

Superfluid ^4He interferometers: construction and experiments

by

Aditya Ajit Joshi

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Richard E. Packard, Chair
Professor Irfan Siddiqi
Professor K. Birgitta Whaley

Spring 2013

Superfluid ^4He interferometers: construction and experiments

Copyright 2013
by
Aditya Ajit Joshi

Abstract

Superfluid ^4He interferometers: construction and experiments

by

Aditya Ajit Joshi

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Richard E. Packard, Chair

This dissertation has two main goals: to highlight some new results in the field of superfluid ^4He interferometry and to provide an in-depth, “hands-on” guide to the physics, design, construction, testing and operation of a continuously operating, fluxlocked ^4He SHeQUID (**S**uperfluid **H**elium **Q**uantum **I**nterference **D**evice). Many of these topics haven’t really been addressed in writing and the hapless new experimenter seeking to develop a SHeQUID is generally forced to reinvent the wheel rather than start at the frontier and push it forward. We would like to prevent that by making this a comprehensive guide to building and operating SHeQUIDs.

We have optimized the fabrication of the nanoscale aperture arrays that are the very heart of the SHeQUID and resolved long-standing issues with their durability and long-term usability. A detailed report on this should assist in avoiding the many pitfalls that await those who fabricate and use these aperture arrays.

We have constructed a new, modular SHeQUID that is designed to be easily adaptable to a wide array of proposed experiments without the necessity of rebuilding and reassembling key components like the displacement transducer. We have automated its working as a continuously operating, linearized (flux-locked) interferometer by using the so-called “chemical potential battery” in conjunction with a feedback system. We have also constructed a new reorientation system that is several orders of magnitude quieter than its predecessors. Together, these developments have allowed us to measure a changing rotation field *in real time*, a new development for this kind of device. We have also developed a module that allows control of the reorientation stage by automated data-taking software for investigating long-term drifts (by safely sweeping the stage back and forth).

We have also investigated the chemical potential battery in further detail and report some fascinating nonlinear mode locking phenomena that have important consequences for practical applications of these devices. We present a crude model that should help in designing and optimizing future devices by giving us at least an initial predictive tool for the critical heater power needed to initiate battery states.

Finally, we analyze some misconceptions about SHeQUIDs regarding what may be considered the logical next step towards improving a double-slit interferometer - the superfluid diffraction grating. We present evidence (experiments, simulations and analytical results) for the somewhat subtle reasons why gratings would be less useful than previously believed and clarifies the proper, limited sense in which such devices *do* improve SHeQUIDs. We also discuss some possible implications of these issues for the field of (electronic) dc-SQUIDs.

To my parents, for nurturing my love for science and for support beyond reason

Contents

Abstract	1
Dedication	i
Contents	ii
List of Figures	vi
List of Tables	xii
Acknowledgments	xiii
1 Introduction	1
1.1 Preamble	1
1.2 The quantum whistle	4
1.3 DC superfluid interferometry	13
1.4 Physical cells	22
1.5 Flow dynamics	22
2 Continuously operating Fiske-enhanced SHeQUID	26
2.1 The chemical potential battery	26
2.2 Cell resonant modes	28
2.3 Interaction of battery with resonances: Fiske-locking and amplification	29
2.4 Mathematical model	32
3 Weak link cell design, components and constraints	40
3.1 Component overview	40
3.2 Constraints	47
3.3 Possible design philosophies	57
3.4 Modular single weak-link cell design	57
3.5 Modular SHeQUID design	60
4 Fabricating nanoscale aperture arrays	64

4.1	Introduction	64
4.2	Fabrication outline	65
4.3	Issues	69
4.4	Wafer considerations	74
4.5	Conclusions	75
5	The displacement sensor	76
5.1	Persistent current type	76
5.2	Magnet type	92
6	Independent component tests	96
6.1	Aperture arrays	96
6.2	Superconducting diaphragm	115
7	The Cryostat	118
7.1	Overview and broad issues	118
7.2	Construction	118
7.3	Thermometry	123
7.4	Wiring	124
7.5	Structural issues	130
8	Cryogenic valves	136
8.1	Introduction and history	136
8.2	Fabrication	138
8.3	Testing the cryovalve	151
9	Laboratory infrastructure	159
9.1	Acoustic and vibration isolation	159
9.2	An improved, quiet rotation system	163
9.3	Automating rotation sweeps (with redundant safety features)	166
10	Operation	173
10.1	Layout and overview	173
10.2	Preliminary tests	178
10.3	Cooldown	183
10.4	The vacuum resonance (simple and mass-loaded)	185
10.5	Thermometry calibrations	186
10.6	Locating the lambda point	190
10.7	Filling and emptying the cell	192
10.8	Cell calibrations	201
10.9	Summary: calibration sequences	212
10.10	Transient analysis	215
10.11	Frequency response of cell	218

10.12	Chemical potential battery	220
10.13	Interferometry	221
11	New results	222
11.1	Gyroscopy with continuous cryostat reorientation	222
11.2	Flux-locked and linearized gyroscope for measuring continuously changing rotation fields	224
12	Noise and drift	229
12.1	Phase drifts	229
12.2	Performance figures of the flux-locked SHeQUID	232
12.3	The resonant landscape and frequency-dependent Fiske gains	234
13	The superfluid diffraction grating	235
13.1	Introduction	235
13.2	Linear grating theory	237
13.3	Experimental evidence	246
13.4	Simulations	250
13.5	The dc SQUID: misconceptions and clarifications	254
	Bibliography	259
	Appendices	268
A	Engineering drawings, fabrication and assembly	268
A.1	Fabrication: SHeQUID structural parts	268
A.2	SHeQUID assembly	279
A.3	Displacement sensor drawings	298
A.4	Single weak-link cell drawings	301
B	Matlab scripts	305
B.1	Interference curves	305
B.2	Fountain calibration	313
C	Flow tests: further analysis	315
C.1	Clausing's factor: $K[u]$	315
C.2	Mathematica code	318
D	On the superposition of phase-coherent oscillations	323
D.1	General problem	323
D.2	In retrospect	325
D.3	Example: The heat current experiment with multiple weak-links	329
D.4	Heat Current Experiment: closed form solution for equal amplitudes	333

D.5	Plots and code listings	334
E	Effect of finite fluid compressibility	337
E.1	Introduction - the zero order approximation	337
E.2	Compression with constant density	338
E.3	Estimates	340
E.4	Compression with changing density	341
F	Inertial effects and resonant frequencies	344
F.1	Inertial effects	344
F.2	Diaphragm resonant frequency	347
F.3	Helmholtz resonance	348
F.4	Cavity acoustic resonances	349
G	Cell dynamics derivations	355
G.1	Normal and total currents	355
G.2	Supercurrent and inductance	358
G.3	Chemical potential difference	362
G.4	The temperature equation	362
G.5	Summary	365
H	Flow simulations for a superfluid grating	367
H.1	Introduction	367
H.2	Single-axis rotation	367
H.3	Matlab code	371

List of Figures

1.1	He-II healing length as a function of temperature.	6
1.2	Two different coupling regimes for two superfluids separated by a wall with a hole.	7
1.3	Solid wall vs. weak-link.	7
1.4	Circuit representation of the Deaver-Pierce model	9
1.5	Current-phase relations for He-II weak-link (calculated using Deaver-Pierce model).	10
1.6	Current-phase relations for He-II weak-link (measured).	10
1.7	2π phase-slippage in a strong-link.	11
1.8	Phase slip oscillations in a strong-link.	11
1.9	Phase slip oscillation amplitude vs. temperature.	13
1.10	Cartoon summary of various coupling regimes.	13
1.11	A 2-slit SHeQUID.	14
1.12	Interference patterns for several asymmetry parameter values.	14
1.13	Vector identity visual proof.	17
1.14	Cartoon representation of Sagnac interferometry using a SHeQUID.	18
1.15	Two extreme orientations of the SHeQUID loop area vector.	18
1.16	Simulated interference pattern for Sagnac effect.	19
1.17	Cartoon depiction of heat-current driven counterflow along a channel.	20
1.18	Depiction of the predicted Aharonov-Bohm (AB) effect for neutral matter.	21
1.19	Schematic of the Aharonov-Bohm (AB) effect for neutral matter.	21
1.20	Single weak-link cell.	22
1.21	Double weak-link cell (SHeQUID).	23
2.1	Diaphragm (velocity) response as a function of drive frequency for 3 different temperatures.	29
2.2	Zoomed version of Fig. 2.2.	30
2.3	Example of a typical battery ramp-up and recorded whistle frequencies during ramp.	31
2.4	Critical power vs. ramp rate: $T = 2.170$ K	37
2.5	Critical power vs. ramp rate: $T = 2.171$ K	37
2.6	Critical power vs. ramp rate: $T = 2.172$ K	38

2.7	Critical power vs. ramp rate: $T = 2.173$ K	38
2.8	Critical power vs. ramp rate: $T = 2.174$ K	39
2.9	Critical power vs. ramp rate: $T = 2.175$ K	39
3.1	AB experiment conceptual sketch and counter-wound sense arm.	47
3.2	Electrical circuit analogue of the SHeQUID with finite loop inductance.	49
3.3	Simulated current vs. time plots for varying α (front view).	52
3.4	Simulated current vs. time plots for varying α (back view).	53
3.5	First 15 integrated peaks in FFT of Figs. 3.3–3.4.	54
3.6	First 15 integrated peaks in FFT of Figs. 3.3–3.4 (normalized in quadrature for each α).	55
3.7	First harmonic amplitude vs. α and φ_a	56
3.8	Partially modular single weak-link cell design.	58
3.9	Proposed modification in the design of Fig. 3.8.	59
3.10	Modular SHeQUID showing V-tunnel plane.	61
3.11	Modular SHeQUID showing breakout plane (Fig. 3.10 rotated 90 deg.)	62
4.1	Wafer processing steps for aperture array fabrication.	66
4.2	SEM image of aperture array.	70
4.3	SEM image of window showing defects.	70
4.4	Membrane wrinkling due to stress.	72
4.5	Before and after images of KCl contaminated nitride membranes.	73
5.1	Persistent current (PI) circuit.	77
5.2	Pancake coil former for PI-style sensor.	81
5.3	Pancake coil winder.	82
5.4	Pancake coil former for PI-style sensor (photo).	83
5.5	Spark welding setup for superconducting joints.	87
5.6	Persistent current circuit shielded box (“PI box”).	89
5.7	Pancake coil former for magnet-style sensor. Made of Stycast 2850FT.	94
5.8	Photo of pancake coil former for magnet-style sensor.	94
5.9	Magnet-loaded diaphragm (magnet being glued).	95
6.1	Flow test cell.	97
6.2	Electrode holder	99
6.3	Diaphragm holder: chip facing side	100
6.4	Diaphragm holder: electrode facing side	101
6.5	Chip holder	102
6.6	Electrode holder (mounted on probe flange spacers) with electrode glued on and wire screwed on to tab using nylon screw.	103
6.7	Diaphragm holder with diaphragm glued on and wire connected.	103
6.8	Diaphragm holder from Fig. 6.7 flipped over and screwed onto electrode holder.	104
6.9	Disposable chip-holder.	104

6.10	Assembled cell before being covered by vacuum can.	105
6.11	KF 4-way breakout with labels for BNCs and pumping port.	106
6.12	Vacuum can	106
6.13	Typical capacitance transient and best-fit curve.	110
6.14	Knudsen number (Kn) as a function of ambient pressure and hole diameter (D) at 77K for ^4He gas.	112
6.15	Knudsen number (Kn) as a function of ambient pressure for different hole diameters (D) at 77K for ^4He gas.	113
6.16	Aspect ratio dependence of the conductance in the molecular flow regime.	114
6.17	A simple jig to test critical fields of superconducting diaphragms.	115
6.18	The control test with no superconducting diaphragm to gauge the bare coupling level.	117
6.19	Example of a “good” diaphragm.	117
6.20	Example of a “bad” diaphragm.	117
7.1	Annotated overview photo of the cryostat.	119
7.2	Close-up of volume elements in the generated mesh of the model cryostat.	132
7.3	Lowest bending mode of (present) cryostat model (bending is exaggerated). $f_1 \sim 4.2 \text{ Hz}$	132
7.4	Example of a raw time-series signal from geophone.	133
7.5	Power spectrum of geophone signal from Fig. 7.4 showing the modal peak.	134
7.6	Lowest bending mode of cryostat model (bending is exaggerated), with more baffles added throughout its length to make it stiffer. $f_1 \sim 9.3 \text{ Hz}$ – more than twice that of the original model.	134
8.1	Comparison of cryostat filling line without and with a cryovalve (CV).	137
8.2	Simplified schematic of pneumatic cryovalve.	138
8.3	Scale model of assembled cryovalve.	138
8.4	All parts after machining.	139
8.5	Servometer FC-14-L bellows. Gets soldered to guide and plunger.	140
8.6	Brass housing can.	140
8.7	Brass guide (plunger slides inside this).	141
8.8	Brass plunger (tip is screwed on to this).	142
8.9	Torlon tip.	142
8.10	Stainless steel seat.	143
8.11	Photo of a stainless steel seat ready for polishing.	145
8.12	Photo of a plunger/guide/bellows subsystem being soldered.	145
8.13	An inline filter. See text for details.	151
8.14	Plumbing circuit for testing cryovalve under various conditions (see text).	152
8.15	Photo of testing jig with cryovalve mounted.	153
8.16	Cryovalve # 1 closing curve.	156
8.17	Cryovalve # 2 closing curve.	157

8.18 Cryovalve # 4 closing curve.	158
9.1 Pneumatic spring suspension system.	161
9.2 Example of battery states switching between two Fiske modes.	164
9.3 Dewar on bearing bolted to air springs (annotated).	165
9.4 photo of working setup for auto-rot (annotated).	169
9.5 Tachometer close-up and wood sandwich closeups, etc. (annotated).	170
9.6 Photo of control room side rot ckt. (annotated)	170
9.7 Full circuit used for auto-rot.	171
9.8 Detail view of the relay box (RB) module.	171
9.9 Home-made optical switch.	172
9.10 LDR circuit used in the optical switch.	172
10.1 Measurement block diagram (comprehensive).	174
10.2 Plumbing setup for cell evacuation.	181
10.3 Vacuum resonance	186
10.4 Calibrating the Sensym pressure gauge.	187
10.5 Calibrating the GRT.	188
10.6 HRT calibration (raw signal vs. temperature).	189
10.7 HRT (salt) calibration (sensitivity vs. T): high T.	190
10.8 HRT (salt) calibration (sensitivity vs. T): low T.	191
10.9 HRT (PdMn) calibration (sensitivity vs. T).	192
10.10 Locating T_λ of the bath.	193
10.11 The lambda line.	194
10.12 Plumbing setup for filling cell.	196
10.13 Plumbing setup for closing the cryovalve.	198
10.14 Normal flow ($T = 2.1658K$)	203
10.15 Fountain series fit.	204
10.16 Thermal boundary resistance.	205
10.17 γ_1 calibration using whistle.	206
10.18 (D-E capacitance) C_x vs. V_b circuit setup.	207
10.19 C_x vs. V_b raw data from bridge output.	208
10.20 C_x vs. V_b data and quadratic fit.	208
10.21 SQUID voltage vs. V_b circuit setup.	209
10.22 SQUID voltage vs. V_b data.	209
10.23 Helmholtz series: frequency f_H	210
10.24 Helmholtz series: decay time τ_H	211
10.25 Helmholtz series: quality factor Q.	211
10.26 Sample battery state timeseries.	213
10.27 Sample battery state power PSD.	213
10.28 Josephson frequency relation.	216
10.29 Current phase relation: strong coupling.	217

10.30	Current phase relation: weak coupling.	218
10.31	Whistle amplitude vs. frequency (strong coupling).	219
10.32	Whistle amplitude vs. frequency (strong coupling) - zoomed in.	219
10.33	Setup to perform resonant frequency sweeps with capacitive drive.	220
11.1	Sagnac effect interference curve (continuous reorientation).	223
11.2	Heat-pipe power interference curve.	225
11.3	Raw data: Whistle amplitude (kept constant) and feedback output power needed to do so.	227
11.4	Raw data: cryostat angle and feedback output power.	227
11.5	Example of dynamic feedback demonstration in a continuously operating SHe- QUID.	228
12.1	(a) Example drift run and (b) summary of drift runs (for phase drift).	230
12.2	Proposed cell modification for reducing phase drift.	231
12.3	A cartoon depiction of a broad resonant peak.	234
13.1	Toroidal superfluid grating design proposed by Sato, <i>et al.</i>	236
13.2	(a) Two slit interferometer and (b) N-slit grating interferometer.	237
13.3	Simulated modulation amplitudes (vs. θ) for various N-slit interferometers.	241
13.4	Simulated modulation amplitudes (vs. Ω) for various N-slit interferometers.	245
13.5	2 slit interferogram.	248
13.6	4 slit interferogram.	248
13.7	75 slit interferogram.	249
13.8	Arrangement for simulation (schematic only).	250
13.9	Sample phase texture in the inertial frame.	252
13.10	Velocity field and speed in rotating frame.	253
13.11	Sample interference pattern vs. Ω for a 10 slit grating.	254
13.12	Maximum sensitivity with respect to Ω (vs. number of slits).	255
13.13	Maximum sensitivity with respect to θ (vs. number of slits).	255
13.14	Interference curves for 6 different gratings from flow simulations.	256
13.15	A schematic SHeQUIG (or SQUIG).	257
A.1	Inner cell piece side view	269
A.2	Inner cell piece (ICP) top and bottom views	270
A.3	D-ring	271
A.4	E-ring	273
A.5	Cell can	274
A.6	Wiring breakout	274
A.7	Sense arm tube flange and side-arm	275
A.8	Sense arm	276
A.9	Sense arm mold	277
A.10	ICP detail view: chip slots	279

A.11	ICP detail view: coordinates	280
A.12	ICP detail view: more coordinates	281
A.13	D-ring detail view	282
A.14	Lead shield cutouts	283
A.15	Three different custom-made mandrels used for off-axis seal polishing.	286
A.16	Cutting indium wire just right.	286
A.17	Thin film chip resistor, leads unsoldered vs. soldered.	288
A.18	Pancake coil glued into E-ring.	293
A.19	Bottom side of D-ring.	294
A.20	Top side of D-ring.	294
A.21	Bottom side of ICP showing triple seal.	295
A.22	E-ring, D-ring and ICP fully assembled.	296
A.23	Fully assembled cell with sense arm.	297
A.24	Tank inductor former for PI circuit.	298
A.25	Inductor former for PI circuit injection line chokes (need 2).	298
A.26	Shielded box for injection line chokes.	299
A.27	Plastic (Lucite or polycarbonate) base for pancake coil winder.	300
A.28	Evaporation mask for diaphragms and electrodes.	300
A.29	Single weak-link cell: pancake coil holder	301
A.30	Single weak-link cell: inner cell ring piece	302
A.31	Single weak-link cell: Modular cell can	303
A.32	Single weak-link cell: wiring breakout (single set of leads)	304
A.33	Single weak-link cell: wiring breakout (double set of leads)	304
C.1	Clousing constant calculation: α vs. u and v	316
C.2	Clousing constant calculation: α vs. v	317
C.3	Mathematica screengrab of code snippets C.2 and C.3.	321
C.4	Example of Mathematica procedure usage and (full=1) output.	322
D.1	Phasor diagram for superposition of oscillations with unequal amplitudes.	327
D.2	Uniform phase gradient set up using superfluid counterflow due to heat current.	329
D.3	Illustration of the “Feynman diagram” approach to writing the asymmetric amplitude for the case of 4 chips.	332
D.4	Theory plots for SHeQUID with 4 chips.	334
D.5	Theory plots for SHeQUID with 10 chips.	335
D.6	Theory plots for SHeQUID with 10 chips (zoomed in).	335
D.7	Mathematica code for applying results in this chapter	336
E.1	Simplified schematic of single weak link cell.	337
G.1	Simple cell schematic for cell dynamics and force diagram for diaphragm.	356
H.1	A snapshot of the system at equilibrium showing the two reference frames.	368

List of Tables

1.1	Glossary of symbols	23
2.1	Mapping JLTP paper to dissertation sections	26
6.1	Some results from four flow test runs.	111
F.1	Cavity resonant modes: SHeQUID 3	349

Acknowledgments

In experimental physics, I have found that there is an insane number of people behind every experiment and I have always wanted to thank these people in some fashion for all their help. I decided long ago that this would be the venue and that my acknowledgments section would *not* be short. My sincere apologies if I forget someone – I really do appreciate all the help and guidance I have received over the years.

My parents, while they have (of course) supported me in everything I did, are also singularly responsible for my seeking a life in science to begin with. To my dear mother, who got me addicted to science fiction and hard science as a child and instilled a deep respect for the profundity (not just the coolness) of it all and to my dear father, who taught me how to tinker (essentially the only seed skills I had coming into experimental physics), who tolerated my many technological failures around the house and who taught me above all to forge my way pragmatically and rationally to my goals (abstract as they may be) – my deep and abiding thanks. I could not have done this without knowing that you were there behind me, every step of the way. I find myself fortunate in having an extended family that is deeply intellectual and whose members have taught me many things over the years, especially in math, science and programming. While naming everyone would be a bit much, I would like to express my gratitude to my sister Smita, who supported me in every way in my freshman year and remains to this day a close friend and respected mentor – thank you for showing me an America that was far more interesting than I'd ever imagined and for never letting me feel homesick.

I owe a deep debt of gratitude to Richard Packard. Thank you for your patience as well as your constant drive for results. I have enjoyed and learned immensely from the many discussions we have had over the years on all sorts of things. Thank you for always being there with encouragement, advice and a steady stream of intriguing ideas; for bringing me into one of the most interesting and challenging fields I could ever have dreamed of working in; and for forging me – hammer and fire – into an Experimentalist. Thanks to Yuki Sato, who trained me in my early years in many of the lab techniques detailed in this dissertation, some of which have been handed down over the years and some of which he himself pioneered. I have enjoyed working with him and discussing physics, philosophy and movies during the tedious times in the lab and I remain in awe of his technical prowess and his penchant for getting results in what appears to be an almost effortless manner (though it obviously isn't). Emile Hoskinson wrote a lot of the software and lab doctrine on which we have built since then and I have learned a tremendous amount from just dissecting his code and poring over his lab notebooks. *In absentia*, he has taught me a lot, just from these notebooks, about what it means to be a physicist and our sacred duty to document our work and even our thought patterns during an experiment as thoroughly as possible. It is safe to say that without his exhaustive notes, a lot of the technology behind creating and using the quantum whistle developed in our group would have simply been lost to the void.

Michael Ray, who worked with me as a post-doc, has contributed a lot toward some of the results presented in this dissertation. At a time when things weren't working so well, I

am grateful that I had him to lean on and keep myself sane with his cheery optimism and can-do attitude. I still marvel sometimes that we essentially dismantled an entire lab, moved our (running) experiment over to it and rebuilt an entirely new setup around it in just under a week – I could never have done it without him. Thank you for the extremely helpful physics discussions and brainstorming sessions – especially in regards to our joint work on the chemical potential battery, and for building and testing the new HRTs, without which we really couldn't have run the experiments in the first place. In more recent times, I have had the pleasure of working with Satoshi Murakawa, who, while visiting our lab from Keio to learn about making SHEQUIDs, taught me quite a few new lab tricks and was extremely helpful in clarifying several concepts for me in the course of some stimulating discussions. I would also like to thank Yashwant Gowda for assisting me quite a bit in our final few experiments during assembly and operation and for keeping things cold while I was busy with other things.

I am extremely grateful to Joseph Kant, who taught me almost everything I know about machining. I do not exaggerate when I say that none of the work here would have been possible without his tutelage in the shop, his extraordinary patience and his proactive attitude in lending his assistance in figuring out how to make a complicated piece with the limited tools at our disposal. A big thanks to all the guys in the main machine shop – Marco, Pete, Warner, Steve, Dave (both of them) and Abel and the guys in the electronics shop – Dave, Gerry and LeVern for always letting me bug you with questions and letting me mooch tools and expertise from you. I've had some of my most instructive discussions at Berkeley just shooting the breeze with these amazing individuals. I cannot imagine doing the numerous taps in my cell pieces without Pete's proselytization of his beloved form taps – they literally changed my life¹.

Our research has required me to reach out to several people for help. It is a testament to the highly collaborative nature of science and the open culture therein, that a large number of people in the department (and beyond) have generously assisted us with crucial things like SEMs, evaporators and other expensive pieces of equipment that we have desperately needed at times. My sincere thanks (and apologies for bugging them so much) go out to Chloe Baldasseroni, Shane Cybart, Stephen Wu and Brian Kessler. My thanks to Aidin, Chris and Benji in the Zettl lab for working with us to cut down on noise above our soundroom² and the folks from the Clarke group (John, Jeff, Steve, Jed, Mike and Sean) and Siddiqi group (Irfan, Vijay, Ned, Andy, Natania and Eli) for numerous discussions, and arrangements for mutual assistance over the years. A special thanks to Prof. Gabor Somorjai (and to Prof. Birgitta Whaley for guiding us to him, among many other things) for helping us resolve some critical issues with our aperture arrays and for teaching us a method for effective cleaning of such devices. I have had the (as I understand it, somewhat rare) pleasure of being able to tap my thesis committee (Irfan and Birgitta) for advice and help over the years. Ty Volkoff and Prof. Alexander Fetter have been extremely helpful for helping me clarify

¹and I'm only partly joking when I say that

²This is not a trivial thing, as we will see throughout this dissertation.

several important concepts relating to the diffraction grating chapter in this dissertation. Richard's former students – Ray Simmonds, Keith Schwab and Steve Garrett, have been very helpful over the years with technical advice (in fact, we owe a lot to Ray's thesis, as our frequent citations demonstrate) and Seamus Davis generously provided us with some very hard to find Cu-Ni capillaries. While our feasibility studies on laboratory detection of enhanced gravitomagnetism (the Lense-Thirring effect) did not make it to this dissertation, I had the distinct pleasure of several meetings and discussions with Martin Tajmar in my early years at Berkeley and came away with a renewed appreciation for the sheer beauty of general relativity.

Making and optimizing aperture arrays has been a significant part of my work here, and I am extremely fortunate to have had the chance to work at the Cornell Nanoscale Facility in “gorges” Ithaca. The highly skilled and helpful staff-members that I had the pleasure of personally working with – John Treichler, Alan Bleier, Rob Ilic, Mike Skvarla, Jerry Drumheller, Garry Bordonaro and Meredith Metzler – have all been invaluable in helping us get our process going and obtaining usable chips. Deirdre Olynick has been a constant source of advice and support in handling nanofab issues on the Berkeley side (at the Molecular Foundry) for all the time that I've worked here.

The Berkeley physics department has some extraordinary staff members who kept us researchers largely insulated from the ravages of bureaucracy. Anthony Vitan, Anne Takizawa, Donna Sakima and all the good folks in the administrative offices in LeConte remain to this day, the most helpful and understanding bunch of people I have had the pleasure of working with. It really helps when my research is the only thing I need to worry about and these are the people responsible for allowing me that focus without unnecessary distractions. A special thanks to Steven from purchasing for always expediting my really urgent and last-minute (aren't they all?) POs.

I have had the privilege of being trained and mentored by a most remarkable set of teachers³ in my undergraduate days at Cleveland State. I am deeply grateful to Dr. Miron Kaufman for taking me under his wing, making sure I received the best opportunities available and guiding me with his matter-of-fact dedication and conscientiousness. Dr. James Lock taught me everything I know about mathematical physics and I find myself, even now, using his (hand-written and photocopied) notes as one of the best reference works in my collection. A special note of thanks to my advisor, Michael Guidry from a summer research internship at UT, Knoxville. This was my first real research project and it was a huge learning experience. His guidance, encouragement and recommendations undoubtedly helped me get to this point. Besides physics, I have benefited greatly from professional relationships with the math department at Cleveland State – specifically, Dr. Keith Kendig for a very early introduction to nonlinear dynamics and chaos in what remains to this day, the most unusual and intellectually stimulating seminar course I have ever seen; Dr. Pratibha Ghatage for a small project we did together that certainly slaked my thirst for rigorous, abstract math; and Dr. Barbara Margolius for mentoring a bunch of us and taking us to math conferences

³I simply cannot, even now, omit their honorifics. Force of habit.

where I first learned how to create an effective talk and present it with a measure of grace. I am grateful to all the wonderful educators in the humanities who strongly encouraged my long time love for writing (without whom, this dissertation might not have been so long and incidentally, useful) – Drs. Carnell (English), Price (Political science) and McIntyre (Philosophy). Thanks also to Dr. Alan McCormack at Portland State, who took a (somewhat bewildered) freshman in a strange new country and showed him a world that was so much bigger and interesting than he had ever imagined.

A huge thanks to all my friends who kept me sane over the years – especially Vivek, Samir, Amin and June and my cheerily insane compatriots from my gaming days. You know who you all are and I thank you from the bottom of my heart.

Chapter 1

Introduction

1.1 Preamble

1.1.1 History and motivations

Here, we present a brief history of the present era of dc-SHeQUIDs that began with Emile Hoskinson's 2005 discovery of the "quantum whistle" in He-II [1]. The concepts relevant to this dissertation will be explained in the introductory material in the latter parts of this chapter.

Continued work by Hoskinson, *et al.* led to elucidation of the current-phase relation of this quantum whistle [2] (that differentiated between whistles caused by the Josephson effect and by phase slippage) and the discovery of various temperature regimes with differing synchronicity in the phase-slip oscillations [3]. Eventually, the first ^4He dc-SHeQUID was demonstrated [4] by measuring the interference pattern due to the Sagnac effect (phase-shifts due to the Earth's rotation) [5] [6]. This may be thought of as a first generation SHeQUID. The first glimpses of resonant amplification were caught during the initial investigations into the so-called "chemical potential battery" [7] - a technique for exciting sustained, stable quantum oscillations rather than the transient techniques used for much of the early whistle work (this would later become a key element in creating a continuously operating SHeQUID).

Sato, *et al.* [8] used this first generation SHeQUID to probe the fundamental link between the hydrodynamic two-fluid model and the condensate order parameter picture for He-II (see Section 1.2) by directly measuring the phase drop due to a heat current generated superflow. The use of a heat current to control the phase in a SHeQUID was a significant breakthrough because it finally provided a *second* phase-shifting influence to accompany the Sagnac effect. A heat current could now be used to cancel out the Sagnac phase-shift via negative feedback to obtain a linearized phase-measurement device. This feedback technique was eventually demonstrated in a static (i.e. manually operated) fashion [9].

Finally, the chemical potential battery (mentioned earlier), paired with a nonlinear amplification using the Fiske effect was used to demonstrate a continuously operating SHeQUID [10]. That leads us to one of the works described in this dissertation where we combine most

of these features together to achieve continuous operation in a SHeQUID while using flux locking to linearize this intrinsically non-linear device in an automated fashion to obtain a practical device that is used to measure time-varying rotation fields [11].

To round out this historical survey, we mention two other parallel developmental paths to the one discussed above:

- A large sense area multi-turn device [12] (with important implications for the design of any SHeQUID, discussed in Section 3.2.2) that can enhance SHeQUID sensitivity or (configured a different way) reduce sensitivity to rotation fields if we are trying to measure some other phase-shifting influence.
- A superfluid *diffraction grating* [13] [14] that can increase phase sensitivity as compared to the conventional two-chip SHeQUID (with limitations and conceptual issues discussed in Chapter 13).

1.1.2 Road map

Chapter 1 will briefly explore the theoretical underpinnings of superfluid ^4He interferometry. That includes the so-called quantum whistle and two of the physical phenomena that couple to the phase of the whistle and which can therefore be measured using SHeQUIDs. We also briefly touch upon a heretofore unobserved phenomenon that could, in principle, be detected using a SHeQUID - the Aharonov-Bohm effect for neutral matter. Most of the information presented in this chapter is already available in several detailed sources, including previous dissertations and review articles that are referenced at the beginning of the relevant sections. Therefore, we will briefly mention old results (when needed) and go into depth only when required for understanding the work presented in this dissertation.

Chapter 2 deals with the chemical potential battery (Section 2.1), its interactions with cell resonances (locking and amplification: Section 2.3) and an initial predictive model for the onset of criticality and the generation of battery states. Together, they form the basis of the *continuously operating*, Fiske-amplified SHeQUID.

The next several chapters delve into the details of designing, building and operating SHeQUIDs (and single weak-link cells). This material may be broadly divided into four overarching categories:

- **Cell**

Chapter 3 focuses on the design of single-weak link cells and SHeQUIDs, including overviews of individual components and discussions about various usable design philosophies as well as design constraints. Chapter 4 delves deep into the fabrication of the most critical ingredient in these experiments - the nanoscale aperture arrays, since there are several important issues and new recipes not discussed elsewhere. Chapter 5 discusses two different designs for and the construction of the superconducting displacement transducer used to detect the quantum oscillations in these experiments. Some (optional) test protocols (independent of the full experimental apparatus) for the

critical components that are especially susceptible to failure (aperture arrays and the flexible diaphragms for the displacement sensor), are described in Chapter 6.

- **Cryostat**

Chapter 7 describes the construction of the (home-built) experimental cryostat. This includes main components, thermometry, wiring (and breakouts) and also structural and acoustic aspects using finite element analysis (FEA) and vibration measurements. An in-depth guide to building, testing and using cryogenic valves (to assist in acoustically isolating the cell from the environment after its filled) appears in Chapter 8.

- **Lab infrastructure/experiment specific installations**

Chapter 9 describes issues and components external to both the cell and cryostat such as: acoustic isolation, design and construction of an ultra-quiet rotation stage¹ and new electronic controllers for interfacing stage motion to data-taking in an automated fashion.

- **Operation**

Chapter 10 contains a detailed description of the operating procedure of a weak link experiment. This includes, among other things, cooling down, cell-filling and cryovalve actuation, testing, calibrations and obtaining data in various experimental scenarios.

Given this background, we then describe some new, recently published results, including a demonstration of real-time tracking of time-varying rotation signals using a computer-driven feedback system in Chapter 11. Noise and drift considerations are discussed in Chapter 12.

A standalone Chapter 13 on the superfluid diffraction grating describes analytical results, numerical simulations and evidence from past experiments to highlight problems with this approach, which was once believed to be a promising road towards improving gyroscope sensitivity.

Appendix A contains engineering drawings for cell components (including descriptions of the fabrication of the structural components and the assembly of the modular SHeQUID and some drawings for a modular single weak-link cell). Appendix B contains selected Matlab code listings that do not appear elsewhere. Appendices C, D, E, F, G and H contain further analytical work (and programming code) behind results used throughout this dissertation.

1.1.3 Control systems

Over the years, group members have created several Labview Virtual Instruments (VIs) to collect data in our experiments, analyze it and present results immediately and also to control the various instruments and systems used to run the experiments. Listings of Matlab code used to do further data analysis (or numerical simulations) have been provided in this dissertation when appropriate. We have endeavored to describe in this dissertation, the

¹This new rotation stage was a critical ingredient in obtaining the main results described in this dissertation.

concepts and techniques embedded within the Labview VIs. However the VIs themselves cannot be properly documented via a print medium such as this. We have therefore freely provided these VIs, accompanied by user manuals, for download from our group website [15]. This website also contains various Matlab and Mathematica scripts written by the author beyond just the excerpts included in this dissertation.

1.2 The quantum whistle

1.2.1 Superfluidity

A more detailed history of theoretical and experimental progress in superfluid helium physics may be found in textbooks like [16, 17, 18]. In this section, we only describe the theoretical aspects that will be required to understand the physics of the quantum whistle at an intuitive level. We will also lay the foundation for the slightly more complicated theoretical framework we will require for a detailed discussion of the chemical potential battery in Chapter 2.

The heavier of the two stable isotopes of helium - ${}^4\text{He}$ liquefies at 4.2 K at atmospheric pressure. Liquid ${}^4\text{He}$, when cooled further down to a temperature of 2.172 K, undergoes a phase transition to a *superfluid* state, which is commonly known as He-II (to distinguish it from the normal liquid above this temperature, which is termed He-I). This transition temperature is known as “the Lambda point” (named after the characteristic shape of the specific heat anomaly that accompanies this transition) and henceforth denoted as T_λ .

He-II is characterized by a flow viscosity that is several orders of magnitude lower than He-I. However, experiments performed to measure this viscosity yielded different results depending on the measurement method used. These anomalous properties of He-II can be explained phenomenologically by the “two-fluid model” put forward by Tisza [19] in 1938 to explain He-II transport phenomena. This model regards He-II as a mixture of two interpenetrating fluids. The “normal fluid” possesses an ordinary viscosity (and obeys the usual Navier-Stokes equations for viscous fluids), while the “super fluid” can flow without viscosity through channels and past obstacles. The normal fluid carries all the entropy in He-II while the superfluid carries none.

Since the two fluids of this model cannot be physically separated, it is not precisely valid to think of some helium atoms as belonging to the superfluid and some belonging to the normal fluid [20, p. 515]. Therefore, the model is most accurately described by considering He-II as being capable of two different motions at the same time; each motion having its own local velocity (\mathbf{v}_s and \mathbf{v}_n) and effective mass density (ρ_s and ρ_n), where the subscripts ‘s’ and ‘n’ denote super and normal components respectively. If \mathbf{j} and ρ are the total current (per unit area) and total density of He-II, we then have: $\mathbf{j} = \rho_s \mathbf{v}_s + \rho_n \mathbf{v}_n$ and $\rho = \rho_s + \rho_n$. In passing, we note that the supercurrent through a channel of cross-sectional area a is thus $I_s = \rho_s v_s a$ and the normal current is $I_n = \rho_n v_n a$.

Viscosity experiments by Andronikashvili [21] (among others) validate this two-fluid model as long as the fluid velocities remain small. At velocities above a *critical velocity*,

the superflow exhibits dissipation by nucleating vortices. The normal fluid can also become turbulent at high enough velocities. This introduces the possibility of interaction between the two fluids [16]. In any case, we assume the two-fluid model holds for the phenomena of interest here.

In parallel with the above hydrodynamic picture, the superfluid component can be described as a macroscopically coherent quantum state using a condensate order parameter ψ that extends over the entire macroscopic volume occupied by the superfluid:

$$\psi = \sqrt{\rho_s} e^{i\phi} \quad (1.1)$$

Here, ϕ is the quantum phase of the order parameter, ρ_s is the superfluid density and both ϕ and ρ_s (and therefore ψ) are functions allowed to vary in space and time. Assuming a spatially uniform density (but *not phase*) for now, we can apply the probability current density operator $\hat{\mathbf{j}}$ to the order parameter in Eq. (1.1) to obtain an eigenvalue \mathbf{j} for this state:

$$\hat{\mathbf{j}}\psi \equiv \frac{-i\hbar}{2m_4} (\psi^* \nabla \psi - \psi \nabla \psi^*) = \left(\rho_s \frac{\hbar}{m_4} \nabla \phi \right) \psi \equiv \mathbf{j}\psi$$

where $\hbar = h/(2\pi)$ is the reduced Planck's constant and we have used Eq. (1.1) to get $\psi^* \psi = |\psi|^2 = \rho_s$. We can identify the eigenvalue \mathbf{j} with the (hydrodynamic) current density of the superfluid: $\mathbf{j} = \rho_s \mathbf{v}_s$, where \mathbf{v}_s is the velocity of the superfluid component described previously. The condensate is thus identified with the superfluid component of He-II and this melding of two viewpoints yields an important relationship between the phase of the order parameter and the superflow velocity (by inspection of the above equations):

$$\mathbf{v}_s = \frac{\hbar}{m_4} \nabla \phi \quad (1.2)$$

This relation was experimentally verified in [8] and an important practical implication for interferometry is that order parameter phase differences (which produce interference patterns) can be manipulated by changing superflow velocities in controlled ways. We will see two examples of this in Section 1.3.

1.2.2 Coupling regimes: the healing length

Near a hard boundary, the condensate order parameter does not die off to zero abruptly, instead decaying gradually over a characteristic length scale known as the healing length ξ_4 (or the coherence length). More generally, the healing length can be defined as the length scale over which perturbations in the order parameter die out and it approaches the bulk condensate value². Superfluidity therefore gets suppressed when He-II is confined in

²Hence the term “healing length”, which may therefore be very crudely thought of as the distance over which the superfluid “heals its wounds” (in a manner of speaking).

geometries with dimensions of the order of the healing length. The behavior of superfluid properties in such confined geometries near the lambda point has been (and continues to be) extensively studied and critical exponents for the characteristic power law divergences of these properties are available in the literature. Fig. 1.1 shows the variation of the healing length over temperature according to the expression shown in the figure, with the critical exponent (0.6717) obtained from recent measurements by Burovski, *et al.* [22].

