB DAQ (a MCCDAQ USB-1608G) and save the data to a file.

All control software is written in Python, and the devices are controlled using the QDGBus protocol described in section 6.1.1.

7.1.1 AOM Throughput Optimization

The fiber-coupled AOMs are rated for a maximum RF power of 1.5W, and have a frequency-dependant diffraction efficiency. The DDS and amplifier chain used on

the apparatus can produce up 1.5 W to 3 W depending on the output frequency and temperature of the amplifiers. To protect the AOMs, we implemented a hard upper limit to the DDS amplitude of any DDS connected to an AOM.

We also mapped out the insertion loss of each AOM for a variety of DDS amplitudes and frequencies and fitted the surface to a 2D spline fit (see fig. 7.2 for an example). This allowed the software to fit and select the optimal DDS amplitude for any set frequency, as well as attenuate the DDS amplitude. The default mode in the software is to treat an amplitude of 1 as requesting the maximum throughput for the set frequency ('relative' mode). However, the software can also be set to 'absolute' mode, where setting the amplitude to the same number maintains the same throughput regardless of frequency.




Figure 7.2: A grid of AOM throughput efficiency at different DDS amplitudes and frequencies

7.2 Measurement Sequences

7.2.1 MOT Pressure Measurement

We can quickly obtain a rough estimate of the pressure in the chamber by recording the loss rate of atoms from a MOT. $\langle \sigma_{\text{loss}} v \rangle \approx 1.3 \times 10^{-15} \text{ m}^3/\text{s}$ for H_2 colliding




Figure 7.3: Timing of different amplitudes and frequencies in order to measure pressure using a MOT

with Rb in a trap depth of 1 K [43], so

$$P \approx (2.3 \times 10^{-8} \text{ Torr s}) \Gamma \quad (7.1)$$

where Γ is the loss rate of atoms from a MOT.

This technique provides a relatively quick measure of the pressure in the chamber to ensure that no major problems occur with the more precise but longer measurements using magnetic traps

Sequence

1. We record a baseline level of fluorescence with the 3D light on to provide a reference level to subtract from all subsequent values. The pump detuning is set to on resonance to prevent accidental trapping of residual rubidium.
2. We load the MOT by turning on the 2D MOT and push beam, and red-

detuning the pump. We found it took ~ 10 s for the MOT to fully load, so we only loaded for 5 s to avoid saturation (which would cause the photodiode signal to become nonlinear in the atom number at high atom numbers). We also remove points where the nonlinearity is present from our curve fits. See Ref [25] for a technique that can be used to verify the photodiode linearity with atom number.

3. We stop loading the MOT by turning off the 2D MOT and push beam. We allow the background collisions to eject atoms from the MOT.
4. We turn off all light to allow Rb to leave the MOT and stick to the glass cell walls.

7.2.2 MAT Pressure Measurement

One of the advantages of the optical system in this apparatus is its flexibility in the frequencies it can produce—it can switch between being able to trap rubidium-85 and rubidium-87 without needing to relock its lasers. While the vacuum chamber baked out, we took advantage of that functionality to connect the Prometheus locking optics and electronics into a pre-existing experiment (MAT) in our lab. Data collection is ongoing, but we have been able to successfully rapidly switch between trapping the two different isotopes of rubidium without human intervention.

This experiment loads its MOT directly from dispensers without using a 2D+ MOT and differential pumping hole, so the 2D MOT optics and push beam were not used. We used the following measurement sequence shown in fig. 7.4

Sequence

1. We take the same baseline measurements as with the MOT.
2. We load the MOT as before
3. We red-detune the pump light, allowing the MOT to expand and cool.
4. We transfer the atoms to the magnetic trap by extinguishing the repump light, transferring ^{87}Rb atoms to the $|F = 1, m_F = -1\rangle$ state and ^{85}Rb atoms to the $|F = 2, m_F = -1, -2\rangle$ states.




Figure 7.4: Timing of different amplitudes and frequencies in order to measure pressure using Prometheus optics and electronics with a different experiment in the lab.

5. This step is only taken when taking measurements with ^{85}Rb . We reduce the magnetic field gradient so that the confinement felt by the $|F = 2, m_F = -1\rangle$ atoms is not strong enough to overcome the pull of gravity, and they fall out of the trap.
6. We repeatedly ramp the frequency of the RF Knife to remove high-energy atoms from the trap, reducing the effect of heating within the trap
7. We hold the atoms in the trap, allowing collisions with the background to eject the atoms
8. We ramp the frequency of the RF Knife again to set the trap depth of the system
9. We reload the trap to see how many atoms remained in the magnetic trap after the hold time.

10. We eject atoms from the trap by shining resonant light.
11. We repeat the baseline measurement to ensure it is the same as before.

While a single decay curve is enough to determine the pressure using a MOT, every MT measurement sequence only produces one point on the decay curve, so the sequence needs to be repeated a multiple times at different hold times.

To extract the number of atoms in the magnetic trap after the hold, we can fit

$$V(t) = At + B \quad (7.2)$$

for a short time after turning on the MOT and use the intercept B to extract the atom number.

7.2.3 MT Pressure Measurement

We plan to use a combination of the two above sequences to measure pressure using a magnetic trap in Prometheus:

We took some preliminary magnetic trap measurements using a modified version of the above sequence with no RF Knife and a fixed magnetic field gradient. These data are shown in section 9.2.3.




Figure 7.5: Timing of different amplitudes and frequencies in order to measure pressure using a magnetic trap.

Chapter 8

Assembly

8.1 Carts

Although the subsystems have been tested, the apparatus has not yet been assembled onto a cart and fully tested.

The original plan was to assemble everything on a 19"x30" server rack enclosed within a 23.5"x37"x58" frame on 4" casters. However, this design of the system used air-cooled coils, which offered inadequate thermal performance. This meant that the rack does not have enough space to house water chiller. A new 26"x48"x36.5" cart on 6" casters was built to house all of the components, including the water chiller.

Currently, the locking optics and control electronics are on the rack, while the water cooling, coils and coil power supply, vacuum optics, and vacuum chamber are mounted on the cart. One of the remaining tasks on the project is to integrate the entire system onto the single cart by removing the section of the rack that the electronics and optics are mounted to and installing it and the boxes in the rack (as shown in fig. 8.1)

The vacuum chamber and vacuum optics will be bolted to the top of the rack. The control electronics will be transferred to a sliding tray (for ease of troubleshooting) and will be located on the bottom of the rack. The locking optics will sit on a sliding tray above the control electronics. Both will remain attached to a section of the 19"x30" server rack used in the previous cart.




Figure 8.1: A mockup of where all of the components in Prometheus can fit on the new cart. The lower two shelves house the coil controller and the hydraulics (not depicted). The acrylic and metal panels in the vacuum optics cover have been made transparent to show the internals - the panels are actually opaque.

A small shelf will be placed on top of the Vescent lock boxes and bolted to the shelf below it, holding the vescent boxes in place. The ion pump controller, dispenser controller, TA power supply, shutter power supplies, and photodiode power supplies will be zip-tied to the small shelf.

The gap between the rack and the side of the cart will be used for routing electrical cables, as well as the coil power cables and water cooling tubing.

8.2 Light Shielding

It is important to cover the optics for two reasons:

- Protecting the optics from being damaged: Damage can occur from stray objects scratching the optical surfaces, but also from dust accumulating then burning.
- Laser safety (protecting the humans from being damaged): The high-power near-infrared (NIR) light we use extremely dangerous because the human eye is not very sensitive to it, so it can cause retinal damage without triggering a person's blink reflex.

8.2.1 Vacuum Optics

We purchased 1" slotted T-rods and acrylic sheets to form an enclosure around the vacuum chamber and optics. The acrylic sheets were laser-cut to fit inside of the slotted rods and around the mounts. This allows the panels to be easily removed to realign the locking optics. Because the acrylic is reflective, we will cover the interior surface with nonreflective photographer's tape.

The panel that the inline valve passes through will be cut from 1/8" aluminum using a waterjet. It has a circular cutout for the inline valve and a horizontal slot that passes through the center of the cutout. This will allow the top part of the panel to be removed without moving the tee slotted frame. This panel is metal so that the optics can remain enclosed during a bakeout of the interconnect between the apparatus and the chamber under test.

8.2.2 Locking Optics

The currently layout (fig. 8.2) of the locking optics is not ideal for light shielding. The original plan had been to cover the entire tray with a shield, but this would enclose the TA without allowing for air circulation. The fiber optics are also very messy - more space will allow us to mount them more cleanly. We plan to cover the free space optics in a small opaque acrylic box that is screwed onto the optical table using L brackets. We will then place the fiber optics on a single acrylic sheet above the free space optics box. The TA will be moved to the back of the optical table where it can cool more efficiently. The proposed new design is shown in fig. 8.3, and will still fit under the 6'' of height allowed for the locking optics in the new cart.

8.3 Cart Walls

The cart does not currently have walls to protect the user from the high-power electronics inside. We will cut walls out of perforated aluminum sheets from McMaster Carr (<https://www.mcmaster.com/9232T341/>) and bolt them to the sides of the cart. The perforated aluminum will provide enough airflow to keep the electronics cool while providing enough strength to protect the interior.




Figure 8.2: Current layout of the locking optics. The fiber optics are spread through two levels - the lower level houses the pre-TA optics and the upper level houses the post-TA optics.




Figure 8.3: Proposed layout of the locking optics. Support structures have been suppressed and acrylic has been made transparent to allow a clearer view of the optics within.

Part III

Results

Chapter 9

Results

Although we have not completed testing of the fully-assembled apparatus, we do have some preliminary results from tests made with the old vacuum chamber. The deviations from the final design are as follows:

- We used the leaky vacuum chamber shown in fig. 3.1 rather than the new chamber with a higher conductance.
- The coil driver and coils had also not been completed at that point, so a benchtop power supply powered old air-cooled coils that produced the magnetic field with a gradient of 11 G/cm when driven at 1 A (see fig. 9.1).
- Intensity stabilization of the 3D MOT light was not installed.
- The RF knife had not been installed so magnetic traps were cell wall limited.
- We had not yet installed the polarizing fiber that should improve the load rate stability.
- The cover for the vacuum optics had not been completed so a cardboard box was used to block external light.
- The vacuum chamber and optics were on a table rather than on the rack.




Figure 9.1: Magnetic field vs axial position of the temporary coils when driven at 1 A

9.1 Analysis

9.1.1 MOT Analysis

The number of atoms in a MOT follows eq. (9.1) [43]:

$$\dot{N}(t) = R - \Gamma N(t) - \int (\beta + \beta_C) n(\vec{r}, t) dV \approx R - \Gamma N(t) - \beta N(t)^2 \quad (9.1)$$

where R is the loading rate of the trap, $b = \beta + \beta_C$ is the two-body collision rate, Γ is the loss rate of atoms from the trap, $n(\vec{r}, t)$ is the density of atoms in the trap, and $N(t)$ is the number of atoms in the trap.

Note that when holding atoms in a MOT, $R = 0$. Solving eq. (9.1) produces

$$N(t) = C + \begin{cases} 0 & t < t_s \\ \frac{2R(1-e^{-\Gamma(t-t_s)})}{\Gamma + \gamma + (\gamma - \Gamma)e^{-\gamma(t-t_s)}} & t_s \leq t < t_l \\ \frac{(n(t_l) - C)(1 - b\epsilon)}{e^{-\Gamma(t-t_l-t_s)} - b\epsilon} & t \geq t_l \end{cases} \quad (9.2)$$

where $\varepsilon = \frac{n(t_{load})}{\Gamma + (n(t_{load}) - C)b}$, $\gamma = \sqrt{\Gamma^2 + 4\beta R}$, C is an offset caused by background light and reflections of the MOT light, t_s is when the MOT starts loading, and t_l is when the MOT stops loading.

We assume that $V(t)$ (the photodiode voltage) is proportional to the atom number $N(t)$. However, for very large values of $N(t)$, this is no longer true - the outer atoms in the MOT deplete the trapping light, preventing the inner atoms from being fully illuminated [25]. We made a crude estimate of this critical photodiode voltage by manually identifying the voltage at which the loading and decay curves deviated from the expected form, and multiplying it by 0.8 to ensure that no nonlinear effects were identified. This rendered a critical photodiode voltage of ~ 0.5 V above which we ignored data points. Once we complete reassembly of the apparatus, we plan to use the alternate output from the push beam shutter to as a probe beam to implement the technique described in ref [25]. This will allow us to more precisely determine this critical value.