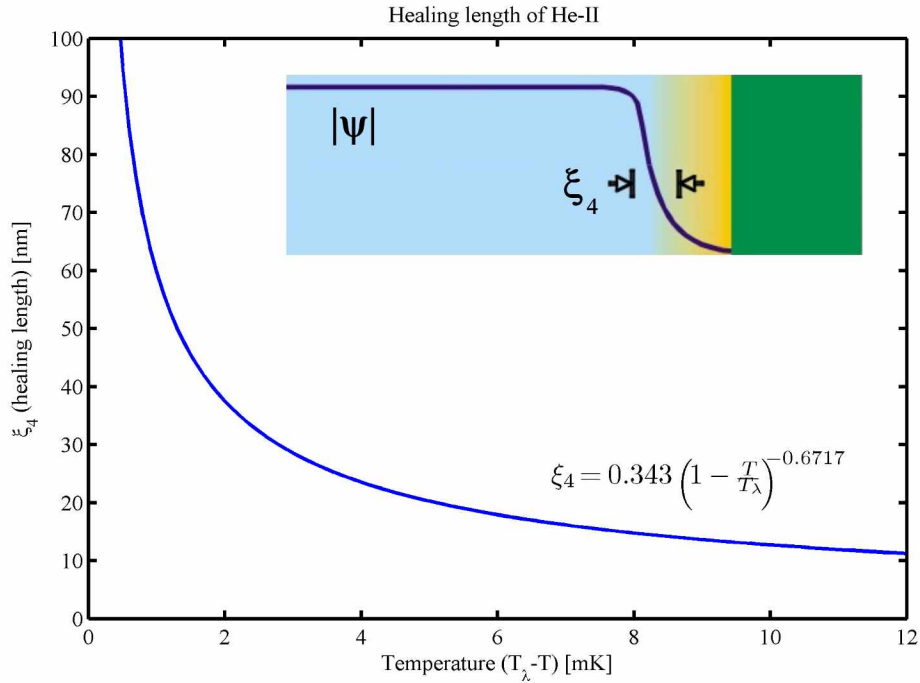


Figure 1.1: He-II healing length as a function of temperature. The expression for ξ_4 gives the value in nm for a given temperature T (where T_λ is the superfluid transition temperature). Inset: Physical interpretation of healing length as the length scale over which superfluidity gets suppressed near a hard wall (see text).

Two coupled superfluids

When two superfluid volumes are separated by a hard wall with an aperture (with dimensions $\sim d$) drilled in it, the condensates in the two volumes can couple to each other in interesting ways. As we saw in Fig. 1.1, the healing length diverges as we approach the lambda point from below. Therefore, at very low temperatures (where $\xi_4 \ll d$), the two volumes are strongly coupled and form one bulk superfluid (and we call this a “strong-link”). At higher temperatures (closer to T_λ), the healing length increases and at some point becomes comparable to the aperture dimensions. The condensate order parameters describing the two volumes now “leak” into each other only weakly since superfluidity is suppressed within

the aperture (and we call this a “weak-link”). This dichotomy is illustrated in Fig. 1.2. We note that the transition from the strongly coupled regime to the weakly coupled regime is gradual and can be parametrized by the ratio ξ_4/d , which in turn varies from ~ 0 to ∞ .

The chief observable consequence of the different coupling regimes is a variation in the relation between the phase difference between the two superfluid volumes and the mass current flowing through the aperture. This relation (called the current-phase relation: $I(\phi)$), changes gradually from linear (strong link) to sinusoidal (weak link), the precise experimental details of which are described in Ref. [2]. In the next sections, we briefly describe the physics of these two regimes.




Figure 1.2: Two different coupling regimes for superfluid volumes separated by a wall containing an aperture with dimension (d) of the order of the healing length. (a) Strongly coupled regime for colder temperatures farther away from T_λ (where $\xi_4 \ll d$) and (b) Weakly coupled regime for warmer temperatures closer to T_λ (where $\xi_4 \sim d$).

Weak coupling: Josephson oscillations

In the weakly coupled regime, the order parameter describing the left side volume (see Fig. 1.3) does not decay to zero until part of it has “leaked into” the right side volume and vice versa.




Figure 1.3: (a) Solid wall. Order parameters go to zero at the wall. (b) A superfluid “weak” link - order parameters “leak” through the aperture.

We can write the Schrödinger equations for the left and right volumes in Fig. 1.3 with a weak coupling term that mixes the behavior of the two sides (see Section G.2 for details) and solve the resulting coupled equations with the ansatz in Eq. (1.1) to obtain the Josephson³ equations (from Eqs. (G.11) and (G.12)) describing the behavior of the supercurrent:

³After Brian Josephson, who first predicted this effect for superconducting weak-links [23]

$$\boxed{I = I_c \sin \Delta\phi} \quad (1.3)$$

$$\boxed{\frac{\partial \Delta\phi}{\partial t} = -\frac{\Delta\mu}{\hbar}} \quad (1.4)$$

The second of the two equations is also known as the Josephson-Anderson phase evolution equation and it can be shown to be valid even in the strongly coupled regime [24]. Here, the chemical potential difference $\Delta\mu$ can be created by a combination of pressure and temperature differences, ΔP and ΔT as [16, pp.79-80]:

$$\Delta\mu = m_4 \left(\frac{\Delta P}{\rho} - s\Delta T \right) \quad (1.5)$$

where ρ and s are the density and specific entropy (per unit mass) of ^4He , respectively. More details on how this $\Delta\mu$ is established in practice are provided in Section 2.1).

If a constant chemical potential difference $\Delta\mu$ is imposed across the aperture, the phase-evolution equation (1.4) can be trivially integrated to get a phase difference evolving linearly in time (with an initial phase difference φ_0):

$$\Delta\phi = -\Delta\mu \cdot t/\hbar + \varphi_0 \quad (1.6)$$

which, in conjunction with the Josephson equation (1.3), implies an oscillating mass current in response to a constant energy difference:

$$I = I_c \sin(-\Delta\mu \cdot t/\hbar + \varphi_0) \equiv I_c \sin(-\omega_J t + \varphi_0) \quad (1.7)$$

where we have defined the Josephson frequency:

$$\boxed{\omega_J \equiv \frac{\Delta\mu}{\hbar} \Rightarrow f_J = \frac{\Delta\mu}{h}} \quad (1.8)$$

We therefore expect to see mass current oscillations sinusoidal in time in this regime at a frequency proportional to $\Delta\mu$.

The Deaver-Pierce model: crossover regime

In direct analogy to the Deaver-Pierce model for superconducting weak-links [25], we can model the current-phase relation as dependent on the ratio of the parasitic (hydrodynamic) inductance⁴ to the ideal Josephson inductance. In other words, we can think of a real weak-link as being composed of an ideal weak-link with a characteristic Josephson inductance L_J

⁴The idea of hydrodynamic inductance is discussed in Section G.2.2. For our purposes here, it is sufficient to think of it in analogy with the inductance in an electrical circuit as the dynamic inertia of the system (how hard it is to change the current).

and a perfect sinusoidal current-phase relation (with phase ϕ , as derived in the previous section) placed in series with a parasitic hydrodynamic inductance L_p associated with flow through the apertures (see Fig. 1.4), so that the total phase across the combination is θ , which is the variable we have access to experimentally.




Figure 1.4: Circuit representation of the Deaver-Pierce model (after [26, p. 92]).

The ratio $\alpha \equiv L_p/L_J$ can be shown⁵ to dial the behavior of the aperture array from a near-ideal sinusoidal Josephson current-phase relation ($\alpha \approx 0$) to an effectively linear current-phase relation where the phase “slips” by 2π from one branch of a multi-valued function to another ($\alpha \gg 1$). The crossover point is at $\alpha = 1$, which is the last point at which the relation is single-valued and roughly sinusoidal. See Figs. 1.5 and 1.6.

Strong coupling: Phase-slip oscillations

For temperatures far below T_λ (where the healing length is much smaller than the aperture dimensions), the two superfluid volumes are strongly coupled into one bulk superfluid. In that case, we can obtain the time rate of change of the superfluid velocity for bulk superflow using Newton’s 2nd law for superfluids (see Eq. (G.17) in Section G.2) over an effective⁶ channel length l_e of the aperture:

$$\dot{\mathbf{v}}_s = -\frac{\nabla\mu}{m_4} \Rightarrow \boxed{\dot{v}_s \approx -\frac{\Delta\mu}{m_4 l_e}} \quad (1.9)$$

which tells us that the superfluid undergoes a constant acceleration for a constant $\Delta\mu$ applied across the aperture. Integrating this equation for a constant $\Delta\mu$ therefore gives a velocity that increases linearly in time.

⁵See Ref. [26, pp. 90-93] for a detailed analysis of the application of this model to superfluid weak-links and Ref. [2] for experimental validation of this model for He-II (reproduced here in Figs. 1.5 and 1.6 respectively). Here, we merely quote and use these results. We use a similar analytical technique in Section 3.2.2 to estimate the effect of the parasitic sense loop path inductance on modulation depth in a SHeQUID.

⁶The geometric channel length is on the order of the thickness of the wall in which the aperture exists. However, changes in the flow field in the vicinity of the aperture should be taken into account since it prevents the phase gradient from being perfectly linear. This information can be folded into an “effective” channel length l_e , which may be very weakly temperature dependent due to perturbations in the flow field.




Figure 1.5: Current-phase relations (for $\alpha = 0, 0.1, 0.3, 0.5, 1, 3, 4.7$) calculated using the Deaver-Pierce model for a single weak-link (from Fig. 3.12 in Ref. [26]). The two distinct regimes (based on whether the function is single-valued or multi-valued) can be seen. The phase θ is the total phase across the combination of the two inductances in series as shown in Fig. 1.4




Figure 1.6: Measured current-phase relations (solid points) for a He-II weak-link and model fits (dotted lines) with fit values of α displayed next to fit curves (from Fig. 4c in Ref. [2]).

However, the superfluid can flow without dissipation only up to a maximum “critical” velocity⁷ beyond which, a quantized vortex (a whirlpool of superfluid surrounding a normal fluid core) is stochastically nucleated. This vortex is pulled orthogonally (see Fig. 1.7) to the flow direction by the Magnus force, whence it bleeds off some of the flow kinetic energy resulting in a velocity drop for the superflow by a fixed, discrete amount (hence the epithet “quantized” for the vortex). A parallel picture describing this event in terms of phase is presented in Ref. [24] where the motion of the vortex across the aperture flow field results in a phase drop along the flow path (across the aperture) of 2π . Given this, we can calculate the drop in the flow velocity (called the “slip size”: v_{slip}) after such a phase-slip event using Eq. (1.2) over a flow channel of effective length l_e :

$$v_{slip} = \frac{\hbar}{m_4} \nabla \phi \approx \frac{\hbar}{m_4} \frac{2\pi}{l_e} = \frac{\kappa_4}{l_e} \quad (1.10)$$

⁷See Ref. [24] for details regarding this as well a more in-depth explanation of how the flow-induced motion of a nucleated vortex during a phase-slip leads to a 2π phase-drop across the aperture.

where $\kappa_4 \equiv h/m_4 \sim \mathcal{O}(10^{-7}m^2/s)$ is called the quantum of circulation. Assuming an effective channel length $l_e \sim \mathcal{O}(10^{-7}m)$, which gives a slip size $v_{slip} \sim \mathcal{O}(1m/s)$ within the aperture.

In terms of superfluid current ($I_s = \rho_s v_s a$) through an aperture channel of cross-sectional area a , the slip size can be written as:

$$I_{slip(1aperture)} = \rho_s a \frac{\kappa_4}{l_e} \quad (1.11)$$

For an aperture array with N apertures, I_{slip} to first approximation would just be N times the above expression, but because of corrections due to more complicated flow fields, is better described by:

$$I_{slip} = \rho_s \beta_s \kappa_4 = \frac{\kappa_4}{L_{ps}} \quad (1.12)$$

where we have used the array inductance definitions described in Eq. (G.21) of Section G.2.3.




Figure 1.7: A strong-link. Bulk superflow through the aperture generated by a force proportional to a chemical potential difference. Superfluid accelerates up to a critical velocity v_c at which point a vortex is nucleated, siphoning off a discrete bit of energy from the flow, thus slowing it down. The vortex moves transverse to the flow and creates a 2π phase drop across the aperture - this event is called a 2π phase-slip.

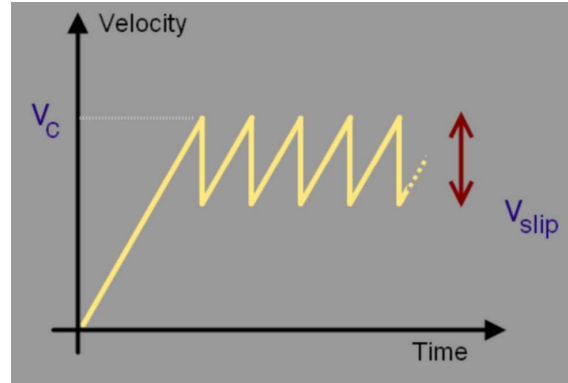


Figure 1.8: Velocity vs. time plot (schematic) for strong-coupling superflow in an aperture. 2π phase-slips occur periodically (see text) as the superfluid continuously accelerates up to a critical velocity v_c and slows down by an amount v_{slip} .

Assuming a constant driving force (i.e. chemical potential difference) is maintained across the aperture, the superfluid continues to be accelerated (linearly in time) up to the critical velocity, undergoes a phase-slip where the velocity drops by v_{slip} and this process repeats indefinitely (see Fig. 1.8 for a cartoon illustration of this process). The supercurrent I_s (proportional to the velocity v_s) is thus a periodic, sawtooth function of time.

We can deduce the frequency f with which 2π phase-slips occur by using the Josephson-Anderson phase evolution equation (Eq. (1.4)) to calculate the time τ needed for the phase-difference $\Delta\phi$ to evolve by 2π :

$$\frac{2\pi}{\tau} = \frac{\Delta\mu}{\hbar} = \frac{2\pi\Delta\mu}{h} \Rightarrow f = \frac{1}{\tau} = \frac{\Delta\mu}{h}$$

Comparing this with the frequency of Josephson oscillations from Eq. (1.8), we see that phase-slip oscillations occur at exactly the same frequency (albeit with a different waveform shape). A sawtooth waveform at frequency $f = f_J$ will have a strong Fourier peak at f_J in the frequency domain - in practice, this is what we can measure.

Conclusions: whistle regimes and operating temperatures

The main consequence of all this for the purposes of interferometry is that in all coupling regimes, we can obtain phase-coherent quantum oscillations (‘whistles’) with a strong Fourier component at the Josephson frequency $f_J = \Delta\mu/h$.

Although we have been assuming a single aperture connecting the two volumes, in practice we use arrays of thousands of apertures. One reason is to amplify the mass current signal to be able to measure it with available sensors. A further reason for using aperture arrays instead of single apertures has to do with thermal fluctuations, which, according to Chui, *et al.* [27] can be strong in a single aperture (and thereby drown out the whistle), but may be suppressed in an array of many apertures. It is not clear at this point of time why such arrays of apertures act synchronously as weak-links. Further, Sato, *et al.* [3] have observed that such synchronicity does not extend indefinitely into the strongly coupled regime, with the result that phase-slip oscillations are *not* always synchronous. Fig. 1.9 (reproduced from Ref. [3]) illustrates this asynchronicity between apertures in an array as an observed drop in the whistle amplitude compared to what it should be for completely synchronized whistling. The predicted rise in whistle amplitude with decreasing temperature comes from increased superfluid density ρ_s (see Eq. (1.12)) but this is eventually overwhelmed by the drop due to loss of synchronicity. An ‘avalanche’ model has been proposed to explain these synchronicity observations by Pekker, *et al.* [28], but will need further testing (of its additional predictions) before it can be validated.

This issue has important implications for the optimum operating temperature of a practical interferometer. It would appear that the top of the broad peak in Fig. 1.9, where the asynchronicity drop just begins to overwhelm the ρ_s rise, is the ‘sweet spot’ with the largest whistle amplitude (higher S/N) and lessened sensitivity to temperature fluctuations. However, there are additional criteria for ‘optimum’ to consider beyond merely these two, including (but not limited to): reasonable transient whistle duration (i.e. lower whistle dissipation, if whistle feedback is used - see Section 2.1) and/or lower drift (implies lower chemical potential battery powers - see Section 12.1). We will highlight such considerations as they arise.

Finally, Fig. 1.10 summarizes the different regimes discussed in this section. We turn now to the process of utilizing these oscillations in an interferometer.




Figure 1.9: Measured phase slip current oscillation amplitude I_{slip} (for $f_j < 300\text{Hz}$) and the expected value for a fully synchronous case I_{slip}^N . The lines are a guide to the eye. Reproduced from Fig. 4 of Ref. [3].




Figure 1.10: Cartoon summary of coupling regimes. Note that the transitions between the different regimes are not sharp as the figure might imply.

1.3 DC superfluid interferometry

1.3.1 A SHeQUID

Given a source of coherent mass current oscillations (henceforth referred to simply as ‘whistles’), we can use two (or more) of them as the coherent sources in an interferometer (analogous to an optical interferometer). The interference pattern in this case is not projected on a screen, but is the result of a coherent superposition of the whistles (each having a well-defined quantum phase) emanating from the aperture arrays. The magnitude of the superposed whistle depends on the relative phase between the interfering whistles. If this relative phase is swept over time, the resultant magnitude will sweep out the interference pattern in time⁸.

To make a 2-slit SHeQUID, we place two aperture arrays in a loop as shown schematically in Fig. 1.11. When the arrays act as weak-links, the currents in the two arrays are given by Eq. (1.7): $I_{c1} \sin(\omega_j t + \varphi_1)$ and $I_{c2} \sin(\omega_j t + \varphi_2)$. The result of their coherent superposition (see Appendix. D for derivations) is given by another sinusoidal oscillation: $I = I_t \sin(\omega_j t + H)$, where I_t and H are respectively, a time-independent amplitude and an overall phase, which depend only on the static quantities I_{c1} , I_{c2} and the phase difference $\Delta\varphi \equiv \Delta\phi_1 - \Delta\phi_2 = \varphi_1 - \varphi_2$ between the two aperture arrays (which is also time-independent). The overall phase H is experimentally unimportant, but I_t is measurable (as the amplitude

⁸Section 13.2.4 has a fuller discussion (“A deeper puzzle” on p. 244) on similarities and crucial differences between optical and superfluid interferometers.




Figure 1.11: A 2-slit SHeQUID. $\Delta\phi_1 = \omega_J t + \varphi_1$ and $\Delta\phi_2 = \omega_J t + \varphi_2$ are phase differences across each aperture array (the phase decreases in the direction of the arrows). X's denote the aperture arrays.

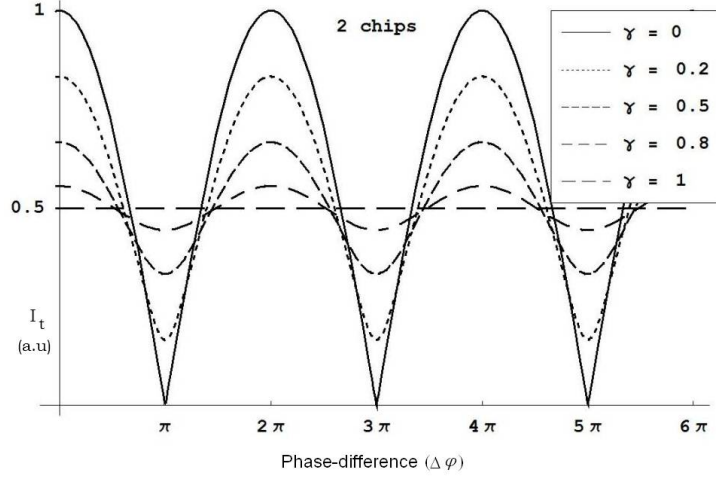


Figure 1.12: Example interference patterns for several values of the asymmetry parameter γ . Without loss of generality, I_{c1} is set to 1, I_{c2} is calculated for each value of γ , and then I_t from Eq. (1.13) is plotted against $\Delta\varphi (\equiv \Delta\phi_1 - \Delta\phi_2)$ after normalizing by the maximum value ($= 2$) in the $\gamma = 0$ case.

of the whistle peak in the frequency spectrum of the oscillation signal) and is given by Eq. (D.25):

$$I_t = I_0 \left[\cos^2 \left(\frac{\Delta\varphi}{2} \right) + \gamma \sin^2 \left(\frac{\Delta\varphi}{2} \right) \right]^{\frac{1}{2}} \quad (1.13)$$

where, $I_0 \equiv (I_{c1} + I_{c2})$ is the maximum current amplitude and $\gamma \equiv \left(\frac{I_{c1} - I_{c2}}{I_{c1} + I_{c2}} \right)^2$ is a parameter that describes the asymmetry between the two aperture arrays (inevitable because it is not practically feasible to fabricate two absolutely identical arrays)⁹.

Asymmetry

As shown in Fig. 1.12, the asymmetry is manifested in a shallowing of the interference pattern (that we frequently refer to as a ‘reduction in the modulation depth’) for γ larger than 0. The ideal case is $\gamma = 0$, which is a perfectly symmetric SHeQUID (with identical aperture arrays) and the other extreme is $\gamma = 1$, which shows no interference at all. In practice,

⁹Notation alert: γ is used in this dissertation in two unrelated places - here, as the asymmetry factor and later, in more technical chapters regarding cell operation and analysis, as a calibration parameter (with a subscript of 1 as a feeble attempt at distinction) used to convert raw signal voltages to mass currents. We do this deliberately to stay consistent with notations used in prior papers, theses and data acquisition and analysis software. Since these quantities do not come into direct contact, it should be clear from context.

this (or a γ very close to 1) is usually a signature of one aperture array completely blocked (by contamination, for example). We have consistently observed that γ (during a given experimental run) increases with temperature. A plausible explanation for this [29] is that closer to T_λ , the qualitative behavior of the aperture array (in terms of which coupling regime it is in) can change by a lot for small variations in aperture dimensions in an array, since the healing length diverges near T_λ (see Section 1.2.2). So, the effect of these small differences can be magnified as subsets of apertures could behave differently (by virtue of being in different coupling regimes). Of course, this idea has not been experimentally verified since one would have to impose a known aperture size variation in the arrays in a very controlled way to induce a predictable asymmetry that could then be measured. Based on the current uncertainties in aperture fabrication, it is unclear whether this is feasible.

Quantization of circulation

The superfluid order parameter (Eq. (1.1)) must be single-valued at every point in space. This implies that going around a closed loop in the superfluid must return the phase to its original value modulo 2π . In other words, the *change in phase* accumulated from going around a closed loop must be an integer multiple of 2π . This phase change can be deduced by computing the path integral of the phase gradient around a loop to obtain the condition:

$$\Delta\phi_{\text{accumulated}} = \oint_{\text{loop}} \nabla\phi \cdot d\mathbf{l} = 2\pi n \quad (1.14)$$

for integer n .

In an inertial reference frame, the phase gradient above can be related to the superfluid velocity via Eq. (1.2) to give the circulation quantization condition:

$$\oint_{\text{loop}} \mathbf{v}_s \cdot d\mathbf{l} = \frac{\hbar}{m_4} 2\pi n = n\kappa_4$$

where (the previously defined) κ_4 is called the quantum of circulation (for reasons that are now abundantly clear). We are careful to inject the requirement of an inertial frame here because in non-inertial frames, velocities must be transformed, thus adding boost terms that do not have to be quantized¹⁰.

We can apply the phase continuity condition in Eq. (1.14) to the SHeQUID loop in Fig. 1.11 while noting that the phase integral will pick up phase contributions ($\Delta\phi_1$ and $\Delta\phi_2$) from each aperture array as well as contributions ($\Delta\varphi_{\text{ext}}$) from any additional physical influence that couples to the phase of the order parameter. Performing the phase-integral this way yields (with $\Delta\phi_1 - \Delta\phi_2 \equiv \Delta\varphi$ as before):

$$\Delta\phi_2 - \Delta\phi_1 + \int_{\text{loop}'} \nabla\phi \cdot d\mathbf{l} = 2\pi n$$

¹⁰see Ref.[30] for a more detailed discussion. We will touch on this issue briefly in the next section (on the Sagnac effect) and return to this issue again in Chapter 13)

where the integral is now taken over the loop length (labeled as loop') excluding the (small) sections that contain the aperture arrays. For low enough superflow velocities, any trapped circulation (non-zero values of n) stays constant while changing the relative phase difference between the aperture arrays ($\Delta\varphi \equiv \Delta\phi_1 - \Delta\phi_2$ as before) and can be folded into a constant phase offset φ_{offset} that merely shifts the interference plots in Fig. 1.11 along the phase axis. $\Delta\varphi$ can be therefore be written as:

$$\boxed{\Delta\varphi = \Delta\varphi_{\text{ext}} + \varphi_{\text{offset}}} \quad \text{with} \quad \boxed{\Delta\varphi_{\text{ext}} \equiv \int_{\text{loop}'} \nabla\phi \cdot d\mathbf{l}} \quad (1.15)$$

can be used to calculate the phase-shift due to the additional physical influences. We know of two such effects that have been measured and one that has been predicted to exist and we briefly discuss them in the following section.

1.3.2 Physical influences

The Sagnac effect

The Sagnac effect is a phase-shift in a quantum-mechanical order parameter as a consequence of being confined to a non-inertial (e.g. rotating) reference frame, such as one attached to the rotating Earth. First postulated for and observed in conventional optical interferometers¹¹, it has since been extended to matter wave interferometers using neutrons, Bose-Einstein condensates and superfluid ^3He and ^4He . A fully relativistic treatment of the Sagnac effect for matter waves can be found in Ref. [32], where we also learn that due to the large rest energies for the particles used in such cases (compared to photon energies in the optical case), a non-relativistic treatment suffices so that we may continue to use the simple order parameter used thus far.

For the superfluid in SHEQUID loop rotating at an angular velocity $\boldsymbol{\Omega}$, its velocity field \mathbf{v}'_{s} at a position \mathbf{r} as seen by an observer in the rotating reference frame is related to the velocity field \mathbf{v}_{s} seen by an inertial observer (in the lab frame) by: $\mathbf{v}'_{\text{s}} = \mathbf{v}_{\text{s}} - \boldsymbol{\Omega} \times \mathbf{r}$. As we saw in the previous section, the superfluid circulation is quantized only in the inertial frame, but the order parameter phase must always return to its original value (modulo 2π) going around a loop. The unquantized contribution to the loop integral therefore comes from the boost term and the external phase-shift (from Eq. (1.15)) becomes:

$$\Delta\varphi_{\text{ext}} = \frac{m_4}{\hbar} \int_{\text{loop}'} (\boldsymbol{\Omega} \times \mathbf{r}) \cdot d\mathbf{l} = \frac{m_4}{\hbar} \int_{\text{loop}'} \boldsymbol{\Omega} \cdot (\mathbf{r} \times d\mathbf{l}) \quad (1.16)$$

where we have used a vector identity to switch the order in the vector products.

We see from Fig. 1.13 that $\mathbf{r} \times d\mathbf{l} = 2d\mathbf{A}$. We also note that for each line-element $d\mathbf{l}$, we have a corresponding (shaded, triangular) $d\mathbf{A}$ and that these area slices together cover

¹¹A full survey is out of place here – a nice historical review with further references may be found in Ref. [31].

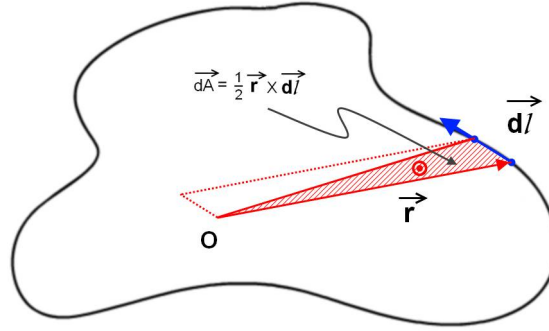


Figure 1.13: Planar loop with arbitrary shape. By inspection, we see that the shaded triangular area is half the area of the full parallelogram, whose area is simply the cross-product of its two sides. Therefore, we must have $\mathbf{r} \times \mathbf{dl} = 2\mathbf{dA}$.

the entire loop area. Therefore, an integral over all such line elements \mathbf{dl} will correspond to a surface integral over all area slices and Eq. (1.16) becomes:

$$\Delta\varphi_{\text{ext}} = \Delta\varphi_{\text{rot}} = 2\frac{m_4}{\hbar} \int_{\text{loop}'} \boldsymbol{\Omega} \cdot \mathbf{dA}$$

Assuming further that the angular velocity “field” is uniform over the loop area (which is the case for the Earth’s rotation field over a sufficiently small SHeQUID loop), we finally have:

$$\Delta\varphi_{\text{rot}} = 2\frac{m_4}{\hbar} \boldsymbol{\Omega} \cdot \mathbf{A} \quad (1.17)$$

where $\boldsymbol{\Omega} \cdot \mathbf{A}$ is informally known as the rotation flux “passing through the loop” (in direct analogy to magnetic flux passing through superconducting SQUID loops, even though in the case of rotation, there are no physically present “rotation fields”).

Figs. 1.14 and 1.15 illustrate the implementation of a SHeQUID configured as a gyroscope to detect the angular velocity of the Earth ($\boldsymbol{\Omega} = \boldsymbol{\Omega}_{\mathbf{E}}$ above). Changing the orientation of the SHeQUID loop area vector \mathbf{A} with respect to $\boldsymbol{\Omega}_{\mathbf{p}}$ (the component of $\boldsymbol{\Omega}_{\mathbf{E}}$ parallel to the ground) changes the rotation flux passing through the loop, thereby changing the phase-shift seen by the SHeQUID. The angular position θ of the experimental cryostat can be changed over a full circle through 360° . As the figures show, $\boldsymbol{\Omega}_{\mathbf{p}}$ always points due North. So, \mathbf{A} pointing North (or South) will result in the greatest magnitude of flux through the loop.

Setting an arbitrary angular position of the cryostat as a zero reference and with θ_{NS} measured to be either of the two positions for which \mathbf{A} lies along the North-South line, the angle between $\boldsymbol{\Omega}_{\mathbf{p}}$ and \mathbf{A} is $\theta - \theta_{NS}$. The rotation flux can then be written as $\boldsymbol{\Omega}_{\mathbf{E}} \cdot \mathbf{A} = \boldsymbol{\Omega}_{\mathbf{p}} \cdot \mathbf{A} = \Omega_p A \cos(\theta - \theta_{NS})$, since the component normal to the surface doesn’t contribute. However, for the sole purpose of staying consistent with previously published papers and theses and all software programs included in this dissertation, we instead define θ_0 as the

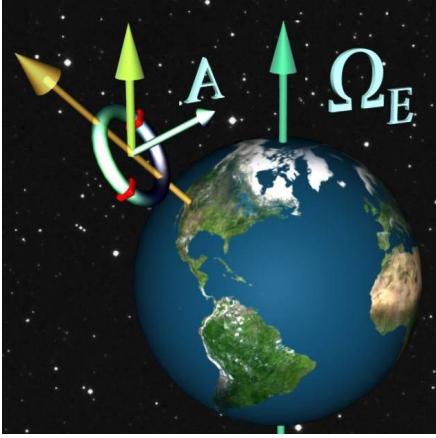


Figure 1.14: Using a superfluid loop (torus) with two aperture arrays (square chips shown in red at opposite ends of a diameter in the torus) as a gyroscope to detect the Sagnac phase-shift. $\Omega_{\mathbf{E}}$ is the Earth's angular velocity vector and \mathbf{A} is the area vector of the loop.

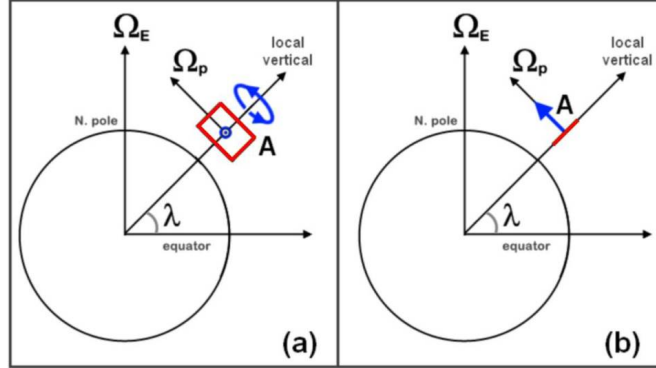


Figure 1.15: Two extreme orientations of the SHeQUID loop area vector \mathbf{A} relative to the Earth's angular velocity vector $\Omega_{\mathbf{E}}$. $\Omega_{\mathbf{p}}$ is the component of $\Omega_{\mathbf{E}}$ parallel to the Earth's surface at the location of the experiment (set by the latitude λ). In a typical reorientation experiment, this parallel component ($\Omega_p = \Omega_E \cos \lambda$) and the area vector \mathbf{A} always lie in the same (horizontal) plane for different orientations of the loop vector. (a) Zero rotation flux: loop vector pointing due East/West ($\mathbf{A} \perp \Omega_{\mathbf{p}}$) (b) Maximum rotation flux: loop vector pointing due North/South ($\mathbf{A} \parallel \Omega_{\mathbf{p}}$)

angular position(s) for which \mathbf{A} lies along the **East-West line**. In that case, the rotation flux becomes: $\Omega_p A \sin(\theta - \theta_0)$. Further setting $\Omega_p = \Omega_E \cos \lambda$ (where λ is the latitude of the experiment location), we finally obtain the Sagnac phase-shift:

$$\Delta\varphi_{\text{rot}} = 2 \left(\frac{m_A}{\hbar} \Omega_E A \cos \lambda \right) \sin(\theta - \theta_0) \equiv 2 c_{\text{rot}} \sin(\theta - \theta_0) \quad (1.18)$$

We can use this phase-shift in the expression for the interferometer amplitude in Eq. (1.13) (with zero asymmetry for visual simplicity) to plot the amplitude against the Sagnac phase-shift (which would look exactly like the $\gamma = 0$ plot in Fig. 1.12) or against the cryostat angular position θ (which is shown in Fig. 1.16 for a typical value¹² of the loop area $A \sim 8.6\text{cm}^2$ and for a latitude $\lambda = 37.9^\circ$). Note that any constant phase offset φ_{offset} in Eq. (1.15) will simply shift the interference pattern along the $\Delta\varphi$ axis in Fig. 1.12 but will change the curve's shape in non-trivial ways in Fig. 1.16 (which is the raw data in an experiment). Therefore, analysis

¹²This value is from an older cell. For the loop used in most of the experiments described in this dissertation, the area is more like $\sim 10.7\text{cm}^2$. See Section 11.1 on “New results” for details. Different areas change the periodicity of the pattern and also affect how many modulation cycles are observed in a full reorientation. The area must be chosen appropriately to be able to observe full cycles.

needs a certain amount of care to allow for this issue. See Ref. [26, pp. 193-197] for a more detailed discussion of the implications of this kind of pattern. Note that $\theta = 0, 180^\circ$ denote an East-West orientation of the loop vector and 90° denotes a North-South orientation. Irrespective of φ_{offset} , the curves are mirror-symmetric around the N-S axis: this fact can be used to determine the local true North direction using the SHeQUID.

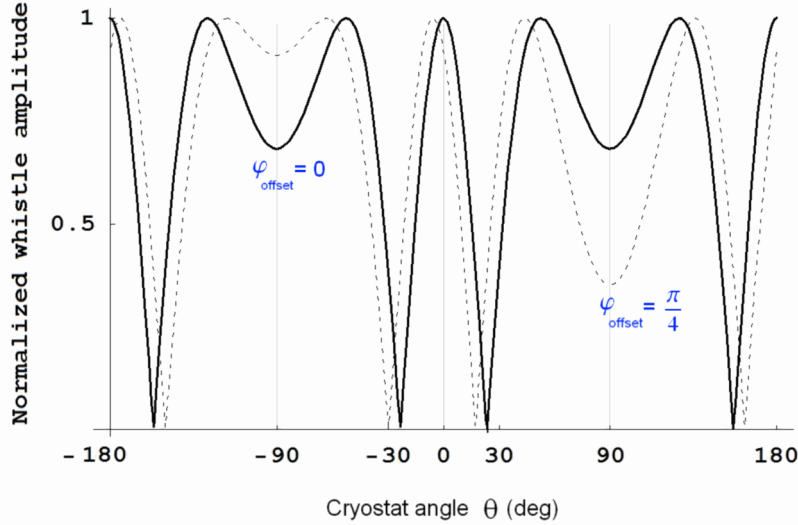


Figure 1.16: Simulated interference pattern for Sagnac effect. Interferometer amplitude plotted against cryostat angular position, with $\theta_0 = 0$ and zero asymmetry ($\gamma = 0$) for simplicity (so $\theta = \pm 90^\circ$ denotes the N-S axis). The solid curve shows a case where the phase offset in the SHeQUID loop is an integer multiple of 2π , while the dotted curve shows the non-trivial change in the curve's shape for other phase offsets ($\pi/4$ in the example shown here).

Heat current driven superflow

According to the two-fluid model mentioned in Section 1.2.1, it is the normal fluid that is responsible for entropy transport (the superfluid does not possess any entropy). In a channel such as the one shown in Fig. 1.17 insulated from the surroundings and filled with He-II, the heater power \dot{Q} dissipated by a resistive heater is transported away from the heater by a normal flow with velocity v_n . The heat flux per unit area (q) from the heater is given by [20, p. 516]:

$$q \equiv \frac{\dot{Q}}{\sigma} = \rho T s v_n \quad (1.19)$$

where σ is the cross-sectional area of the channel (see Fig. 1.21), T is the temperature and s is the specific entropy (per unit mass) of ^4He . The net current in He-II (in one dimension

along the channel) is $j = \rho_s v_s + \rho_n v_n$ and when steady state is reached, this must go to zero. Further, v_s can be written in terms of the order parameter phase using Eq. (1.2) to give:

$$|v_n| = \frac{\rho_s}{\rho_n} v_s = \frac{\rho_s}{\rho_n} \frac{\hbar}{m_A} \nabla \phi \quad (1.20)$$

where $\nabla \phi$ is the phase-gradient induced along the channel because of the heat current.




Figure 1.17: Cartoon depiction of heat-current driven counterflow along a channel. This apparatus is henceforth referred to as a “heat-pipe”. See Fig. 1.21 for a schematic of a SHeQUID containing such a heat-pipe.

Putting Eqs. (1.19) and (1.20) together and integrating the phase-gradient (according to Eq. (1.15)) over a length l of the channel, we obtain an expression for the phase-shift due to a heat current:

$$\Delta\varphi_{heat} = 2 \left[\frac{l}{\sigma} \frac{\pi m_A}{h} \frac{\rho_n}{\rho_s \rho T_s} \right] \dot{Q} \equiv 2 c_h \dot{Q} \quad (1.21)$$

Note that this phase shift (unlike the Sagnac case) is *proportional* to the physical parameter being changed (the heater power \dot{Q} here). The change in heater power required to create a 2π phase-shift across the channel is (by inspection of the above equation): $\dot{Q}_{2\pi} = \pi/c_h$

The Aharonov-Bohm effect for neutral matter (proposed)

The original Aharonov-Bohm (AB) effect [33] predicted that an electron beam split in the classic two-path interferometer scheme in to two beams traveling around a perfect solenoid would exhibit interference effects even though the region accessible to the electrons does not contain any actual magnetic fields. The experimental observation of this effect [34] resulted in a re-evaluation of the “reality” of the magnetic vector potential (which, until then, had been thought of as a purely abstract mathematical construct). After this, a variety of such effects were predicted for charged particles and also neutral particles (and primarily observed with neutron interferometers). The story becomes relevant to our SHeQUID in 1994, when Wilkens [35] predicted a phase-shift for electrically polarized neutral particles traveling in a *radial* magnetic field. Radial magnetic fields are difficult to realize in practice, so the first

practical suggestion came soon after from Wei, *et al.* [36], where an axial magnetic field could be used instead (see Fig. 1.18 for a cartoon depiction of this setup).




Figure 1.18: Depiction of the predicted Aharonov-Bohm (AB) effect for neutral matter. Superfluid helium is confined to the torus. A radial electric field E in the toroidal plane polarizes a helium atom (red-green dumbbell) so that it obtains a dipole moment. Motion of a dipole in a magnetic field B along the axis of the torus creates the AB phase-shift mentioned in the text. See Fig. 3.1 for a model of the proposed experimental cell to test this prediction.




Figure 1.19: Schematic of the Aharonov-Bohm (AB) effect for neutral matter based on Ref. [37]. Radial electric field E is created by putting a voltage V across a cylindrical capacitor with an outer shell of radius b and an inner shell of radius a . This field polarizes the helium atom, whose subsequent motion in a magnetic field B (normal to the plane) creates the AB phase-shift in Eq. (1.22).

A short summary of these developments, a derivation of the AB phase-shift for He-II and a feasibility analysis of detecting such a phase-shift using current SHeQUID technology may be found in Ref. [37]. In this reference, Sato, *et al.* derive the AB phase-shift for a proposed experiment illustrated in Fig. 1.19 and Fig. 3.1, where a voltage V applied across a pair of concentric cylindrical plates (inner plate radius a and outer plate radius b) create a radial electric field that polarizes the helium atoms (where ${}^4\text{He}$ has a polarizability α_{pol}). The dipoles move in the external magnetic field B orthogonal to the plane and this induces a phase-shift (for a single turn torus) given by:

$$\Delta\phi_{AB} = \frac{2\pi \alpha_{pol} B V}{\hbar \ln(b/a)} \quad (1.22)$$

For practically achievable values of the parameters chosen in Ref. [37], with $b/a \sim 1.1$, $B \sim 7 \text{ T}$ and $V \sim 5 \text{ kV}$ (and where $\alpha_{pol} = 2 \times 10^{-41} \text{ F m}^2$), we obtain a phase-shift of $\sim 0.5 \text{ rad}$ with a single turn. This is around 15 times larger than the phase-resolution of our typical SHeQUIDs. Since that feasibility paper was published, multi-turn SHeQUIDs have been successfully tested [12], so that this experiment becomes ever more feasible to perform.

1.4 Physical cells

Up to this point, we have dealt with the superfluid cells in an abstract manner, ignoring the manner in which chemical potential differences are imposed and the methods used to detect the quantum whistles. In this section, we present the archetypal cell schematics for both a single chip cell (Fig. 1.20) and a double chip SHeQUID (Fig. 1.21). These model cells represent (topologically speaking) the actual cells considered (for the respective cell species) in this dissertation. Chapter 3 contains more detailed cell component descriptions.




Figure 1.20: Cell with single aperture array (X) connecting inner and outer cell fluid. Inner cell (pink) is capped by flexible plastic diaphragm (D) coated with superconducting metal. Fixed normal metal electrode (E) is used to exert electrostatic force on diaphragm (also, E-D forms a parallel plate capacitor, which is useful during cell evacuation, filling and calibrations). A superconducting spiral-wound “pancake” coil (PC) is part of a SQUID-based circuit used to detect changes in magnetic flux caused by motion of the diaphragm. R_{in} is a resistive heater used to inject heater power into the inner cell. Refer to Chapter 3 for a more detailed description.

1.5 Flow dynamics

The dynamical equations listed in this section (derived in Appendix G) are valid for both kinds of cells. These cell diagrams and dynamical equations will be useful in the next chapter, which deals with the dynamics of the chemical potential battery.

Table 1.1 contains a list of physical quantities used in these equations and elsewhere in this dissertation (excluding some of the self-contained chapters and appendices).




Figure 1.21: See Fig. 1.20 caption first. Fixed electrode and pancake coil are not shown here for clarity (but are present and identical to those in the single chip cell). R_{in} is (just like the single chip cell) a resistive heater used to inject heater power into the inner cell. Two apertures arrays (X) couple the inner cell to the sense arm, which includes a heat-pipe (see Section 1.3.2) with cross-sectional area σ . Spacing between the vertical side-arms is l . Heat dissipated by a resistive heater R_{sense} flows towards a thin, roughened copper sink (S), creating a superflow towards R_{sense} .

Table 1.1: Glossary of symbols

Symbol	Units (SI)	Description
ρ, ρ_s, ρ_n	kg/m ³	Total, super and normal density
I_t, I_s, I_n	kg/s	Total, super and normal current
v_s, v_n	m/s	Super and normal velocity
η	Pa·s	He-II Viscosity
h, \hbar	J·s	Planck's constant (h) and reduced constant (\hbar)
T	K	Temperature ¹³
$\Delta\mu$	J	Chemical potential difference ¹⁴

... Continued on next page

¹³see Section G.4 for subtleties regarding temperature

¹⁴all differences are [inner cell – outer cell]

Table 1.1 – continued from previous page

Symbol	Units (SI)	Description
ΔP	Pa	Pressure difference between inner and outer cells
ΔT	K	Temperature difference between inner and outer cells
s	$\text{J kg}^{-1}\text{K}^{-1}$	Specific entropy of He-II per unit mass
m_4	kg	Atomic mass of ^4He (6.64647285E-27)
M_4	kg/mol	Molar mass of ^4He (4.0026E-3)
κ_4	m^2/s	^4He quantum of circulation (9.96929973E-8)
κ	1/Pa	He-II compressibility
c_p	$\text{J mol}^{-1}\text{K}^{-1}$	Molar specific heat capacity of inner cell fluid
$c_{p,v}$	$\text{J m}^{-3} \text{K}^{-1}$	Volume specific heat capacity of ^4He ($c_{p,v} = c_p \rho/M_4$)
C_p	J/K	Total heat capacity of inner cell fluid ($C_p = c_{p,v}V_{in}$)
α_p	K^{-1}	Isobaric coefficient of linear expansion of ^4He
R_K	K/W	Effective thermal resistance of inner/outer cell boundary (empirical)
β_n	m^3	Normal flow conductance (empirical)
β_s	m	Superflow conductance (empirical)
\dot{Q}_{in}	W	Heater power dissipated by the inner cell heater
\dot{Q}	W	Generic heater power symbol (depends on context)
A	m^2	<i>Movable</i> ¹⁵ area of flexible diaphragm
k	N/m	Effective spring constant of flexible diaphragm
V_{in}, V_{out}	m^3	Inner and outer cell volumes
V_{SQ}	V	Output voltage of SQUID displacement sensor
Δx	m	Displacement of flexible diaphragm from equilibrium position
α	V/m	Displacement calibration $\Delta x \equiv \Delta V_{SQ}/\alpha$
γ_1	Pa/V	Pressure calibration $\Delta P \equiv \gamma_1 \Delta V_{SQ}$
γ	–	Asymmetry factor in a two-chip SHeQUID
Ω_E	rad/s	Angular speed of Earth's rotation

1.5.1 Dynamical equations

See Appendix G for a more detailed discussion about the physical significance and implications of these equations as well as their derivations.

¹⁵Design area may be different from this because when diaphragm is glued down, the glue can render an unpredictable, roughly annular area at the edges immobile.

The normal current from Eq. (G.1):

$$I_n = \rho_n \frac{\beta_n}{\eta} \left(\frac{\rho_n}{\rho} \Delta P + s \rho_s \Delta T \right) \quad (1.23)$$

The total current from Eq. (G.6) and Eq. (G.7):

$$I_t = \rho \frac{A^2}{k} \Delta \dot{P} = \rho \frac{A^2}{k} \gamma_1 \Delta \dot{V}_{SQ} \quad (1.24)$$

$$I_t = I_s + I_n \quad (1.25)$$

The chemical potential difference from Eq. (1.5):

$$\Delta \mu = m_4 \left(\frac{\Delta P}{\rho} - s \Delta T \right) \quad (1.26)$$

The temperature equation from Eq. (G.30):

$$C_p \Delta \dot{T} = sT \left(\frac{\rho}{\rho_n} I_n - I_t \right) + \dot{Q}_{in} - \frac{\Delta T}{R_K} \quad (1.27)$$

Chapter 2

Continuously operating Fiske-enhanced SHeQUID

The new results described in this dissertation have been compiled into an article that has been published in the Journal of Low Temperature Physics under the title “A continuously operating, flux locked, superfluid interferometer” [11]. Separate sections of this publication are reproduced in this dissertation in an appropriate sequence of chapters. Table 2.1 describes how its contents map to sections in this dissertation.

Table 2.1: Mapping Ref. [11] to dissertation sections

Dissertation section	Publication section	Description
2.1	3	The chemical potential battery
2.3	4	Resonant locking - attractors and repulsors
11.1	5	Results: Continuously operating SHeQUID as a gyroscope
11.2	6,7	Results: Flux locking and linearization; Feedback
12.2, 12.1	8, 10	Noise and drift
9.2, 9.3	9	Low-noise rotation stage, automatic reorientation runs

2.1 The chemical potential battery

As we saw in the introduction (Eq. (1.5)), the chemical potential difference $\Delta\mu$ depends on both pressure and temperature differences, ΔP and ΔT as $\Delta\mu = m_4(\Delta P/\rho - s\Delta T)$. Therefore, unbalanced ΔP or ΔT terms can both result in a non-zero $\Delta\mu$ (and thus a whistle¹ with frequency $f_J = \Delta\mu/h$). In previous versions of the SHeQUID (used as a

¹We have described in some detail the physics of how a non-zero $\Delta\mu$ gives rise to a quantum whistle in Section 1.2. We therefore take that as a given and focus in this chapter on the behavior of $\Delta\mu$ over time.

rotation sensor), the quantum whistle was generated by the application of a pressure step ΔP [1] or a temperature step ΔT [38]. However, the frequency of these whistles decays to zero over time, mirroring a decay in $\Delta\mu$. More detailed descriptions of the decaying whistles using the two different excitation methods can be found in their respective publications cited above and a compiled version may be found in Ref. [39].

The dissipation mechanism is briefly described below for two different whistle excitation methods. We show here how the second method can be modified to obtain continuous oscillations.

2.1.1 Pressure step

When a pressure step is applied across the aperture array by electrostatically pulling on the diaphragm, a quantum whistle is observed with a frequency that decays over time (due to a $\Delta\mu$ that relaxes over time). $\Delta\mu$ relaxes both due to net fluid flow into the cell causing ΔP to relax as well as due to net *superfluid* flow inducing a ΔT that reduces the magnitude of $\Delta\mu$. We can delay this relaxation for a few seconds by continuously increasing the pull on the diaphragm (applying more and more ΔP) to keep f_J constant (i.e. using a feedback routine on f_J).

2.1.2 Temperature step

When a constant heater power step \dot{Q}_{in} is injected into the inner cell (as a sudden step), the temperature of the fluid inside the inner cell begins to rise, creating a temperature difference ΔT across the aperture array (and therefore a $\Delta\mu$). Josephson oscillations are observed, beginning at a low frequency, which begins to increase together with ΔT . Heat is carried out of the inner cell by the normal current I_n and conduction through the cell walls. $\Delta\mu$ drives a net DC supercurrent I_s into the cell, causing a pressure difference ΔP to build, which counteracts the ΔT term in the expression for $\Delta\mu$. The Josephson frequency ($\propto \Delta\mu$) thus rises to a maximum and drops again as the ΔP term catches up to the ΔT term. If this process is sufficient bring $\Delta\mu$ down to zero, equilibrium is reached when ΔP reaches a steady “Fountain pressure” given by $\Delta P = \rho s \Delta T$. At steady state, the net current is $I_t = I_n + I_s = 0$ (so that $I_n = -I_s$ where these are mean, DC values of the currents) and the heater power injected into the cell is balanced by heat flowing out of the cell via the normal flow and wall conduction [38].

2.1.3 Continuous whistling

The whistles created this way are therefore transitory and typically decay in a few seconds (in temperature regimes where the signal is high enough to provide good sensitivity). This transient method of monitoring a phase difference is not optimal since it involves a low duty cycle and requires measuring the amplitude of a continuously changing oscillation frequency. The feedback technique described above (for the pressure-step excited whistle) works in

prolonging the whistle transient but the limited feedback dynamic range is such that it has to be reset often (see Ref. [10] for limit details). Due to the low duty cycle of the transient methods, the signal to noise and the response time are compromised (in practice, thousands of averages were required per point in previous interference experiments that used this feedback method and interference curves could really only be obtained by a point-by-point data acquisition, while sitting at point of constant phase).

A continuous whistle needs a different technique – the so-called chemical potential battery [7]. If the cell heater power \dot{Q}_{in} (in the previously described scenario) is now increased to a new fixed value, the quantities I_s , I_n , ΔP and ΔT all increase towards trying to maintain $\Delta\mu = 0$ in the final steady state. But the supercurrent I_s cannot exceed the critical current I_c . This puts an upper bound on the DC supercurrent and consequently (in steady state) on the normal flow as well, both of which govern the whistle dissipation. Increasing \dot{Q}_{in} beyond this point (which defines a critical value for $\dot{Q}_{in,c}$) therefore leads to a steady state where $\Delta\mu > 0$ and the Josephson oscillations occur continuously without any further decays in frequency. This constitutes a chemical potential “battery” and this state will be henceforth referred to as a “battery state”. In an ideal case, we would be able to conveniently change $\Delta\mu$ by adjusting the heater power in order to set the whistle frequency to arbitrary values. The reality is more complicated and involves phenomena that we describe in following sections.

2.2 Cell resonant modes

The cell supports several hydrodynamic resonant modes, which are essentially standing waves in various cavities. We can calculate these resonant modes for all such cavities that we have been able to identify (using concepts described in Section F.4) for the SHeQUID cell used in this work.

We can try to excite these resonances by driving the diaphragm capacitively (putting an AC voltage with different frequencies across the electrode and diaphragm as described in Section 10.11) and observing the response of the diaphragm (via the displacement sensor). The results are shown in Fig. 2.1 for our SHeQUID #3 (of the form of Fig. 1.21).

Zooming in (Fig. 2.2) shows clearly that the higher temperature curves are shifted lower in frequency. This is consistent with the first sound speed decreasing (see Eq. (F.15) in Section F.4) with increasing temperature, going from $\sim 220m/s$ at 2.15 K to $\sim 218m/s$ at 2.17 K. This is a fractional decrease of $\sim 1\%$. Finding the approximate peak locations in the zoomed plot by eye, we see that they too shift down by $\sim 1\%$.

Note that we observe many more modes (below 6 kHz) than those we predicted from cavity mode calculations in Table F.1 of Section F.4. These remain unidentified. We will return to this sweep data at the end of the next section.

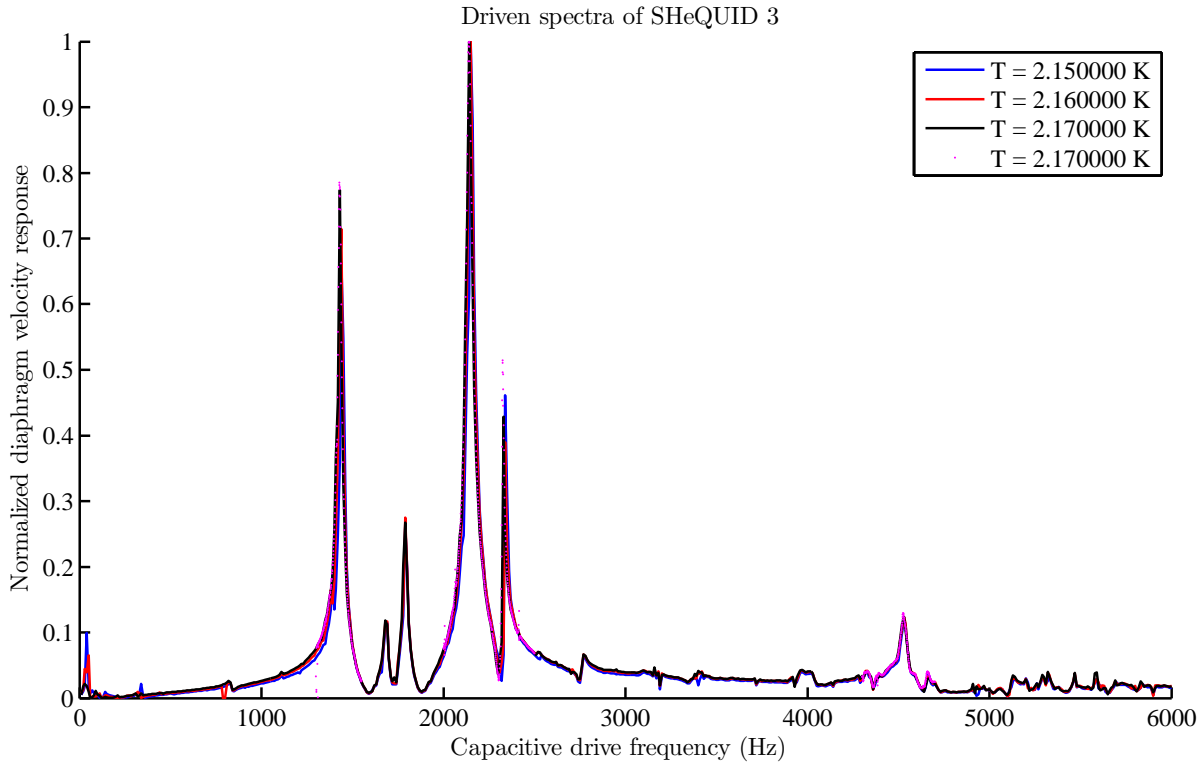


Figure 2.1: Diaphragm (velocity) response as a function of drive frequency for 3 different temperatures.

2.3 Interaction of battery with resonances: Fiske-locking and amplification

The interaction of a battery state with cell resonances involves physics analogous to the Fiske effect in superconducting weak links [40]. During the cell heater ramp-up, when the battery frequency approaches a cell resonance, homodyne mixing of the resonant mode with the whistle provides an additional DC current [41], which could be flowing either into or out of the inner cell, depending on the relative phases of the two oscillations. This is because the DC current enhancement is proportional to the sine of the relative phase [26, p.139], which can be positive or negative. This current combines with the battery driven flow to (respectively) accelerate or retard the battery state's progress in frequency space towards the resonant mode during the heater ramp. The two cases then lead to either an attractor or a repulsor in frequency space.

We have investigated the resonant modes of our cell below 6 kHz. As we saw in the previous section, we can determine some of the cell resonances by exciting the cell with an electrostatically applied harmonic force while monitoring the displacement transducer response. The resonant frequencies we observe with this method are battery repulsors and do

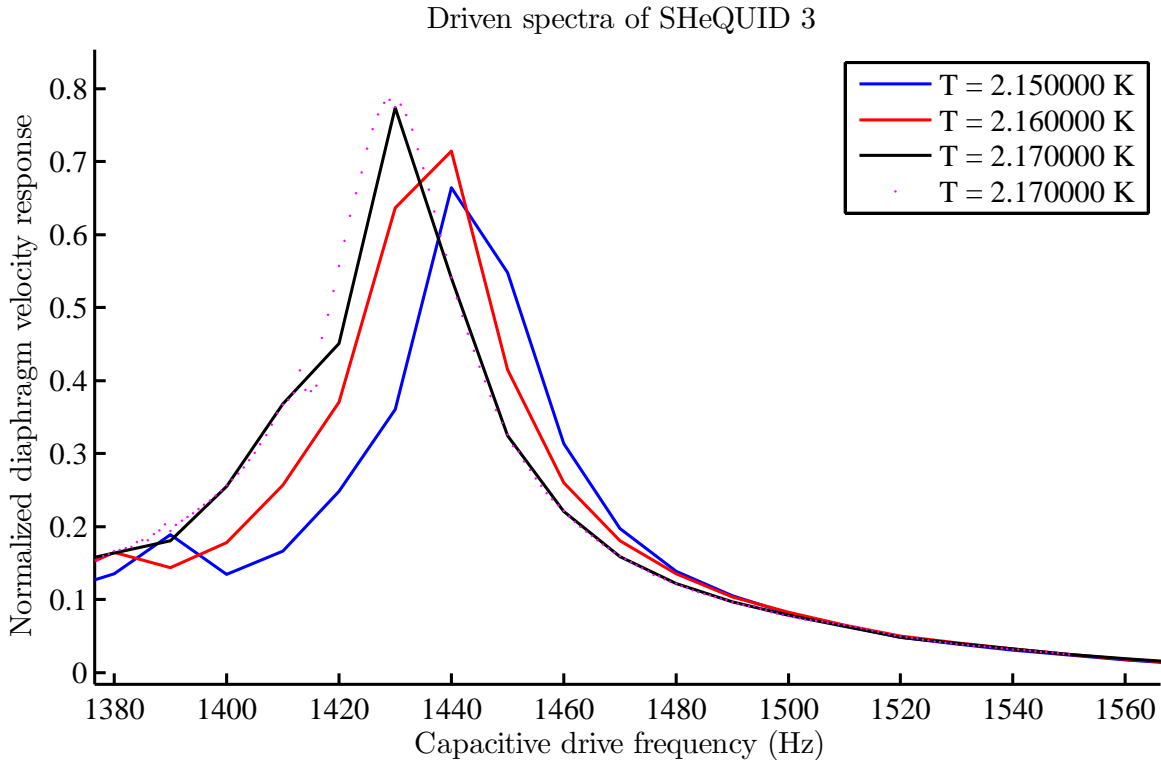


Figure 2.2: Zoomed version of Fig. 2.2.

not seem to match the many battery attractors seen². This may be strongly cell dependent, as Sato, *et al.* found no repulsors and observed attractors at a mix of identified/unidentified cell resonant frequencies³. This is not surprising since the relative phases that determine the sign of the extra DC superflow would depend on cell dimensions and specific details of node/anti-node locations of the resonances.

The attractor cell resonances have the twin advantages of locking the battery at a resonant state with excellent stability and significantly amplifying the whistle amplitude. The unfortunate aspect of the cell resonances is that the resonant amplification (referred to henceforth as the “Fiske gain”) is strongly frequency dependent. For example, the interference patterns in Fig. 11.1 and Fig. 11.2 of Chapter 11 were taken at the same temperature but used different battery states (1080 Hz and 2507 Hz respectively) and have maximum amplitudes of 29 and 55 ng/s owing to the differing Fiske gains. We explore the implications of this issue further in Section 12.3.

These resonant behaviors also lead to complex hysteretic behavior, where the equilibrium battery frequency attained depends not only on the inner cell heater power (\dot{Q}_{in}) but also

²One exception to this rule is the so-called Helmholtz mode (the fundamental mode for small oscillations of the superfluid in the apertures), which does behave as an attractor.

³Yuki Sato, personal communications (pertaining to experiments in Ref. [10]).

on the *details of the path* followed to get to that value. We are thus far unable to predict the spectrum of battery frequencies $f(\dot{Q}_{in})$ as such an analysis requires a more detailed understanding of the cell dynamics. We note that none of these nonlinear mixing and hysteresis effects prevent one from actually *using* these phenomena for practical purposes, though an improved understanding would undoubtedly enable even greater ease of operation. For instance, being able to engineer a sharp, resonant attractor in an otherwise clean frequency regime would simplify the techniques used and improve stability.

2.3.1 Battery state data




Figure 2.3: An example of a typical battery ramp-up (variable rate) and corresponding recorded whistle frequencies (taken with SHeQUID#3 – 7/14/12 at 2.175 K). Black scatter plot is the whistle frequency data and blue line is the battery heater power. The stable, thick bands are the battery states (the FFT routine sometimes loses lock because it is programmed to find the biggest peak in a given frequency range and any noise or acoustic spikes can temporarily kick up nearby resonances and confuse the routine).

Fig. 2.3 shows example whistle frequency data from a cell heater ramp-up. Looking at the data, the system only seems to allow certain discrete frequency values for the whistle in a stable or metastable battery state⁴. We can compare the discrete frequencies that the whistle will stay at with the resonant peaks from the sweep shown earlier in Fig. 2.1. The band near ~ 1450 was essentially the only discernible battery state matching one of the peaks (the lowest frequency large peak) observed in the resonance sweep (and it was not stable over time). After that point, the battery state frequencies *consistently* miss the resonant peaks

⁴even though we can observe the whistle transitioning continuously between states.

by a large distance in frequency space (the small temperature difference between the sweep run and battery run shown is negligible in terms of frequency shift so that doesn't explain the misses). Further, the frequencies that the whistle *is* attracted to (and sticks to) do not match any of the peaks found in the resonant sweep.

One speculation is that the battery sticks at frequencies that are between two first sound resonances, trapped because of the direction that the extra DC currents flow for the two flanking Fiske modes (in that they both push away and the state finds an uneasy equilibrium in between). Another possibility is that there are other kinds of resonant modes that we are just not detecting with the capacitive drive technique and this is what the whistle is more strongly influenced by for some reason. This is a heretofore unresolved issue about the battery.

2.4 Mathematical model

We turn now to a crude mathematical model of the battery that can predict (transient) critical powers with some limited success. Predicting steady-state criticality awaits the development of a more sophisticated model that includes a detailed description of the cell resonances and their interaction with the battery state. However, as we shall see, when the battery state is obtained with sufficiently slow heater power ramps, this distinction tends to disappear and even transient battery states get locked onto cell resonances in extremely long-lived bound states. The lowest frequency attractor, which is invariably the Helmholtz mode, is sufficiently strong that hitting criticality at ramp rates of a few nW/s or less tends to be essentially permanent (at least over several hours). The primary utility of this predictive model is discussed in the final paragraph of this chapter.

2.4.1 Strategy

Since the whistle interacts with cell resonances as soon as it is born, we start with the assumption that the system is sub-critical (so that the chemical potential difference $\Delta\mu \approx 0$) and the inner cell heater power \dot{Q}_{in} is ramped up slowly enough (as described in Section 2.1.3) that the process is quasi-static and $\Delta\mu$ stays approximately zero during the heater ramp.

Given this assumption, we are willfully cutting ourselves off from the critical regime where the whistle exists. However, we can at least solve the flow equations to make a definite prediction of the heater power $\dot{Q}_{in,c}$ for which the supercurrent equals some critical current I_c .

We will find that this predicted critical power depends linearly on the heater power ramp rate. We will check this trend against actual data to see whether our crude model is of some value. It is important to note here that the advent of criticality makes our main assumption ($\Delta\mu \approx 0$) invalid, so that we can say nothing about whether this critical state is stable or merely a transient phenomenon (since we will, at some point need to stop ramping the heater and maintain it at a stable value).

The reason we find this predictive model to be of some use (in practice), is that given sufficiently small ramp rates⁵, the critical state remains stable. This may be related (as mentioned previously) to the Helmholtz attractor, since that is the initial frequency of the whistle when the system goes critical.

Our strategy in the following section is to eliminate dynamical variables in favor of the pressure difference $\Delta P(t)$, which is easily related to the diaphragm mean position and hence the displacement sensor mean voltage.

2.4.2 Pressure equation

Refer to Section 1.5.1 or Appendix G for the set of dynamical equations that govern this system and Table 1.1 for symbol definitions.

With the quasi-static, subcritical assumption ($\Delta\mu \approx 0$), Eq. (1.26) simplifies to:

$$\Delta T = \frac{\Delta P}{s\rho} \quad (2.1)$$

Using this in Eq. (1.23) for the normal current I_n , we obtain:

$$I_n = \frac{-\rho_n \beta_n}{\eta} \Delta P \quad (2.2)$$

Substituting Eq. (2.1) for ΔT , Eq. (2.2) for I_n and Eq. (1.24) for the total current I_t into the temperature Eq. (1.27), we finally obtain a differential equation for ΔP as a function of time:

$$\boxed{a\dot{\Delta P} + b\Delta P - \dot{Q}_{in} = 0} \quad (2.3)$$

where $a = \left(\frac{C_p}{s\rho} + sT\rho\frac{A^2}{k}\right)$ is a capacitance-like term (units of m^3) and $b = \left(\frac{1}{s\rho R} + sT\frac{\rho\beta_n}{\eta}\right)$ is a conductance-like term (units of m^3/s). We see that it is a linear, inhomogeneous differential equation with the (user-controlled) inner cell heater power function as the inhomogeneous term.

2.4.3 Solution for linear ramp

We can solve this equation exactly⁶ for a heater power that is ramped linearly in time at a rate r (with $\dot{Q}_{in}(t) = \dot{Q}_0 + rt$). The solution, for an initial pressure $\Delta P = \Delta P_0$ at $t = 0$, is:

⁵In practice, it is sufficient to ramp up faster until you gets to within a few μW of the critical power and then ramp more slowly to avoid metastable states.

⁶Laplace transforms work easily here, especially for the more generalized heater power functions that we might want to apply to the system. Alternately, we can recognize this to be in the form of the Riccati equation, which has a standard solution that can be looked up in mathematical tables. The solution is a closed-form for (at least) the special case of a ramp function that is polynomial in time.

$$\Delta P(t) = \left(\Delta P_0 - \frac{\dot{Q}_0}{b} + \frac{ra}{b^2} \right) e^{-\frac{b}{a}t} + \frac{\dot{Q}_{in}(t)}{b} - \frac{ar}{b^2} \quad (2.4)$$

where the exponential term dies off fairly quickly with a time constant $\tau = a/b$, which is on the order of half a second for the cell parameters used here. Note that this looks like an RC time constant with $R \sim 1/b$ and $C \sim a$, which is why we called b the conductance term and a the capacitive term.

Note that this is not an equilibrium solution since the ramp is still ongoing. We can construct various piecewise continuous (or even piecewise smooth) ramp functions that level off after a chosen duration (after hitting criticality), but this doesn't help much since the resonant interactions still need to be modeled before we can have an equation that is also valid in the critical regime.

From Eqs. (1.25), (1.24) and (2.2); the supercurrent I_s can be written as:

$$I_s = I_t - I_n = \rho \frac{A^2}{k} \dot{\Delta P} - \left(\frac{-\rho_n \beta_n}{\eta} \Delta P \right) \quad (2.5)$$

Using our solution from Eq. (2.4) in the above equation, we obtain:

$$I_s = \frac{r}{b} \left[\rho \frac{A^2}{k} - \frac{\rho_n \beta_n a}{\eta} \right] + \left[\frac{\rho_n \beta_n}{\eta} \frac{\dot{Q}_{in}(t)}{b} \right] + \left[\left(\frac{\rho_n \beta_n}{\eta} - \frac{b}{a} \rho \frac{A^2}{k} \right) \left(\Delta P_0 - \frac{\dot{Q}_0}{b} + \frac{ra}{b^2} \right) e^{-\frac{b}{a}t} \right] \quad (2.6)$$

The first term is a constant in time, while the second term goes exactly as the heater power function (not universally, just for the specific case of a linear ramp). The third term is a transient that dies off quickly (as mentioned before, the time constant is on the order of half a second for the parameters used in the experiment whose data we will compare this result to) just after the ramp first begins at $t = 0$ (since the t that appears in the exponential is the time measured from the beginning of the ramp - for typical ramps, this can be on the order of several hours). We will therefore ignore the transient term for the analysis that follows. In retrospect, we could have dropped the term prior to this, but it makes sense to ensure that the time derivative does not contribute anything from the transient term.

Now, suppose that the flow hits the critical velocity so that $I_s = I_c$ (for some critical current I_c). We can estimate I_c from the critical velocity measured from single phase-slips. However, there is a geometric factor owing to the array inductance corrections so that this is strictly only an order of magnitude estimate. Regardless, we can make some tests of our predictions even without knowing the exact I_c . With $I_s = I_c$ in the solution above, we can solve for the heater power at which the flow goes critical and call it the critical power $\dot{Q}_{in,c}^d$.

$$\boxed{\dot{Q}_{in,c}^d = \left[I_c b \frac{\eta}{\rho_n \beta_n} \right] - \left[\rho \frac{A^2}{k} \frac{\eta}{\rho_n \beta_n} - \frac{a}{b} \right] r} \quad (2.7)$$

where the superscript d is to clarify that it is a dynamic quantity and might not (and indeed, *should* not⁷) necessarily be *the* stable critical power; in the sense that if we stopped the ramp when the flow went critical (i.e. when we heard a whistle at the Helmholtz frequency), the whistle might not be in a stable battery state and might decay into nothingness after a while. For fast ramp rates, this might just be a truly transient whistle (the limiting case of our usual heater step transients⁸, which are just ramps with arbitrarily large ramp rates). For slower ramps, if the battery state goes high enough to get locked on to an attractor before the ramp is turned off, a relatively long-lived metastable battery state can form, which decays slowly (and in steps as it hits other resonances on the way back down) over several hours. The truly insidious aspect about the battery is that when we are near criticality, a small injection of energy into the system (most commonly, a sharp acoustic spike from outside the Dewar, but sometimes even a faint structural cracking from thermal relaxations in the cryostat) can push the system over the threshold and excite a metastable battery state. It is very difficult to distinguish it from a stable state because the decay times (including the effect of resonances) can be quite long. As stated earlier, the system looks quite simple to model, but only if we ignore the resonances.

Given all these caveats, we can think of the first term in Eq. (2.7) as the “true” critical power (were it possible to ramp up the power at very very slow rates). This is therefore a (very crude) model for predicting critical powers for a given cell and at a given temperature, which should be useful while designing the cell and choosing values for the inner cell resistor.

2.4.4 Experimental tests

Despite the caveats mentioned in the previous section, we can check the prediction in Eq. (2.7) against data from a previous experiment⁹ (when this prediction was not available). The cell power was ramped up to a value a few μW short of the critical value and the cell pressure allowed to stabilize. After this point, the power was ramped up linearly at a fixed rate (r) and the critical power $\dot{Q}_{in,c}^d$ noted (manually) as the power at which a reasonably stable whistle is observed (usually at the Helmholtz frequency). There is some error involved in this because the determination of whether it is “reasonably” stable is a subjective one. Nevertheless, we observed that this ambiguity is considerably less pronounced for temperatures farther from T_λ than for those closer to T_λ . This is reflected in the comparison between prediction and data.

The power is ramped back down to the subcritical intermediate value from before and ramped up again at a different rate. At each stopping point, the cell is allowed to relax to stable pressures (signaled by stable DC displacement readings). The critical powers for

⁷This is because at equilibrium, without resonances, there should be a unique critical power that is determined by the balance of heat flows into and out of the cell, by simple conservation of energy. However, the resonantly pumped metastable states create the illusion of multiple stable battery states and hide the “true” one.

⁸see Section 2.1.2

⁹We are grateful to Michael Ray for collecting a large part of the ramp rate data used in this section.

various ramp rates (for a fixed temperature) are plotted against the rates and a straight line fit to the data. The slope of this line should be the factor multiplying r in Eq. (2.7) and the vertical intercept, the first (constant) term proportional to I_c . We show the results of this analysis in Figs. 2.4–2.9. The black squares are the values calculated from the linear (in r) function of Eq. (2.7) with the intercept replaced by the fit intercept to the data. The reason for not using the predicted intercept is the lack of a precise value for I_c for the intercept calculation.

We note that the slope prediction is spot-on for colder temperatures but gets worse as we go farther from T_λ . As mentioned earlier, we attribute this to the increased hysteresis closer to T_λ . An important reason for all this is that the change in power required to change the battery frequency (and hence $\Delta\mu$) is observed to get smaller as we get closer to T_λ . This is reflected in the cell power needed to hit criticality (observed as a function of temperature). This implies that smaller and smaller disturbances (acoustic energy spikes, temperature spikes, etc.) can push the system past critical where it can get stuck on resonant attractors to give long-lived metastable battery states (as we get closer and closer to T_λ).

Finally, we note that the data is sparse because the goal of those ramp rate runs was simply to find the limiting value of the critical power if the ramp rate were extrapolated to zero (for a cruder, steady-state model). If the goal were a more thorough test of the predictions in this chapter, more data would be needed for a larger set of ramp rates. As it stands, all we can safely say is that our model can give us at least a good idea of what to expect as the critical power for a given temperature. This is useful in at least two ways: choosing the correct inner cell heater resistance optimized for the power ranges we will be needing, and knowing how far one can ramp the power at a fast rate before one needs to slow down the rate to avoid metastability. From a practical standpoint, this is not trivial, as ramping up at “safe” rates all the way from zero to the final battery power can take up to ten hours, a significant chunk of available time between helium bath transfers. We suggest starting a power ramp as soon as possible after a transfer up to the intermediate point so that the cell can stabilize and the cryostat quiet down at around the same time.




Figure 2.4: Critical power vs. ramp rate: $T = 2.170$ K




Figure 2.5: Critical power vs. ramp rate: $T = 2.171$ K

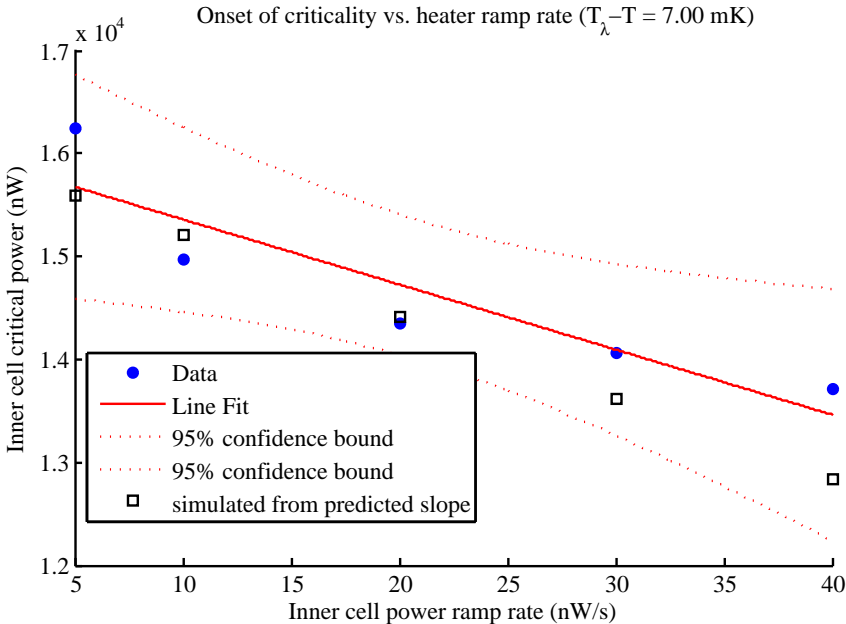


Figure 2.6: Critical power vs. ramp rate: $T = 2.172$ K

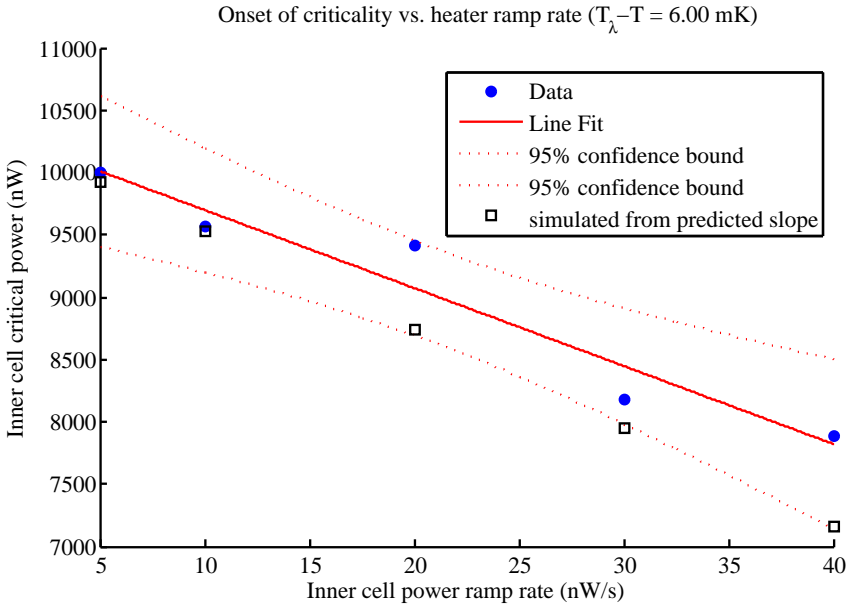


Figure 2.7: Critical power vs. ramp rate: $T = 2.173$ K

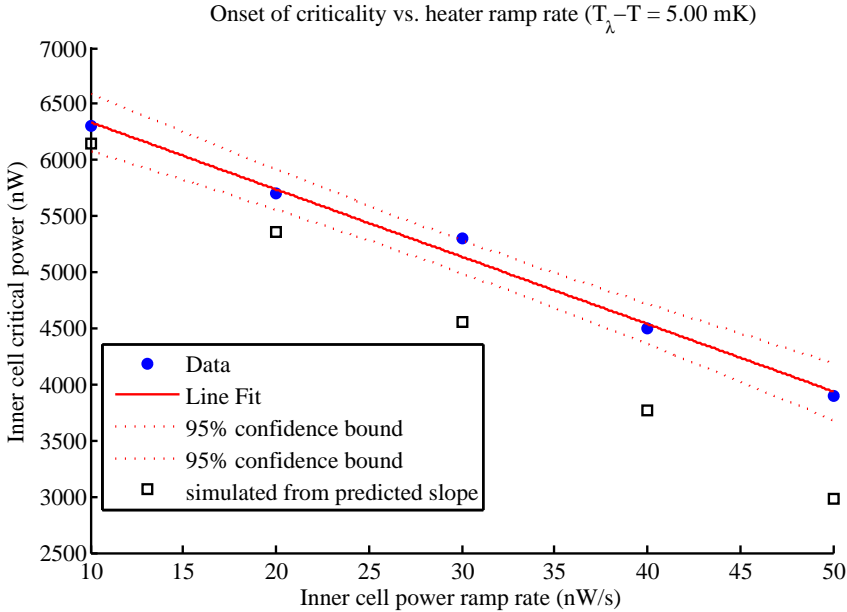


Figure 2.8: Critical power vs. ramp rate: $T = 2.174$ K




Figure 2.9: Critical power vs. ramp rate: $T = 2.175$ K

Chapter 3

Weak link cell design, components and constraints

Superfluid weak link experiments generally require a common set of modules that can be re-purposed towards various goals. This chapter attempts to explain their individual designs in a way that might help the reader abstract the overall design of the experiment in terms of these basic modules. The chapter ends with a *specific* design for the modular double weak link SHeQUID used in the interferometry experiments in this dissertation. Unless specified otherwise, the term SHeQUID will henceforth refer exclusively to double weak link interferometers.

Figs. 1.20 and 1.21 should be used as schematics for the descriptions in this chapter. Scale drawings of the actual cells can be found in Fig. 3.8 and Figs. 3.10–3.11 for the single weak-link cell and SHeQUID respectively.

3.1 Component overview

Recall (Section 1.2) that the quantum whistle requires two volumes of superfluid helium connected only by a weak link. These two volumes are the inner and outer cell volumes. Since the whistle can be thought of as being generated *in* the weak links, these volumes can be connected together by multiple paths, with each path having one weak link. The two path case was explained in Section 1.3.

Now, we need a way to excite the whistle and a way to detect it. A microphone with a flexible (metallized) diaphragm that is driven by the whistle performs the detection. The excitation is provided by an electrostatic force between the flexible diaphragm and a fixed electrode when a voltage is applied between them. The whistle can also be excited by heater power injected into the inner cell via a resistive heater. This is the essence of a single weak link cell used to explore the dynamics of the quantum whistle.

In the case of two (or more) weak links used to create an interferometer, the whistles generated at the two weak links interfere in the inner cell, which forms one common end of

the paths. The resultant oscillation is detected by the microphone. An additional module (the “sense arm”) is used to define the closed sense loops on the other end of the paths and to expand out the “sense area” defined by the loops. Devices to impose phase changes in the interferometer (such as a resistive heater to create the phase gradient in Section 1.3.2) can also be part of this module.