For large MOTs, we can fit the first 10% of the loading curve to a linear slope to extract a value r proportional loading rate R of the MOT:

$$V(t) = r(t - t_s) \quad (9.3)$$

where t is the time since the 2D MOT was turned on.

However, if the load rate too small ($V < 0.05$ V), it can be difficult to make this linear fit. In this case, we can add R to the fit of eq. (9.2).

We also tried fitting to a simplified version of eq. (9.2) by removing the two-body collision term (eq. (9.4)). However, it did not accurately describe the MOT signal for large MOTs.

$$N(t) = C + \begin{cases} 0 & t < t_s \\ \frac{R(1 - e^{-\Gamma(t-t_s)})}{\Gamma} & t_s \leq t < t_l \\ \frac{R(1 - e^{-\Gamma(t-t_l)})}{\Gamma} & t \geq t_l \end{cases} \quad (9.4)$$




Figure 9.2: Load rate (in units of V/s) for different push beam parameters (push power is in μW and push detuning is in MHz from resonance with the pump transition). The MOT loads for 10 s, the pump light detuning is 12 MHz, and the load rate is in units of “normalized photodiode voltage” as discussed in eq. (9.6).

9.2 Measurements

9.2.1 Optimal Push Beam Detuning

The push beam’s frequency can be set independently from the pump light’s frequency, but moving its frequency away from the pump light’s detuning reduces its maximum power because it involves running the push AOM away from its optimal frequency. However, we knew that the push beam works optimally for small blue detunings. We tried loading the MOT at a variety of different push beam powers and detunings to determine what produced the best load rate. Because the MOT load rate also depends on the pump detuning, we repeated the scan for three different pump detunings.

The steady-state photon scattering rate of an atom in a rubidium-87 MOT is




Figure 9.3: Load rate (in units of V/s) for different push beam parameters (push power is in μW and push detuning is in MHz from resonance with the pump transition). The MOT loads for 10 s, the pump light detuning is 15 MHz, and the load rate is in units of “normalized photodiode voltage” as discussed in eq. (9.6).

[48]

$$R_{sc}(\Delta) = \frac{\gamma}{2} \frac{\frac{I}{I_{sat}}}{1 + 4 \left(\frac{\Delta}{\gamma} \right)^2 + \frac{I}{I_{sat}}} \quad (9.5)$$

where γ is the natural linewidth of the transition (6.06 MHz), Δ is the detuning of the light from the pump transition, I_{sat} is the saturation intensity of the pump transition (this is dependant on the polarization of the light as well as the alignment of the atom but it is at least 3.58 MHz), and I is the total intensity of the light going into the MOT (we estimated it was $\sim 360 \text{ mW/cm}^2$). Therefore, we normalized the fluorescence signal from the MOT by

$$R_{sc}(0)/R_{sc}(\Delta) \quad (9.6)$$

so that the atom number would be linear in the fitted signal.




Figure 9.4: Load rate (in units of V/s) for different push beam parameters (push power is in μW and push detuning is in MHz from resonance with the pump transition). The MOT loads for 10 s, the pump light detuning is 18 MHz, and the load rate is in units of “normalized photodiode voltage” as discussed in eq. (9.6).


We determined that the best loading rate could be achieved by using a maximum power push beam of the same detuning as the pump light, and that the best pump light detuning to use was -12 MHz.

9.2.2 Dispenser Performance


Load Rate for different dispenser currents

Using the above parameters, we monitored the loading rate from the dispenser at different dispenser currents. We also monitored the decay rate from the MOT as a function of dispenser current to see if changing the dispenser current affected the pressure in the measurement chamber. We conducted tests at the optimal detunings discovered in the prior section.

The loading rate clearly peaks at 3.3 A and drops off steeply afterwards. We



(a) Load rates with the push beam turned on.




(b) Load rates with the push beam turned off.

Figure 9.5: Comparison of load rates (in units of V/s) across different dispenser currents with the push beam on compared to off.


hypothesize that this is due to background collisions reducing the 2D MOT size, reducing the number of atoms than can be pushed through the differential pumping hole.

We also saw a small increase in the fitted decay rate near the optimal loading conditions, however manual inspection of the decay curves revealed it was due to a tendency for fitting errors to occur at higher MOT fluorescence (see fig. 9.7b for an example). If the poorly fitted curves are discarded, actual decay rate remained at $0.65(2)$ MHz during the entire test. This tendency for fitting errors along with the nonlinearity due to MOT saturation illustrates the need to load MOTs at smaller amplitudes.

A second reason to load the MOT to smaller amplitudes is the risk of affecting the chamber's pressure. ColdQuanta promises atomic fluxes of up to 1×10^{10} atoms/s if the 2D+ MOT is operated at full power. If all rubidium atoms make it to the chamber and none are captured by the MOT or stick to the test chamber walls, this could increase the pressure in the test chamber by $>4 \times 10^{-8}$ Torr/s. In the real world, most atoms will be captured by the MOT or will stick to the chamber walls, but once we can perform tests at lower base pressures, this is an important effect to characterize. It may be necessary to hold the MOT prior to conducting a measurement at low pressures to allow the excess Rubidium to dissipate.




(a)




(b)

Figure 9.6: Comparison of loss rates (in units of Hz) at different dispenser currents. Re-analyzing the data revealed that fitting errors caused the apparent increase in the decay rate near the maximum load rate. Efforts to improve the loss rate fitting are ongoing.



(a)



(b)

Figure 9.7: Comparison of well-fitted and poorly-fitted MOT decay curves. Points that are in red are ignored due to MOT saturation while points in blue were used to generate the fit. Curves such as the one on the right were manually rejected when re-analyzing the MOT decay curve data.

Rubidium lifetime after being dispensed

In case we discovered that the source leaked contaminants into the chamber once activated, we examined the possibility of loading the 2D and 3D MOTs from the residual rubidium vapour that remained in the 2D MOT section after the dispensers were turned off. To study this, we recorded the load rate for a period of time after activating the dispenser (we ran it at 5 A for 10 s then 3 A for 60 s before turning it off). Whether we took measurements continuously (loading 50% of the time, as seen in fig. 9.8) or spaced five minutes apart (loading 1.7% of the time, as seen in section 9.2.2), we found we could load enough atoms to get a reasonable signal for about fifteen minutes before the signal amplitude grew too small.

Reloading After Releasing

We investigated whether the residual rubidium vapour might collect and remain in the 3D MOT region after loading the 3D MOT. To study this, we would load a 3D MOT from the 2D MOT using the parameters that maximized our loading rate, hold the atoms for five seconds, then release the atoms by blue-detuning the light. After one second, we would then try to reload the 3D MOT, this time without activating the 2D MOT. Fluorescence in the 3D MOT would indicate that residual rubidium vapour remained in the 3D MOT region. However, we observed no change in the scattered light from the 3D MOT region.

This means that

- The released Rb atoms quickly dissipate, likely sticking to the chamber walls.
- The source cell does not appear to leak rubidium at a noticeable rate into the main chamber.

This means that the apparatus is unlikely to contaminate the chamber under test with a large influx of rubidium atoms.




Figure 9.8: Loading rate (V/s) and fitted MOT loss rate (Hz) over time with the dispenser turned off and measurements being taken continuously. It is clear that after 10-20 minutes, the fitted loss rates cannot be trusted.




Figure 9.9: Loading rate (V/s) and MOT loss rate (Hz) over time with the dispenser turned off and measurements being taken every five minutes.




Figure 9.10: Comparison of loading amplitudes (as a proxy for loading rate) over time when measuring continuously as compared to measuring infrequently. It illustrates how taking frequent measurements does not significantly deplete the supply of free rubidium in the source cell.

9.2.3 Magnetic Trapping




Figure 9.11: A sample curve from the magnetic trap measurement

We tried taking a preliminary set of magnetic trap measurements with the magnetic field gradient set to 11 G/cm. A sample trace from the measurement sequence is shown in fig. 9.11. The exponential decay in the magnetic trap recapture fraction is shown in fig. 9.12. The non-exponential behaviour high recapture fraction at small hold times (hold times less than 10 ms) is from the atoms not having enough time to leave the magnetic trap before the MOT is turned on again.




Figure 9.12: A preliminary decay curve generated from a series of magnetic trap measurements. The fitting code is not fully functional so the fit parameters were estimated by hand.

Chapter 10

Conclusions

This thesis has demonstrated significant progress towards constructing a portable version of a quantum pressure standard. We have a vacuum chamber that can dispense rubidium in a manner that minimizes contamination of the chamber under test. It takes up roughly 400 mL of volume¹ when conducting measurements, and provides reasonable conductance (1.3 L/s) to the chamber under test². It has a base pressure of at least 5×10^{-10} Torr, which should be sufficient for conducting measurements at the high range of UHV, and may be able to reach much lower pressures.

We have a magnetic field coil that can ramp to maximum gradient or back in under 1 ms, allowing us to explore deep magnetic traps. We have a cooling system that can maintain the surface temperature of the coils within 3 K under normal operating conditions. It is adjustable in three dimensions which should speed up realignment after we transport the apparatus without significantly compromising alignment stability.

We have a system of electronics and optics that have been proven to be able to operate the 2D+ and 3D MOTs in the vacuum chamber. They produce over 70 mW of pump and repump light with a linewidth of under 3 MHz that can make small adjustments within microseconds and can switch to trapping an entirely different

¹This number does not include the volume of the interconnect from the apparatus to the chamber under test

²Not including the conductance of the interconnect

species within milliseconds. We can keep the standard deviation of the 3D MOT laser's amplitude within 0.3% of its fullscale value, and fully settle after a step response within $200\ \mu\text{s}$. Finally, we found that after adding vibration damping mounts, we could hit our rack with a screwdriver without unlocking our lasers.

There is still some work that needs to be done. We need to finish reassembling the apparatus in the new cart. We need to characterize the RF Knife's performance, then measure the base pressure in the apparatus using magnetic trap measurements. We also should compare with our previous data to see if the optimal loading parameters have changed after exposing the Rubidium source to an atmosphere of Argon when we rebuilt the chamber.

Once this work has been done, we plan to ship the apparatus to Physikalisch-Technische Bundesanstalt (PTB) in Germany to directly compare against their orifice flow standard. We have reason to be optimistic that this measurement can happen soon!

Bibliography

- [1] S. Bali, K. O’Hara, M. Gehm, S. Granade, and J. Thomas. Quantum-diffractive background gas collisions in atom-trap heating and loss. *Physical Review A*, 60(1):R29–R32, 1 1999. ISSN 2469-9926. [doi:10.1103/PhysRevA.60.R29](https://doi.org/10.1103/PhysRevA.60.R29). → pages 16, 17, 18
- [2] D. G. Bills. Causes of nonstability and nonreproducibility in widely used bayard–alpert ionization gauges. *Journal of vacuum science & technology. A, Vacuum, surfaces, and films*, 12(2):574–579, 1994. → pages 4, 5
- [3] J. L. Booth, P. Shen, R. V. Krems, and K. W. Madison. Universality of quantum diffractive collisions and the quantum pressure standard. *New Journal of Physics*, 21(10):102001, oct 2019. [doi:10.1088/1367-2630/ab452a](https://doi.org/10.1088/1367-2630/ab452a). → pages vi, xiv, xv, 5, 6, 14, 18, 19, 20, 21, 22, 23, 24
- [4] R. Calder and G. Lewin. Reduction of stainless-steel outgassing in ultra-high vacuum. *British Journal of Applied Physics*, 18(10):1459–1472, oct 1967. [doi:10.1088/0508-3443/18/10/313](https://doi.org/10.1088/0508-3443/18/10/313). URL <https://doi.org/10.1088/0508-3443/18/10/313>. → page 36
- [5] P. Chiggiato. Outgassing properties of vacuum materials for particle accelerators, 2020. → page 36
- [6] M. S. Child. *Molecular collision theory / M. S. Child*. Academic Press London ; New York, 1974. ISBN 0121726509. → pages vi, xv, 13, 14, 15, 23
- [7] ColdQuanta. *Cold Atom Source Cell User’s Manual*. ColdQuanta, June 2015. → page 40
- [8] ColdQuanta. *Photonically Integrated Cold Atom Source Fixed PICAS User Manual*. ColdQuanta, 2021. → pages vii, xx, 80, 101