The following is a summary of the various components needed for a weak link cell, with references to fabrication details and other issues (provided in later chapters). Important constraints and other issues will be fleshed out in more detail in later sections.

3.1.1 Aperture arrays

See Chapter 4 for fabrication details and a more rigorous analysis of issues that arise and the resulting constraints on usable fab processes and dimensions. The heart of a weak link cell is the actual weak link in the form of a 3 mm square silicon chip with a thin (~ 60 nm), freely suspended, silicon nitride (just “nitride” henceforth) membrane in which an array of nanoscale (~ 70 nm) holes is shot using various clean-room techniques. These are called “aperture arrays” henceforth, because the term “weak link” refers to just one possible temperature regime. We can (and do) use the interferometer even in the “strong link” and “crossover” regimes mentioned in Section 1.2.2.

Aperture size

The size of the apertures and the thickness of the nitride membrane is determined by the healing length (see Fig. 1.1) of superfluid helium near T_λ (~ 1 mK away, this length is ~ 60 nm, dropping to ~ 10 nm as we go colder to ~ 12 mK away). For aperture dimensions larger than this, one would need to work at temperatures even closer to the lambda point to enter the weakly coupled Josephson regime and temperature regulation (with tens of nK stability or better needed for weak link experiments) becomes more and more difficult, even with the simple pump-bath cryostat we use. For interferometry, being in the Josephson regime is not necessary since we can obtain coherent whistles even in the strongly coupled regime. However, the larger the apertures, the greater the normal flow through the apertures, which leads to faster decays of generated whistles in any kind of transient methods. This makes observing the whistle quite difficult because the whistle frequency will be changing too rapidly for there to be even one pure frequency cycle in a given transient.

In the chemical potential battery used to build a continuously operating SHeQUID (Section 2.1), we run into different (but related) problems, where significantly more heater power would need to be injected into the inner cell (the larger the aperture diameter) to reach criticality and get a stable whistle. These larger cell powers lead to larger phase drifts in operation, so it is imperative that the inner cell heater powers be kept as low as humanly possible (see Section 12.1 for a discussion about this issue). For a practical SHeQUID, this inevitably leads to a need for fabricating apertures with as small a diameter as feasible, while still keeping the total exposed area of the apertures about the same as (or higher than) the

present values (to prevent the total signal current from dropping too low and being drowned by noise). The values used in our lab are ~ 73 nm diameter apertures in a 50×50 grid spaced $3 \mu\text{m}$ apart on a single chip. The typical current amplitudes observed for these arrays (in a two-chip SHeQUID) are $10 - 100 \text{ ng/s}$ with best-case signal to noise ratios of $\sim 10:1$. This baseline can be used to scale down to the single aperture and further, to calculate the approximate signal sizes for different aperture arrays.

Aperture spacing

The array spacing is not terribly crucial, with spacings of 1, 2 and $3 \mu\text{m}$ successfully used in the past with no observed effect on SHeQUID characteristics. However, for spacings less than $1 \mu\text{m}$, we do run into issues with fabrication (not fatal, but proximity effects become important for e-beam lithography and must be taken into account). It is also possible that when spacings start being comparable to hole sizes, coupling between apertures may become important when it comes to the coherence of the phase-slip oscillations described in Section 1.2.2 (see Ref. [3] for the experiments and Ref. [28] for a model that explains the experiments based on coupling between neighboring apertures). However, this is largely speculative at this point of time and could be verified by testing the even more extreme array spacings and distributions described in Ref. [28].

Number of apertures

Having set the typical aperture dimension (diameter and depth), we run up against limits on nitride membrane size, which further sets limits on how many apertures we can have on a chip. This limit on membrane size stems from the breaking stress of a freely suspended thin film for impressed pressure differences across the membrane. The pressure differences arise during evacuation of the cell before cooling down and while handling the chips during fabrication processes. Membranes should typically be strong enough to withstand at least an atmosphere of differential pressure with ease. In practice, we have found that membranes $\sim 200 \mu\text{m}$ square are extremely robust and quite difficult to break using differential pressures of several bar. Membranes $\sim 400 \mu\text{m}$ square, *when processed correctly* (details in Chapter 4), are likewise robust. However, above this membrane size, they start becoming weak enough to break easily during fabrication steps, in addition to other problems like wrinkling, which we discuss in more detail in Chapter 4. Suffice to say that we have found $\sim 400 \mu\text{m}$ to be a practical limit on membrane size.

The spacing and membrane size limits defines a limit on the total number of apertures we can have on a single chip. Based on the above discussion, this also limits how small the apertures can be and still obtain sufficient signal to detect the whistle. The total exposed area with a 50×50 grid of $\sim 73 \text{ nm}$ diameter holes at $3 \mu\text{m}$ spacing is $\sim 1 \times 10^{-11} \text{ m}^2$. As an extreme case, if we wanted to make the apertures as small as possible (while keeping the signal size of the same order of magnitude as our present value - i.e. keeping the total exposed area the same), we could fill a $400 \mu\text{m}$ membrane “to the brim” ($\sim 350 \times 350$ for practical

reasons) with apertures spaced $1\mu\text{m}$ apart. The aperture diameter in that case would be $\sim 10\text{nm}$, which is a value we have observed for certain aperture arrays that we suspected were closed up due to contamination, but is otherwise very difficult to obtain reliably and stably using our present fabrication methods.

Another important criterion is uniformity of average aperture size and shape between different chips on the wafer. Non-uniformity induces higher asymmetry factors and reduces the modulation depth in SHeQUIDs. See the sub-section on asymmetry in Section 1.3.1 for more details and simulated plots.

That summarizes most of the practical constraints placed on the aperture arrays. Of course, it is possible that more sophisticated fabrication methods than the ones we use might be used to go beyond these constraints. We talk about some of these possibilities in Chapter 4 (but the literature shows that they come with their own set of problems). It is likely that E-beam lithography, at this point of time, remains the only *scalable* fabrication method to make many chips in a single fab run at a reasonable cost. At the risk of stating the obvious, given sufficient time and money, many things become quite practical.

3.1.2 Displacement sensor (Chapter 5)

The displacement sensor can come in at least two flavors, both making use of a SQUID¹ to measure small changes in the magnetic flux in a superconducting spiral-wound coil (called a “pancake coil”). These flux changes are caused by movements of a flexible diaphragm (coupled to mass currents in the inner cell) placed close to the pancake coil.

Original design (Section 5.1)

The original sensor, developed by Paik, *et al.*[42][43], uses a flexible diaphragm with a thin superconducting film (so far, lead and niobium have been used successfully for this purpose) located next to the pancake coil. The coil and a commercial SQUID are part of a superconducting circuit in which a persistent current is circulated. Motion of the superconducting diaphragm changes the flux in the circuit. This flux change is detected by the SQUID and is proportional to the displacement of the diaphragm.

Magnet based design (Section 5.2)

The magnet based design, developed by Sato, *et al.*[44], uses a flexible diaphragm with a small rare-Earth magnet (commercially available neodymium magnets have been used successfully) glued in the center. The diaphragm is still metallized, but with a normal conducting metal (not a superconductor), which is not used as a part of this sensor². A superconducting pancake coil is placed next to the diaphragm as before and connected directly to the SQUID,

¹Superconducting Quantum Interference Device

²The normal metal coating in this case is just used for capacitive coupling to the fixed electrode as described in the introduction to this section.

with no other circuitry needed. Motions of the diaphragm+magnet change the flux seen by the pancake coil proportionally to the displacement.

Comparison

- The major advantage in the original design is that the displacement sensitivity can be changed during experiments (simply by changing the injected persistent current in the circuit). This can be useful because the SQUID has a limited dynamic range, so that (for example) measuring DC position values can be tricky for a high sensitivity sensor (forced to keep track of fast resets in the SQUID). There are a lot of different experiments one can do once cooled down and having an easily adjustable dynamic range/sensitivity can come in quite handy. This is not so in the magnet-based design, where the sensitivity is set during assembly, with the magnet-pancake coil spacing and the magnet's field strength.
- A possible problem with the magnet-based design is the mass of the magnet (greater than the typical mass of the bare, metallized diaphragm), which tends to suppress high frequency oscillations (see Section 3.2.1 for details) and may reduce the sensor's effectiveness.
- The major advantage in the magnet-based design is the lack of the persistent current (PI) circuit, which is a common point of failure (especially over several thermal cycles) since it contains several superconducting joints, either spark-welded or screwed together. The former tend to be robust but the latter can loosen due to thermal cycling. Any normal joints instantly kill the ability of the circuit to maintain a persistent current. Historical/anecdotal surveys suggest that such joints tend to be pretty binary in that they either work or don't, with very little (if any) middle ground (in the sense of the decay times of the trapped current being either essentially infinite or of the order of seconds or minutes). However, any problems in the PI circuit (which is mostly outside of the cell) can be fixed without much hassle by warming up and cooling back down relatively quickly ($\lesssim 1$ week) so this is not a huge argument against using such a design.
- A somewhat practical disadvantage in the original design has to do with depositing superconducting metal (Pb or Nb) on the diaphragm. Both these metals have issues with deposition that, depending on available facilities, can range from trivial to fatal. Further details on this issue can be found in Section 5.1.2. Since the magnet-based sensor can be realized with any metal that stays normal near T_λ – typically Aluminum – deposition is not an issue as Aluminum deposition can be done easily in most basic evaporators.

3.1.3 Cell body (Section A.1)

The cell body defines the inner and outer cell volumes and is the framework holding the other modules together. We can think of the cell body as composed of four distinct parts:

Inner cell piece (ICP) (Section A.1.1)

The ICP holds the aperture array chip(s). For SHeQUIDs, it also defines two distinct paths for the sense loop. Constraints on the path dimensions derive from hydrodynamic inductance constraints and are discussed in Section 3.2.2.

D-ring (Section A.1.2)

The D-ring holds the flexible diaphragm from the displacement sensor. It also holds a resistive heater for the inner cell and together with the ICP, defines the inner cell volume. In less modular designs, the ICP and D-ring are typically combined into one part. Important constraints on the inner cell volume dimensions are discussed in Section 3.2.1. A resistance of $1\text{ k}\Omega$ is found to be optimal for the kinds of heater powers we need in the inner cell.

E-ring (Section A.1.3)

The E-ring holds the fixed metallized electrode used to exert electrostatic forces on the diaphragm. It also holds the pancake coil from the displacement sensor close to the diaphragm.

Cell can (Section A.1.4)

The cell can defines the outer cell volume and is used to mount the cell to the cryostat. It is also lead plated to shield the sensitive displacement sensor components from stray magnetic fields. A copper-nickel tube stuck through the wall and soldered in place provides a fill-line to evacuate and fill the cell with helium during the experiment. Depending on the design requirements, electrical leads are brought out of the cell through a sealed breakout on either a separate cap for the cell can or the ICP (which, in most versions of the SHeQUID created so far, itself serves as a cap for the cell).

3.1.4 Wiring breakouts

Bringing the various electrical leads out of the cell is a crucial (and altogether easily overlooked) aspect of designing a cell. An important issue is adequate shielding from electromagnetic noise in the cryostat environment and cross-talk (inductive or capacitive coupling between nearby leads). Another issue is the mechanical difficulty of working with the quite fragile superconducting wires used nearly everywhere in the cell (which dictates large turn radii for tunnels). All wires must be mechanically anchored to prevent motion-induced stray fields. Finally, all wiring paths and anchors (epoxy or in rare cases, cryogenic tape) must

be planned out during the design phase because once assembly begins, everything must be kept as clean as possible, which makes it very difficult to modify structural parts without contaminating the parts. The electrode, diaphragm, pancake coil and aperture arrays are all highly intolerant of dust/debris created during such late stage fine-tuning operations. One might even find late in the game that a given wire simply cannot be safely maneuvered through a poorly planned breakout.

The possibility of scratching wires while feeding them through breakouts and thereby shorting them to the cell body is very likely and all paths, edges and holes must be thoroughly deburred and smoothed. Small diameter teflon (PTFE) tubing can be used to safely guide wires through tunnels as an extra precaution. Screw joints to connect wiring to the electrode and diaphragm tabs (thin plastic with evaporated metal) are made over embedded epoxy washers in the E-ring and D-ring so that accidental rips through the fragile plastic tabs do not short the leads to the body. All these possibilities must be thoroughly considered and designed for unless one wishes to waste several cooldowns tediously tracking down mysterious shorts. A healthy sense of (constructive) paranoia and (cheerful) pessimism can come in very handy for this purpose³.

Another, particularly insidious issue is bends and kinks in superconducting wires (specifically, the Cu-Ni clad, NbTi wire we use for all cell wiring, including the pancake coil and the persistent current circuit). Anecdotal evidence suggests that such defects in superconducting wire can significantly reduce the critical current supported by it. Physically, this makes sense, because sharp bends tends to produce larger fields nearby, making it easier to reach the critical field for the superconductor.

3.1.5 Sense arm (Section A.1.6)

The sense arm can be made in several different forms, depending on the goal of the experiment. Its most basic job is to close the loop defined by the two paths leading out of the aperture arrays and define the “sense loop”. It may also have a resistive heater in one end and a copper sink in the other to create a superflow as described in Section 1.3.2. Other forms have been constructed for various purposes. For example, Fig. 3.1 shows a vertical sense loop with counterwound clockwise and counterclockwise turns to cancel out any coupling to rotational fields while measuring other (non-rotational) novel interactions such as the Aharonov-Bohm effect for neutral matter [36][45]. This kind of sense arm was recently demonstrated experimentally by Narayana, *et al.* [12].

Constraints on the sense arm dimensions derive from hydrodynamic inductance considerations and are discussed in Section 3.2.2. We have found a resistance of $1\text{ k}\Omega$ to be optimal for the kinds of heater powers we need in the heat-pipe.

³The author has found these techniques useful in combating the dreaded *planning fallacy*.




Figure 3.1: (a) Conceptual sketch of proposed Aharonov-Bohm effect experiment. (b) “Reciprocal configuration” (i.e. counter-wound CW and CCW turns) sense arm from Narayana, *et al.* [12].

3.2 Constraints

Many of the constraints relevant to SHeQUID design are scattered throughout this dissertation and are explicitly pointed out in such cases (such as the ones for aperture arrays, discussed in Section 3.1.1). Here, we discuss issues that require special attention because they impose somewhat more global constraints on design and must be considered before designing even the kernel of the SHeQUID.

3.2.1 Helium compressibility and diaphragm mass

The fluid mass current generated by the quantum oscillations deforms the flexible diaphragm. This effective motion is measured by the displacement sensor. When we relate the velocity of the flexible diaphragm to the current in the aperture arrays, we make the assumption that the fluid in the inner cell is incompressible. This allows us to assume that the amount of fluid entering or leaving the cell is exactly equal to the volume change in the inner cell due to deformation of the diaphragm.

This assumption, however, ignores the fact that the volume change (and hence the sensor’s effective sensitivity to current) is reduced by the compression or expansion of the inner cell fluid. We define and derive the “signal loss” ϵ due to finite compressibility as the fractional

amount by which the displacement of the diaphragm is reduced as compared to the case where the fluid is assumed to be incompressible in Eq. (E.7) (see Appendix. E).

However, the analysis of Appendix. E assumes quasi-static behavior and consequently, does not consider the effect of mass current oscillation frequency as it pertains to this signal loss. It does however, give us an inside look at the dynamics of how compressibility enters the picture. We can see the consequences of both the “capacitive” effects of finite compressibility and the “inductive” or inertial effects of finite diaphragm mass on the displacement signal caused by a time-varying mass current, by writing out the equation of motion of the diaphragm and taking its Fourier transform. We see then that the mass of the diaphragm (including magnet for the alternate style displacement sensor) serves to suppress high frequency oscillations. A full analysis of this is presented in Appendix F (which is meant to accompany Appendix. E). The reader is invited to explore those chapters for more detailed discussions about these issues, including numerical estimates on the consequences of these effects.

3.2.2 Hydrodynamic inductances

The ratio of the hydrodynamic inductance⁴ of the individual segments of the sense loop (L_t) to the inductance of a weak-link (L_J), termed $\alpha \equiv L_t/L_J$ affects the modulation depth of the SHeQUID interference pattern and therefore the sensitivity. The following analysis quantifies this effect and is based on the reasoning described in Ref. [26, pp. 161-2]. In essence, we find that increasing α leads to decreasing modulation depth (and sensitivity), which is borne out by experiments performed by Narayana, *et al.*[12]. This lends us a quantitative tool for designing the sense loop within these constraints.

The SHeQUID with finite loop inductance L_t is modeled by the electrical circuit in Fig. 3.2. To retain maximal generality, we do not assume any particular regime (Josephson or phase-slip) at the moment. Referring to the circuit diagram, we only assume that the current in each branch is given by the current passing through each aperture array, with some well-defined current-phase relations $I_1 \equiv I_1(\theta_1)$ and $I_2 \equiv I_2(\theta_2)$, where θ_1 and θ_2 are the phase-differences across the aperture arrays.

We note that the path inductance segments L_1 and L_2 are always in bulk superfluid (strongly coupled). We can therefore use our inductance results from Eq. (G.20) of Section G.2.2 to write the phase-drops across the paths L_1 and L_2 as $m_4 L_1 I_1 / \hbar$ and $m_4 L_2 I_2 / \hbar$, respectively.

The total phase-drops ($\Delta\phi_L$ and $\Delta\phi_R$) along the left and right branches can now be written as the sum of the phase-drops across the aperture array and the path inductance:

$$\Delta\phi_L = \theta_1 + \frac{m_4}{\hbar} L_1 I_1 \quad (3.1)$$

$$\Delta\phi_R = \theta_2 + \frac{m_4}{\hbar} L_2 I_2 \quad (3.2)$$

⁴See footnote 4 on p. 8 for a note on the hydrodynamic inductance.




Figure 3.2: Electrical circuit analogue of the SHeQUID with finite loop inductance. The X's are the aperture arrays (indexed as 1 and 2 for the left and right arrays). L_1 and L_2 are the parasitic inductances of the paths that make up the sense loop from each aperture array to the top junction. θ_1 and θ_2 are the phase-differences across the aperture arrays and μ_o and μ_i are respectively the chemical potentials at the outer and inner cells. All phase/chemical potential differences are defined: (lower value – upper value). $I_{tot} = I_1 + I_2$ is the total current, where I_1 and I_2 are the individual currents entering each branch from the junction.

Even though the end points for these branch total phase-drops are the same (the top and bottom junction), they are not, in general, equal. That is because we have not yet included any externally imposed phase-drops (like the Sagnac phase-shift) that can contribute asymmetrically to each branch. We can see this more generally (and in a different way) by noting that the Anderson phase-evolution equation (G.12) governs the evolution of $\Delta\phi_L$ and $\Delta\phi_R$ as:

$$\dot{\Delta\phi}_L = \dot{\Delta\phi}_R = \frac{-\Delta\mu}{\hbar}$$

where we have defined a total chemical potential difference $\Delta\mu \equiv \mu_i - \mu_o$ between the inner and outer cells (see Fig. 3.2). For a constant $\Delta\mu$, we can integrate these equations separately with the initial conditions that at $t = 0$, $\Delta\phi_L = D_{L0}$ and $\Delta\phi_R = D_{R0}$ to obtain:

$$\Delta\phi_L = D_{L0} - \omega_J t \quad \text{and} \quad \Delta\phi_R = D_{R0} - \omega_J t \quad (3.3)$$

where we have identified the usual Josephson frequency $\frac{\Delta\mu}{\hbar} = \omega_J$. Note that the 2 constants encode information about the initial phase-difference between the two branches. Using these

expressions in Eqs. (3.1) and (3.2), we obtain:

$$D_{L0} - \omega_J t = \theta_1 + \frac{m_4}{\hbar} L_1 I_1 \quad (3.4)$$

$$D_{R0} - \omega_J t = \theta_2 + \frac{m_4}{\hbar} L_2 I_2 \quad (3.5)$$

We can finally address the somewhat *ad hoc* assumption that is fundamental to the calculation procedures used in Refs. [26, p. 161] and [12], where an average phase is defined as half the sum of the right-hand sides of the above equations and this phase is claimed to be proportional to the time variable. Adding the two equations above and dividing by 2, it is immediately clear that this is indeed the case (but it is hardly obvious from just physical intuition).

This diversion taken care of, we subtract the two equations and define a new constant $D \equiv D_{L0} - D_{R0}$ to obtain:

$$\theta_2 + \frac{m_4}{\hbar} L_2 I_2 - \theta_1 - \frac{m_4}{\hbar} L_1 I_1 + D = 0 \quad (3.6)$$

This entire exercise was done towards determining this important constant D in terms of known quantities. As discussed in Section 1.3.1, the single-valuedness of the phase of the superfluid order parameter around the interferometer loop demands that $\oint \nabla \phi \cdot d\mathbf{l} = 2\pi n$. We can assume here that changing the applied (external) phase-influence φ_a does not change the circulation number n so that the constant can be folded in as a constant offset to φ_a . Using the phase drops above and integrating the phase-gradient counter-clockwise around the sense loop, we get:

$$\theta_2 + \frac{m_4}{\hbar} L_2 I_2 - \theta_1 - \frac{m_4}{\hbar} L_1 I_1 + \varphi_a = 0 \quad (3.7)$$

Eqs. (3.6) and (3.7) imply that the phase offset D between the left and right branches of our circuit is precisely equal to the externally applied phase-shift φ_a (obvious but now rigorously shown). With no loss in generality⁵, we therefore set $D_{R0} = 0$ and $D_{L0} = D = \varphi_a$ in Eqs. (3.4) and (3.5) to obtain the final equations that we will use to compute the phases numerically:

$$\boxed{\theta_2 + \frac{m_4}{\hbar} L_2 I_2(\theta_2) = -\omega_J t} \quad (3.8)$$

$$\boxed{\theta_1 + \frac{m_4}{\hbar} L_1 I_1(\theta_2) = \varphi_a - \omega_J t} \quad (3.9)$$

The total current passing through the SHeQUID is

$$\boxed{I_{tot} = I_1(\theta_1) + I_2(\theta_2)} \quad (3.10)$$

⁵Think of it as starting our time axis when $\Delta\phi_R$ happens to be 0. Therefore, $D_{R0} = 0$ in this coordinate choice.

Procedure

The first two of the three boxed equations above can now be used to numerically compute the phases θ_2 and θ_1 (separately) vs. time values t for a fixed external phase φ_a and the last equation used to compute the total current measured at the inner cell. The inductances L_1 and L_2 , and the current-phase relations must of course be known. The whistle frequency is unimportant in this analysis and any convenient frequency may be chosen without loss of generality⁶.

This process gives us a family of numerically generated timeseries functions $I_{tot}(t)$; one for each (fixed) value of φ_a . We can numerically find the Fourier transforms of these functions and integrate the whistle peak (the first harmonic I_0 at ω_J). Repeating this for different values of φ_a , we finally obtain the interference curve $I_0(\varphi_a)$.

This allows us to see how the path inductances affect the modulation depth (and hence sensitivity) of the SHeQUID.

Josephson regime

In the Josephson regime, the current-phase relation is sinusoidal and using Eq. (1.3), the master equations (3.8), (3.9) and (3.10) become:

$$\begin{aligned}\theta_2 + (m_4/\hbar)L_2I_{c2} \sin \theta_2 &= -\omega_J t \\ \theta_1 + (m_4/\hbar)L_1I_{c1} \sin \theta_1 &= \varphi_a - \omega_J t \\ I_{tot} &= I_{c1} \sin \theta_1 + I_{c2} \sin \theta_2\end{aligned}$$

Now, we observe (from Eq. (G.15)) that the Josephson inductance of a weak-link at zero phase is $L_J(0) = \hbar/(m_4I_c)$. Note that this is only the *minimum* inductance (the maximum is ∞). We therefore define two new parameters $\alpha_1 \equiv L_1/L_{J1}(0)$ and $\alpha_2 \equiv L_2/L_{J2}(0)$, which determine the relative dominance of the parasitic (path) inductance over the intrinsic Josephson inductance. The trio of equations can now be written more clearly as:

$$\theta_2 + \alpha_2 \sin \theta_2 = -\omega_J t \quad (3.11)$$

$$\theta_1 + \alpha_1 \sin \theta_1 = \varphi_a - \omega_J t \quad (3.12)$$

$$I_{tot} = I_{c1} \sin \theta_1 + I_{c2} \sin \theta_2 \quad (3.13)$$

where the two alphas can be directly dialed from 0 (no parasitic inductance \Rightarrow ideal weak-link SHeQUID) to arbitrarily large values to see how the interference curves are transformed. For simplicity, we set $L_1 = L_2$ (a symmetric path SHeQUID) and $I_{c1} = I_{c2} = 1$ (identical aperture arrays), so that $\alpha_1 = \alpha_2 \equiv \alpha$.

We present the results of this numerical analysis below:

⁶Here, we used the same Labview VIs for windowing the timeseries data and computing the frequency spectrum that we used during experiments (for consistency). Therefore, we used realistic frequencies to simulate realistic current amplitudes so that we could gauge the effects on signal to noise as well.

Time-series: Notice that the total current above can, in principle, be a complicated function of time. So, the shape of the time-function for the total current in the SHeQUID should vary with the parameter α (recall that α rises linearly with the loop inductance for fixed weak-links). This in turn means that the frequency spectrum (Fourier transform, which uses a sinusoidal basis) will no longer be just a single peak at ω_J and can, in general, have large contributions from higher harmonics of ω_J . Sample time-series and FFT for a non-zero α are shown in Figs. 3.3–3.4 and 3.5.




Figure 3.3: Front view (back view in next figure) of a set of simulated current vs. time plots (of Eq. (3.11)) for different values of α . Time is displayed in units of the Josephson period ($2\pi/\omega_J$), where ω_J is the Josephson angular frequency used to generate the data in Eq. (3.11). There is one timeseries shown for each α , where φ_a is chosen to give the largest first harmonic in the FFT for that α .

Spectral composition: We find the size of the first 15 integrated peaks in the FFT of the timeseries (i.e. the first 15 harmonics of the Josephson frequency). This data is displayed in Fig. 3.5

Now, this current function can no longer be analytically separated into a time-dependent “carrier” oscillation (at ω_J) and a time-independent “modulation envelope” as we could do in the case with negligible loop inductance (Section 1.3). Instead, we need to compute it




Figure 3.4: Back view of timeseries plot in Fig. 3.3.

numerically. Figure 3.3 shows the results of such a numerical computation of the timeseries function. The figure shows a series of plots of the total current vs. time from Eq. (3.11) for varying values of α . Note that the current function gets progressively more distorted from that of a pure tone at ω_J as α increases. We should therefore expect its decomposition in the frequency domain to yield nontrivial contributions from higher harmonics of ω_J . Fig. 3.5 shows us precisely such an outcome. As α increases, the spectral power gets spread out over more and more harmonics in a nontrivial manner. This decreases the size of the first harmonic contribution, *which is the quantity that is actually measured*. So, if the first harmonic amplitude is defined as the “total whistle amplitude”, we will interpret this chain of events as a reduction in the modulation depth of the interference pattern, accompanied by a corresponding reduction in the maximum phase sensitivity of the SHeQUID. Note also that it is not merely a matter of measuring more and more harmonics as the spectrum does not, in general, converge within a few harmonics.

First harmonic size vs. α and φ_a : We then plot the magnitude of (integrated power under) the first harmonic of the FFT against α and φ_a in Fig. 3.5 and note that it decreases with increasing α .




Figure 3.5: As described in the text, the first 15 integrated peaks in the FFT for each timeseries plot from Fig. 3.3 are shown for each value of α . These are the raw harmonic amplitudes.

3.2.3 Conclusions: modulation depth summary

The expected modulation depth (and therefore the phase-sensitivity) depends on several factors, including:

- the individual current amplitudes in each aperture array
- the asymmetry between the two (or more) aperture arrays that make up the SHeQUID:
- the inductance ratio α : Calculations in the previous section and results from [12] suggest that increasing the ratio of sense loop path inductance to the weak-link inductance to more than $\mathcal{O}(1)$ can severely degrade the phase-sensitivity of the SHeQUID.




Figure 3.6: For a fixed value of α , the set of harmonic 15 amplitudes from Fig. 3.5 is normalized (in quadrature) to determine the relative contribution of each harmonic to the interfered whistle and displayed here. This bar graph should be read as a set of relative amplitudes, one for each α . Unlike Fig. 3.5, amplitudes do not compare across α values. We can see more clearly here that the nature of the oscillation goes from sinusoidal to sawtooth-like for increasing α . Note the anomalies showing up for larger α - they are due to numerical errors in computing the phases that start becoming more and more important as the overall current decreases.




Figure 3.7: Magnitude of the first harmonic in Fig. 3.5 plotted against α and φ_a

3.3 Possible design philosophies

The base components described above tend to be present in one form or another in a SHeQUID. The time and effort that goes into making each one can vary enormously (we will get some idea of this when we discuss their fabrication in later chapters). Historically however, the least reliable component has been the aperture array, with problems ranging from weak membranes (that break easily during cell evacuation or filling) to apertures that have a limited shelf-life and close up over time. No matter what the problem with the aperture arrays, it is a given that if they suffer from such problems, it is extremely unlikely that one will obtain a usable quantum whistle.

Therefore, having a cell that has to be completely rebuilt with all new components in the event of aperture array failure or contamination is probably not a very good strategy in terms of wasted time and effort. This is the major reason why the base structural components are typically made of metal (aluminum or brass) and the chips (and other finicky parts like diaphragms, pancake coil, cell heater, wiring) are glued on to them. This way, any defects in these active parts (prone to failure) can be fixed by simply heating the pieces, thereby undoing the glue joints and cleaning out any debris/residues using sandblasters or other finishing tools and gluing in a new component. Time savings on complex (but completely passive) structural parts like all the parts falling under “cell body” (Section A.1) can be quite considerable, possibly adding up to a week for each iteration for a skilled machinist⁷.

3.4 Modular single weak-link cell design

This is a cell we designed to quickly be able to test different aperture array chips in full superfluid experiments. However, it stopped short of full modularity in the sense that the inner cell heater, flexible diaphragm and the wiring breakouts for both had to be remade for each new chip tested. This was tolerable because the most (relatively) difficult part to make reliably and reproducibly is the pancake coil. Based on the recent successful test of the modular SHeQUID design, we can make some small changes to this design to have a fully modular cell where one has to simply replace extremely simple metal holders with chips glued in them for each new run. This sort of design should be very useful for someone trying to test the superfluid dynamics of different aperture arrays (which was the original dissertation goal of the author).

Before making something like this for the first time, it probably makes more sense to build a very simple single-use cell with little to no modularity⁸ in order to validate the processes

⁷Of course, with access to relatively long CNC times or with the rapidly evolving new technologies of rapid prototyping using additive instead of subtractive machining (3D printers such as the Makerbot), these considerations may no longer be relevant after a few years (it would be rather disappointing if they were). Of course, such techniques are as yet very material specific and it may take a while before cryogenically suitable materials are available in this context.

⁸Cells like these may be found in several early publications on the ^4He quantum whistle coming out of our research group and details may be found in dissertations such as Refs. [39] and [46].

and components described in this dissertation. The complexity of the cell design does go up somewhat with increasing modularity as will the initial time investment in making the parts. However, the main advantage of modularity is being able to reuse the bulk of the cell and shorten overall turnaround time between test runs. This, however, is useful only if one is reasonably constant in time when it comes to research goals.

3.4.1 Original (tested) design: partially modular

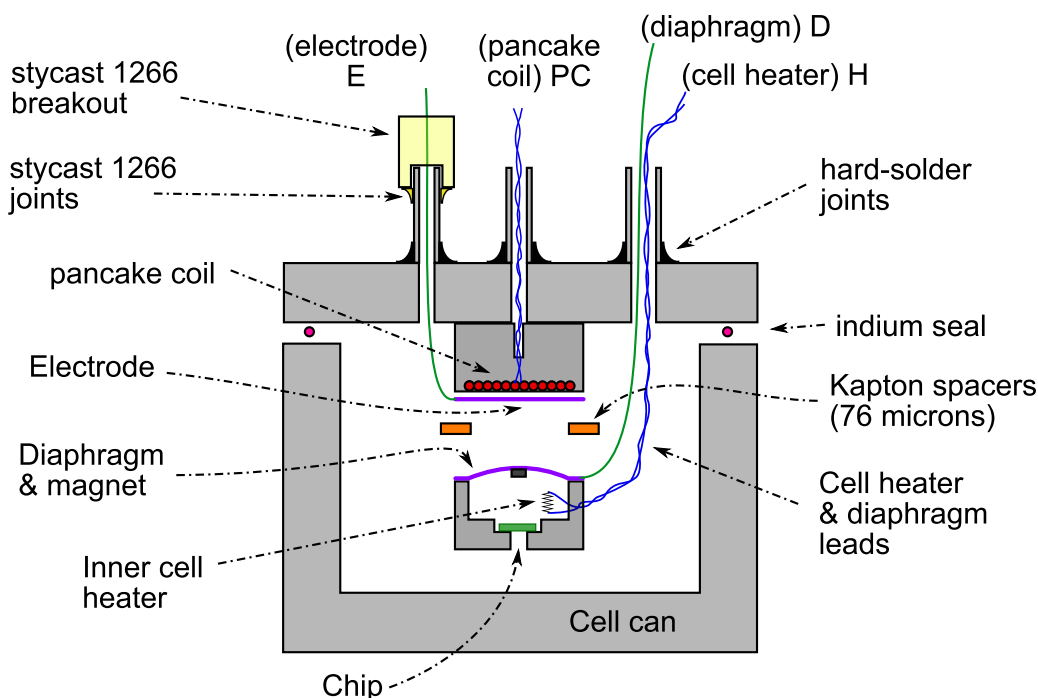


Figure 3.8: Partially modular single weak-link cell design. Used in Weak-link cell 6, which was used for the demonstration results in Chapter 10.

As shown in Fig. 3.8, the pancake coil (PC) in its holder is mounted to the main flange that caps the cell can. The central breakout has the most electrical shielding (with the twisted-pair PC leads inside a lead tube and further screened by the holder metal). The electrode (E) lead, which is connected to the fixed electrode (glued over the PC) by a screw-joint, is deliberately given its own breakout to capacitively decouple it from the (PC) leads. This forms the permanent part of the cell and should never have to be replaced under normal circumstances. The inner cell piece (ICP) has the chip, cell heater (H) and flexible diaphragm (D), with a twisted-pair for the heater and a single lead wire for the diaphragm (again, screw-fit onto the diaphragm tab). This entire piece (with a new chip, heater and diaphragm) is screwed on to the permanent (PC) assembly with thick Kapton spacers in

between to separate the fixed and flexible diaphragms and the leads brought out together through the third (larger) breakout tube.

The breakouts are machined clear stycast⁹ 1266 forms designed to minimize wiring stress and provide a superleak-tight joint (they are sealed with more (fresh) stycast 1266). Now, the way the breakouts are setup here, a multiply folded piece of aluminum foil (or a similar heat-shield) can be used to cover the E and PC breakout tubes to protect it, while the (D+H) breakout can be unsealed using a small, commercial butane torch and blasting the seal edge with the flame for just 1-2 sec. The seal should peel off cleanly. The wires can be cut and the ICP removed. A new ICP can now be put in its place once the breakout tube is gently cleaned with alcohol wipes. We have done this successfully several times without disturbing the PC+E assembly. One must be careful not to expose the sealing surface of the breakout tube to a direct flame to avoid charring or other deposits. It is only the stycast seal that needs to be heated (it is alright if it momentarily catches fire as long it's put out immediately) and the seal must be pulled off while the stycast is still soft.

3.4.2 Suggested changes for full modularity




Figure 3.9: Proposed modification in the design of Fig. 3.8.

The semi-modular design in the previous section falls short of full modularity because the inner cell heater and diaphragm need to be replaced each time. In the case of the conventional (Paik) displacement sensor, the lead-coated diaphragms are not easily available (albeit for sociological, not technological reasons discussed in Section 5.1.2). In the case of the magnet sensor, the new magnet's field properties, its precise centering on the diaphragm and the precise centering of the diaphragm on the ICP might all be different if the ICP is rebuilt with new parts, which can change the displacement sensitivity. If this change is too large, it will necessitate a detailed recalibration of the sensor, which adds time to the run.

Since the main issue here is the wiring that needs to be re-done each time, a simple fix is to break the ICP into two parts as shown in Fig. 3.9. The upper part now contains the inner

⁹Stycast 1266 and 2850FT epoxies are used extensively in this work. They are manufactured by Emerson & Cuming and distributed within the USA by Ellsworth Adhesives (<http://www.ellsworth.com>)

cell heater and the diaphragm, and is (optionally) permanently affixed to the PC assembly. The D and H wires are taken out as before but now do not need to be replaced each time. The chip now has its own holder (which contains nothing else and is a very uncomplicated piece that can be mass-produced quickly). The chip holder is now indium-sealed onto the new ICP for each new run involving a new aperture array chip. The chip-gluing platform is raised above the indium seal surface to protect the latter during gluing and to minimize the inner cell volume (see Section 3.2.1 for reasons why the volume should be minimized).

This method has already been used in well-tested cells such as the modular SHeQUID (to be discussed in the next section and subsequent chapters) and the apparatus to conduct 77 K flow tests on aperture arrays (discussed in Section 6.1).

3.5 Modular SHeQUID design

The primary layout in the case of the modular SHeQUID¹⁰ is necessarily reversed as compared to the single weak-link cells discussed previously. The reason is the sense arm, which is large (and can be even larger for certain experiments like the Aharonov-Bohm effect mentioned Section 1.3.2) and would require a very large cell can to contain if we went with the same layout as before.

As shown in Figs. 3.10 and 3.11, we build the cell around the inner cell piece (ICP) rather than the pancake coil (PC). The ICP defines the sense loop paths for the SHeQUID, which are closed by the sense arm. The sense arm remains independent, and can be replaced anytime without even breaking the cell can seal. The tube flanges (the metal base of the sense arm that is otherwise made of stycast 1266) can also be reused by unsealing the tubes from the stycast using a butane torch as described in Section 3.4.1. The chips are glued onto the ICP, which has three separate indium seals on the same flange. The D-ring (which holds the cell heater (H) and a flexible diaphragm (D)) is sealed onto the ICP to define the inner cell volume. The two off-axis seals are for wiring breakouts, which work in much the same way as the ones in the single weak-link cell (except that these breakouts are now inside the cell can rather than outside). The (H) and (D) leads are taken out through one of the breakouts (mutually shielded by lead tubing). The (PC) and fixed electrode (E) are glued into the E-ring and this assembly is screwed onto the D-ring with a Kapton spacer (76 μm thick) in between as usual to define the E-D capacitance. The (PC) and (E) leads are taken out the other remaining breakout tube and sealed in place as before with machined stycast breakout forms. All 4 sets of wiring are inside lead tubes and freely slide through matching holes in the ICP (this explains the reason for the off-axis indium seals around these holes).

The modularity within the D-ring/E-ring assembly is somewhat limited because the breakout tubes are soft-soldered onto the D-ring and the butane torch removal method is just a little riskier than in the single weak-link cell (where the tubes are hard-soldered). The reason is that the D-ring is a considerably more delicate piece and hard-soldering seems

¹⁰The base designs upon which we have built, come from previously made cells by Emile Hoskinson[29] and Yuki Sato[46], which were further based on cells made by Ray Simmonds.

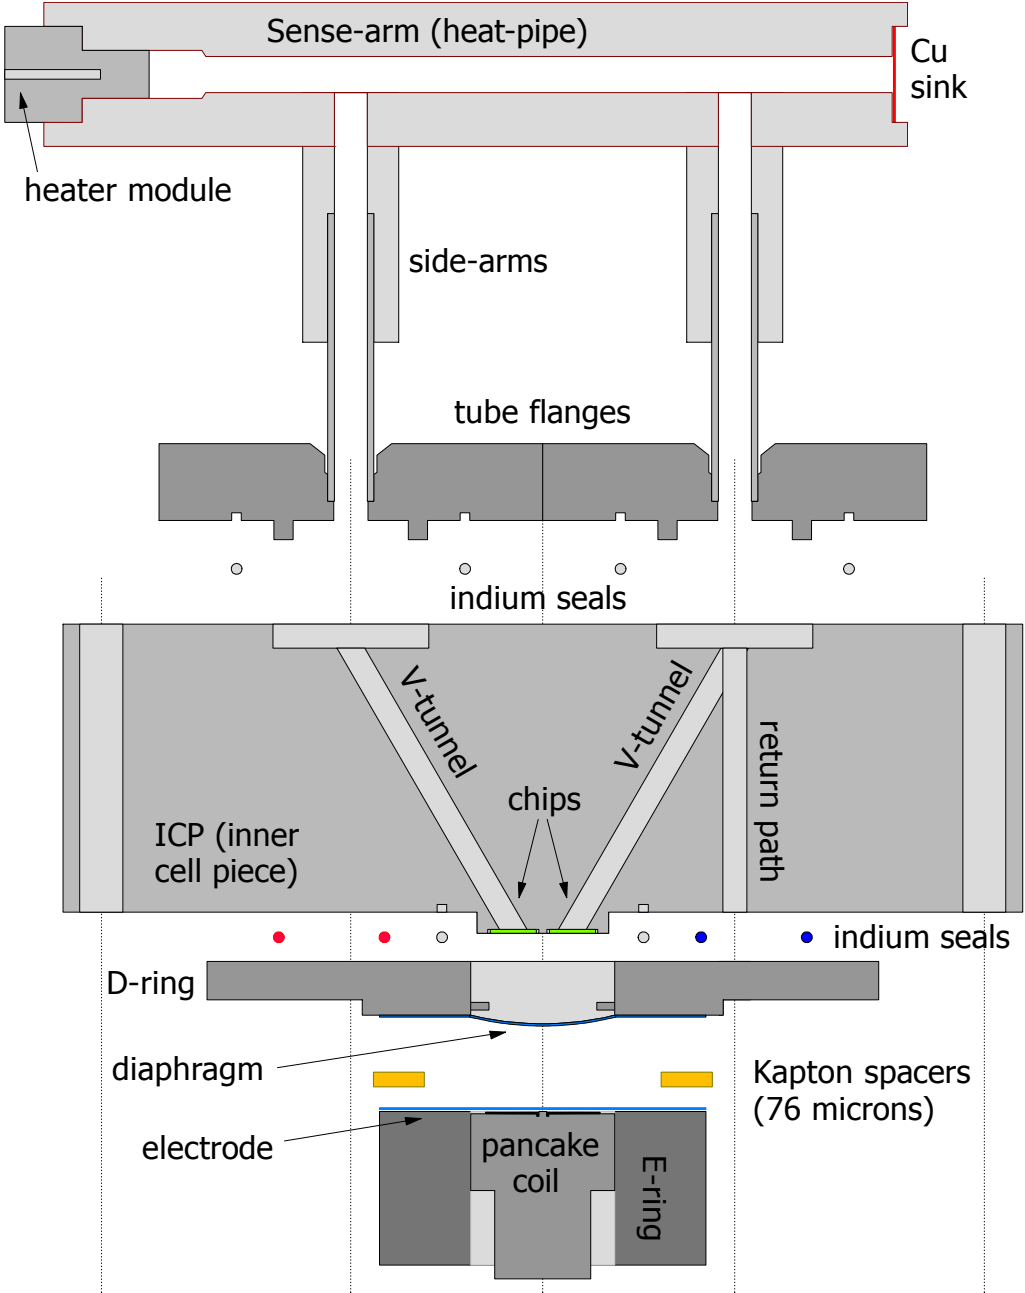


Figure 3.10: Modular SHeQUID showing V-tunnel plane.

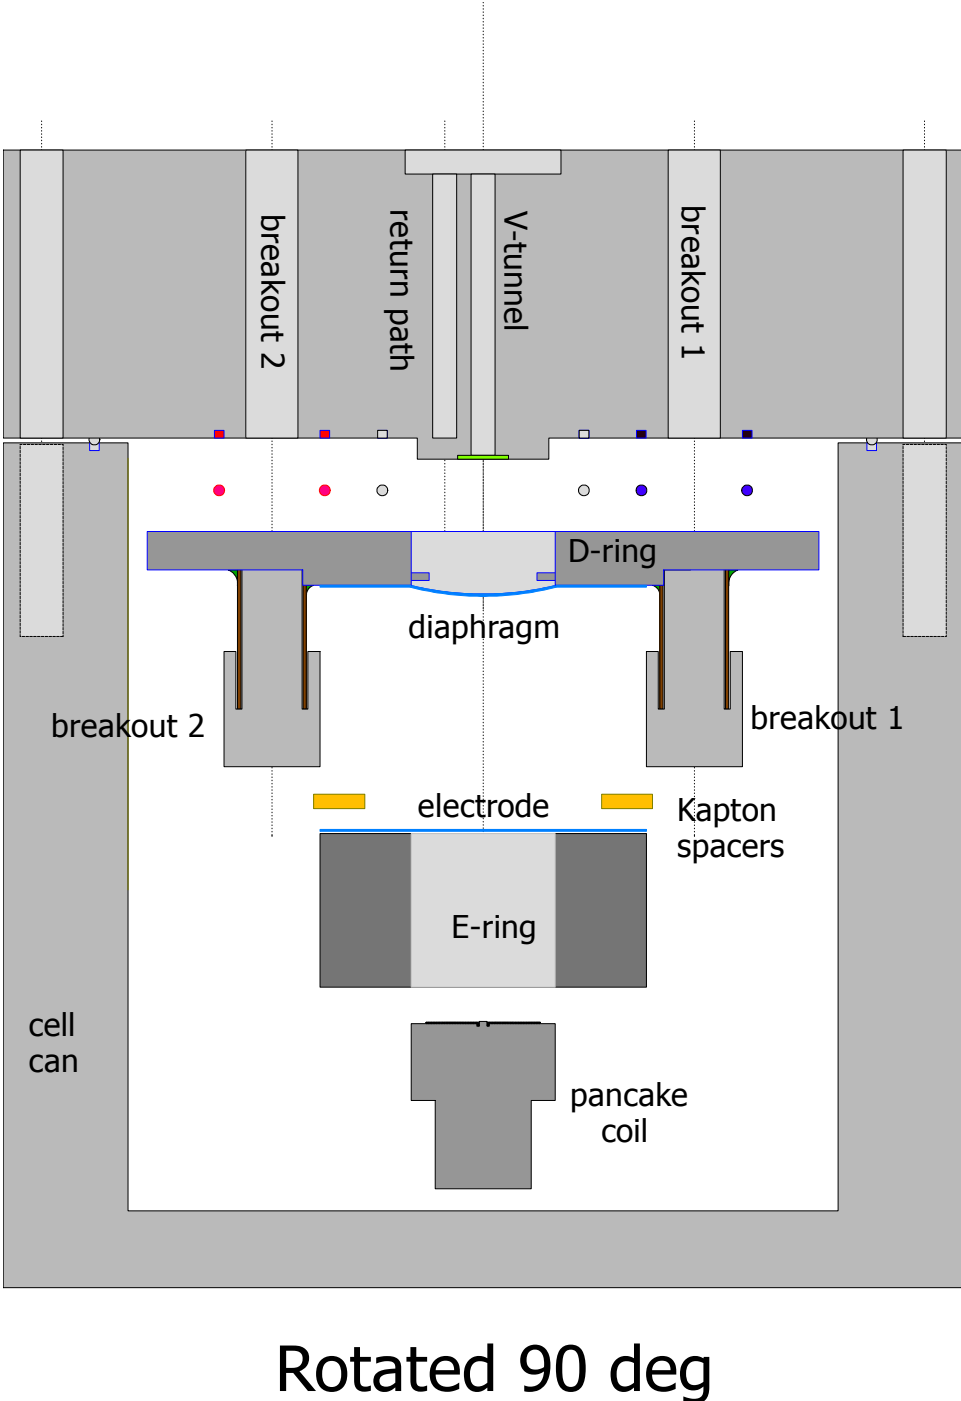


Figure 3.11: Modular SHeQUID showing breakout plane (Fig. 3.10 rotated 90 deg.)

to warp the surface a bit (bad for indium sealing). These are tractable problems but do need more work to get around. In any case, the entire D-ring/E-ring assembly (once validated in use) can be unsealed from the ICP and removed as a whole (with leads intact). The chips can be replaced with minimal finishing work on the ICP (polishing the seal surfaces if required) and the sense arm can be replaced (without disturbing the chips if need be). We can continue to reuse the D-ring/E-ring assembly indefinitely, thus removing the need for any rewiring or remaking of complicated components within the cell.

Even in extreme circumstances where cell wiring is suspect, all the metal parts can be cleaned of epoxy and reused indefinitely. We will discuss the fabrication and assembly of this modular SHeQUID in some detail in subsequent chapters. Note that most of the fabrication and assembly guidelines are rather universal when it comes to performing superfluid weak-link experiments, regardless of design paradigms followed.

Chapter 4

Fabricating nanoscale aperture arrays

Parts of this chapter were previously published in the conference proceedings of the 25th international conference on low temperature physics [47].

Arrays of nanoscale apertures have been found to exhibit fascinating quantum phenomena such as the Josephson effect and collective quantized phase-slippage. To be in the Josephson regime, the aperture size must be comparable to the healing length of the ^4He order parameter. For 50 nm apertures this regime is attained $\sim\text{mK}$ below T_λ . Collective phase slippage occurs at considerably lower temperatures. Fabricating aperture arrays with appropriate properties (strength, temporal stability and reproducibility) at these length scales has been a long-standing goal for our group. Here, we present the techniques used thus far, based on recent work performed at the Cornell Nanoscale Facility. We discuss some issues that arise and their possible solutions. This work was supported by the NSF and the ONR.

4.1 Introduction

We have fabricated large arrays of nanoscale apertures with diameters $\sim 90\text{nm}$ and less using electron beam lithography (EBL) on freestanding silicon nitride membranes $\sim 70\text{nm}$ in thickness. These aperture arrays, when used to separate two reservoirs of superfluid ^4He near the superfluid transition temperature T_λ , act interchangeably as Josephson junctions (weak-links) or phase-slip centers (strong-links) depending on the degree of coupling between the reservoirs [2]. The coupling strength decreases with the ratio of the coherence length ξ (see Fig. 1.1) of the superfluid order parameter to the aperture linear dimension. To have a superfluid ^4He weak-link in an accessible temperature regime (a few mK away from T_λ) therefore requires channels with dimensions comparable to the healing length (90nm or smaller). Submicron holes/slots have been fabricated for superfluid studies by several groups [48],[49], in the past using ion milling techniques in thin foils. Nanofabrication of such apertures via EBL has also been reported in the literature [50],[51]. The recipes developed

here are based partly on techniques described in [52].

Here, we present the methods used in our group to fabricate such aperture arrays at the Cornell Nanoscale Facility (Ithaca, NY) over the past few years. We discuss the salient features of the processes involved as well as the difficulties encountered - both solved and unsolved - and suggest possible ways to improve the techniques in future iterations. The aperture arrays must satisfy some very stringent requirements to be useful in superfluid experiments - the nitride membranes must be strong enough to survive pressure differentials of several hundred mPa or higher as well as thermal cycling between 300K and 2K and the aperture sizes must be robust over time and fairly uniform across the array. Also, the fabrication process must be scalable so that dozens of chips containing arrays with several thousand apertures each can be made in a reproducible and cost-effective way¹. Focused Ion Beam (FIB) lithography is ideal for making nanoscale holes but is not scalable in this sense. EBL is therefore the default method.

4.2 Fabrication outline

The process can be broken up into three major parts: making freestanding silicon nitride membranes using photolithography, shooting apertures in them with EBL, and proper storage and cleaning before use. Fig. 4.1 shows the fabrication procedure (step numbers in the list match up with part numbers in the figure).

4.2.1 Making freestanding membranes

1. Double side polished (DSP) 100mm silicon wafers with 1,0,0 surface orientation and 400mm nominal thickness² are coated on both sides with $\sim 70\text{nm}$ low stress ($\sim 200\text{MPa}$ tensile) silicon nitride (Si_3N_4) in an LPCVD (low pressure chemical vapor deposition) furnace. A chrome-on-glass contact photolithography positive mask with clear areas for the windows is made using an optical pattern generator (the GCA/Mann 3600F).
2. The backside of the wafer is spin-coated with $\sim 2\mu\text{m}$ thick Shipley 1818 positive photoresist and the pattern is transferred onto it from the mask using a contact exposure tool (EVG 620).
3. The developing step removes the resist from the regions to be etched.
4. The exposed nitride is then etched away with a CHF_3/O_2 plasma in a reactive ion etch (RIE) tool (the Oxford Plasmalab 80) leaving the underlying silicon exposed.

¹With unlimited funds, this is not an important consideration.

²This wafer thickness is standard for DSP wafers. Standard single side polished (SSP) wafers however, are about 500 microns thick. Given the relatively insignificant cost fraction of DSP wafers in these processes, SSP wafers (essentially a cost-cutting measure) are unnecessary given the small quantities of wafers needed.

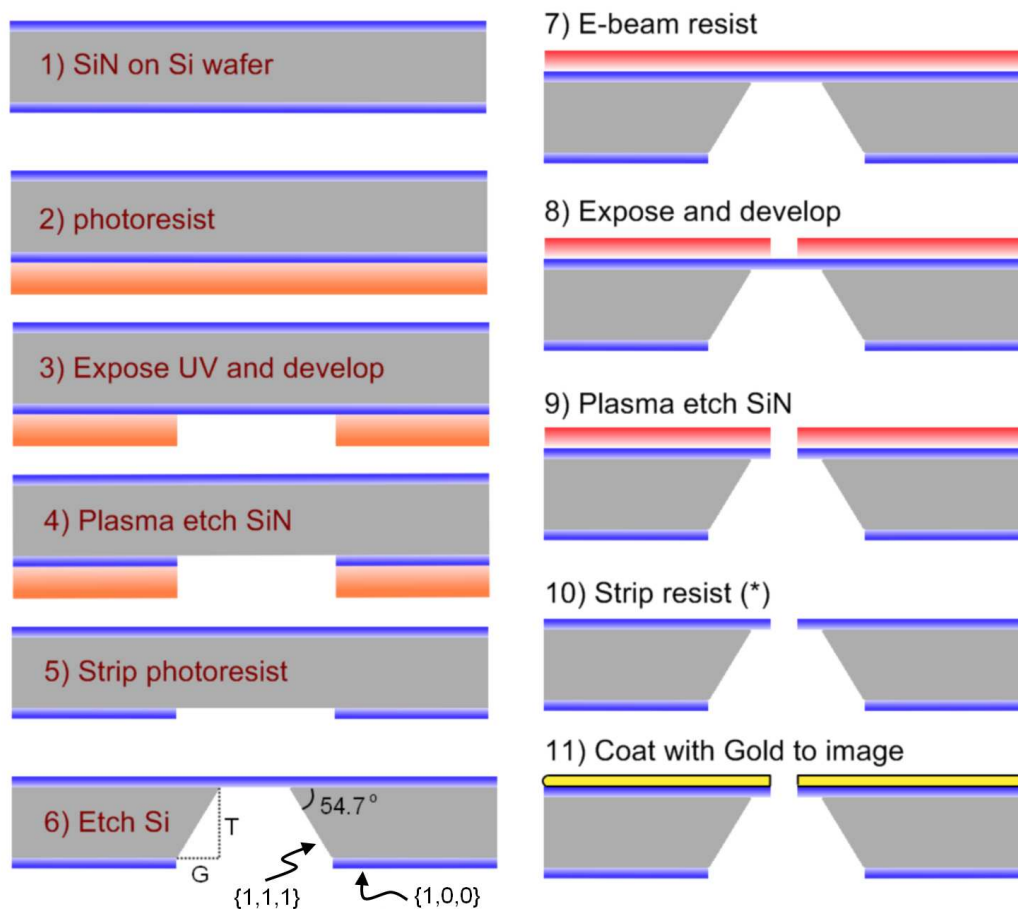


Figure 4.1: Wafer processing steps for fabricating free-standing Si_3N_4 membranes (Section. 4.2.1) and shooting apertures on them using electron beam lithography (EBL) (Section. 4.2.2).

- The resist is then stripped away in a hot chemical bath (70°C) containing propylene glycol, NMP (N-Methylpyrrolidone) and TMAH (Tetramethylammonium hydroxide). TMAH has been recently found to be more toxic than previously thought, so it may not be available in a given cleanroom. However, plasma-based strippers should be avoided whenever possible to avoid stressing/wrinkling the nitride films (more on this later).
- The silicon is then etched anisotropically from the backside in a hot KOH bath at $\sim 92^\circ\text{C}$, giving an etch rate of $\sim 110\text{mm/hr}$ and almost infinite selectivity to the nitride (used as a hard mask in this step). This leaves the wafer with freestanding nitride membranes on the front side and cleave lines and coordinate codes etched into the backside. The former enable us to snap single chips off the wafer with great precision

simply by pressing a sharp edged tool (blade, glass slide, etc.) on the other side of the cleave line while the latter are invaluable in cataloging the chips by type (the marks are visible under a powerful optical microscope), thus making it possible to shoot several array variants on the same wafer with impunity.

4.2.2 Shooting the aperture arrays

7. The wafer is then spin-coated with $\sim 140\text{nm}$ of 4% PMMA³ 495K and baked at 170°C for 15 min. As a rule of thumb, the PMMA thickness should be around twice the smallest feature size for accurate sizing. Since the selectivity of PMMA to the CHF_3/O_2 plasma etch is not much better than 1:1, this lets us over etch by at most a factor of two. This is necessary, as we have found the published etch rate ($\sim 54\text{nm}/\text{min}$ for the Oxford Plasmalab 80 etch tool) to be overly optimistic when it comes to ensuring that the apertures are etched all the way through the membrane. In general, RIE is suppressed in constrained geometries [50] so that published rates tend to be more meaningful for negative patterns.
8. The array patterns (150×150 and 300×300 with 1 mm spacing and 50×50 and 100×100 with $3\ \mu\text{m}$ spacing - aperture sizes of 250nm, 150nm, 90nm and 70nm) are shot on entire wafers in a 100keV e-beam tool (the Leica VB6 and the JEOL 9300FS have both been used successfully) by exposing small spots at a time to the electron beam. Exposure doses are calibrated by shooting dose matrices for each array type and aperture size and optimal doses are determined by exhaustive SEM imaging. Electron beam exposure breaks the polymer into fragments that are dissolved preferentially by a developer such as methyl isobutyl ketone (MIBK). MIBK alone is too strong a developer and removes some of the unexposed resist. Therefore, the developer is usually diluted by mixing in a weaker developer such as isopropyl alcohol (IPA). We use a 1:3 solution of MIBK:IPA for high contrast, low sensitivity. Raising the concentration to 1:1 can improve sensitivity significantly with only a small loss of contrast [53], but this will need to be tested if increased sensitivity is felt necessary in the future.
9. A full four inch wafer is etched in a Plasmalab 80 for 2 minutes to ensure proper etch-through - the etch time clearly a variable parameter that depends on the tool used. Wafer pieces tend to etch much faster. **The etch recipe is the single most critical parameter to be adjusted in order to obtain smaller hole sizes reproducibly and must be painstakingly re-calibrated for a new tool, aperture size, resist thickness or membrane thickness.** It is easy, though tedious, to do this etch-calibration. One obtains EBL-exposed, nearly identical chips and etches them for systematically varied etch recipes (chiefly, etch time is varied since the etch chemistry and pressure are pretty well calibrated for a given material in a well-equipped

³PMMA, or Poly(methyl methacrylate) is a positive e-beam resist with very high resolution but low etch resistance.

cleanroom). Anecdotal evidence suggests that the plasma flow is not spatially uniform within the chamber so that sample placement in the chamber can also affect the etch rate. This means that samples from all over a single wafer should be used to calibrate the etch rate. It is possible that some chips may have through holes while others don't, if the chosen etch time is just barely sufficient to cut through the membrane. It may be wisest to etch smaller blocks of chips at the same location in the chamber for the most reliable results since merely over-etching can increase the hole size beyond nominal values.

10. Once etched, the residual PMMA is stripped in a barrel etcher (Branson/IPC P2000), taking care not to let the temperature go over $\sim 150^\circ\text{C}$ (see section 4.3.1). If necessary (it usually is), the resist strip is performed in stages of duration (typically 2 min, but the temperature must be monitored) such that the temperature stays below this limit. Waiting several minutes between steps for the wafers to cool back down to near room temperature ensures that the intrinsic stresses are not changed, thus avoiding weakening of the membranes, which was observed after regular (directional) plasma cleaning processes. Two of these 2 min cleanings (with time in between for cool down) was found to be sufficient to clean out an entire wafer with (exposed, developed and CHF_3/O_2 plasma-etched) PMMA on the order of 140 nm thickness.

4.2.3 Post-processing: cleaving, storage and cleaning

11. The finished wafers should be thoroughly dried and stored on top of a blank, clean wafer in wafer storage boxes with securing springs to prevent motion. For transporting them from the cleanroom to our lab, we pack these boxes in nitrogen filled sealed packaging and thoroughly protect the packages with foam before shipping or transporting.
12. The resulting chips are then cleaved out of the wafer by hand (only when needed and only in a clean environment), sputter-coated with a thin Au-Pd film (5-10nm) and imaged (see Fig. 4.2) in a scanning electron microscope (SEM). A statistical sampling should be performed on several chips from different locations on the wafer and it must be assumed that the imaged chips are essentially unusable in experiments due to the sputter-coating, whose effects in the experiment have not been determined⁴ and more importantly, due to hydrocarbon contamination from the vacuum pump of the SEM.
13. The cleaving can be very difficult. We wished to avoid finer particulate contamination from using a diamond saw, so we cleave chips by scoring the backside of the wafer with a diamond tipped scribe to remove the partitions between preexisting cleave lines (which are just shallow trenches to define the cleaving crack). Then, placing the wafer backside down (with nitride membranes up), we hold down both sides adjoining a

⁴It may well be completely harmless but this would require an actual test to verify. We did not consider it a high enough priority to waste a cooldown on since there are always more variables than we can reasonably tweak (to test) for a given cooldown.

cleave line with opened tweezers and push down on the cleave line with a sharp, clean (and rigid) razor blade or other sharp straight-edged tool. Done properly, this should crack the cleave cleanly and the tweezers should hold the chips down so they don't go flying off. With a full wafer, we start with removing the curved sides first until we have just a 20×20 chip-block in hand. For such a big piece, the razor method doesn't work (it only really works when we get down to blocks of 5-6 chips on a side). Instead, we hold the big block and place it against a secured sharp edge (say a very clean lab jack edge) with the front side touching the edge at 45° such that the cleave line is parallel to the edge. Then, gentle pressure on the block should snap the block cleanly in half. Needless to say, this takes infinite patience and a steady hand and lots of practice (save wafers with badly etched membranes for this purpose or use good spares with no apertures - this is important!). Continue to break things in half until the blocks are small enough to use the razor method on. The easier the wafer is to cleave, the weaker it is during the processing steps. We are, however, not at the optimal point yet and there is room to improve the ease of cleaving while still maintaining wafer strength.

14. Our most reliable runs came after we started to store chips in dessicator jars⁵ under vacuum. Before each run, the aperture array chips are glued in as the very last assembly step after everything else is good to go (to avoid exposing them to atmosphere for too long, and then too are confined to a clean, filtered laminar flow bench). Before gluing them in, they are cleaned for ~ 100 min in a commercial UV/Ozone cleaner⁶ by placing individual chips on two clean silicon wafer pieces with the nitride membrane exposed on top and bottom to allow access to the ozone. This has been sufficient for chips that are already clean to begin with (PMMA stripped using a barrel etcher) and it must be noted that UV/Ozone cleaners can only clean organic contaminants that form volatile oxides⁷.

4.3 Issues

There are 3 overarching issues to be considered: structural integrity of the membranes, throughness of the apertures and reliable characterization.

⁵SPI (Structure Probe, Inc., West Chester, Pennsylvania) makes inexpensive polycarbonate dessicators. We hacked in to these by epoxying in a metal diaphragm valve and putting Apiezon grease on the rubber O-ring provided. Silica gel dessicant cartridges (also sold by SPI) should also be used to ensure a very dry environment as water films could close apertures and concentrate impurities at the holes upon evaporation.

⁶Model UVO 42-220 cleaner (made by Jelight company, Inc., Irvine, California) found in the Molecular Foundry at the Lawrence Berkeley National Labs.

⁷See Ref. [54] for details on the UV/Ozone cleaning process. One can even build an inexpensive cleaner using information in that reference and by searching around on the web.




Figure 4.2: SEM image of 150×150 array of ~ 75 nm apertures spaced $1 \mu\text{m}$ apart at 7kV accelerating voltage. (Inset) Close-up of a single aperture at fast scan speed (Section 4.3.3). The images are of the front side with Au-Pd sputtered on for contrast.




Figure 4.3: SEM micrographs of backside of window showing defects arising from using single side polished silicon wafers. The smallest defects seen are tens of microns in size.

4.3.1 Membrane integrity

Pinhole formation:

The wafer front side (that forms the membrane) must be protected as much as possible during the photolithography step when it is most roughly treated. A chemical similar to photoresist but with no photosensitivity (FSC-M) is spun on this side and baked for several hours so that any subsequent hotplate steps do not melt the layer nor does it outgas. Also, a 1/4 inch perimeter including the edges has to be recoated and baked just before the plasma etch (section 4.2.1) to prevent subsequent KOH etching of the wafer edge as this can deposit nitride scraps on to the windows and generally cause the wafers/windows to be more fragile and susceptible to cracking.

Membrane relaxation/wrinkling

This relates to the general issue of whether to use dry or wet etch tools in a process. While a dry etch is generally much cleaner and thorough, wet etches tend to be much gentler, especially on thin and consequently fragile membranes. The nitride membranes, just after deposition are originally under tensile stress (~ 200 MPa). Thin film stress has two compo-

nents - intrinsic and thermally induced [55],[56]. In attempting to strip the PMMA off the wafer in the last step (section 4.2.2), the intrinsic stress change due to ion bombardment or implantation (a byproduct of low pressure plasmas) was found to non-uniformly relax and thus wrinkle the membranes, thereby making them fragile along the edge at random stress points that can be seen under an optical microscope using phase contrast microscopy (see Fig. 4.4 for an example of such stress points).

This was confirmed by using a chemical plasma instead (as in the Branson/IPC P2000 barrel etcher), which is struck in a denser gas (~ 1250 mTorr as opposed to ~ 50 mTorr for kinetic plasmas) - the resulting smaller mean free path drastically reducing the bombardment on the membranes. There is only a small amount of hysteresis in the membrane stress under thermal cycling [57] so that the stripping could be done in steps, preventing the temperature from rising above $\sim 150^\circ\text{C}$, where thermal stress might permanently set in. We cannot over-emphasize the importance of this step as we have found it to be the only way to thoroughly clean off the resist while at the same time not weakening the membranes. Resist residues (migrating over time) are the single most likely suspect in the hole-closing affair⁸ to be discussed in the next section.

The precise recipe followed for PMMA thicknesses and plasma etch recipes⁹ used here is to strip in the Branson etcher for about 2 min (which heats up the chamber to around 150°C - the temperature must be monitored and the process stopped earlier if needed). At this point, the PMMA should be nearly stripped away, but the process should be repeated (after first waiting long enough for the chamber to cool back down close to room temperature, or at least $\sim 30^\circ\text{C}$) to ensure there are no trace residues, especially within the apertures. Smaller chip blocks must be balanced on glass slides or similar objects to ensure that there is a clear path on both sides of the chip for plasma to flow without obstruction.

4.3.2 Aperture throughness

The etch recipe used is the most sensitive parameter affecting the throughness of the apertures. If this is done properly, all care must be taken to further prevent the apertures from closing up. It is still an open question whether we have discovered and addressed all the factors responsible for this, but the conclusions so far are as follows:

Exposure and etching:

While the minimum critical exposure dose determined from dose matrices (usually $\sim 1000 \mu\text{C}/\text{cm}^2$) during EBL is a useful number, we have found that any dose greater than this minimum tends to work well (up to $\sim 3000 \mu\text{C}/\text{cm}^2$). This is because we are working with $\sim 60\text{nm}$ membranes so that substrate backscatter is negligible and proximity effects are not evident since the array spacings are usually more than ~ 10 times the aperture size. The

⁸Gabor Somorjai, personal communications.

⁹Plasmas like CHF_3/O_2 tend to polymerize the PMMA and make it difficult to lift it off using chemical etchants like Methylene Chloride.




Figure 4.4: A membrane that was stripped of photoresist using an asher (RF driven oxygen plasma) observed under a phase contrast (Nomarski) microscope after first sputtering with Au-Pd. The wrinkles and edge stress points can be clearly seen. Such membranes were found to be too weak for superfluid experiments as the minimum differential pressure applied during cell evacuation, filling, etc. of ~ 1 bar was observed to break the membranes with little exception. Visible wrinkling was not observed without the sputtering for the weak membranes but it is still a strong test for membrane integrity as the membranes stripped in a hot chemical bath showed no wrinkling regardless of sputtering.

etch recipe on the other hand is extremely delicate and the optimal etch time usually lies within a tiny window to avoid underetch (blocked holes) and overetch (much bigger holes).

Process contamination:

The nitride etch in section 4.2.2 can be done by a CHF_3/O_2 plasma or a CF_4/H_2 plasma. The former gives greater selectivity, which is crucial for this process but it also results in polymer deposits at the end of the etch. The oxygen is used to prevent these deposits but the selectivity suffers with increasing oxygen content. The polymer makes it extremely difficult to strip away the residual PMMA after the EBL using even a wet chemical stripper as strong as methylene chloride (in an extreme case - left overnight with mild agitation), which is why we have found dry etch tools much more effective in this step. This is a crucial point as apertures that were through up to this point can easily become clogged due to an imperfect wet strip. While almost any plasma etcher works very well to thoroughly clean the wafers, chemical plasmas (as opposed to kinetic plasmas) should be used for reasons discussed in section 4.3.1. The higher temperature also helps in burning off the polymer deposits within the constraints imposed in section 4.3.1.

The wafers must be cleaned thoroughly after the KOH etch (section 4.2.1) or else the residual potassium ions react with chlorine (probably environmental traces) to form copious crystalline deposits that clog the apertures. This is particularly insidious as the growth is rather gradual and difficult to detect. After our initial samples showed evidence of such hole size reduction, a closer SEM inspection in addition to a spectrographic analysis confirmed the presence of KCl crystals blocking the apertures (see Fig. 4.5). Several rinses with deionised (DI) water, followed by a dilute HCl rinse, followed by DI water again and ending

with acetone and isopropanol washes before blow-drying with nitrogen has since solved this problem. In particular, water must never be allowed to dry on the wafers as that is very efficient at aggregating impurities.




Figure 4.5: (Left) Spectrographic analysis (in an SEM) of deposits found on membranes confirms the presence of KCl. (Right) After cleaning in (just) DI water, the gross contamination lessened considerably, but apertures remained blocked, presumably because water surface tension makes it difficult, if not impossible for it to properly clean inside the apertures. Therefore, this cleaning step must be done prior to shooting the aperture arrays.

In the past, we would selectively coat small sections of the wafers with Au-Pd (by masking the rest with Al foil) and image entire wafers in the SEM. This turned out to be a mistake as SEM imaging tends to deposit hydrocarbons (probably from pump oil fumes, even with filters in place), in some cases at sufficiently high rates as to completely close the apertures¹⁰. It is more difficult to deal with large numbers of small chips during imaging but that has proved entirely too necessary.

Experiment contamination:

As in the SEM, pumps used in the experiment can back stream oil fumes toward the chips - inline filters suffice to prevent this. Since the chips are exposed to atmospheric air (even within a class 100 clean room), water films forming across the apertures and consequently freezing at the cryogenic temperatures in our experiments can also be a problem. Bake-out would be recommended, but is not possible in our assemblies. This is particularly difficult to detect, since the severity of the problem can actually depend on the prevailing weather!

¹⁰Real time SEM videos have been captured of this process, which occurs over half a minute.

4.3.3 Characterization challenges

Charging vs. SNR (signal to noise ratio):

Si_3N_4 is an insulator and therefore quickly gets charged while imaging in an SEM. Au-Pd can be sputtered onto the surface to provide a grounding path. The optimal accelerating voltage is then found to be around 7kV for uncharged viewing. Also, slow scans are preferred to increase the SNR and get crisp images but this results in more hydrocarbon deposits and decreased apparent sizes of the apertures. Surprisingly, we have found that the contrast provided by fast scans (fewer averages) is quite adequate as far as measuring hole sizes is concerned (see Fig. 4.2(inset)). Also, if the etch is insufficient, it leaves a thin remnant nitride film on the backside of the chips which can be highlighted quite clearly by the sputtered Au-Pd film. The stress of the sputter coated film must be matched to the nitride stress to prevent wrinkling due to stress competition. This can be tricky as the sputtered film stress can vary dramatically and even change from tensile to compressive over tiny pressure ranges in the argon chamber (~ 1 mTorr). In general, we find that higher pressures (~ 50 mTorr above base vacuum) tend to give more tensile Au-Pd films that match well with the nitride.

Backside imaging:

To sputter coat and SEM image the backside of the windows, the chips have to be mounted with the delicate window side down; this can be done by cutting a small slit in double-sided copper tape and sticking the chips with the window suspended over the slit. This should be done in any case to provide a path for evacuating air during the processes without developing a pressure differential across the window (that might break it).

4.4 Wafer considerations

Our masks contain 400 chips (3mm square in a 20x20 square array) with windows for 200 mm and 400 mm square final membrane sizes, cleave lines and row/column coordinate code grids (visible under a strong optical microscope) for record keeping. The KOH etch process used to make the membranes is anisotropic - KOH attacks the 1,0,0 plane through the back side of the wafer preferentially to the 1,1,1 plane - leaving cavities in the silicon substrate with the characteristic 1,1,1 sloping walls and the freestanding membranes on the front side. The etch angle (54.7°) is fixed by the silicon crystal geometry (see step 6 of Fig. 4.1), which makes the final membrane size strongly dependent on wafer thickness (T). The grow size (G) for making the windows is simply the length added to each edge (parallel or orthogonal to the major flat) of a feature. From elementary trigonometric analysis, G is nominally 282 mm for standard double-side polished (DSP) wafers and 354 mm for standard SSP wafers. For example, a mask designed to make 200 mm square membranes in 0.4 mm thick wafers (standard DSP) would make 56 mm square membranes in 0.5 mm thick wafers (standard single side polished: SSP). Cleave lines have to be similarly resized to ensure that chips can

be easily cleaved out for individual use. The upshot of all this is that a mask can only be used for one thickness of wafer. In practice, DSP wafers are so easily/cheaply available in recent times that this shouldn't be an issue.

Further, for SSP wafers, the rough side¹¹ is used as the backside for the photolithography step and due to diffuse reflections off this surface during contact photolithography UV exposure of the resist, may give unsatisfactory results for the smaller (~ 5 mm) coordinate codes unless exposed for a longer time. The windows themselves are relatively unaffected but longer exposures do tend to enhance contrast and give crisper edges. An SEM micrograph of some of the defects arising post-KOH etch due to the coarseness of the surface is shown in Fig. 4.3. We have found exposure times of 6 sec and 15 sec and nitride etch times of 3 min and 6 min optimal for smooth and rough sides respectively.

4.5 Conclusions

Based on our work thus far, we seem to have reached a point where getting apertures smaller than 70 nm consistently and reproducibly (and not as a mere accident of contamination) is limited by the etch step. To get to smaller aperture sizes, we may have to attempt several different schemes; for instance, using a carbon hard mask instead of PMMA to improve the selectivity to nitride and find an effective way to strip off the carbon [58]. The newly emerging technique of nanoimprint lithography is also a promising avenue to be explored.

Finally, in addition to imaging, we would recommend testing aperture arrays using the gas flow tests described in Section 6.1 to obtain an independent measure of the aperture size. This test is much quicker and simpler than the full superfluid experiment and can help reliably determine if the aperture arrays are within spec.

¹¹This side remains rough even after nitride deposition as the roughness (tens of microns) is of a much larger length scale than the nitride thickness (~ 60 nm).

Chapter 5

The displacement sensor

Brief overviews of the two types of SQUID-based displacement sensors that we have used for superfluid experiments, and a list of their pros and cons were provided in Section 3.1.2.

The superconducting displacement sensor (persistent current type) is the most structurally complicated component of weak-link cells. The physics and optimization of these displacement sensors, specifically for the kind of experiments described here has been covered in some detail in dissertations of previous students from our group¹, so we avoid reproducing those details here.

The magnet type sensor is covered in some detail by Sato, *et al.* in Ref. [44] and the only new issues not discussed in that reference (resonant frequencies of magnet-loaded membranes) have been treated thoroughly in Sections 10.4 and F.2.

In this chapter therefore, we will present only a brief description of the physics and proceed directly to the practical aspects of fabricating, assembling and using these devices. Except for the pancake coil, engineering drawings are included in Section A.3 of Appendix A.

5.1 Persistent current type

This type of sensor relies on a superconducting diaphragm placed next to a spiral wound pickup coil (“pancake coil”), which is an integral part of a so-called “persistent current circuit”. Referring to the entirely superconductive circuit shown in Fig. 5.1 and Ref. [59, pp. 139-148] for equations, a persistent current is circulated in both loops. The bulk of the current is confined to the lower loop, which contains the pancake coil (bare inductance L_{PC}) and a (much larger) tank inductor (L_T). The upper loop contains the input loop (L_S) of a commercial dc SQUID magnetometer². Two sections of the circuit wire (marked L_1 and L_2) are wound around resistive heaters (H_1 and H_2) in order to switch those sections between

¹Ref. [59, pp. 139-148] is particularly comprehensive and useful. Refs. [60, pp. 93-96] and [26, pp. 249-252] are briefer synopses, although there is an error in Fig. A3 of the last reference listed, where current injection lead (b) should actually be placed at the other end of the inductor L_T .

²Quantum Design (San Diego, California) model 550 SQUID with model 50 controller.

normal and superconducting states at will. These persistent current switches (PCS) are used to inject current into the circuit (or change the injected current) in a manner described later in this chapter.

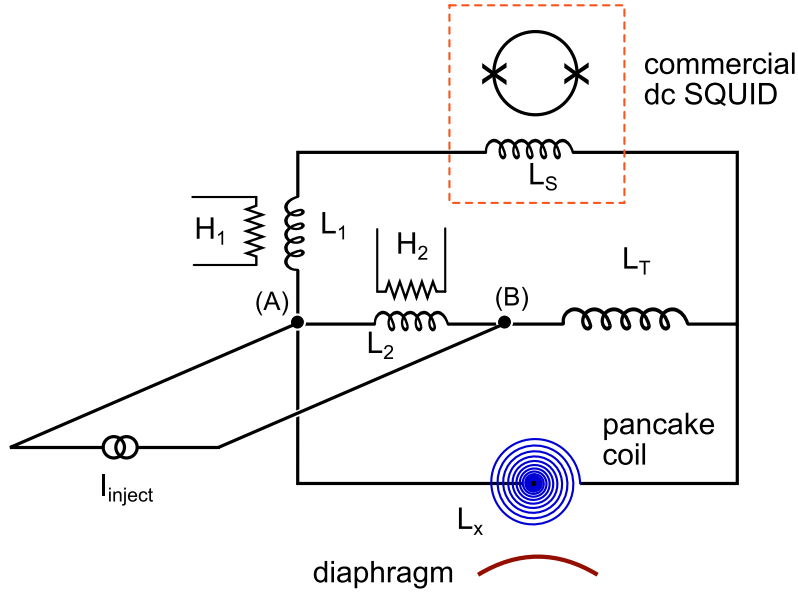


Figure 5.1: Persistent current (PI) circuit.

The *effective* inductance of the pancake coil, including the effect of the nearby superconducting diaphragm plane is

$$L_x = \mu_0 n^2 A x \quad (5.1)$$

where μ_0 is the vacuum permeability, n is the radial turn density of the pancake coil, A is the coil area (assuming $A \leq A_{diaphragm}$) and x is the spacing between the coil and diaphragm. If the diaphragm moves, the effective inductance L_x changes, which in turn changes the inductance of the lower loop (and hence the magnetic flux in the lower loop). Since flux is a conserved quantity in a closed superconducting loop, the lower loop current changes in response to a change in L_x . This in turn adds or removes some current from the upper loop to conserve charge. The SQUID is therefore used in this application, as a sensitive ammeter to measure this small change in input coil current. Changes in the SQUID input current show up as changes in the magnetic flux incident on the SQUID loop, which correspond to changes in the feedback signal applied by the SQUID electronics to maintain a stationary flux. This feedback signal (voltage) is the final output of the SQUID that we can measure.

From the sequence of the events just described, we can see that this output signal is proportional to the initial diaphragm displacement that produced it. We simply define an effective proportionality constant α such that $\Delta x \equiv \Delta V_{SQ}/\alpha$ (see Eq. (G.8)) and determine this constant empirically as described in the various calibration sections of Chapter 10.

5.1.1 Typical parameters

We choose values for the various dimensions to obtain optimal values for the SQUID sensitivity and dynamic range. The (flux) sensitivity can be shown [59, p. 141] to be:

$$\frac{d\phi_{SQ}}{dx} = -\frac{M I_1}{x} \left(1 + \frac{L_S}{L_x} + \frac{L_S}{L_T} \right)^{-1} \quad (5.2)$$

where M is the (internal) coupling between the SQUID input coil and the loop containing the Josephson junctions and I_1 is the current in the lower loop of the circuit (in most cases, this is nearly equal to the total current injected in the circuit). We would like to keep the (potential) sensitivity easily adjustable to high values merely by adjusting the injected current. Towards this end, we ensure that the rest of the parameters are kept optimal so they do not decrease the sensitivity. Therefore, $L_T \gg L_S$ ($L_T \sim 200 \mu H$ for ~ 400 close-packed turns in 4 layers filling the tank inductor former shown in Fig. A.24 and $L_S \sim 1 - 2 \mu H$ from the SQUID specs). While we would like $L_x \gg L_S$; we find in practice that $L_x \sim 4 \mu H$ (for the pancake coil former shown in Fig. 5.2 with ~ 36 turns of the 4.2 mil superconductive wire mentioned in the next section and with $\sim 150 \mu m$ spacing between the coil and diaphragm³), so that the ratio is just about 1/2. Still, with an injected current of $\sim 1 A$, we find ourselves limited by ambient vibrational noise rather than the SQUID sensitivity so that this is not something that overly concerns us at the present level of SHeQUID development.