- [9] A. Devices. *Microwave Wideband Synthesizer with Integrated VCO (ADF5356)*. Analog Devices, 2021. → page 113
- [10] A. Devices. *EV-ADF5356SD1Z User Guide*. Analog Devices, 2021. → page 107
- [11] M. R. Dietrich and B. B. Blinov. Use of a microcontroller for fast feedback control of a fiber laser, 2009. URL <https://arxiv.org/abs/0905.2484>. → page 116
- [12] J. V. Dongen. *Study of background gas collisions in atomic traps*. Phd thesis, University of British Columbia, 2014. → pages 4, 5, 37
- [13] D. Edwards. An upper bound to the outgassing rate of metal surfaces. *Journal of Vacuum Science and Technology*, 14(4):1030–1032, 1977. doi:10.1116/1.569315. → page 35
- [14] L. H. Ehinger, B. P. Acharya, D. S. Barker, J. A. Fedchak, J. Scherschligt, E. Tiesinga, and S. Eckel. Comparison of two multiplexed portable cold-atom vacuum standards. *AVS Quantum Science*, 4(3):034403, sep 2022. doi:10.1116/5.0095011. URL <https://doi.org/10.1116/5.0095011>. → pages 18, 36
- [15] A. Elkatmis. An analysis of the stability of spinning rotor gauges. *Measurement: Sensors*, 18:100244, 2021. ISSN 2665-9174. doi:<https://doi.org/10.1016/j.measen.2021.100244>. URL <https://www.sciencedirect.com/science/article/pii/S2665917421002075>. → page 3
- [16] D. Fagnan. *Study of Collision Cross Section of Ultra-Cold Rubidium using a Magneto-Optic and pure Magnetic trap*. Bachelor’s thesis, University of British Columbia, 2009. → pages 10, 15, 16
- [17] D. E. Fagnan, J. Wang, C. Zhu, P. Djuricanin, B. G. Klappauf, J. L. Booth, and K. W. Madison. Observation of quantum diffractive collisions using shallow atomic traps. *Phys. Rev. A*, 80:022712, Aug 2009. doi:10.1103/PhysRevA.80.022712. → pages 11, 16
- [18] J. A. Fedchak and D. R. Defibaugh. Long-term stability of metal-envelope enclosed bayard–alpert ionization gauges. *Journal of vacuum science & technology. A, Vacuum, surfaces, and films*, 30(6):61601, 2012. → page 4

- [19] J. A. Fedchak, J. Scherschligt, D. Barker, S. Eckel, A. P. Farrell, and M. Sefa. Vacuum furnace for degassing stainless-steel vacuum components. *Journal of Vacuum Science & Technology A*, 36(2):023201, 2018. doi:10.1116/1.5016181. → page 36
- [20] D. J. Griffiths. *Introduction to electrodynamics; 4th ed.* Pearson, Boston, MA, 2013. doi:1108420419. URL <https://cds.cern.ch/record/1492149>. Re-published by Cambridge University Press in 2017. → page 47
- [21] E. O. Göbel, U. Siegner, W. O. Library, and W. F. A. O. 2019. *The New International System of Units (SI): Quantum Metrology and Quantum Standards*. Wiley-VCH [Imprint], Hoboken, 2019. ISBN 9783527344598;3527344594;. → page 2
- [22] D. M. Hoffman, B. Singh, and J. H. Thomas. Handbook of vacuum science and technology. In *Handbook of Vacuum Science and Technology*, 1997. → pages 4, 5, 29, 30, 35
- [23] K. Huang, H. Le Jeannic, J. Ruau del, O. Morin, and J. Laurat. Microcontroller-based locking in optics experiments. *The Review of scientific instruments*, 85, 09 2014. doi:10.1063/1.4903869. → page 116
- [24] G. T. INC. Magnetic resonance imaging system and method of manufacturing thereof, European Patent Office EP0310212A2, Feb. 1988. → page 48
- [25] K. Jooya, N. Musterer, K. W. Madison, and J. L. Booth. Photon-scattering-rate measurement of atoms in a magneto-optical trap. *Phys. Rev. A*, 88:063401, Dec 2013. doi:10.1103/PhysRevA.88.063401. URL <https://link.aps.org/doi/10.1103/PhysRevA.88.063401>. → pages 129, 142
- [26] K. Jousten. Gauges for fine and high vacuum. In *Gauges for fine and high vacuum*. AS 2006 - CERN Accelerator School: Vacuum in Accelerators, Proceedings, Jan 2007. → page 3
- [27] W. Ketterle and N. V. Druten. Evaporative cooling of trapped atoms. In B. Bederson and H. Walther, editors, *Evaporative Cooling of Trapped Atoms*, volume 37 of *Advances In Atomic, Molecular, and Optical Physics*, pages 181–236. Academic Press, 1996. doi:[https://doi.org/10.1016/S1049-250X\(08\)60101-9](https://doi.org/10.1016/S1049-250X(08)60101-9). URL <https://www.sciencedirect.com/science/article/pii/S1049250X08601019>. → page 8

- [28] Lesker. Conflat (cf) uhv flanges & components technical notes, 2022. URL https://www.lesker.com/newweb/flanges/flanges_technicalnotes_conflat_1.cfm. → page 24
- [29] D. Li, G. Zhao, M. Guo, J. Xu, and Y. Cheng. Static expansion vacuum standard with extended low pressure range. *MAPAN*, 24:95–100, 06 2009. doi:10.1007/s12647-009-0012-z. → page 3
- [30] H. J. Metcalf and P. van der Straten. *Laser Cooling and Trapping of Neutral Atoms*, chapter 1-18, pages 1–261. American Cancer Society, 2007. ISBN 9783527600441. doi:10.1002/9783527600441.oe005. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/9783527600441.oe005>. → pages vi, xv, 8, 9, 10, 11, 12
- [31] T. Meyrath. *Experiments with Bose-Einstein Condensation in an Optical Box*. PhD thesis, University of Texas, 2005. → page 106
- [32] P. W. Milonni, J. H. Eberly, and W. O. Library. *Lasers physics*. John Wiley & Sons, Hoboken, N.J., 2010. → page 84
- [33] M. Neumann-Brosig, A. Marco, D. Schwarzmann, and S. Trimpe. Data-efficient autotuning with bayesian optimization: An industrial control study. *IEEE Transactions on Control Systems Technology*, 28:730–740, 2020. → page 122
- [34] D. Nino, H. Wang, and J. N. Milstein. Rapid feedback control and stabilization of an optical tweezers with a budget microcontroller. *European Journal of Physics*, 35(5):055009, jul 2014. doi:10.1088/0143-0807/35/5/055009. URL <https://doi.org/10.1088/0143-0807/35/5/055009>. → page 116
- [35] C. O’Sullivan. Newton’s law of cooling-a critical assessment. *American Journal of Physics - AMER J PHYS*, 58:956–960, 10 1990. doi:10.1119/1.16309. → page 48
- [36] C. D. Park, S. M. Chung, X. Liu, and Y. Li. Reduction in hydrogen outgassing from stainless steels by a medium-temperature heat treatment. *Journal of Vacuum Science & Technology A*, 26(5):1166–1171, 2008. doi:10.1116/1.2956625. → page 36
- [37] L. Peksa, D. Prazak, T. Gronych, P. Řepa, M. Vièar, J. Tesar, Z. Krajčůek, and F. Staník. Primary vacuum standard for uhv range - standing experience

- and present problems. *MAPAN-Journal of Metrology Society of India*, 24: 77–88, 06 2009. doi:10.1007/s12647-009-0010-1. → page 4
- [38] J. Perez-Rios and A. S. Sanz. How does a magnetic trap work? *American Journal of Physics*, 81(11):836–843, nov 2013. doi:10.1119/1.4819167. URL <https://doi.org/10.1119%2F1.4819167>. → page 47
- [39] R. H. Perry and D. W. Green. *Perry’s chemical engineers’ handbook*. McGraw-Hill, New York, 8th edition, 2008. ISBN 9780071422949;0071422943;. → pages 166, 169
- [40] A. L. Schawlow. Spectroscopy in a new light. *Science*, 217(4554):9–16, 1982. ISSN 00368075, 10959203. URL <http://www.jstor.org/stable/1689577>. → page 83
- [41] R. J. Scherrer. Time variation of a fundamental dimensionless constant, 2009. URL <https://arxiv.org/abs/0903.5321>. → page 2
- [42] J. Scherschligt, J. A. Fedchak, D. S. Barker, S. Eckel, N. Klimov, C. Makrides, and E. Tiesinga. Development of a new uhv/xhv pressure standard (cold atom vacuum standard). *Metrologia*, 54(6):S125, 2017. URL <http://stacks.iop.org/0026-1394/54/i=6/a=S125>. → page 11
- [43] P. Shen. *Development of a Cold Atom Pressure Standard*. Phd thesis, University of British Columbia, 2022. → pages vi, 11, 21, 22, 80, 128, 141
- [44] P. Shen, K. W. Madison, and J. L. Booth. Realization of a universal quantum pressure standard. *Metrologia*, 57(2):025015, mar 2020. doi:10.1088/1681-7575/ab7170. → page 19
- [45] P. Shen, K. W. Madison, and J. L. Booth. Refining the cold atom pressure standard. *Metrologia*, 58(2):022101, feb 2021. doi:10.1088/1681-7575/abe02f. → pages vi, xiv, 19, 23
- [46] P. Shen, E. Frieling, K. R. Herperger, D. Uhland, R. A. Stewart, A. Deshmukh, R. V. Krems, J. L. Booth, and K. W. Madison. Cross-calibration of atomic pressure sensors and deviation from quantum diffractive collision universality for light particles. *Preprint*, sept 2022. → pages vi, xiv, 23, 24
- [47] J. L. Siegel, D. S. Barker, J. A. Fedchak, J. Scherschligt, and S. Eckel. A bitter-type electromagnet for complex atomic trapping and manipulation. *Review of Scientific Instruments*, 92(3):033201, mar 2021.

- doi:10.1063/5.0026812. URL <https://doi.org/10.1063%2F5.0026812>. → page 49
- [48] D. A. Steck. Rubidium 87 d line data, Jul 2021. URL <https://steck.us/alkalidata/rubidium87numbers.pdf>. → pages vii, xviii, 11, 78, 84, 144
- [49] C. V. Sukumar and D. M. Brink. Spin-flip transitions in a magnetic trap. *Phys. Rev. A*, 56:2451–2454, Sep 1997. doi:10.1103/PhysRevA.56.2451. URL <https://link.aps.org/doi/10.1103/PhysRevA.56.2451>. → page 22
- [50] J. Söding, D. guéry odelin, P. Desbiolles, G. Ferrari, and J. Dalibard. Giant spin relaxation of an ultracold cesium gas. *Physical Review Letters - PHYS REV LETT*, 80:1869–1872, 03 1998. doi:10.1103/PhysRevLett.80.1869. → page 47
- [51] Thorlabs. Polarization-maintaining fc/apc fiber optic patch cables, 2022. URL https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=3345. → page 95
- [52] E. Toolbox. Colebrook equation, 2021. URL https://www.engineeringtoolbox.com/colebrook-equation-d_1031.html. → page 166
- [53] E. Toolbox. Convective heat transfer, 2021. URL https://www.engineeringtoolbox.com/convective-heat-transfer-d_430.html. → page 168
- [54] E. Toolbox. Darcy-weisbach pressure and major head loss equation, 2021. URL https://www.engineeringtoolbox.com/darcy-weisbach-equation-d_646.html. → page 166
- [55] E. Toolbox. Poiseuille’s formula, 2021. URL https://www.engineeringtoolbox.com/poiseuille-formula-d_1625.html. → page 165
- [56] W. Umrath. Fundamentals of vacuum technology. In *Fundamentals of Vacuum Technology*, 1998. → pages 29, 30
- [57] Vescent. D2-135 offset phase lock servo, 2022. URL https://www.vescent.com/manuals/doku.php?id=d2:offset_phase_lock_servo. → pages xiv, xix, 87, 88

- [58] Vescent. D2-210 spectroscopy module, 2022. URL https://www.vescent.com/manuals/doku.php?id=d2:spectroscopy_module_210.
→ pages xix, 84
- [59] Vescent. D2-250 heterodyne module, 2022. URL https://www.vescent.com/manuals/doku.php?id=d2:d2_250_heterodyne_module.
→ pages xix, 86
- [60] R. Weaver. Numerical methods for inductance calculation, 2021. URL <http://electronbunker.ca/eb/CalcMethods.html>. → page 164

Appendix A

Magnetic Field Coil Modelling

A.1 Inductance

The mutual inductance between two coils is [60]

$$M = -\mu_0 \sqrt{r_1 r_2} \left(\left(k - \frac{2}{k} \right) K(k) + \frac{2}{k} E(k) \right), \quad (\text{A.1})$$

$$k = \frac{2\sqrt{r_1 r_2}}{\sqrt{(r_1 + r_2)^2 + x^2}}$$

where r_1 and r_2 are the radii of the two coils, x is the spacing between them, and E and K are the elliptical integral functions. When finding the self-inductance of the coil, set $x = g = r_w e^{-.25}$ where r_w is the radius of the wire to compensate for the finite diameter of the wire.