The injected current can be as high as some critical value at which the diaphragm is penetrated by flux lines (this saturates the sensitivity). This issue is discussed further in Section 6.2. When everything is properly made, currents of at least $\sim 1 A$ should be easily possible. We turn now to a discussion on how to make the various parts for the circuit.

5.1.2 Fabrication

We use 4.2 mil NbTi wire⁴ with Cu-Ni cladding and an overall formvar coating to make the components and connecting wires for the circuit. Since this wire is delicate and difficult to handle, we sometimes use a thicker (9 mil) wire with similar structure. We will henceforth refer to these wires as simply 4.2 mil SC wire or 9 mil SC wire to signify superconductive wire.

Persistent current switches (PCS')

The PCS' are made by winding 5 – 6 turns of 4.2 mil SC wire around a 100 Ω metal film resistor and gluing it in place with Stycast 1266 epoxy. This can be cumbersome but there is an easy way to do it. We tape the resistor leads onto a Teflon sheet, carefully wind the wire around the resistor body and tap the wire securely to the sheet. Then, a small amount

³This spacing estimate includes a 76 μm 300 HN Kapton spacer, glue and diaphragm thickness and taking into account the finite thickness of the wire.

⁴Supercon Inc., Shrewsbury, Massachusetts

of Stycast 1266 is sufficient to embed the wire turns completely. Too much Stycast can be counter-productive since we want the heat to dissipate quickly once the PCS is de-energized. To confine dissipated heat to only the resistor body, we snip off the resistor leads almost completely and solder 9 mil SC wires (stripped but with Cu-Ni cladding intact) to the ends. Since the initial Stycast gluing will have formed a small puddle, leaving the top resistor body a bit bare, we can do another gluing step with the PCS upside down. At this time, we can also embed the resistor lead solder joints in Stycast.

The inductance owing to the handful of turns around the resistor comes out to around $\sim 0.5 \mu H$ (L_1 and L_2 in the circuit). About 5 V across the resistive heaters should be sufficient to make the wound wire normal under ordinary circumstances.

Tank inductor and current injection chokes

The tank inductor is made by winding ~ 400 close-packed turns (in 4 layers) of 4.2 mil SC wire around a former machined out of Stycast 1266 stock (see Fig. A.24). This gives $\sim 200 \mu H$ of measured inductance. Two more similar (but differently dimensioned) inductors are made to serve as RF chokes inline with the two current injection leads. These inductors (~ 730 turns of 4.2 mil SC wire in 3 layers for an inductance of $\sim 150 \mu H$ for each inductor) sit inside an aluminum box with lead sheets glued on for shielding (we find this to be much more robust than relying on lead-plating as aluminum can be machined with great ease and lead sheets are much more reliable as shields). See Figs. A.25 and A.26 for drawings of the filter inductor former and filter box, respectively.

These formers are difficult to machine out of Stycast stock (the filter inductor former is probably impossible owing to its small diameter⁵) and we recommend making them inside aluminum molds. These molds can be made very easily by taking aluminum rod of appropriate thickness and drilling in from both sides with a flat drill (or end-mill). Fresh Stycast 1266 can be gently poured in (or injected in with a syringe) and the two ends capped off with scotch tape (doesn't stick to Stycast). The molds can be machined a bit to remove as much aluminum as possible without touching the Stycast and then etched in 1M NaOH⁶ to remove the aluminum.

We use a coil-winder (a small lathe with a turn counter) to wind these large coils. A notch is made on one end of the former to enable feeding the wire out before the first turn. An appropriate length of starting lead length is bunched and taped to the coil-winder chuck and the wire guided in over the notch. With a gentle grip on the wire and the spool freely spinning on a fixed rod nearby, we run the coil-winder and guide the wire to obtain close-packed turns in multiple layers over the former. We always end the coil at the start point (along another notch directed such that the two notches guide the free ends of the wire towards each other), tape down the ends temporarily and glue the coil in place using either Stycast 1266 or GE

⁵The dimensions for the molds are chosen to obtain the largest inductance possible with the constraint that the fields generated within them during current injection ($\sim 1 A$) do not exceed the critical field of the lead sheets used to shield them.

⁶see Section A.1.6 for a detailed recipe

varnish. If using varnish, at least a day should pass before mechanically stressing the coil as the solvent used for the varnish (typically toluene) can temporarily craze the wire insulation and cause it to break when stressed. Once dry, the leads are twisted together manually for a bit, then taped securely and the twisting completed using a slow rotary tool (as described in Section 7.4.2).

Pancake coil

The spiral wound pancake coil (PC) can be notoriously difficult to make without some experience at it. This is only because there are several small issues that, if ignored, can needlessly complicate things. We will start with making the coil former from black Stycast (2850FT) and then winding a coil onto it.

Referring to Fig. 5.2, the main body of the former can be either machined from Stycast 2850FT stock or made directly in a mold. The critical feature on the former is the central post, which is just 6 mils high and 20 mils wide. This material is chosen for its well-matched thermal expansion to brass (to prevent stress-induced distortions upon cooling down). Silicon carbide tool bits should be used for working with black stycast as it is infused with quartz powder and is extremely hard on regular tool bits (to the extent that it is difficult to make even a single pass with the tool before abrading a significant part of the cutting edge). Copious amounts of cutting oil or other lubricant must be used to protect the tool (even carbide tools). Clean off debris frequently as the slurry is rich with abrasive quartz.

A good way to make the central post reliably and accurately is to raise the tool bit on the lathe tool post by a small piece of 10 mil shim stock and *then* adjust the vertical position of the tool bit progressively until the tool faces off the former at dead center. Once this is achieved, we gently remove the 10 mil shim without changing the tool vertical setting so that the tool bit drops by exactly 10 mils below center. Now, we feed the tool longitudinally by exactly 6 mils and face off the former again. Due to the lowered tool offset, the facing operation now leaves a central post 20 mils wide and 6 mils high. Of course, it should be possible to get the post through clever use of molds. It is known [61] that aluminum molds work quite well for Stycast 1266 as well as 2850FT, in that we can etch away the aluminum once the epoxy is set using a 1 molar solution of sodium hydroxide (we have confirmed⁷ this for Stycast 1266).

The hard part is getting the 6 mil wide (and deep) groove around the central post to accommodate the first wire loop. We make this by grinding a custom groove tool bit (which, consequently cannot be of carbide and is therefore short-lived) of the right size and going in gently, flush against the central post. A straight groove is cut into the surface *very carefully* by holding the former in a vise and drawing a sharp X-acto knife blade radially from the central groove outward several times (one can hold up a metal ruler as a guide). This straight groove is gently deburred with a miniature file so that the surface remains smooth and the two grooves are ensured to be in contact under a microscope. Performing these finishing

⁷See Section A.1.6 for details.




Figure 5.2: Pancake coil former for PI-style sensor. (Vector graphics can be zoomed in indefinitely on screen).

steps under a low power microscope can make things a lot easier. The top edge is rounded (around the groove escape point) with abrasive tools for the wire to enter the groove without any sharp bends. See Fig. 5.4 for a side view of the former.

The former is cleaned ultrasonically in soap water, acetone and isopropanol baths (about 5-10 min each), blow-dried and kept in a clean box thereafter. Referring to Fig. 5.3 henceforth, the former is mounted on a home-made winding jig consisting of a central plastic platform (drawing provided in Fig. A.27) that is screwed vertically onto an aluminum base with a small piece of Teflon between the platform and the aluminum as a lubricant. The screw is tightened just enough that the platform rotates freely but without any wobble. A small screw tightened into a threaded hole in the side of the base intersecting the vertical screw orthogonally is used to hold the screw in place so it doesn't rotate with the platform. We have designed the former so that the platform *has to* be rotated clockwise, which will not loosen the vertical screw. All surfaces on the jig are carefully deburred so that the wire will not scratch by accident. A few feet starting length of 4.2 mil SC wire is wound around a step on the platform and scotch taped in place and the wire guided into place as shown in the figure into the straight groove on the former (the groove escape and the small slot in the platform well are lined up before tightening the horizontal set screws into the former's




Figure 5.3: Pancake coil winder.




Figure 5.4: Pancake coil former for PI-style sensor (photo).

side). Past the former, the wire is clamped in a miniature vise that has two felt pads glued on the inside of its jaws. This lets us hold the wire with a steady tension without risking breakage or abrasion of the wire. Scotch tape is put on all surfaces that may have sharp edges to protect the wire. The spool is free to rotate on its own stand and the wire is not under tension at the spool.

A thick, clear plastic window is glued onto a simple, commercial ball-bearing, which is glued onto an aluminum plate with a matching hole. This plastic window goes over the former now with a small piece of ink-jet printer transparency⁸ sandwiched between (free to move). Ensuring that all surfaces are properly aligned orthogonal to each other, and with the window tightened gently (not too much) with thumbscrews, we now start slowly turning the platform clockwise so that the wire loops around the central post. While doing so, tiny drops of freshly made Stycast 1266 are placed on the incoming wire from the spool end to fix the coil in place once wound. We also used to put a small drop on the central post before starting the winding but find that it works marginally better without doing so (it may well help as long as it is an *extremely* tiny drop). Stop sending in stycast drops for the last 1 – 2 turns. The window (and bearing) should move with the coil once several turns have built up. We continually inspect the coil through magnifying lenses under bright lights to ensure tightly wound and circular turns. Any distortions in the central post or incorrect winding will make the turns elliptical. Loss of the tension while winding, wobble in the platform or just breathing wrong can destroy hours of prep time. It is therefore prudent to have several formers at the ready in case things go wrong. Formers can be cleaned of undried stycast 1266 and reused but the central post invariably needs to be remade because it is designed to deform slightly during winding (under window pressure) and lock the first coil to some

⁸Laser printer transparency is a bit softer and sometimes leads to distorted coils

extent.

After about 35 turns, we reach close to the former's edge. Most of the displacement sensitivity comes from the hubward (rather than perimeter) turns so the last few turns are not especially crucial. The final turn should be looped down the rounded edge near the straight groove escape point and held down with tape until the epoxy dries (overnight usually but at least 8 hours before handling). This piece of wire should be epoxy free so that it can be twisted around the starting wire once dried.

Once the coil is dry and passes inspection, the wire is cut from the spool to appropriate lead length and the two leads are manually twisted around each other (see Fig. 5.4). Once a long enough length has been twisted, it can be taped to the platform and a slow rotary tool used to finish twisting the rest of the leads. This is not a question of laziness⁹ – machine wound twisted pairs are straighter and fit more easily into (shielding) lead tubes as compared to hand-wound ones. A small section of Teflon tubing is slipped over the leads to protect the fork in the wires from being scratched by the lead tube. The coil must be carefully inspected under a microscope to ensure that the surface is flush and there are no overlapping turns or too much glue thickness or other anomalies (it doesn't take a more than a wire thickness to lower the sensitivity considerably).

Superconducting diaphragm

30 HN Kapton¹⁰ – nominally $7.6 \mu\text{m}$ thick – is clamped between an evaporation mask (see Fig. A.28 and a blank plate, both machined out of $\sim 1/16''$ thick aluminum plate. If the evaporator being used is unduly directional, thinner plate can be used, but that entails a different processing scheme to make the mask. For 5 mil thick brass shim stock, we have successfully machined a mask by sandwiching it between two flat (finely sanded), thick and clear plastic plates and gluing it in place with Plexiglas or PVC cement (clamped tightly). Once dry, the sandwich is machined easily and the plates prevent the shim stock from warping. This is a bit wasteful of plastic, but we do not need to make these masks very often (only if there's a diaphragm dimension change). Alternately, brass shim stock can be processed via photolithography (the home hobby version should work quite well), where a pattern is transferred from a printed transfer paper (or even transparency) to the shim using a household iron (details can be found on the Internet on hobby sites – the printer apparently matters quite a lot) and the brass etched in ferric chloride solution. Using photoresist and a UV lamp and a mask made out of transparencies might be overkill, especially for the kind of large feature sizes and rough resolution we need, but if one is already set up for such things (a PCB setup for instance), it may make sense to just go that route.

Regardless, using this evaporation mask, we need to deposit three layers: $\sim 50 \text{ nm}$ Al on the Kapton to provide a sticking layer for the lead to follow; $\sim 400 \text{ nm}$ Pb for the

⁹Not entirely anyway.

¹⁰Dupont corporation, available from a number of distributors. For the small quantities we need, ordering some sample sheets can be enough to last several years.

superconducting layer and finally, another ~ 50 nm Al as a protection layer for the lead (against oxidation and mechanical flaking).

Issues: Lead evaporation is extremely expensive through industrial sources due largely to regulatory policies in the United States (upwards of \$ 1000 for a single evaporation run, which, however can contain several diaphragms and might well suffice for several years assuming no dimension changes). This may be different in other countries. In academic facilities, lead evaporation is getting more and more difficult to find for reasons having to do with the high volatility of lead and its tendency to contaminate the insides of the evaporator, which cannot subsequently be used for more critical applications without thorough cleaning. Sadly, obtaining lead films has become more difficult over time, due purely to sociological (rather than technological) reasons.

Niobium has been reported to work just as well in displacement sensors. However, it has such a high melting point that niobium deposition tends to be done most easily through sputtering where intrinsic stresses in the film can easily wrinkle the diaphragm. Schwab [62] reports a simple way to prevent wrinkling, where he glued the Kapton sheets to a stainless steel ring, ~ 3 " in diameter, which had 3 holes on the ring for alignment to a larger holder. The film was glued to the ring so that the Kapton was flat, without tension. Then these rings with Kapton sheet glued across, were put into the sputter chamber. For membrane sizes that are small enough, the longer wavelength ripples from this method do not affect anything once the membranes are cold. Schwab also mentions that there are more sophisticated ways to control the tension of the deposited Nb film to be found in the literature.

Note that choosing Pb over Nb merely because the former is a Type I superconductor while the latter is Type II would be misguided because thin films typically behave as Type II regardless of bulk behavior [59, p. 146]. Therefore, problems like flux creep would be present in either kind of film. Evaporating Nb (which tends to avoid the wrinkling problems mentioned earlier) can be done using E-beam evaporators, but even that requires rather large power supplies and might not be readily available.

5.1.3 Assembly

Wire/joint preparation

For normal joints to this kind of wire, we need to strip off the formvar insulation cleanly. Techniques on doing this are described in detail in Section 7.4.2. For superconducting joints, the Cu-Ni cladding must be removed as well. We can do this by dipping the exposed wire in 50% diluted nitric acid. For best results, we have the acid ready in a small (10 mL) glass beaker on a glass petri dish (for safety). We also have a small beaker of water and a beaker of sodium bicarbonate (simple household baking soda) dissolved in water ready. The acid and base should be kept at a distance (for obvious reasons). Some lab wipes¹¹ and isopropanol

¹¹Kim wipes appear to be a near-universal standard.

(IPA) are also kept ready. Alcohols should also be kept far from acids as they can sometimes form explosive mixtures.

The formvar-stripped wire is dipped in the nitric acid, whereupon we observe vigorous bubbling at the (gold-hued) Cu-Ni surface. Shining a light on this can be helpful. The solution starts getting bluer over time. Once the Cu-Ni has been etched away, the reaction stops, leaving behind the somewhat black NbTi surface. If formvar was imperfectly stripped, strands of residual formvar freed by the removal of the Cu-Ni substrate will be seen dangling around the wire (one should re-evaluate the stripping method used in this case). Immediately after this, we neutralize any remaining acid on the wire by dipping in the baking soda solution and subsequently rinsing with water and wiping clean with IPA to remove water residues. Typically, any water cleaning (of any parts) should be followed by an IPA rinse or wipe to avoid deposition of any impurities the water is carrying (unless it is distilled water).

Typically, multiple wire-ends can be etched in a 10 mL beaker. Since the reactivity of the acid goes down over time (as it gets bluer due to dissolved copper salts), it is wise to re-etch all the etched wires in a second acid dip to ensure it is really clean. Remember that we will not find out about the success or failure of these joints until we get to 4 K. We must be over-cautious and extremely paranoid about such things in order to not waste valuable cooldowns tracking down these problems later.

Oxide layers form on all surfaces over time. A good rule of thumb is to make superconducting joints within 1 – 2 hours of nitric acid cleaning of wires. If not, re-dipping in acid is a prudent pre-assembly step. Lightly sanding the etched wires (and any other superconducting pads, etc.) with fine grit (1500 or higher) sandpaper is also a safe thing to do.

Note that acid-etching wires on the cryostat should be done only after covering it with aluminum foil to protect components, wiring and plumbing from acid fumes. This is also necessary while stripping formvar from wires if some especially nasty stripping methods are used (see Section 7.4.2).

Spark-welding

Spark-welding¹² is the more robust of the two types of superconducting joints we discuss here. As such, we try to use it everywhere possible and use screw joints (the other type) only when absolutely unavoidable. Prior to welding joints, wires are stripped of formvar and the Cu-Ni cladding removed as described previously. Gloves must be worn at all times to prevent contamination of wire surfaces. 2 mil thick Nb foil¹³ is scrubbed with fine grit sandpaper and wiped clean with IPA. Small sections ($\sim 1/4'' \times 3/8''$) are cut out with scissors and folded thrice to give a three layer sheath ($\sim 1/4'' \times 1/8''$ in size). The sizes are of course not critical but keep in mind that the sheath opening will be completely welded off spot by spot and keeping it small is much more robust. To enable screw-joints between already

¹²We are grateful to previous group members for passing down the base technique described here. This is sometimes referred to as heliarc welding because it is performed in a helium atmosphere. However, that term has a precise industrial meaning, so we use spark-welding to avoid unnecessary confusion.

¹³99.8% Nb foil from Alfa Aesar, Ward Hill, Massachusetts.

welded wires and a third wire, we sometimes drill screw holes in these Nb sheaths prior to welding. To do this in such small objects, it is helpful to make a jig by bending a thick steel shim in half and drilling a clearance hole in it. Then, the Nb sheath is clamped between the two shims (by hand) and drilled through the hole with a small hand drill.




Figure 5.5: Spark welding setup for superconducting joints.

It is important that the wire surfaces and the Nb sheath surface be clean and the oxide layers removed by either acid etching or fine sanding. We are now ready to spark-weld the joint. Referring to Fig. 5.5, the Nb sheath with the two wires to be welded inserted inside it, is crimped and held between the cross-hatched jaws of a stainless steel surgical hemostat, which is then held securely in an insulated vise. A thick braided cable connects the Nb sheath (via the hemostat) to the negative terminal of a 30 V (1-3 A) power supply. The spark is provided by a thick Nb tip inserted in a simple, commercially available inert gas welding handle hooked up to a helium gas bottle. Like any inert gas welding setup, a shield over the tip ensures that the spark area is flooded with helium so that the metals do not oxidize. The tip connects (through a sealed cable) to the positive terminal of the power supply. A set of 10 electrolytic capacitors ($\sim 3300 \mu F$ each) are connected in parallel with the power supply output.

The (bare) superconducting ends of the wires are gently twisted together and inserted in the Nb sheath so that a little bit of the unstripped wire goes into one end and only the superconducting parts poke out the other end (and the twisted wires are pushed to one edge of the sheath). The capacitors are charged up to $\sim 20 V$ and the tip is brought very close to the wires to initiate the spark. Dark glasses should be worn to protect the eyes during this time. We have found it helpful to not move the tip towards the wires; instead holding the tip offset parallel to the wires and gently move the tip orthogonally past the wires, momentarily coming close enough to initiate the spark. This helps avoid the tip getting stuck on the sheath. If the tip does get stuck, the power supply will overload (it should have sufficient protection against this and this **will** occur at some point). Though it is difficult (without practice), it is important to fight against the instinctual reaction to pull the tip away as this could break the fragile wires¹⁴. Instead, think of the tip getting stuck as a common event

¹⁴Another danger with jerking the tip away is that the wire is dragged out through the hemostat jaws,

(practice helps). Leave it stuck, turn off the power supply and let the capacitors drain out to ground. Then, gently rocking the welding handle back and forth will eventually break the tip free in a safe manner. We can simply continue with the welding after this.

Our goal is to weld the wires together directly at first (as described). Subsequently, we “stitch” the open end of the sheath closed by a continuous series of spot-welds. This ensures a robust and mechanically strong joint.

Before doing a real joint, one should experiment with different voltages to ensure that joints (say, on an empty sheath) do not look charred (too much energy dumped by the spark) and look like a continuous flow of metal under a microscope. Watch for cracks in the joints and adjust the voltages, helium flow and other parameters until things look right. Lots of practice ensures a steady, unflinching hand when working with actual critical components.

Screw joints

Screw joints are significantly easier than spark-welds but still carry most of the same precautions and prep work. Wires must be cleaned and etched as before and all superconducting surfaces (pads, wires, sheaths) should be sanded clean with fine grit sandpaper and then wiped clean with IPA. Sanding small pads embedded in shielded boxes (like in Fig. 5.6) can be done by using a simple jig made by rolling up a small piece of sandpaper, inserting one jaw of a tweezers inside and then folding the roll in half (lengthwise) to insert the other jaw into the other end of the roll. Of course, there are plenty of sanding tools (sanding strips, twigs, etc.) available to do this more elegantly. For single wires to be connected at a screw pad, it is wise to either weld it into a Nb sheath with a screw hole or at the very least cut a small Nb foil washer to place between a steel (or brass) washer and the wire before screwing it down. Care must be taken to ensure that the bare NbTi part of the wire has been properly contacted to the Nb pad and is not simply twisted around the screw. Lock-washers must be used to ensure that thermal cycling does not loosen the screws and undo the joint.

Final assembly guidelines

Converting the circuit from Fig. 5.1 to the real-world circuit in Fig. 5.6 is a non-trivial task for superconducting circuits since it is not merely the connectivity but the actual relative placements that matter (in terms of mutual inductance, noise coupling, etc.). To this end, we try to keep the high current components separated from the low current components as much as possible. The sections in the box are all shielded from each other. Lead sheets can be used for further shielding by simply gluing pieces in with common cyanoacrylate glues (“superglue”), which do maintain mechanical strength down to cryogenic temperatures.

strips the insulation and the high current passing through the sheath and into ground now goes through the stripped wires, burning off the insulation and in general causing a nice little disaster (any connected components are now also suspect for shorts or frayed insulation). Needless to say, this is not a hypothetical scenario and we’ll just leave it at that.




Figure 5.6: Persistent current circuit shielded box (“PI box”). Since the dimensions of the box will depend critically on cryostat space and wiring choices, we simply provide a photo here instead of a schematic in order to show where everything goes. This photo shows the circuit in the process of being assembled. This box was designed (and probably made) by Emile Hoskinson/Tom Haard.

It is very important that the wires going to the SQUID be disconnected from the SQUID prior to any spark-welding as the high current could damage the SQUID. The SQUID is, in general a very delicate device, which is especially sensitive to damage from electrostatic discharge. Therefore, anti-static gloves must be worn at all times¹⁵ when handling a bare SQUID (even if it is inside its shield can). Periodically touch a large (grounded) metal rack or other metal surface with one’s hands and instruments (screwdrivers, tweezers, etc.) before touching the SQUID to be even safer. With all these precautions, the incidence of mysteriously damaged SQUIDS has slowly dwindled to nothing. Storing the SQUIDS in anti-static bags is also a good idea.

The tank inductor and the two persistent current switches are placed in their respective slots and their wires connected appropriately. The leads to the SQUID are connected carefully with the usual prep work on the wires and Nb pads. The leads from the PI box leading to the pancake coil (PC) in the cell are connected to a pair of Nb pads in an intermediate shielded box where they are further connected to the PC leads coming from the cell. This way, the persistent current circuit can be made *once*, with the SQUID and the interface leads from the top plate semi-permanently hooked up and the only thing that needs to be

¹⁵Fleece clothing should be avoided at all costs, especially in dry weather. These SQUIDS die very easily.

connected before each new cooldown is the set of PC leads in the intermediate breakout box. The PI circuit and the SQUID need never be opened up or disturbed. If cared for properly, it should last over dozens of cooldowns without need for maintenance. The SQUID leads and the PC leads are shielded inside lead tubing¹⁶.

After wiring up the circuit, all leads are secured to insulated surfaces (tape helps to insulate surfaces) using GE varnish¹⁷. This is especially important for high current carrying wires as even small relative motions can induce large noise voltages in nearby components/sections.

5.1.4 Current injection procedures

Referring again to the PI circuit in Fig. 5.1, we describe¹⁸ how a persistent current is circulated in the PI circuit, for two different starting points. We assume that all inductances stay constant during the injection and (for estimations only) that $L_x \approx L_{PC}$.

Starting with zero injected current

1. With $I = 0$ everywhere in the circuit, we begin by turning on both heaters: H_1 to protect the SQUID input from seeing large currents¹⁹ and H_2 to make the small wire section L_2 normal. For 100Ω heaters, we use a voltage of $4 - 5 V$ (more generally, a few hundred mW should be sufficient power).
2. The current source is now slowly ramped up from 0 to some final value I_{inj} . Of the three paths across the points A and B, only the lower branch (containing L_T and L_{PC}) is superconducting and this step ends with current I_{inj} flowing in that branch.
3. Now, H_2 is turned off so that the segment L_2 becomes superconducting a few seconds later. H_1 remains on for now, as does the current source (steady at I_{inj}). Let the currents in the L_2 and L_{PC} branches be denoted as I_2 and I_p , respectively. At this point, if any of the currents change, a voltage V_{AB} will be induced across points A and B by Faraday's Law: $V_{AB} = -L_2\dot{I}_2 = -(L_T + L_{PC})\dot{I}_p$. But charge conservation gives us $I_2 + I_p = I_{inj} = \text{constant}$, so that $\dot{I}_2 = -\dot{I}_p$. Since the currents are thus forced to always change in opposite directions and Faraday's law forces them to change together, they end up being forced to not change at all to satisfy both conditions. Therefore, when H_2 is turned off, the current continues to flow from the source only into the L_{PC} branch ($I_p = I_{inj}$) and I_2 stays 0.

¹⁶Lead tubing is 88% Pb/10% Sn/2% Ag (Lead-Tin-Silver) Solder with a hollowed out core from GWR instruments (San Diego, California). Techniques to make this in-house from commercial solder wire can be found in Ref. [61].

¹⁷Should be available online from cryogenic accessory suppliers such as Lakeshore or CMR direct.

¹⁸We are grateful to Satoshi Murakawa for a spirited brainstorming session where we rediscovered the detailed dynamics of the current injection that is presented here.

¹⁹The SQUID can handle a maximum current of $\sim 20 mA$ at its input coil [59, p. 143].

4. The source current is now slowly ramped from I_{inj} down to 0. It is worthwhile to understand what happens during this rampdown. Since charge is conserved at each instant, we still have $I_2(t) + I_p(t) = I_{source}(t)$ during the rampdown. At the beginning of the rampdown, $I_p = I_{inj} = I_{source}$ and $I_2 = 0$. At the end of the ramp, when $I_{source} = 0$, we must have $I_2(t) + I_p(t) = 0$, so that $I_2(t) = -I_p(t)$. This only tells us that the currents in the two branches are equal and opposite, i.e. that there is a net *circulating current* in the entire lower loop of the circuit.
5. The SQUID protection heater H_1 is now turned off. It is interesting to note here that no current enters the upper loop at this point (this can be shown by an analysis similar to that in step 3 above). Only when the effective inductance L_x of the pancake coil deviates from its (fixed) value during the injection (due to diaphragm displacements), does some current get diverted to (or from) the upper loop in order to maintain constant flux in the lower loop.

But how big is the circulating current that we injected? We can answer this question two ways: (i) by integrating the 2 coupled differential equations generated by Faraday's Law and current conservation, or (ii) by remembering that the flux is a conserved quantity in a superconducting circuit. Going the second route (starting at step 4, with H_2 off), we note that the initial flux in the entire lower loop at the beginning of the rampdown was $\Phi_i = (L_T + L_{PC})I_{inj}$ and the flux at any instant during the rampdown²⁰ is $\Phi(t) = -L_2I_2 + I_p(L_T + L_{PC})$. Using the current conservation equation and equating Φ_i and $\Phi(t)$, we finally obtain the instantaneous values of the currents in the two branches:

$$I_p(t) = [I_{source}(t) - I_{inj}] \frac{L_2}{L_T + L_{PC} + L_2} + I_{inj}$$

$$I_2(t) = [I_{source}(t) - I_{inj}] \frac{L_T + L_{PC}}{L_T + L_{PC} + L_2}$$

These equations tell us that if we start with some initial current $I_{source}(0) = I_{inj}$ being provided by the source, the current in the PC branch starts at $I_p(0) = I_{inj}$ and decreases as the source current is ramped down. At the same time, the current in the L_2 branch starts at $I_2(0) = 0$ and increases during the rampdown. Values at any intermediate time during the rampdown are provided by the above equations. At the end, if the source current reaches 0 at some time t_f , the final values of the two currents will be: $I_p(t_f) = I_{inj} \frac{L_T + L_{PC}}{L_T + L_{PC} + L_2}$ and $I_2(t_f) = -I_{inj} \frac{L_T + L_{PC}}{L_T + L_{PC} + L_2}$. These are equal and opposite and denote a circulating persistent current in the lower loop.

²⁰Note here that writing the junction rule the way we do is tantamount to a sign choice that the source current entering junction A (or B) splits into I_2 and I_p where both *go away from* the junction. This means that at the points A and B, the current directions are chosen to be positive away from each other. This further implies that the fluxes for the two branches are opposing each other for this choice of sign convention. So, the total flux in the lower loop will be $-L_2I_2 + I_p(L_T + L_{PC})$, where the (-) sign ensures that we add the fluxes properly.

Note that as long as the combined inductance of the tank inductor ($\sim 200 \mu H$) and pancake coil ($\sim 3 - 4 \mu H$) is much larger than the inductance (L_2) of the wire wound around the heater ($\sim 0.5 \mu H$), the circulating current is nearly equal to the current injected by the source.

Changing a preexisting injected current

We include this procedure for completeness.

1. Assuming an initial circulating current I_{inj} in the circuit, we start by turning on H_1 to protect the SQUID. This quenches any (small) current circulating in the top loop. Recall however, that if the diaphragm is the same position as when we injected the current, there shouldn't be anything in there to quench.
2. The source current is ramped up slowly to match I_{inj} .
3. H_2 is now turned on. Matching the current in the previous step ensures that nothing much will happen to any other currents.
4. After waiting a few seconds, the current is now slowly ramped up or down to the desired new value: $I_{inj, new}$.
5. H_2 is turned off.
6. After a few seconds, the source current is turned down to zero. This leaves a circulating current $\approx I_{inj, new}$ in the lower loop.
7. H_1 is turned off.

5.2 Magnet type

The magnet type sensor is covered in some detail by Sato, *et al.* in Ref. [44]. The essence of this type of sensor can be stated almost trivially: a magnet is mounted on the flexible diaphragm and when it moves, it changes the magnetic flux seen by a fixed pancake coil next to it, which is connected directly to a commercial dc SQUID. The SQUID reads the changes in magnetic flux, which are proportional to the diaphragm displacement (and we can calibrate this in the same way that we described for the PI sensor). We discuss here (very briefly) the fabrication of the pancake coil (which is actually easier than that used in the PI type sensor) and the (normal) diaphragm with glued magnet.

The pancake coil former is machined from Stycast 2850FT in a similar way as we did for the PI sensor. The brass holder piece (cell E-ring) for the coil is machined first so that the height of the smaller diameter stub can be adjusted on the fly during machining. The goal of this adjustment is to ensure that the stub with the central post sticks out a little past the surface of the coil holder piece. This way, we can glue the coil into the holder with the

coil mated flush with the surface. The coil can be wound in the same winding jig that we used for the PI sensor with no change in procedure. Since only 5 turns are needed (for the magnet mentioned below, in order to get around the same sensitivity as we get for ~ 750 mA of injected current for the PI sensor), the winding is considerably easier and ellipticity and other distortions (unless wildly exaggerated) are largely inconsequential. See Fig. 5.7 for an engineering drawing and Fig. 5.8 for a photo of the finished pancake coil.

The diaphragm is now easier to make as we only need a normal metal (aluminum works well). See Ref. [44] for a note on Al deposition (ensuring that it remains normal near T_λ). Around 200-300 nm of Al is sufficient for our purposes. The magnet used here is an N50 grade neodymium (NdFeB) disk magnet²¹ (1/16" diameter and 1/32" thickness), axially magnetized with a nominal weight of 0.0118 g, and a specified surface field strength of 0.18 T. As shown in Fig. 5.9(a), the diaphragm is laid (Al face down) on a transparency (for stiffness) with a scale drawing of a radial grid printed on it. This grid is used to sight the center of the diaphragm with some accuracy and paint a single dot of fresh Stycast 1266 at that point using a blunt toothpick (slightly *less* than the quantity shown in the photo would be better). The magnet is cleaned by running it gently between the folds of a Kimwipe wetted with IPA. It is then placed gently on the Stycast dot using plastic tweezers, being careful to hit it dead center so as not to spread the Stycast around. Fig. 5.9(b) shows the result. One should see a *very* thin ring of Stycast around the magnet to know that it's secure.

One must be careful with storage and handling of these magnets as shocks and exposure to stronger magnets can distort their fields. The nominal sensitivity was manually attained by Sato, *et al.* [44] by trying different magnets and number of turns in the coil, assuming nominal behavior for the magnets.

²¹K&J Magnetics, Inc, Jamison, Pennsylvania




Figure 5.7: Pancake coil former for magnet-style sensor. Made of Stycast 2850FT.




Figure 5.8: Photo of pancake coil former for magnet-style sensor.




Figure 5.9: Magnet-loaded diaphragm. (a) Kapton diaphragm (Al-coated side face down) centered using printed transparency with single dot of Stycast 1266 in the center. (b) Magnet has been glued to the diaphragm.

Chapter 6

Independent component tests

The more critical parts of the cell can be tested independently of the full experiment to save time. This is especially true for new designs/recipes of the aperture arrays since they have historically been prone to failure and being able to test several samples very quickly and at liquid nitrogen temperatures in a simple table-top apparatus can be very useful.

The other critical component is the displacement sensor (original design from Section 5.1), whose persistent current circuit contains several superconducting joints which require some experience to make reliably. It is helpful to test out the circuit separately if possible to ascertain that it can sustain a persistent circulating current. However, this can be impractical and one usually tests for persisted currents by quenching the current and measuring the decay of the resulting voltage step in time. The superconducting flexible diaphragm (that is the main sensing element in the sensor) can have significantly lowered critical fields so that it is penetrated at persistent current levels that are too low to afford sufficient displacement sensitivity. This happens if the superconducting film is too thin or patchy or oxidized. There are certain signatures for both these components that can signal whether they will work optimally and it is worthwhile to perform these tests to avoid significantly longer downtimes with the complete superfluid experiments. Also, in the full superfluid experiments, it may so happen that with the added complexity, problems with individual components cannot always be easily tracked down.

In this chapter, we will discuss suggested independent tests for the components described above.

6.1 Aperture arrays

6.1.1 Introduction

Experiments with superfluid ^4He and a single nanoscale aperture array on a silicon chip performed by our group [1, 2, 3, 7, 38, 44, 63] involve two reservoirs of superfluid separated by the aperture array. In practice, this is realized (see Fig. 1.20) as an inner cell capped

with a flexible metallized diaphragm on one end and the aperture array chip as the sole means of entry on the other end. A fixed electrode is placed next to the diaphragm and the duo define a capacitor whose value changes with the distance between them (the curvature is exaggerated in the picture - the actual movement of the diaphragm is much less than the equilibrium separation). In actual experiments, the motion of the magnet glued to the diaphragm results in a changing flux that is picked up by the (superconducting) pickup coil and read off the commercial dc SQUID connected to this coil.

While there already exist techniques [26, 39, 46] to measure the aperture size near $\sim 2K$ by measuring flow transients in the normal regime, this requires filling the cell, closing it with a cryo valve, and in general, having the machinery of the full experiment on hand. This chapter describes a much simpler setup that is (a) modular, so that multiple chips may be tested with quick turnaround time, and (b) operated at 77K, thus requiring a simpler dewar (the setup is essentially just a small, short probe dunked in an open mouth dewar and held on a lab stand). The experiment and theory are described in the following sections.

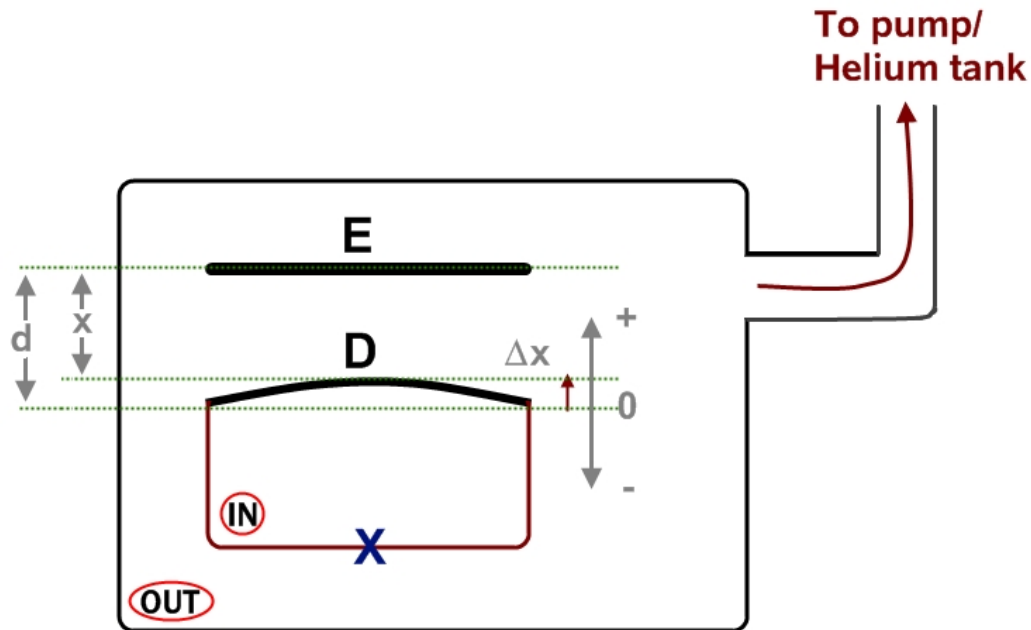


Figure 6.1: Flow test cell. Fixed electrode (E) and movable diaphragm (D) form capacitance $C[x]$ where x is the instantaneous D-E spacing. $x = d$ defines the equilibrium position (very stable at 77K). Δx is the mean instantaneous displacement of the diaphragm away from equilibrium (+ or -) as shown. See section 6.1.3 for more details.

6.1.2 Flow test experiment

As shown in Fig.6.1, the inner cell of the test apparatus is composed of an aluminum body with the aperture array chip (marked X) epoxy-sealed at one end and the other end sealed by a metallized diaphragm (D). A fixed electrode (E) next to the diaphragm defines a variable capacitance that depends on the position of the diaphragm, which in turn varies according to the pressure difference between the inner and outer cells, the outer cell being defined simply by a vacuum can enclosing the inner cell.

Apparatus

This setup consists of 4 pieces (engineering drawings are referenced for each part):

1. Probe (Figs. 6.11, 6.12): this is a thin wall 3/8 inch OD stainless steel tube hard-soldered onto a 1/2 inch thick brass flange with a central through hole that connects to one end of the tube. The other end of the tube has a standard KF hard-soldered on with a 4-way KF adapter clamped on it. Two of the ports have KF plugs with hermetic BNC connectors screwed on - these provide electrical access to the diaphragm and electrode. The third free port is used for gas handling (pumping out for filling with helium). It is important that all tube inductances be large compared to the aperture arrays so that they do not contribute to the pressure decay times. The main brass flange has a larger bolt circle with clearance holes and jacking screw taps for indium sealing to an overall vacuum can. It has a smaller diameter circle of 8X #4-40 blind-taps for mounting the experiment.
2. Electrode holder (Fig. 6.2): This is just a flat brass disk with mounting holes. A fixed electrode is glued onto this and the piece is mounted onto the probe over a set of commercial electrical spacers (to allow easy access for the wires to the central hole).
3. Diaphragm holder (Figs. 6.3, 6.4): Made of aluminum, this is (relatively) the most complicated part to make and has two distinct sides with the through hole in the center. One side has the flexible diaphragm glued on to it while the other side has an O-ring groove (and finely polished surface) and screw taps for an indium seal. This piece screws onto the electrode holder with the diaphragm facing the electrode and with a thin ($76\mu m$ thick) Kapton spacer between them. This assembly defines the capacitance that we will measure as a function of time.
4. Chip holder (Fig. 6.5): This is (by design) the simplest piece to make so that it can be mass-produced for single use with the aperture array to be tested. If needed, the aperture array chip, which is epoxied in place with Stycast 2850FT, can be easily removed by heating the piece (made of aluminum) and cleaning it up with sandpaper or a sandblaster gun. This piece indium-seals onto the diaphragm holder, so its seal surface has to be polished fine on a lathe. It has a bolt circle to match the indium taps on the diaphragm holder with two or more jacking screw taps to aid in removing the

seal. Note that a small amount of vacuum grease (Apiezon M or N has worked well) should be dabbed onto the indium before sealing so that it comes off cleanly.




Figure 6.2: Electrode holder

Engineering drawings (and fabrication guidelines) of parts are shown in Figs. 6.2, 6.3, 6.4 and 6.5. Photos of individual parts being assembled are shown in Figs. 6.6, 6.7, 6.8, 6.9, 6.11 and 6.12; and the fully assembled cell is shown in Fig. 6.10. The apparatus is designed to be highly modular so that any piece may be replaced independently if needed. However, the diaphragm and electrode holder assembly should ideally never need to be touched. After a chip test, one need only remove the chip holder indium seal, put on a new seal with a new chip holder (and a different chip to be tested), re-seal the vacuum can and be cold again within an hour or two. All holes are designed with a standard dividing head (15° increments) in mind for machining ease.

Procedure

1. Once completely assembled, the setup is tested at room temperature by measuring the capacitance between the diaphragm and electrode. It should be close to the calculated capacitance for a parallel plate capacitor with the chosen dimensions. If the chip being tested has too much flow impedance, it may happen than the diaphragm bulges




Figure 6.3: Diaphragm holder: chip facing side

out enough to touch the electrode after assembly, in which case the capacitor will be shorted.

2. Once the vacuum can is sealed, the probe is evacuated slowly (over a few hours at least to avoid putting too much differential pressure across the aperture array) through a metering valve. Once the probe pressure is low enough, it is switched over to a diffusion pump and pumped out to a high vacuum (a few mTorr at least at room temperature). At this point, if the capacitance is not nominal, it should be allowed a chance to relax some more (as gas flows out of the inner cell). If, after several hours, it still hasn't unshorted, it implies a near total blockage of the aperture array and there is little point in proceeding with the cooldown.
3. If however, the capacitance is nominal, a simple test can be conducted when the probe pressure is still high where the pumping is halted suddenly and the capacitance relaxation monitored. If some relaxation is visible, things look good. One can even use the analytical machinery in this chapter to obtain a very crude hole size from such rough transient tests.
4. Once all these things have been verified and the probe (and cell) is at high vacuum and the diaphragm has relaxed to an equilibrium position so that starting or stopping




Figure 6.4: Diaphragm holder: electrode facing side

pumping doesn't affect it anymore, the probe is gently lowered into a small dewar of liquid nitrogen (LN₂). A small amount of helium gas (just a squirt) can be injected into the probe at this point to help speed up thermalization of the parts. In any case, we will need to have a small pressure of helium¹ in the probe for the tests so we might as well put it some use. 2-3 hours should be good enough for stability.

- At this point, the capacitance will become much quieter and more well-defined. Rampant drifts observed at room temperature will suddenly disappear. This is the entire reason for doing this test at 77 K. The diaphragm becomes taut and less susceptible to static cling and floppy motion.
- Now, we need to generate a pressure step and allow it to relax via aperture flow so that the resulting decay time in capacitance (i.e. diaphragm position and hence pressure) can give us information about the flow impedance of the aperture array. There are two ways to do this: inject a bolus shot of helium into the outer cell or open the pumping valve suddenly to evacuate the outer cell. We find that the latter is more reliable as the gas injection takes a while to diffuse down into the outer cell. The evacuation step is much more instantaneous. The idea is to balance the capacitance bridge being

¹We need to be in the molecular flow regime with high Knudsen number, so pressures less than a few inches of Hg are optimal for hole sizes ranging from 15 to 100 nm. See Figs. 6.14 and 6.15.




Figure 6.5: Chip holder

used for the measurement, decide on the maximum dynamic range of the bridge and evacuate the outer cell only to the point where the capacitance being measured is still meaningful (i.e. on scale). The step should also not be so large that the capacitor shorts or the diaphragm bulges so much that the parallel plate paradigm is threatened. Again, the experimenter should run the numbers to decide these issues as they will affect the accuracy of the measurement.

7. We note in passing that another way to excite these transients (that is probably more optimal than the ones discussed before) is by putting an electrostatic force step (i.e. a voltage step) between the fixed electrode and flexible diaphragm as we do in our superfluid experiments. A circuit like the one shown in Fig. 10.18 can be used to do this since we will be measuring the capacitance of the same parallel plate pair across which we will be putting a sizable DC voltage. Despite everything, the lock-in analyzer used for the capacitance bridge will get overloaded for a few seconds just after the step is applied. If the transient decay times are on the order of hundreds of seconds (as they should be for the sort of aperture dimensions and diaphragms we use), this is irrelevant as all we need is a relatively exponential-looking section of the transient decay to fit to in order to obtain the decay time. In extreme circumstances, a tactic we have used before might come in handy, where we included a relay in the circuit just before the




Figure 6.6: Electrode holder (mounted on probe flange spacers) with electrode glued on and wire screwed on to tab using nylon screw.




Figure 6.7: Diaphragm holder with diaphragm glued on and wire connected.

lock-in that was initially open (to isolate the lock-in from the voltage step) and was closed after a short (~ 100 's of ms), user-specified duration after the step was applied (all computer controlled).

8. Once the desired capacitance step is induced, we stop pumping and allow the capacitance to relax on its own, recording the capacitance vs. time. This data will be used in later sections to determine the aperture array conductance and thence the average hole size.

6.1.3 Analysis

The capacitance is obtained in terms of the diaphragm displacement. The displacement is related to the pressure difference across the chip. The pressure difference as a function of time is then related to the gas (mass) flow conductance of the aperture array on the chip assuming the gas is in the molecular flow regime (see section 6.1.5 for a discussion of different flow regimes). The flow conductance can be obtained from the time constant found by fitting this model to the data. Of course, this analysis is valid only in the molecular flow regime (again, see section 6.1.5).

Capacitance

Approximating the setup as a parallel plate capacitor, the D-E capacitance is:

$$C[x] = \frac{A'\epsilon}{x} \quad (6.1)$$




Figure 6.8: Diaphragm holder from Fig. 6.7 flipped over and screwed onto electrode holder. The inner cell cavity and indium seal groove, surface and taps can be seen.




Figure 6.9: Disposable chip-holder. The square milled slot in the center is where the 3 mm square chip is placed and glued in place with 2850FT (black) Stycast. The lower surface is where the indium seal sits and is finely polished. Bolt holes (8X) for indium seal screws and jacking screw threads (2X) can be seen.

where ϵ is the permittivity of the medium ², A' is the metallized area of the diaphragm (or electrode) and x is the instantaneous spacing. Then, from Fig. 6.1,

$$x = d - \Delta x \quad (6.2)$$

so that,

$$C[x] = \frac{A'\epsilon}{d - \Delta x} = \frac{A'\epsilon}{d(1 - \frac{\Delta x}{d})} = \frac{C[d]}{1 - \frac{\Delta x}{d}} \quad (6.3)$$


If the spring constant of the diaphragm is k and the movable area of the diaphragm is A , and assuming that the diaphragm is in mechanical equilibrium at all times with the gas in the cell, the pressure difference (ΔP) across the chip that causes a displacement Δx is given by

$$\Delta P \equiv P_{in} - P_{out} = \frac{force}{A} = \frac{k\Delta x}{A} \quad (6.4)$$

so that,

$$\Delta x = \frac{A}{k} \Delta P \quad (6.5)$$

² $\epsilon \approx \epsilon_0$ here but the exact permittivity is not needed in this analysis.



**Quick change
chip-holder**




Figure 6.10: Assembled cell before being covered by vacuum can. Chip holder has been sealed onto the partly assembled setup from Fig. 6.8. Simple schematic shows basic assembly and indium seal. Chip-holder with chip glued on can be seen on the left.




Figure 6.11: KF 4-way breakout with labels for BNCs and pumping port.



Figure 6.12: Vacuum can

From Eqs. (6.5) and (6.3),

$$C[x] = \frac{C[d]}{1 - \frac{A}{kd}\Delta P} \quad (6.6)$$

For future reference, the above equation can be solved for ΔP :

$$\Delta P = \frac{kd}{A} \left(1 - \frac{C[d]}{C[x]} \right) \quad (6.7)$$

Note from Eq. (6.5) that a positive pressure difference ($\Delta P > 0 \Rightarrow P_{in} > P_{out}$) gives (as it should) a positive diaphragm displacement ($\Delta x > 0$).

A model for the time evolution of the pressure difference is now presented.

Pressure evolution

The following quantities are defined (all quantities in SI units unless specified otherwise):

- m_4 : Mass of a ${}^4\text{He}$ atom (or other gas at 77K)
- V_{in}, V_{out} : Inner and outer volumes as defined in Fig. 6.1
- N_{in}, N_{out} : Number of atoms in the inner and outer volumes
- Q_m : Mass flow through the aperture array (kg/s)
- k_B : Boltzmann's constant
- T : Absolute temperature

The mass flow conductance of a channel that admits a flow Q_m for an impressed pressure difference ΔP is defined as

$$G \equiv \frac{Q_m}{\Delta P} \quad (6.8)$$

We may express the mass flow in terms of the rate of change of the number of atoms in the inner volume. To stay consistent with prior notation ($\Delta P \equiv P_{in} - P_{out}$), the flow should be out of the inner cell for a positive pressure difference, resulting in a decrease in the number of atoms in the inner cell. With this convention, we have

$$\boxed{Q_m = m_4 \dot{N}_{in} = -G\Delta P = -G(P_{in} - P_{out})} \quad (6.9)$$

In what follows, we assume that the system is in the molecular flow regime where the conductance is independent of pressure and therefore constant during the transient. This is justified for reasons discussed in section 6.1.5.

Using the ideal gas law and differentiating once with respect to time,

$$\begin{aligned} m_4 \dot{N}_{in} &= -G \left(\frac{N_{in} k_B T}{V_{in}} - \frac{N_{out} k_B T}{V_{out}} \right) \\ m_4 \ddot{N}_{in} &= -G k_B T \left(\frac{\dot{N}_{in}}{V_{in}} - \frac{\dot{N}_{out}}{V_{out}} \right) \\ \ddot{N}_{in} &= -\frac{G k_B T}{m_4} \dot{N}_{in} \left(\frac{1}{V_{in}} + \frac{1}{V_{out}} \right) \end{aligned}$$

where we have assumed no leaks in the system so that $\dot{N}_{out} = -\dot{N}_{in}$. The inner and outer volumes do not change by much³ when the diaphragm moves so that they may be considered constant during the transient. Since $V_{out} \gg V_{in}$ by design, we may write:

$$\ddot{N}_{in} \approx -\frac{G k_B T}{m_4 V_{in}} \dot{N}_{in}$$

Using our master equation (Eq. (6.9)), we can rewrite this in terms of ΔP :

$$\frac{d}{dt} \left[-\frac{G\Delta P}{m_4} \right] = -\frac{G k_B T}{m_4 V_{in}} \left[-\frac{G\Delta P}{m_4} \right]$$

Finally,

$$\Delta \dot{P} = -\frac{G k_B T}{m_4 V_{in}} \Delta P$$

whose solution is an exponentially decaying pressure difference from an initially imposed pressure drop of ΔP_0 :

$$\boxed{\Delta P[t] = \Delta P_0 e^{-t/\tau}} \quad (6.10)$$

with a characteristic time

$$\tau \equiv \frac{m_4 V_{in}}{G k_B T} \quad (6.11)$$

³the maximum diaphragm displacement is $\Delta x \sim d \sim 76\mu m$, giving a volume change of less than $\sim 5 \times 10^{-9} m^3$ while the inner cell volume is $\sim 10^{-6} m^3$

where we recall that G is the conductance of the entire array.

Since the aperture spacing in the array is usually much larger than the aperture size, the array may be considered to be n equal conductances in parallel⁴, where n is the number of apertures in the array. From Eq. (6.11) therefore, the experimentally obtained conductance of a *single aperture* is:

$$G_{fit} = \frac{m_4 V_{in}}{\tau k_B T n} \quad (6.12)$$

Fit model

Using Eqs. (6.10) and (6.6), the full model becomes:

$$C[x] = \frac{C[d]}{1 - \frac{A}{kd} \Delta P_0 e^{-t/\tau}} \quad (6.13)$$

with 2 additional fitting parameters ($C[d]$ and $\frac{A}{kd} \Delta P_0$).

However, note that the last parameter is not an independent one. From Eq. (6.7), we find that the initial pressure drop is related to the starting value of $C[x]$ (i.e. $C[x_0]$) in the data used for fitting as follows:

$$\Delta P_0 = \frac{kd}{A} \left(1 - \frac{C[d]}{C[x_0]} \right)$$

Substituting this in Eq. (6.13), our final fitting model becomes:

$$C[x] = \frac{C[d]}{1 - \left(1 - \frac{C[d]}{C[x_0]} \right) e^{-t/\tau}} \quad (6.14)$$

where $C[x_0]$ is simply the first datapoint in the fitted data (may be used as a parameter if desired) and where $C[d]$ and τ (defined in Eq. (6.11)) are the fitting parameters.

The idea then is to measure the capacitance $C[x]$ during a pressure step transient, fit it to the above model and obtain the time constant τ , which (according to Eq. (6.12)) yields the experimentally measured conductance of a single aperture in the array.

To obtain the aperture size from this information requires additional theoretical input, which we discuss next.

⁴Conductances add in parallel. Intuitively, in the molecular flow regime, where the mean free path is greater than the channel diameter, I don't expect the proximity corrections for neighboring apertures to be very significant.

Aperture size

In the molecular flow regime, the conductance for a single channel takes the following form (see Section. 6.1.5 and [64, 65]):

$$G_{th}[u] = u^2 K[u] \left(\frac{3L^2}{8} \sqrt{\frac{32\pi m_A}{9k_B T}} \right) \quad (6.15)$$

where $K[u]$ is Clausing's factor, which depends on the aspect ratio of the channel $u \equiv R/L$ (where R is the radius of the channel and L is the length of the channel. The length here is simply the thickness of the nitride membrane, which we assume is known⁵). An important thing to note here is that some of the aperture arrays (specifically, the CNF apertures) tested were square in shape and the aperture sizes thus obtained must be viewed as approximations that nevertheless can be examined for trends. Besides, the etch inevitably rounds out the corners of the apertures so that the shape actually obtained is more of a rounded square.

Eq. (6.15) above will be finally set equal to the experimentally obtained value of the single-hole conductance (Eq. (6.12)) and the resulting equation solved for u (and therefore R - the hole radius). It is therefore extremely important here to possess a valid functional form⁶ for $K[u]$.

$K[u]$ exists as a definite integral to which analytical approximations have been constructed ([66, 67]). Clausing's approximation [66] was constructed as an improvement over that of Dushman [67] and that is what we will use in this chapter. A usable functional form for $K[u]$ is provided in Appendix C as Eq. (C.1).

The aperture size is now obtained by solving the equation

$$G_{th}[u] = G_{fit} \quad (6.16)$$

for u and extracting the aperture radius R (knowing the nitride thickness L).

While deviations from Clausing's conductance formula have been observed (e.g. [68, 69]), Srekanth [69] notes that for high Knudsen numbers ($Kn \gtrsim 1$, which is true here), the conductance was observed to obey the Clausing relation closely.

It is more efficient to solve Eq. (6.16) numerically (one look at Eq. (C.1) in Appendix C tells us why). To aid in this, we first plot $G_{th}[u]$ in Eq. (6.15) and find the approximate value of u that gives the observed conductance (G_{fit}). This is then used as a guess value for a numerical (iterative) equation solver to get a refined value for u (which is easily verified by applying Eq. (6.15) and checking it with G_{fit}). This iterative procedure has been coded up in Mathematica [70] and the code is provided in Section C.2.

It is observed that the solution for u depends only weakly on L (see the last three rows of Table 6.1), so that a given uncertainty δ in L results in a much lesser uncertainty in the diameter (at least for the runs we have studied so far).

⁵A good assumption - we measure the nitride thickness during the fabrication with an optical tool so it is a known parameter.

⁶A dense table of K values vs. u might also be sufficient as a crude function in the form of a lookup table

6.1.4 Some results

A typical capacitance transient is shown in Fig. 6.13. Results from four flow test runs (on different chips) are summarized in Table 6.1. Silicon nitride film thickness (L) was measured to be very close to 60 nm using the FilMetrics F40 system at the Cornell Nanoscale Facility, Ithaca, NY. The additional values in the table are provided merely to show the lack of sensitivity of the calculation to inaccuracies in L .

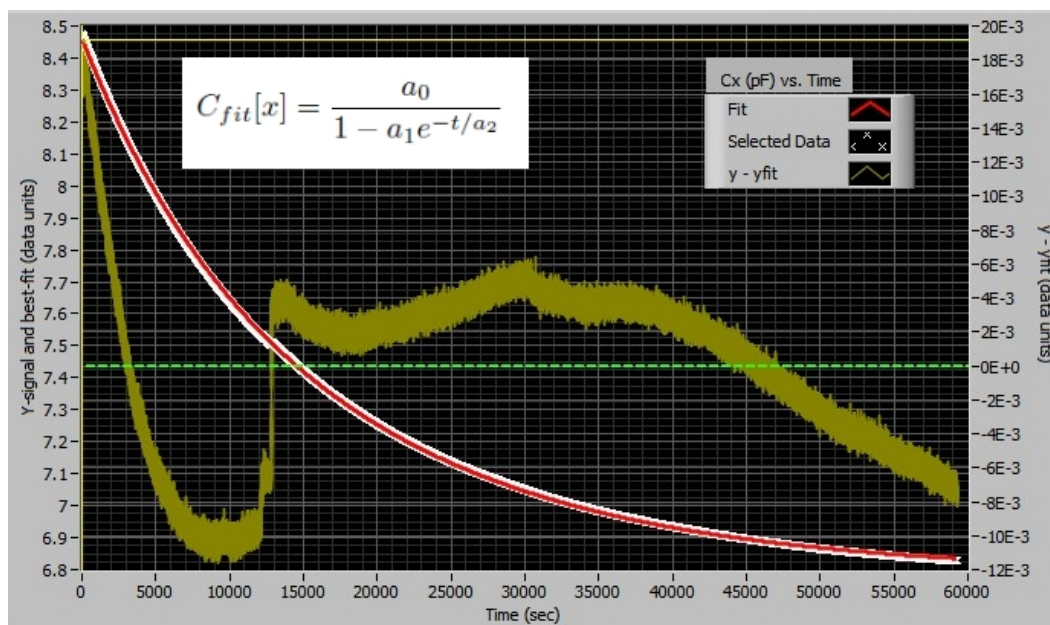


Figure 6.13: Typical capacitance transient and best-fit curve (model described on p. 108) - Eq. (6.13). The capacitance data and fit are in units of pF (left axis) and the fuzzy, background curve (right axis) is the difference between the fit and the data (also in pF). We can see that the model fits the data quite well.

Note that the two LBL chips tested had different number of holes (100x100 vs. 300x300). Also, the last one was tested in the full experimental cell on the cryostat where the inner cell volume was ~ 500 times less than the one in the jig I'm using now. The time constant scales accordingly.

6.1.5 Flow regimes

Flow through a channel behaves differently depending on the relative sizes of the mean free path (λ) of the gas and the dimensions of the channel (diameter D). The relevant parameter here is the so-called Knudsen number:

$$Kn \equiv \frac{\lambda}{D} \quad (6.17)$$

Wafer Array	CNF-J23 3a100w65	CNF-J23 3a100w65	LBL-1 3a100w40	LBL-1 1a300w40
Cleaning	UV/ozone Somorjai lab	UV/ozone Foundry	UV/ozone Foundry	O ₂ plasma Microlab
Time	~26hr	~ 99min	~ 99min	~ 5min
Measured τ (sec) in Eq.(6.11)	117	178	17811, 45049	70, 19
for nitride thicknesses:	Hole diameters (nm)			
L = 60 nm	109	91	15, 11	15, 25
L = 45 nm	105	88	14, 10	14, 24
L = 75 nm	113	95	16, 11	16, 27

Table 6.1: Some results from four flow test runs. Two values are provided for the final hole diameters when there is significant disagreement in the values obtained during filling and evacuation transients. In that case, the first value provided is the average during the evacuation transient and the second value that from the filling transient. See Fig. C.4 in Section C.2 for an example of how these numbers are calculated. Code listings are provided in the same section.

where the mean free path for a gas with effective molecular size (hard sphere radius) d at absolute temperature T and absolute pressure P is

$$\lambda = \frac{k_B T}{\sqrt{2}\pi d^2 P} \quad (6.18)$$

where k_B is Boltzmann's constant.

Since we are working at $T = 77K$ ⁷, helium-4 is best suited as a working gas and its hard sphere radius is $d = 2.2 \text{ \AA}$ [65] so that we have:

$$\boxed{Kn = \frac{1.46008 \times 10^6}{DP}} \quad \dots \text{ for } ^4\text{He at } 77K \quad (6.19)$$

where D is entered in nm and P in inches of mercury.

Two extreme regimes are readily identified with a Knudsen number much less than 1 (viscous - Poiseuille flow) or much greater than 1 (molecular flow - effusion). It is desirable to work in the effusive regime as the flow conductance there is a constant (with respect to pressure) so that during a transient, it does not change. As opposed to this, the viscous

⁷for reasons of stability, low noise and lack of drift in the diaphragm. It is stretched taut at 77K and behaves better as a spring with less floppiness and less susceptibility to (sometimes considerable) static cling to the electrode.

conductance depends on pressure and to make things worse, requires even more corrections for the case where the aspect ratio for the hole is near 1 (which is the case for us).

In practice (as seen in [64] for nanopores with sizes similar to the ones described here), $Kn \lesssim 0.1$ and $Kn \gtrsim 10$ suffice to separate the two regimes. Further, the conductance appears to stay fairly constant even near $Kn \sim 1$, deviating from its high-Kn value by just $\sim 5\%$.

Typical Knudsen numbers in these experiments

Since we would like to be in the effusive regime, it was decided to stick with low pressures in setting up the transients (typically, around 5 inches of Hg absolute). The hole sizes we deal with range from $\sim 15\text{nm}$ to 100nm . The Knudsen numbers for these ranges are shown in Figs.6.14 and 6.15.




Figure 6.14: Knudsen number (Kn) as a function of ambient pressure and hole diameter (D) at 77K for ${}^4\text{He}$ gas.

The molecular flow regime ($Kn \gg 1$)

Here, the mean free path is greater than the hole size so that the flow is a series of individual molecules passing through the tube described by the kinetic theory in statistical mechanics. The mass flow conductance for tubes of any aspect ratio (u) in this regime is given by Eq. (6.15).

Clausing's correction ($K[u]$) is deemed necessary for channels with aspect ratios any bigger than ~ 0.1 (below which a much simplified expression for the conductance holds where $G \sim u^3$). Since we routinely deal with aspect ratios near 1, we cannot afford to




Figure 6.15: Knudsen number (Kn) as a function of ambient pressure for different hole diameters (D) at 77K for 4He gas.

neglect it since it affects the calculation in a significant way. See Fig6.16 for an illustration of this issue.

From the typical working conditions in this experiment and Figs.6.14 and 6.15, we can say that using the molecular flow approximation with the Clausing correction for short tubes is a valid approach.

The viscous regime ($Kn \ll 1$)

The viscous regime can be described by the continuum dynamics of the Navier Stokes equations. The flow conductance of a channel (as defined in Eq. (6.8)) is in this case described by the Poiseuille equation (see [65]):

$$G_{viscous} = \frac{m_4}{k_B T} \frac{\pi L^3 P_{avg}}{8\eta} u^4 \quad (6.20)$$

where $u \equiv R/L$ as defined previously (p. 109), η is the gas viscosity and P_{avg} is the average pressure in the neighborhood of the aperture. Since P_{in} is impossible to measure absolutely without knowing the spring constant of the movable diaphragm, we can rewrite the above equation in terms of P_{out} , which can be known with some accuracy if a large buffer volume




Figure 6.16: The aspect ratio (u) dependence of the conductance in the molecular flow regime (arbitrary units). Inset: close-up of region $u \in [0, 0.15]$

is used outside ⁸ so that the pressure doesn't change significantly during the transient. We have

$$P_{avg} \equiv \frac{P_{in} + P_{out}}{2} = \frac{P_{in} - P_{out} + 2P_{out}}{2} = P_{out} + \frac{\Delta P}{2} \quad (6.21)$$

Using this in Eq. (6.20) we finally obtain

$$G_{viscous} = \frac{m_4}{k_B T} \frac{\pi L^3}{8\eta} \left[P_{out} + \frac{\Delta P}{2} \right] u^4 \quad (6.22)$$

This conductance would replace the differential equation (6.1.3) with a more complicated one since now the conductance would change during the transient. In full weak link experiments near 2K, the pressure differences are so small compared to the ambient pressure that the ΔP term can be neglected and the conductance stays approximately constant during the transient (this condition is used for finding the hole size using normal flow transients in a liquid helium-filled cell just above T_λ).

The transition regime

The in-between regime ($0.1 \lesssim Kn \lesssim 10$) - spanning 2 orders of magnitude - is known as the transition regime[64, 65], where an empirical model is used that interpolates between

⁸as has been done in this experiment

effusive and viscous flow.

6.2 Superconducting diaphragm

We have conducted several tests of the superconducting diaphragms used here. The essence of the test is to sandwich the diaphragm between two pancake coils⁹ with a Kapton spacer to provide a similar spacing to the one used in actual experiments. One of the coils (the “sense coil”) is connected to a SQUID and the other (“source coil”) is connected to a current source. In experiments, the diaphragm (of the displacement sensor) is exposed to a magnetic field generated in the pancake coil because of nearly $\sim 1A$ of persistent current circulating in the coil. The arrangement shown in Fig. 6.17 mimics this scenario to find out whether the diaphragm can continue to be a near-perfect magnetic shield under such field conditions. If the diaphragm is penetrated (and vortices form), the sensor sensitivity is compromised. Further, vortex drift and acoustically-induced motion can increase the base noise of the sensor. This is something to be avoided.




Figure 6.17: A simple jig to test critical fields of superconducting diaphragms. Insets show the individual components. The test diaphragm is sandwiched between a spacer Kapton sheet and a thinner Kapton sheet to mimic the thickness and spacing in experiments and is then held snugly between the two coils by screw joints. Source coil is fed by a current source and sense coil is read by the same kind of commercial SQUID magnetometer that we use in our displacement sensor.

The test is quite simple. The coil sandwich is immersed in a liquid helium storage Dewar on a probe (essentially just a wooden stick). The source coil should have an inline filter (simple RC filter is fine) to suppress line noise. The source coil current is slowly ramped

⁹These are crude tests so any imperfect, even kinked coils rejected for the main cell can be used here.

up from 0 and the sense coil magnetic pickup is read by the SQUID. It is observed that for properly made diaphragms (with thick enough lead coating), the sense coil reads nothing (within noise). Noise levels do go up due to leakage from the sides.

We can define a coupling parameter between the two coils as the ratio of the magnetic flux measured by the SQUID in units of flux quanta¹⁰ divided by the source coil current (typical sizes are $\phi_0/\mu A$ so this is a good unit to use). This coupling parameter is computed by measuring the SQUID response over time for slow ramps of the source coil current. Plotting the flux vs. source current gives us a sawtooth waveform, since the output of the SQUID fluxlocked loop resets to 0 after it exceeds its dynamic range. A recently created Labview VI automatically detects the continuous sections between resets and fits straight lines to each section to obtain the slope (which is the desired coupling parameter). Of course, this is a bit more sophisticated than necessary because all we need is a few data points between resets to find the coupling at a given source current level and then move the current up significantly and repeat.

Improperly made diaphragms are fine up to a critical current, at which point they show distinct signs of penetration, with the coupling rising rapidly past this point. We performed a null test (with blank Kapton and no diaphragm to keep the spacing the same) to verify that the bare coupling is on the order of $\sim 2 - 4 \phi_0/\mu A$. A properly screened coil typically shows a coupling several orders of magnitude smaller than this control value up to the point when it is penetrated. In practice, the sense coil SQUID signal is a series of increasing ramps followed by resets (as the SQUID reaches the end of its dynamic range). Figs. 6.18, 6.19 and 6.20 show, respectively, the results of this test for a blank, “good” and “bad” diaphragm, where the last two are defined simply by how well they screen the sense coil from the source coil. Note that there is some scatter in Fig. 6.19 - a consequence of our analysis technique where noise or overloads in the SQUID signal in some sections causes bad fits and yields invalid slopes.

The reason we perform the experiment and do the analysis this way is to also figure out if (in addition to a maximum value of the source field) there is also a maximum ramp rate beyond which we will create trapped vorticity in the diaphragm. It is possible that ramping the current up gently in the persistent current circuit might prevent the creation of trapped flux lines, whose motion (vibration induced or just drift - the latter of which is known as “flux creep” in the SQUID literature) might contribute to displacement sensor noise. We have found that to a limited extent, ramping slowly *might* delay the onset of penetration, but this is not a clear conclusion (since the ramp rate seems not to make a difference in “good diaphragms”). Practically speaking, based on experience, we would suggest ramping the injection current (up or down) at no more than $\sim \text{mA/s}$.

In conclusion, this constitutes a simple, yet powerful test of the superconducting used for the displacement sensor.

¹⁰The magnetic flux quantum is $\phi_0 = \frac{h}{2e}$, where h is Planck’s constant and e is the electron charge




Figure 6.18: The control test with no superconducting diaphragm to gauge the bare coupling level.




Figure 6.19: Example of a “good” diaphragm. Coupling stays 0 within noise level and is at least 2 orders of magnitude lower than the bare coupling. Note that various ramp rates (for the source coil current) have been tried and the results all shown here.




Figure 6.20: Example of a “bad” diaphragm. Coupling is already high to begin with (compared to the good diaphragm) and starts increasing arbitrarily past ~ 250 mA in the source coil. Past 500 mA, the penetration and noise is too high for the SQUID and the signal is just garbage.

Chapter 7

The Cryostat

7.1 Overview and broad issues

The cryostat is a pumped-bath design based largely on one previously made by Emile Hoskinson. The main difference is that the one described here was made to be compatible with two entirely different structural solutions in mind. We will touch on this briefly in the final section of this chapter, where we will also explore the vibrational properties of the cryostat. Helpful guidelines on building such cryostat inserts may be found in practical textbooks on cryogenics such as Refs. [61] and [71].

Fig. 7.1 shows an overview photograph of the assembled cryostat. The main components are the top plate, structural frame and the two bottom experimental stages. Wiring and thermometry are especially important subjects that will be covered in their own sections. The cryostat insert is sealed onto a neck ring (a brass flange that raises it by about 6”), which is further sealed onto a cryogenic Dewar. The Dewar has a narrow neck and tail and a wide belly to increase the time during which the experimental stages can remain cold. To this end, the insert length is engineered to put the cryogenic components as low as possible in the tail. Using a 4.25” mouth Dewar with a ~ 20 L capacity, we have been able to stay cold for about 2.5 days between transfers (while still maintaining a liquid level that keeps components like the persistent current circuit and the cryogenic valve submerged). The hold time can of course be increased quite a bit by using a larger Dewar.

7.2 Construction

Note that more details on the various components used during experiments are provided in Chapter 10 (“operation”) and in preceding chapters dealing with the individual components. In this chapter, we focus on the cryostat itself and its various interfaces to the outside world. Any component referenced here that has heretofore not been introduced, may be found in the chapter on operation.

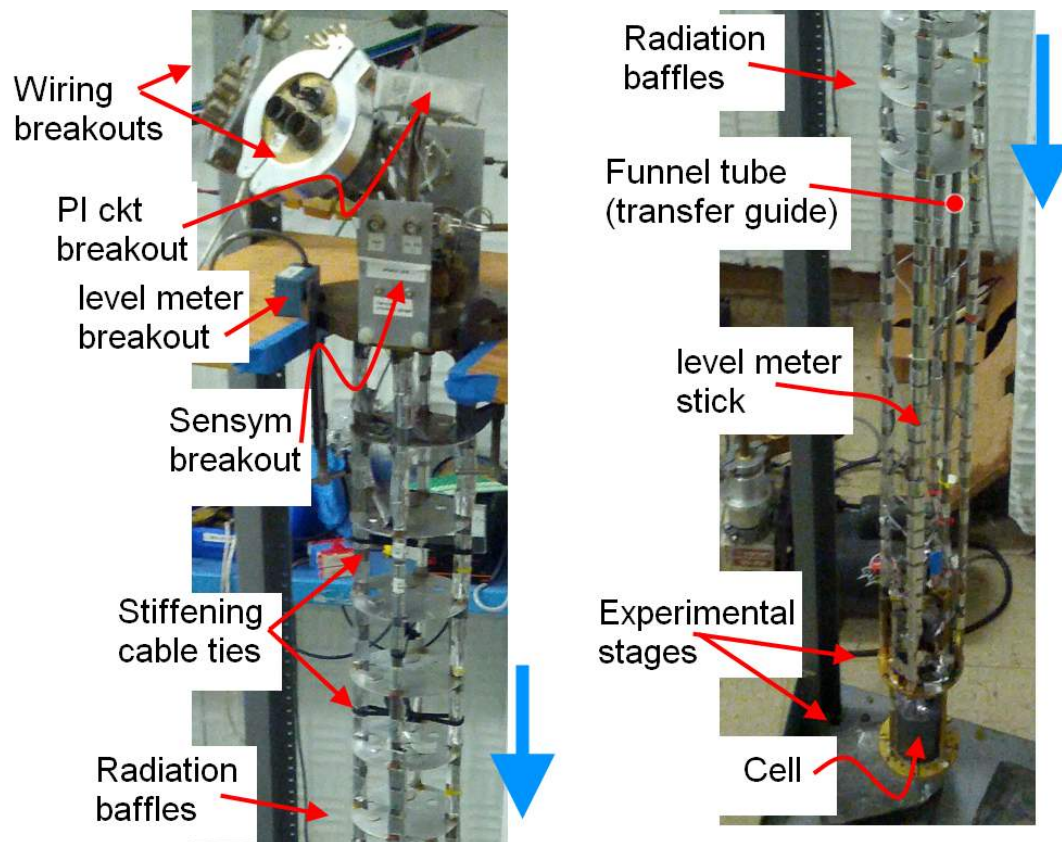


Figure 7.1: Overview of the cryostat. Major components are labeled. The Amphenol breakout flange can be seen in the foreground, with the BNC breakout flange off to the left. The PI circuit breakout is hidden by the RF shield cap over it. The top plate valves and pressure gauges are on the side facing away from the camera.

7.2.1 Top plate

The top plate is a 1/2" thick brass plate with a polished underside used to make the O-ring seal to the brass neck ring. It has several quick-connect ports hard-soldered in for the various wiring and plumbing breakouts. We use two blank KF flanges with drilled holes and affix several bulkhead BNC jacks with hermetic seals on one (referred to as the "BNC breakout") and 3 Fisher (Amphenol) multi-pin connectors on the other. Of course, as long as the connectors can be sealed, the particular brands are arbitrary. The essential thing is that we have a bunch of shielded coax connectors for the 3 capacitance bridge leads and a bunch of shielded multi-pin connectors to breakout the several 4-wire leads for the resistive heaters, bath level-meter and thermometer. Shielding the wiring will continue to be a major theme in this chapter.

These (large) breakout KF's are sealed onto conical reducer nipples, which themselves seal onto custom-built KF adapters that can be secured to 2 permanently mounted stainless

steel tubes on the top plate. This way, the entire breakout can be removed for repairs or rewiring and simply sealed back on and the top plate maintains a degree of versatility¹.

Two commercial SQUID cables seal directly into other quick-connects and we allow for future changes by providing for as many quick-connects as will fit on the plate (sealed off with plugs until necessary). It is wise to always provide for such extra breakouts as well as have a bunch of extra, unused leads already wired up. A solid state pressure gauge is connected to the bath via a small plastic tube.

A free quick-connect precisely aligned over matching holes down the rest of the cryostat provides for a transfer port, through which various tubes (LN2 transfer and blowout, LHe transfer, etc.) can be inserted during the cooldown process. We also use this port (occasionally) to insert a dipstick heater for rapid warmup, a simple thermocouple for temperature measurements from 77 K up to room temperature and (rarely) to insert a dipstick level meter² if the electronic level meter becomes inoperative.

7.2.2 Structural frame

The top plate has an adapter bolted underneath (with 4 large 3/8" bolts), which holds the actual structural frame. This is done to enable use of the top plate with different kinds of frames, because an outstanding goal of our research has been to try to stiffen the cryostat in different ways to drive its resonant frequency higher in order to reduce its acoustic coupling to low frequency noise sources that can inject rotation noise into the SHeQUID. We discuss this some more towards the end of this chapter.

Stiffness

The main supports are thin wall (0.028"), 3/8" OD stainless steel tubes hard-soldered onto threaded adapters at both ends that are securely bolted to the top plate adapter and the experimental stages. A set of Aluminum radiation baffles reflects incident radiation from 300 K and reduces the impact of that heat leak. They also provide much-needed rigidity to the entire structure. Traditionally, the baffles are brass or steel and are hard-soldered to the support tubes. However, we have made them removable as they are just tied on with wire threaded through small holes in the support tubing. The baffles have open clearance holes for the tubing to allow easy removal. Since this is not very rigid, we have found that tying on some simple (thick) cable ties diagonally, as shown in Fig. 7.1, significantly helps stiffen up the cryostat. We expect this to improve at cryogenic temperatures as the plastic cable ties shrink more than the metal and grip the tubing tighter.

¹It is a large and complicated part and hard-soldering all the quick-connects at once can be tricky, so it is desirable to not have to remake the piece for every little change in the experiment

²This is merely a thin stainless steel tube with a larger hollow chamber capped by a rubber glove at one end (for use as an acoustic amplifier), which is used to detect a helium bath free surface using the phenomenon of Taconis oscillations.

Straightness

The structure must be extremely straight since we typically make the experimental stages as large as possible (space is always at a premium) so that we have only about a 1/8" clearance between the stage OD and the tail section ID. It is generally understood³ that an experimental stage that touches the cryostat is an acoustic short for cryostat vibrations to propagate to the cell. This may be checked for via electrical methods. We did this recently by measuring the resistance between an isolated sharp studded band⁴ wrapped around the experimental stage and the Dewar body (which must obviously be metal). The studs are made sharp to break through any oxide layers on the tail metal. Sometimes, the Dewar belly and tail section are suspended from the Dewar mouth by a fiberglass (G10) neck to improve thermal isolation. This is sub-optimal for other reasons⁵ but the issue here is simply that the Dewar outer body might not be connected to the inner tail section and this would invalidate our contact test. A simple workaround is to attach a spring-loaded needle or rod to the bottom of the cryostat insert with a lead that provides electrical contact to it coming out of the Dewar. This is designed to provide a gentle press-contact to the Dewar tail bottom, which can be used with the studded band leads to do the contact test.

7.2.3 Plumbing

Fill lines and other plumbing capillaries are made out of thin wall⁶ stainless steel capillaries hard-soldered to small brass end-rods with drilled holes. These hollowed 1/4" end-rods are used to connect to valves via (Swagelok) compression fittings. The lower terminus of the plumbing lines are typically hard-soldered into similar (but smaller) brass end-rods onto which cupronickel (Cu-Ni) capillaries can be soft-soldered. This is done to ensure that the cryostat plumbing connections can be removed and installed when needed without having to apply the high heat needed for hard-soldered joints. Cu-Ni capillary (to brass end-rod or cell) joints can be made with a small butane torch, while (telescoping) capillary-capillary joints can be made with just a simple soldering iron! This makes installing and removing

³We only have anecdotal evidence to this effect because the noise sources can be difficult to distinguish unless one specifically conducts experiments geared towards that goal.

⁴An easy way to make this is by cutting out a small strip of metal shim stock and punching it at regular intervals with a hammer-driven, blunt hand punch over a wooden base. This provides beak-like protrusions on the other side that are reasonably sharp, yet short enough not to touch the tail metal if the insert is reasonably straight.

⁵The G10 neck can take a long time to thermalize and continues to relax in a noisy fashion for several hours after a bath helium transfer. This forces us to wait until things quiet down for our acoustically sensitive experiments and wastes valuable cold time. A Dewar whose innards are made entirely of stainless steel would be preferable acoustically, but we have been made aware of a possible problem with stray magnetic fields from the steel that could be detrimental to our SQUID based measurement systems. SQUID researchers use shields made of Cryoperm alloy that is placed inside the Dewar. We have not tested these possibilities but simply provide it here for future researchers to consider.

⁶Except for the cryovalve actuation line, which is made of thick wall capillary to withstand the higher pressures.

the cell and cryovalve much easier than otherwise. All plumbing lines are fed through snug holes in a small brass insert (that can seal into a quick-connect) and carefully hard-soldered in place to make a leak-tight joint.

The neck ring has ports in its side for pumping on the bath space and bleeding gas in. The cryostat should also be equipped with a relief valve (typically set at 5-10 psig) to prevent blowing out any feed-throughs or quick-connect O-ring seals while pressurizing the bath. As a general rule of thumb, we never exceed a differential pressure of ~ 1 bar between the bath space and the outside. A large port is used for maintenance-level pumping (initial cooldown, post-transfer pumpdown, etc.) and a valve is used to seal off the Dewar neck during most experimental runs. Bath pumping during runs is done through an axial pumping jig⁷, which is essentially a 3/8" OD thin wall stainless steel tube divided by a valve and inserted into the transfer port when needed. This axial jig comes in particularly handy when reorienting the cryostat as the pumping line can be more easily secured and kept out of the way as compared to the traditional side-port line.

Design of external pumping lines is described in Section 9.1 on vibration isolation, since that is a critically important subject in its own right. Based on our experiments so far, we suspect that (given rudimentary isolation protocols) the bulk of the residual acoustic noise probably comes through the pumping lines.