If there are n axial turns and m radial turns to the coil with radii r_i and axial positions x_i , then the total inductance of the coil is

$$L = n \sum_{i=1}^m M(r_i, r_j, g) + 2n \sum_{i=1}^{m-1} \sum_{j=i+1}^m M(r_i, r_j, 0)$$

$$+ 2 \sum_{i=1}^m \sum_{j=1}^{n-1} (n-j) M(r_i, r_i, jp) + 4 \sum_{i=1}^{m-1} \sum_{j=i+1}^m \sum_{k=1}^{n-1} (n-k) M(r_i, r_j, kp) \quad (\text{A.2})$$

where $p = x_{i+1} - x_i$.

- the first term corresponds to the self-inductance of each coil
- the second term is the mutual inductance of all coils with different radii at the same axial position
- the third term is the mutual inductance of all coils with the same radii but different axial positions
- the fourth term is the mutual inductance of all coils with different radii and different axial positions

A.2 Fluid Flow

A.2.1 Flow Rates

Fluid flow in a pipe is either laminar, turbulent, or mixed. Which regime it is in is determined by its Reynolds number:

$$Re = \frac{\rho v d}{\mu} \quad (\text{A.3})$$

where ρ is the density of the fluid, v is the mean velocity of the fluid, and d is the diameter of the pipe.

If $Re < 2300$, the flow is laminar. If $Re > 4600$, the flow is turbulent. Note that these numbers are approximate; different literature quotes the laminar cutoff as between 2000 and 2300, and the turbulent cutoff between 4600 and 6100. We chose 2300 and 4600 as thresholds, but numbers may produce slightly more accurate results.

Laminar Flow

If the flow is laminar, then the relationship between the pressure drop and the velocity of the fluid is given by the Hagen-Poiseuille law [55]:

$$\Delta P = \frac{32\mu L v}{d^2} \quad (\text{A.4})$$

where L is the length of the pipe. Note that $L/d > 96Re$ for eq. (A.4) to hold; Bernouilli's principle requires that $v < \sqrt{\frac{2\Delta P}{\rho}}$.

Turbulent Flow

If the flow is turbulent, then the relationship between the pressure drop and the velocity of the fluid is given by the Darcy-Weisbach equation [54]:

$$\Delta P = \frac{fLv^2}{2d} \quad (\text{A.5})$$

where f (the friction factor) solves the Colebrook equation [52]:

$$\frac{1}{\sqrt{f}} = -2 \log \left(\frac{2.51}{Re\sqrt{f}} \right) + \frac{e}{3.72} \quad (\text{A.6})$$

where e is the “relative roughness” - the ratio of roughness to tube diameter. Note that $f = \frac{64}{Re}$ for laminar flow.

Mixed Flow

We could not find a consensus solution on how to treat “mixed flow”, where $2300 < Re < 4600$. As a result, we took the weighted average of the two solutions:

$$\Delta P = \frac{(Re - 2300)\Delta P_{turb} + (4600 - Re)\Delta P_{lam}}{4600 - 2300} \quad (\text{A.7})$$

A.2.2 Helical Corrections

There are a few corrections that need to be made to the above equations if the water is flowing through a helical pipe.

Laminar Flow Threshold

Flow remains laminar in a helical pipe as long as $Re < 2300 \left(1 + 12\sqrt{\frac{d_c}{d_t}} \right)$ where d_c is the diameter of the helix and d_t is the diameter of the tube [39].

Updated Friction Factor

eq. (A.5) is used for both laminar and turbulent flow in helical pipes. If flow is laminar, then

$$f = \begin{cases} \frac{64}{Re} & De < 11.6 \\ \frac{64}{Re} \frac{1}{1 - \left(1 - \left(\frac{11.6}{De}\right)^{0.45}\right)^{2.2}} & 11.6 < De < 2000 \\ \frac{7.0144\sqrt{De}}{Re} & De > 2000 \end{cases} \quad (A.8)$$

where $De = Re\sqrt{\frac{d_c}{d_t}}$ is the Dean number.

If flow is turbulent, $Re\left(\frac{d_c}{d_t}\right)^2 < 700$, and $0.0097 < \frac{d_c}{d_t} < 0.135$, then

$$f = 0.336 \left(\frac{d_c}{d_t}\right)^{0.1} Re^{-0.2} \quad (A.9)$$

If the tube is rough, then let f_c be the value calculated using eq. (A.9), f_r be the value calculated using eq. (A.6), and f_s be the value calculated using eq. (A.6) with $e = 0$. Then, use

$$f = f_c \frac{f_r}{f_s} \quad (A.10)$$

A.3 Cooling Theory

The heat generated by the coils is

$$\dot{Q} = \frac{I^2 R_0}{L} \int_0^L 1 + \alpha(T(x) - T_r) dx = I^2 R_0 (1 + \alpha(\bar{T} - T_r)) \quad (A.11)$$

where I is the current flowing through the coils, R_0 is the resistance of the coils at temperature T_r , \bar{T} is the mean temperature of the coils, and α is the temperature coefficient of the conductor the coils are made out of.

A.3.1 Non-liquid Cooling

Convective Cooling

The rate that heat is lost from a plate at temperature T of area A into a fluid of ambient temperature T_a follows Newton's law of cooling (eq. (A.12)):

$$\dot{Q} = hA(T - T_a) \quad (\text{A.12})$$

h , the convection coefficient, is typically experimentally determined, but some efforts have been made to come up with an analytic equation for it. This type of cooling is referred to as forced convection if the fluid is propelled over the surface being cooled rather than simply allowed to convect naturally.

One empirical equation for the convection coefficient of air is [53]

$$h = 12.12 - 1.16v + 11.6v^{1/2} \quad (\text{A.13})$$

where all units are in MKS. It is valid for velocities between 2 and 20 m/s.

Radiative Cooling

$$\dot{Q} = \sigma eA(T^4 - T_a^4) \quad (\text{A.14})$$

where σ is the Stefan-Boltzmann constant, e is the emissivity of the surface, A is the area of the surface, T is the temperature of the surface, and T_a is the ambient temperature of the air.

Typically, radiative effects are less than 1% the size of all other effects and can easily be ignored.

Conductive Cooling

The flow of heat through an object of cross-sectional area A is

$$\dot{Q} = kA \frac{dT}{dx} \quad (\text{A.15})$$

where k is the thermal conductivity of the material the object is made from.

A.3.2 Liquid Cooling

The rate that heat is transferred to a fluid flowing through a pipe is also described by eq. (A.12), where $h = \frac{k}{d}Nu$, where k is the thermal conductivity of the fluid and Nu is the Nusselt number [39].

Laminar Flow

If the flow in the pipe is laminar, then

$$Nu = 1.86 \left(Re Pr \frac{d}{L} \right)^{\frac{1}{3}} \left(\frac{\mu(T_f)}{\mu(T_w)} \right)^{0.14} \quad (A.16)$$

where $Pr = \frac{c\mu}{k}$ (Prandtl number), c is the isobaric heat capacity of the fluid, T_f is the temperature of the fluid, and T_w is the temperature of the wall. At low temperature differences, it is reasonable to say $\mu(T_f) \approx \mu(T_w)$.

Turbulent Flow

If the flow in the pipe is turbulent, $0.7 < Pr < 16000$, and $L/d > 60$ then

$$Nu = 0.027 Re^{\frac{4}{3}} Pr^{\frac{1}{3}} \left(\frac{\mu(T_f)}{\mu(T_w)} \right)^{0.14} \quad (A.17)$$

If the flow is through helical pipe rather than straight pipe and $Re > 10^4$, then multiply eq. (A.17) by $1 + 3.5 \frac{d_t}{d_c}$.

If $2 < \frac{L}{d} < 20$, multiply eq. (A.17) by $1 + (\frac{d_t}{d_c})^{0.7}$. If $20 < \frac{L}{d} < 60$, multiply eq. (A.17) by $1 + 6(\frac{d_t}{d_c})$.

Mixed Flow

Like before, there is no clear solution to the mixed flow regime. We took the weighted average:

$$Nu = \frac{(Re - 2300)Nu_{turb} + (4600 - Re)Nu_{lam}}{4600 - 2300} \quad (A.18)$$

A.4 Cooling Models

A.4.1 Hollow-core Coil Models

Naive Model

The simplest way of modelling the temperature of a hollow fluid-cooled coil is to assume the coil is the same temperature as the fluid and that the temperature of the fluid increases linearly over the length of the coil. Therefore, at steady-state,

$$T_f(L) - T_f(0) = \frac{\dot{Q}}{\pi(d/2)^2 \rho c v}$$

where v is the velocity of the fluid and \dot{Q} is the heat created by the coils.

$$\bar{T}_f = \bar{T}_c = T = \frac{1}{2} \frac{\dot{Q}}{\pi(d/2)^2 \rho c v} + T_f(0) = \frac{1}{2} \frac{I^2 R_0 (1 + \alpha(T - T_r))}{\pi(d/2)^2 \rho c v} + T_0$$

Some algebra gives

$$T = \frac{2I^2 R_0 (1 - \alpha T_r) + \pi d^2 \rho c v T_0}{2I^2 R_0 \alpha + \pi d^2 \rho c v} \quad (\text{A.19})$$

However, this does not take into account

- the temperature difference between the fluid and the conductor.
- the position-dependent power due to the changing temperature along the length of the coil.
- the conductive flow of heat from hot sections of the coil to cold sections.
- heat flow between adjacent windings.
- any cooling effect other than the cooling from the water.

Exponential Model

This model takes into account

- the temperature difference between the fluid and the conductor.
- the position-dependent power due to the changing temperature along the length of the coil.

At steady-state, the infinitesimal heat flow from the conductor to the fluid is

$$d\dot{Q} = hS(T_c(x) - T_f(x))dx = \frac{I^2 R_0}{L}(1 + \alpha(T_c(x) - T_r))dx \int_0^x d\dot{Q} = \frac{\pi d}{4} \rho_{cv}(T_f(x) - T_f(0))$$

where $S = \pi d$ is the circumference of the pipe.

Taking the derivative of both sides of the second equation and solving the ODE produces

$$\begin{aligned} T_c(x) &= \frac{\pi L d h}{\pi L d h - I^2 R_0 \alpha} T_w(x) + \frac{I^2 R_0 (1 - \alpha T_r)}{\pi L d h - I^2 R_0 \alpha} \\ T_w(x) &= \frac{\pi L d h}{\pi L d h - I^2 R_0 \alpha} \left((T_w(0) - T_r + \frac{1}{\alpha}) e^{\kappa x} + T_r - \frac{1}{\alpha} \right) + \frac{I^2 R_0 (1 - \alpha T_r)}{\pi L d h - I^2 R_0 \alpha} \\ P &= \frac{I^2 R_0}{L} \left(\left((\alpha T_r - 1) \left(\frac{\pi L d h}{\pi L d h - I^2 R_0 \alpha} - 1 \right) + \alpha \frac{I^2 R_0 (1 - \alpha T_r)}{\pi L d h - I^2 R_0 \alpha} \right) x \right. \\ &\quad \left. + \frac{\pi L d h}{\pi L d h - I^2 R_0 \alpha} \frac{\alpha (T_w(0) - T_r) + 1}{\kappa} (e^{\kappa x} - 1) \right) \end{aligned}$$

where $\kappa = \frac{4hI^2 R_0 \alpha}{\rho_{cv} d (\pi L d h - I^2 R_0 \alpha)}$ and P is the total power produced by the coil.

Exponential Conduction Model

It is possible to also take into account the effect of conduction along the length of the coil. To do so, use

$$\begin{aligned} hS(T_c(x) - T_f(x)) + kA \frac{dT_c(x)}{dx} &= \frac{I^2 R_0}{L} (1 + \alpha(T_c(x) - T_r)) \\ \int_0^x \frac{I^2 R_0}{L} (1 + \alpha(T_c(x) - T_r)) dx &= \frac{\pi d}{4} \rho_{cv}(T_f(x) - T_f(0)) \\ \left. \frac{dT_c(x)}{dx} \right|_{x=0} &= 0 \end{aligned}$$

where $S = \pi d$ is the circumference of the pipe, and $A = \frac{\pi(D^2 - d^2)}{4}$ is the area of the conductor.

We this to model the behaviour of the hollow-core water-cooled coils. It still neglected the effect of heat being transferred between adjacent windings or any passive cooling, but produced approximately correct results.