7.2.4 Experimental stages

Instrument stage

This is the lower terminus for the support tubes and it houses the various instruments used for the experiment. A superconducting level meter stick is taped to one of the support tubes using cryogenic tape⁸.

The cryovalve (Chapter 8) and persistent current circuit (Chapter 5) are both mounted on this stage. So are the 2 SQUIDs (for the displacement sensor and High Resolution Thermometer [HRT]) as well as the Germanium Resistance Thermometer [GRT] and the actual HRT. Lastly, we have some lead-plated breakout boxes to enable shielded connections to the leads coming from the cell. This last category includes the cryogenic reference capacitor (used in a bridge circuit to measure the D-E capacitance) and breakout boxes for the inner cell and sense arm heaters.

⁷Based on a suggestion by Yuki Sato for the purpose of reducing transverse wobbling of the cryostat, which can add rotational noise that the SHeQUID is particularly sensitive to.

⁸This is aluminum tape made by 3M and distributed by Lakeshore (model C8-105) and works well at cryogenic temperatures. We use this kind of tape extensively for securing wires and other things to the cryostat body. The bright aluminum surface also makes it useful to fill in any gaps in the radiation baffles or even provide a quick RF shield for small components or wire joints. We can make small, cylindrical shield boxes for such joints by taping over plastic tubing with aluminum or lead tape and inserting unshielded joints inside.

Cell stage

This stage typically contains only the (bolted on) cell and can be completely removed from the cryostat. The cell fill line coming from the cell is rigidly secured to the stage supports. The modular design of the cryostat makes it relatively easy to lower either experimental stage by simply adding spacers to the main support tubes. Increasing the cell space is more difficult and involves cutting and re-soldering the supports (a significant undertaking as all components have to be removed). However, a simpler workaround (done recently) is to put in a taller neck ring and then add short support spacers to fine-adjust the heights. Recall that to get the longest experiment duration (between transfers), it is necessary to keep both experimental stages as low into the Dewar tail section as possible.

7.3 Thermometry

Three separate thermometers are used under normal circumstances. During warmup, a fourth (portable) thermocouple junction thermometer may be inserted through the transfer port to monitor the warmup process since the other thermometers are not very useful at temperatures above 4 K. A permanently mounted thermometer like this can also be a useful thing to have if one has the wiring breakouts to spare. These thermocouple junctions are readily available commercially or can be made by taking the two different metal wires, twisting them together and blasting the junction with a blowtorch until they melt together into a bead.

7.3.1 Vapor pressure

A solid state pressure gauge (Sensym/Honeywell model ASCX15AN) is used to measure the absolute pressure in the bath space. Its output is a voltage signal proportional to the pressure and to the input power supply voltage applied to run it. Therefore, it must be calibrated for a given supply voltage against a mercury (or oil) manometer (see Section 10.5).

7.3.2 GRT

Our primary thermometer is a Lakeshore model GR 200A-1500 germanium resistance thermometer (GRT) with a negative temperature coefficient (chosen to be sensitive near T_λ , where we perform our experiments) and resistances on the order of $\sim 10\text{ k}\Omega$ near T_λ (increasing as we get colder). A 4-wire ($I \pm V \pm$) measurement is performed to read this resistance using a commercial AC bridge (the Lakeshore 340 temperature controller), which is transmitted to a computer over a GPIB interface (more details on this in Chapter 10). This GRT resistance is calibrated (see Section 10.5) against a parallel vapor pressure measurement using a standard, published [72] vapor pressure curve for ^4He by letting the temperature drift up from $\sim 1.5\text{ K}$ to slightly above T_λ .

The bridge resolution is on the order of $\sim 1 \Omega$ in the temperature range of interest, which corresponds to $\sim 10 \mu K$ in temperature. Practically, this lets us regulate the bath temperature to a stability of a few tens of μK using just the GRT. This is not sufficient as temperature gradients on the much less than this can cause fountain transients. So, we need a secondary thermometer, with much more sensitivity (which leads to a correspondingly lower dynamic range).

7.3.3 HRT

The HRT stands for High Resolution Thermometer (which is not a particularly descriptive name). Narayana, *et al.* [73] describe in detail the construction of such an HRT, so we provide only a brief description here⁹. It essentially consists of a small cylinder of PdMn alloy whose magnetic susceptibility varies with temperature. The alloy concentrations are precisely engineered to give a large sensitivity to temperature changes in the temperature region of interest (in our case, near T_λ). Two permanent magnets are fixed on the cylinder ends to provide a steady field through the alloy. A pickup coil would around the cylinder is connected to a commercial dc SQUID. As the susceptibility changes with temperature, so does the flux picked up by the coil. The SQUID signal is found to be approximately linear in the temperature changes within sufficiently small temperature domains. We have also used HRTs made using paramagnetic salt pills, such as the ones made by Welander, *et al.* [74].

The HRT (sensitivity) is calibrated (see Section 10.5) by measuring the SQUID output vs. temperature and finding slopes to sectional linear fits. The HRT resolution is less than 5 nK, and its use in temperature regulation allows stability to around 20 nK (even down to 10 nK on a good day). The stability seems limited by acoustic noise and not intrinsic SQUID noise – we have seen that vibrations affect the HRT quite strongly (presumably by causing fluctuating fields due to motion of coil/magnet in the HRT).

7.4 Wiring

7.4.1 Breakouts and wiring choices

The choice of wiring for the components on the cryostat requires careful thought. We have changed the wiring for various reasons over the years and we discuss these choices in the list included later in this section. Note that wire sizes in the United States are typically quoted as either a wire diameter in “mils” or thousandths of an inch ($0.001'' = 1 \text{ mil}$) or in terms of its AWG (American Wire Gauge) number, whose decimal sizes can be easily looked up in online references. Wire resistances are provided in Ω per foot (at 300 K unless mentioned otherwise).

⁹We would like to express our deep gratitude to Michael Ray for building and testing the HRT used in this work.

Note also that the precise top plate connectors used to interface to these wires are not as critical and many possibilities exist (that are probably more durable than the ones we use for compatibility reasons). The main criteria for connectors are their ability to form robust hermetic seals, reasonable heat and water resistance to guard against damage during LHe transfers, ease of making wire joints to connector pins (possibly with the system cold and in awkward positions during ad hoc repairs) and finally, their ability to make secure mechanical joints with mating cables (especially important during Dewar reorientations). We have 3 main wiring breakout clusters on the top plate and they are discussed next.

BNC flange

This is a KF-50 blank flange with 7 hermetic BNC jacks (sealed by rubber O-rings dabbed with Apiezon-M vacuum grease).

- 3 of these (with shields grounded) –labeled 2,3,4 — have stainless steel coax wire¹⁰ for the capacitance leads (E, D and C_{ref} in Fig. 10.1). This is multi-wire inner conductor (nominally $7 \Omega/ft$ at 300 K) with fluorocarbon insulation and outer braid, both made of stainless steel covered with overall fluorocarbon insulation. The outer insulation is stripped off so that the braid is electrically and thermally sunk to the cryostat all the way down. Only the inner conductor is used to carry signals and the braid shield conducts pickup noise to ground.
- 4 more (with shields grounded) are connected to 2 separate twisted pairs of conformally insulated, $\sim 8 \text{ mil}$ “Nico” wire¹¹. This is a Cu-55%, Ni-45% alloy with a 300 K resistance of about $5 \Omega/ft$. One pair (labeled 5,6) is used for the bath heater, which is $\sim 8 - 10 \text{ ft}$ of 36 AWG manganin wire at $\sim 12 \Omega/ft$, while the other pair (labeled 0,1) is an unused spare. Manganin is used for its high resistivity and low temperature coefficient (the latter making it relatively slowly stable over temperature changes). These Nico wire leads are protected from scratches by a small diameter Teflon (PTFE) tube. It can be difficult to insert wires into small tubing and one way to do this is by cutting the tubing into sections and joining the sections after wire insertion using heat-shrink tubing or tape. A better way¹² might be to pull the wires in at one end by applying suction (with a pump or lab vacuum) at the other end. Another way that has worked to some degree is to insert a smaller, sacrificial wire, tie or hook it on to the wires to be inserted and pull them through using the sacrificial wire.

¹⁰Model# AS636-1SSF from Cooner Wire Company (Chatsworth, California).

¹¹From California Fine Wire company (Grover Beach, CA).

¹²We have only recently heard of this from Jeff Birenbaum and have not tested this technique as yet.

Amphenols flange

This is another KF-50 blank flange with 3 breakout plugs soldered into it. Two of these (#2 and #3 in the list below) are actual male (pin) connectors that fit into Amphenol¹³ (model 165-16-1003) cable receptacles, which we install on multi-wire cables for 300 K use made of individually shielded pairs of heavy gauge wire. The third is just a multi-pin feed-through (not a plug).

- **Amphenols-1:** The name is purely historical – this breakout does not plug into an Amphenol connector but has bare pins. This is a 9-pin glass-to-metal feed-through with solderable pins. 5 pins are unwired while the remaining 4 are hooked up to a DB-9 female connector in a breakout box on the top plate for a Lakeshore 241 level meter controller.
- **Amphenols-2:** This is a 9-pin Amphenol (male). 5 pins are unused. 4 are used for the sense arm heater in the (SHeQUID) cell. One pair of wires is the same 36 AWG manganin used for the bath heater while the other pair is larger gauge¹⁴ manganin (~ 0.56 mm). The larger gauge pair is used to run current through the sense arm heater (lead resistance very small at ~ 1 Ω and temperature related changes even smaller because it is manganin), while the smaller gauge pair used for voltage sensing across the resistor (for more accurate heater power measurements while in use).

Both pairs are inserted in a PTFE tube for protection and shielded by 1/8" stainless steel braid¹⁵ We would have used the larger gauge wire for both pairs were it not for space constraints in the breakout tubes, PTFE tubes and sleeving. The 4 wires connect up to the sense arm heater leads on the second experimental stage inside a shielded box.

- **Amphenols-3:** This is another 9-pin Amphenol (male) used for wiring two components:
 - Inner cell heater: we use the same combination (of 2 twisted pairs of manganin wires) connected to 4 pins here that we used for the sense arm heater in Amphenols-2 above. The 4 wires connect up to the cell heater leads on the second experimental stage inside a shielded box.
 - GRT: 4 of the remaining pins are connected to a commercial cryocable (type CYRC¹⁶) with 4 superconductive (32 AWG Cu-Ni clad NbTi) wires quad-twisted

¹³Amphenol corporation, Wallingford, Connecticut

¹⁴We could only find this at GVL Cryoengineering, Stolberg, Germany.

¹⁵Small quantities of this braided sleeving, made of 304 (non-magnetic) alloy can be found at an online store (for, of all things, motorcycle parts!) called 4RCustoms (Rowley, Ma) – most other companies require large minimum orders. Using any other metal braid (such as the more common tin coated copper) would be an intolerable heat-leak down the cryostat. One must also be careful about the magnetic steel alloys as vibrations could force the braids into inducing noise in the leads.

¹⁶Lakeshore Cryotronics, Westerville, Ohio

and coated with Teflon. This is further shielded by a (304) stainless steel braid and again covered by Teflon outer insulation. We remove the outer insulation so that the steel braid touches the cryostat throughout its length. The 4 wires are connected to the 4-wire measurements leads ($I \pm V \pm$) of the GRT.

The above cryostat breakouts plug into their respective (room temperature) multi-wire cables and finally separated into individual BNC jacks at an electronics rack where appropriate inputs and outputs can be connected to the various components. Most recently, we have split the room cable for Amphenols-3 into two separately shielded (with tin coated copper braid) 4-wire cables, one of which (GRT) goes directly to the Lakeshore 340 to be measured, while the other is broken out into BNC jacks on the electronics rack.

DB-25 (persistent current (PI) circuit breakout)

This is a custom-built shield box with a DB-25 connector on one end of a bent stainless steel tube with the other end hard-soldered into a brass adapter that seals into one of the spare quick-connects on the top plate. This was added on recently to isolate the persistent current (PI) circuit wiring from the rest of the wiring. A snug-fitting cover made of bent aluminum sheet metal is snapped on over the DB-25 connector to shield the PI circuit from stray RF noise that can make its way into the displacement sensor. We have observed greater levels of high-frequency “buzz” in the displacement sensor signal with the cover off.

For the two heaters used as PI switches, we use twisted pairs of the same Nico wire used for the bath heater (in the BNC breakout) protected by PTFE tubing. For the current injection leads, we would like something with very low resistance (to avoid heating the bath during current injection, where we might be slowly ramping up the persistent current to around an ampere over several minutes) and a lot of shielding to prevent noise from being injected into the displacement sensor. For these reasons, we use a twisted pair of 28 AWG (Belden 8080) solid copper wire ($\lesssim 1 \Omega$ for about 8 ft of total wire) with an outer coating of Poly-Thermaleze and inserted into PTFE tubing covered by a (304) stainless steel braided shielding¹⁷. The cover is opened and a breakout cable (with banana jacks on the other end for connecting the 3 pairs of wires to power supplies) is connected to the DB-25 plug *only* when we need to change the persisted current in the PI circuit. At all other times, the cable is kept disconnected and the cover kept closed because the displacement sensor is essentially unusable (noise driven continuous SQUID resets) with the cable connected.

As described in Chapter 5, the 2 injection leads terminate in the shielded filter inductor box (one superconducting inductor for each lead) and continue onward to the appropriate superconducting joints in the PI circuit box. Both these boxes are on the first experimental stage. Each pair of heater leads terminates across a separate filter capacitor and continues onward to the heater resistors in the PI circuit box. The importance of properly shielding all these leads cannot be over-emphasized.

¹⁷from Star Cryoelectronics, Santa Fe, New Mexico. See footnote 15 on p. 126 for a note on the steel alloy.

7.4.2 Wiring techniques

Stainless steel wire joints

Since stainless steel wires cannot easily be soldered to¹⁸, we slip any stainless steel wires inside Cu-Ni capillary tubes with part of the insulated wire end also inserted and then crimp the conducting part with cross-hatched pliers. The extra insulated part provides some strain relief and the now roughened Cu-Ni stubs form a nice, solderable surface to make joints to. This technique can also be used for any small wires that are difficult to solder to, such as small gauge superconducting wires, which despite their Cu-Ni cladding may not always make secure solder joints.

Soldering

Also see Section 8.2.4 for an extended discussion on soldering issues. All wiring joints (and plumbing joints) on the cryostat are made with flux-free solder used with Superior # 30 blue liquid solder flux (manufactured by Superior Flux & Mfg. Co. and distributed by several suppliers, such as Amtech, INC., Deep River, CT, USA). While this is an important issue for plumbing joints, as far as wiring joints are concerned, we have found no discernible difference between this and the regular rosin flux that comes embedded in standard, electrical multi-core solders and in recent times, we have drifted more and more towards using these latter solders for wiring.

Stripping conformal wire coatings

Stripping painted-on insulation from fine wires is a daunting proposition due mainly to present-day industrial regulations against certain types of chemicals, particularly those containing Methylene Chloride (which is highly toxic to living beings). Ref. [75] contains a useful discussion about the various methods one can employ towards this goal.

The venerable Strip-X is now all but impossible to find, as is the Conformable Coating Stripper with Methylene Chloride once manufactured by MG Chemicals (Surrey, British Columbia). The latter *name* may still be found attached to a similar product by MG chemicals but we have found it to be largely useless for stripping Formvar insulation from superconducting wires (though, to be fair, this is never promised by the product). As of this writing, only two techniques out of the distressingly many¹⁹ that we have tried have worked for Formvar insulation.

One of these is mechanical – two conical, abrasive rotating wheels with an adjustable gap for the wire. Construction of such a device is described in the previous reference and

¹⁸Except with highly corrosive zinc chloride flux, which can damage the cryostat and its components and possibly cause plumbing leaks and electrical shorts over time with exposure to its fumes. It is recommended that this type of flux be avoided entirely if possible.

¹⁹Limited, of course, to those legally available in the United States. At the risk of sounding bitter, we have found that any commercial chemical that works well is promptly banned soon after we discover it.

commercial implementations (such as the RT2S by the Eraser Company, Syracuse, New York) exist.

The other is a chemical method²⁰ that is somewhat hazardous. A common aspirin pill²¹ is placed on a glass slide (or other heat-resistant surface) and a soldering iron is pressed into it with the wire to be stripped sandwiched between the two. The melting aspirin is very corrosive and it strips the Formvar sheath off cleanly. However, the corrosiveness immediately goes away, leaving a largely useless re-solidified puddle. Also, noxious fumes arise during this process, which should not be inhaled (a respirator or at least a dust mask should be used as well as chemical goggles to protect the eyes – they do sting quite a bit with exposure to the fumes). For table-top use, we have had some success with table-fans blowing the fumes away from the user (a common vacuum cleaner might also help here, especially if one wets the air filter a bit to block the fumes from exiting the cleaner). Care should be taken to shield the other sensitive components from the fumes if this is done for wires on the cryostat. We cover the cryostat with aluminum foil in such cases. We would recommend experimenting with variations on this technique, such as making a custom attachment tip (for soldering irons with threaded tips) that is essentially spoon-shaped. Aspirin pills can be powdered and placed in the spoon and the wire-end simply dipped in the powder prior to activating the stripping by heating the iron. A flat, hollow, hinged attachment with holes for the wire might work better to contain the fumes. The mechanical method is the safer and more robust method and a commercial stripping tool (or a home-built one) is not a bad investment if such wires are used regularly.

An alternative method is simply scraping the wires with a razor or xacto blade (a bit risky to the wires, but gets easier with practice) can also be used, especially in tight spaces where the above methods are impractical. To make this a bit easier, we put a small piece of scotch tape near the wire end with only a small end-length uncovered. Instead of a blade, a slow rotary tool with a cylindrical sanding head²² can be used on this taped up wire by just holding the wire between one's fingers and brushing away and towards the end. Doing this a few times while rotating the wire occasionally can be sufficient to strip off the insulation. Removing the tape is made trivial by wetting a pair of tweezers with some isopropanol (IPA) or ethanol and taking off the tape with it. Simply pressing the tape between IPA-wetted wipes for a few seconds and then gently pulling it off also works.

All these stripping methods are typically followed by a final cleaning with fine ($\gtrsim 1500$ grit) sandpaper and an IPA wipe.

Etching Cu-Ni cladding from superconducting wires is described in Section 5.1.3.

²⁰We are grateful to Mark Kimball for informing us of this technique.

²¹Banning aspirin seems like a formidable enough challenge for regulators that this method would likely remain feasible (at least for a while).

²²Such tools, essentially small abrasive drums, are available as accessories for Dremel tools. Using a Dremel tool here would be inappropriate and a low speed motorized screwdriver is much more controllable and useful in this case.

Making twisted pairs

Making twisted pairs of wires is remarkably simple, yet equally easy to get wrong. After trying several ways to do this, we have found an optimal method that involves a slow-speed cordless screwdriver (or similar rotary tool), a common rubber band and a small, smooth (deburred) hollow cylinder (about 1" long and 1/4 – 3/8" in diameter – a standard electrical spacer works fine). The wire to be twisted can be two equal length pieces or one piece folded in half with one end remaining continuous after twisting (this is needed, for instance, when making a bath heater out of manganin wire). In either case, the free ends are knotted together and one end is looped over a thin, smooth, fixed object (such as a toothpick held in a vise) with a tapering end so that this end can be safely slid off after twisting. The other end-loop is slid around the spacer and taped in place. The rubber band is cut and inserted through the spacer hollow and the cut ends secured together in the rotary tool chuck. Many variations of this setup are possible of course. The important thing is that anything the wire touches should be deburred so it doesn't scratch the wire and that the rotary tool end of the wire loop be spread open a bit (hence the spacer) to allow the twists to propagate freely down the wire instead of bunching up at the tool end. Given these basic ideas, there are several ways to get it done. In the end, once the twisting is done, it is important to not simply cut the wire out or let go of the tool as the stored torsion can tangle the wire up. Instead, we hold on firmly to the rotary tool end of the wire and gently release the rubber band from the chuck and let the wire unwind freely to release the excess torsion. We have even slipped out the chuck and let the wire unwind a bit by letting it swing freely over a finger to maintain tension on the wire. After this, it is useful to spread the wires out (or just tape them) to "lock" the twists in place and prevent it from slowly unwinding in storage.

7.5 Structural issues

The main issue to be discussed here is the question of resonant modes of the cryostat frame. As we will see in a later chapter (Section 9.1.2), the fundamental frequency of the pneumatic springs that isolate the experimental platform from ground vibrations should be ~ 0.7 Hz. Our goal is to make the cryostat insert stiff enough that *its* resonant modes are pushed higher in frequency so that the greatest noise leakage that makes it past the springs does not drive the cryostat on resonance. Towards this end, we describe some simple calculations for the cryostat resonant modes, simulate them using a finite element analysis (FEA) package and attempt to measure them using an accelerometer (in order to validate our predictive models). Further, we discuss some ideas on how to enhance the cryostat's stiffness. Given these analytical tools, it should be possible to model changes in the mode frequencies for any proposed changes in the cryostat structure.

7.5.1 Cryostat normal modes

Analytical

The cryostat insert undergoes normal mode oscillations at certain well-defined frequencies. We can estimate the lowest mode by approximating the structure (which is 4 steel support tubes with a fixed top plate, several radiation baffles and most of the mass concentrated on the experimental stages) as simply a single equivalent cantilever fixed at one end with a point mass on the other (free) end. The frequency of this toy model is found using the including Matlab script (based on expressions found in Ref. [76]) (all physical quantities are defined therein). The lowest mode is estimated to be $\sim 2.4 \text{ Hz}$.

Note that this model *cannot* include the additional stiffness imposed by the radiation baffles. The FEM simulations in the next section *do* take the baffles into account, but they do so for baffles welded in place, while our actual cryostat has them tied to the support tubes with wire (with additional glue and cable ties to make them stiffer). This is not as stiff as welding them in place. This means that the analytical method should underestimate the stiffness (and hence the frequency) while FEM simulations should overestimate it. The actual measured value should lie somewhere in between for these results to be consistent.

Listing 7.1: Matlab script for analytical estimation of cryostat modal frequency.

```
%% Formula
% Beam clamped at one end and free at other (with point mass m loading the
% free end).
% L = length (m) n = number of parallel tubes (spacing between tubes << L)
% A = cross-section area (m^2)
% rho = volume density (kg/m^3)
% E = Young's modulus (N/m^2 = Pa)
% IO = area moment of inertia (m^4)

% Fundamental frequency in Hz
f = inline('(0.5/pi)*sqrt(3*E*n*IO/((0.2235*rho*A*L_+_m)*L^3))', 'L', 'm', 'A', 'rho', 'n', 'E', 'IO');
% Usage: f(L,m,A,rho,n,E,IO)

%% unit converters
pound = 0.45359237; %in kg
in = .0254; %in m

%% Cryostat (steel tubes, 3 lb point mass at bottom)
% All units SI
L = 44*in;
n = 4; %4 tubes in parallel. Spacing ~ 3.5" << 44"
r2 = (0.375/2)*in; %OD = 3/8"
r1 = (0.319/2)*in; %ID = 0.319"
A = pi*(r2^2 -r1^2);
IO = pi*(r2^4 -r1^4)/4; %Wikipedia (area moments of inertia -annulus)

m = 3*pound; % 3 lbs cell and experimental stage contents (approx.)
rho = 8000;
E = 193e9;
% 316 stainless steel props from
% http://www.azom.com/article.aspx?ArticleID=863#_Mechanical_Properties

f_cryo = f(L,m,A,rho,n,E,IO);
['normal_cryostat:_f_=_ ' num2str(f_cryo) ' Hz']
```

```
%%
%Output is:
%normal cryostat: f = 2.4036 Hz
```

FEM

We can model the cryostat in some detail (structurally at least), by virtually welding together components made of different materials like steel and brass to better estimate its behavior. We use commercial software²³ for this purpose. The model is subsequently meshed, i.e. divided up into small volume elements that are modeled as solid volumes connected with neighboring elements through spring-like interactions. The elements can also distort under stress according to their material properties. An example of such a mesh is shown in Fig. 7.2 and the results of a modal analysis performed on this model shown in Fig. 7.3. The lowest mode frequency is found to be $f_1 \sim 4.2 \text{ Hz}$, consistent with the baffle discussion in the previous section. Note that while we display the models horizontally here (to save space), the FEM software knows that the cryostat longitudinal axis is along the vertical and also knows what the gravitational acceleration of the Earth is.




Figure 7.2: Close-up of volume elements in the generated mesh of the model cryostat.




Figure 7.3: Lowest bending mode of (present) cryostat model (bending is exaggerated). $f_1 \sim 4.2 \text{ Hz}$

²³Autodesk Simulation Multiphysics 2012 [77], graciously made available for free educational use by Autodesk.

Measurements

We mounted a dummy cell – essentially a heavy metal block to approximate the weight of the cell – embedded with two HS-J-L1 geophones²⁴ used as accelerometers. The cryostat was mounted on a thick wooden support on a wall such that the support rod mounting points were square with the X and Y axes in the FEM model from the previous section. In other words, we could excite the cryostat along the same X and Y axes as the model. This excitation merely involved pulling the bottom experimental stage and letting it go into free oscillations that died out over time. An example of the geophone signal during such an excitation is shown in Fig. 7.4 and its power spectrum shown in Fig. 7.5. We find a modal frequency of $f_1 \sim 2.7Hz$.

The example shown is for a mode excited along the direction of maximum stiffness of the mounting support. This ensures that the oscillation being monitored is as close to being the actual bending mode as we can get, rather than a mixed mode composed of cryostat bending in addition to the mounting support bending. The error arising from this issue tends to push the measured frequency lower than the bare bending frequency that we are actually trying to measure. This is consistent with the predictions previous discussed.




Figure 7.4: Example of a raw time-series signal from geophone.

7.5.2 Stiffness enhancement ideas

Since the radiation baffles (when rigidly mounted on the support tubes) do raise the stiffness and modal frequencies of the FEM model relative to the non-baffled analytical estimates, we can try to add more baffles to the model to see how the effect scales. The result is shown in Fig. 7.6 – an effective doubling of the 1st mode frequency upon adding 5 more baffles spread over the cryostat length.

Taking the increasing number of baffles to its ultimate conclusion – we can test a model with the support rods replaced by a thin walled, stainless steel jacket around the entire cryostat. We find that the modal frequency in this simulation comes out to around 60 Hz!

²⁴Geospace Technologies, Houston, Texas.




Figure 7.5: Power spectrum of geophone signal from Fig. 7.4 showing the modal peak.




Figure 7.6: Lowest bending mode of cryostat model (bending is exaggerated), with more baffles added throughout its length to make it stiffer. $f_1 \sim 9.3 \text{ Hz}$ – more than twice that of the original model.

This seems like a promising avenue to explore. We have attempted in the past to build a cryostat with a PVC jacket as a support, but the distributed thermal mass throughout the rather heavy PVC pipe (we suspect) made the hold times for our Dewar intolerably low. Essentially, if the majority of the cryostat specific heat is not all the way at the bottom, we are wasting helium by having to cool down the upper parts of the thermal mass after every single transfer (since that part warms up to near room temperature). However, a stainless steel jacket might work differently, if it is thin and light enough. Other materials can also be explored for this purpose.

Other ways to stiffen a cryostat include arrangements of Kevlar wire²⁵ under great tension (see Refs. [78] and [79] for more details).

7.5.3 Conclusions

The techniques and results presented in this section give us usable tools (approximate but not terribly so) to predict the normal modes of cryostats modified in different ways. We can use these tools to test any proposed modifications before investing the rather large amounts of time and labor needed to modify most cryostats. The goal (presently) is simply to drive up the cryostat normal mode frequency away from the resonant frequency of the suspension system (again, see Section 9.1.2 for details).

²⁵Our thanks to Bryan Steinbach for discussing this issue and providing references. This kind of suspension was used primarily for thermal isolation in helium fridges used for infrared bolometry for near-Earth space experiments.

Chapter 8

Cryogenic valves

This chapter is based on an internal report on the construction and testing of three new cryogenic valves (henceforth called cryovalves) in 2007-2008. The techniques described here owe much to extensive discussions with Yuki Sato and (as with any other piece of laboratory lore passed down through the ages) anecdotal/apocryphal information is largely inevitable. We include such information for completeness, with appropriate labels when needed.

For guidelines on proper usage of these valves in superfluid experiments (cell-filling, closing and opening the valve below 4 K), see Section 10.7.

8.1 Introduction and history

The designs described here are based on previously constructed valves by Bruckner, et al. [80] that were modified by Sato [46, pp. 51-56]. As we have seen in earlier chapters, these quantum whistle experiments involve filling a small cell with helium at temperature close to the lambda point ($\sim 2K$) and observing small displacements of a flexible diaphragm using a very sensitive microphone. As such, isolation of the experiment from ambient acoustic vibrations is a very important requirement¹. After the cell is filled and the fill valve on the top plate is closed off, liquid helium still remains in the fill line approximately up to the bath level during filling, above which helium exists as a gas at the vapor pressure at ambient temperature (see Fig. 8.1). It is observed that this column of liquid+gas channels a large amount of acoustic noise down the fill-line and into the cell, which interferes with the measurements². The main purpose of the cryovalve is therefore to acoustically isolate the fluid in the cell from ambient noise outside the cryostat. A useful side-effect is the reduction in the risk of over-pressure in the filling line when the bath liquid level drops and during bath helium transfers (when warm gas is used to pressurize the bath space).

¹Other aspects of vibration isolation discussed in Section 9.1.

²Even with a cryovalve, this noise is easily observable every time after filling the cell and before closing the cryovalve




Figure 8.1: Comparison of cryostat filling line without and with a cryo valve (CV). Hashed area is the bath liquid. Solid blue is the liquid inside the fill line after filling the cell. Pale blue is the vapor column above the free liquid surface in the line without a cryo valve. If a cryo valve is available, the cell can be filled, the cryo valve closed off and the liquid and gas above the cryo valve can (and is) pumped out to enhance the acoustic isolation. This has the further, highly desirable consequence of preventing over-pressure in the fill line (and running the risk of popping a seal somewhere) when the bath level drops and during bath transfers (when warm gas is used to pressurize the bath space).

Here we describe a pneumatically actuated valve, which has worked well at sub-4K temperatures in making superleak-tight plugs. The main idea (see Fig. 8.2 for a simplified schematic and Fig. 8.3) is to pressurize a volume defined by a flexible bellows using helium gas, thereby expanding the bellows and driving a conical plastic tip into a highly polished (conical) steel hole. The plastic deforms to form a tight seal against the steel, to the extent that even superfluid helium cannot leak through. The only design change in this version of the valve has been to raise the relaxed position of the tip by 7 mils by making the outer housing can slightly longer (0.907" as compared to 0.900" for the older design) so that the relaxed position of the tip is slightly above the seat - this makes it more likely that the tip won't get jammed in too easily at usual closing pressures. The tip diameter is reduced from 0.125" to 0.100" keeping the tip angle 90 deg. This makes the equilibrium position of the tip far enough away from the seat that flow is relatively unimpeded in the open position. Also, when valve is actuated, the tip is smashed into the seat but not too tight so that merely pumping out the actuation line is enough to pop the seal and open the valve. This design

(as an improvement over the Bruckner design cited earlier) includes a non-chamfered seat hole (that is subsequently polished to a small chamfer). This design has been demonstrated (over dozens of cooldowns) to be less likely to get the tip jammed in to the seat hole.




Figure 8.2: Simplified schematic of pneumatic cryo-valve. The use of the off-axis port to send in the helium for filling the cell is not an arbitrary choice. In the event that the valve tip gets stuck in the seat, the much larger area of the plunger “seen” by the gas entering the off-axis port can allow popping the tip out by pressurizing this port.




Figure 8.3: Scale model of assembled cryo-valve. Brass tubes to be hard-soldered to actuation, off-axis, and axis ports are not shown. Bellows are shown at relaxed length. Note that the seat where the tip will touch upon actuation is not shown chamfered. This is because it is not chamfered during machining and is polished directly. The resultant chamfer will be too small to see at this scale.

8.2 Fabrication

8.2.1 Parts

The various parts (Fig. 8.4) that go into the final assembled cryo-valve shown in Fig. 8.3 are listed below with engineering drawings included. Note that these parts must be machined to very strict tolerances (~ 1 mil in most places), especially with respect to the axial alignments of all pieces (also known as concentricity of all central axis bores and holes and tip, which must be within 0.5 mil). Axial misalignment is the most tricky problem when constructing valves and can crop up at several stages of fabrication (we highlight this at the relevant stages).




Figure 8.4: All parts after machining.

1. Bellows (Fig. 8.5) are electrodeposited nickel, model FC-14-L, purchased from Servometer Corporation, Cedar Grove, NJ, USA.
2. Brass housing can (Fig. 8.6) encloses all parts. Flat edges are indium seal surfaces and should be kept smooth³. A flat may be milled on the curved surface and blind threads tapped into it for convenient mounting (see the brass housing can used in Fig. 8.11).
3. Brass guide (Fig. 8.7) seals one end of the cryo valve chamber and forms the main reference for axial alignment of all parts. The bellows slides onto the outer diameter and the plunger slides into the inner diameter.
4. Brass plunger (Fig. 8.8) is soldered to the bellows and guide. A plastic tip is screwed onto it and the plunger moves the tip to make or break the main valve seal. It is important that the tip be screwed in so that it sits flush against the plunger surface. To ensure this, a 6-32 clearance step is drilled to allow the unthreaded part of the tip's shank to fit inside.
5. The tip (Fig. 8.9) is made of Torlon 4203L (purchased from E. Jordan Brooks, Fremont, CA, USA). There is no significant difference between the 4203 and 4203L grades except that the former is extruded and the latter is injection-molded. Chemically, they are the same and 4203L is slightly tougher at room temperature⁴, which makes it a more desirable alternative since it is less likely to get too badly deformed during 300 K testing.
6. Stainless steel seat (Fig. 8.10) has two ports (off-axis inlet and axial outlet) which are connected or disconnected from each other by the seal formed at the central seat hole

³See Appendix A.2 for more notes on making reliable indium seals.

⁴According to Solvay Advanced Polymers, private communication.

by the tip. The seat hole should not be chamfered by the machinist (the whole inner surface should be finely polished on the lathe however - this deburrs the central hole and gets it ready for the polishing procedure described below).




Figure 8.5: Servometer FC-14-L bellows. Gets soldered to guide and plunger.




Figure 8.6: Brass housing can. Flat should be milled on outer curved surface and appropriate blind-taps sunk into it for mounting. Obviously, this part of the design can be customized to the end-user's needs.

8.2.2 Preliminary preparation

1. Get all parts machined by main shop to very tight tolerances. The Torlon tip especially **must** be machined by a highly competent machinist. Marco Ambrosini has done this in the past with excellent results so plan early and beg him to take it on himself.
2. Clean all parts in standard methanol-acetone-isopropanol⁵ ultrasonic clean sequence preceded by a detergent rinse (Simple Green or the like). Be very careful while using simple green. Dilute with 10 parts water. Clean for less than 10 minutes and rinse with water immediately thereafter. Do not use it concentrated as brass is especially prone to tarnishing in concentrated Simple Green if kept too long. If there's the slightest sign of tarnishing (or do it anyway), dip the pieces in very dilute Hydrochloric acid. A concentration of about 10% works vigorously and restores shine before you even

⁵Isopropanol = IPA




Figure 8.7: Brass guide (plunger slides inside this).

finish dipping it in! So, lower concentrations might be better. 5-6% is also fine in my experience. Immediately rinse with water afterwards as the acid does etch the surface. The etch rate is not too high but if acid is allowed to stay on after removing from bath, it will evaporate non-uniformly and therefore etch unevenly. For a dip of less than 10 seconds in 6% HCl, I saw that the finish is brilliant while retaining the machining grooves from the lathe. These grooves are useful for the indium seal. Torlon is fairly resilient to these baths (not the acid) so there shouldn't be any chemical damage as is the case with Stycast. Nevertheless, about five minutes in each bath should be sufficient. Also, it may be wise to skip the stronger solvent baths for the Torlon tips and instead just do a soap water, water and IPA clean sequence.

- Note that once the tip is machined properly, it should NOT be polished by hand as this will ruin the finish and may introduce scratches. The lathe should be the last cutting tool the tip sees. Also be sure to properly store the cleaned parts and never allow dust to settle on them.




Figure 8.8: Brass plunger (tip is screwed on to this).




Figure 8.9: Torlon tip. The side-hole dimensions are non-critical as it is used solely to insert a small rod and screw the tip into the plunger.

4. Immediately after machining, brass tubes ($3/32$ " OD, thin wall) are hard-soldered into each of the two ports in the seat and the actuation port on the guide. Around 1-2" length is sufficient. once the bellows/plunger/guide subsystem has been assembled, further capillaries can be soldered onto the brass tube for use during testing. High heat should be avoided on the seat after polishing (flux fumes and heat might mess with the polished surface), however soft-soldering a capillary to the end of a long brass tube should not be a problem. If extra caution is desired, a small wet wipe can be placed on the hard-solder joint to dissipate the heat and prevent it from reaching the polished seat hole area.

8.2.3 Seat polishing

Do not perform this on the clean bench - this only contaminates the bench as the process is messy.

1. The stainless steel seat is cleaned thoroughly as described above and mounted carefully on a convenient stand (an old or unused housing can makes for a perfect stand - see Fig. 8.11).
2. A wooden applicator is sharpened with a *clean*⁶ knife to pyramidal shape (a *clean* pencil sharpener might do things more symmetrically). A much better tool is the humble toothpick (since it is made of hardwood). A pyramidal-head toothpick with a large opening angle (shallow slopes) works best in this case since it tends to polish (and define a chamfer on) the upper surface of the hole, where the tip will touch, rather than lower inside the hole. It doesn't have to be conical or perfectly symmetric, just shallow

⁶Note that all tools used must be sparkling clean and free of humidity to prevent contamination of the parts

5. Thoroughly mix the powder (starting from the large size - 15 microns and moving down to 0.3 microns⁸) and oil to form a thick paste with no lumps in a clean petri dish (several in the dark room).
6. Dab some on the sharpened wooden applicator and affix to some rotary device (a motorized screwdriver works best as the speed can be adjusted by lightness of touch rather than dialed to rather large values as is the case in a drill or a dremel tool).
7. As symmetrically as possible polish the seat ensuring that no one side gets preferential treatment. Precessing the tool in a circle should help avoid non-uniformities. The tool should be held at an angle from the vertical rather than fully vertical.
8. For smaller grain sizes (1 micron and below), the polishing action is very fast and murder on the toothpick. Do not run it for too long (\sim half a minute to be safe) or the tip gets polished to a long pointed saber which easily breaks off in the seat hole. If this happens, run some wire gently from the outside to push the nub out. This could scratch the area so it's not the ideal scenario. Even after taking this precaution, the tip can sometimes stay stuck. In that case, pull vertically to avoid breaking the tip off. Better yet, remove the tool while the motor is still running.
9. The toothpick can be used in the motorized tool as above (just break it in half - each half can be used and the lengths are perfect). Alternately, it can be spun by hand in the seat. I have found this to produce comparable results, at least when viewed under a microscope (haven't done any cryovalve tests yet). The advantage is exquisite control and stability compared to the bulky and inconvenient power tools. **Update:** In retrospect, polishing by hand is simply not fast enough. Several iterations may have to be done to attain the same finish as a single iteration of motorized screwdriver polishing.
10. After polishing with each powder size, the seat must be cleaned thoroughly (simple green and tap water rinse - apply some force with water spray, then ultrasonic cleaner with all 3 treatments as before) before proceeding to the next smaller size. Any larger particle leftover from the previous polish step will only serve to scratch the surface during the next step so this is crucial. It must be noted that fresh solvents and clean beakers⁹ must be used for cleaning each time as leftover particles from previous cleans can aggregate on the pieces. Existing solvents may be used only if the same particle size is being repeated. I found that when the polishing paste is blackened (metal has

⁸Yuki reports that he went from 12 to 0.3 in 3 steps. Bruckner, *et al* [80] also report a terminal polish size of 0.3 microns. What matters is that the order be strictly observed else the process is rendered counter-productive.

⁹This might lead to using a lot of solvent. A nice way to conserve solvent is by putting in 3 or 4 #4-40 screws in the seat screw holes the wrong side up so that the seat piece can rest on the screw heads (use socket head screws) leaving a small gap between the beaker bottom and the seat surface for solvent to flush. Then, it only takes a tiny bit of solvent for each cleaning step.

come off), it is best to just clean it instead of using a fresh stick on the goeey mess - the metal slurry might scratch?

11. In addition, after every polishing step, it is necessary to check the seat under a powerful optical microscope (near the clean bench) to ensure the circular uniformity of the seat and the lack of scratches and remnant polish powder particles. Be sure to use each applicator only once as wood debris tends to form and in extreme cases also tends to get charred (especially when power tools are used).

Safely store the polished seat on the clean bench in a closed box. Practise polishing on bad seats first to ensure a mastery of the process.



Figure 8.11: Photo of a seat mounted upside down on to an old housing can and ready for polishing. Be careful to use plastic washers (cut out of transparencies is fine) under the screws to avoid ruining the indium seal finish.



Figure 8.12: Photo of a plunger/guide/bellows subsystem being soldered (solder is still wet in the photo). The stand on which the guide is mounted, is a copper block with an axial hole to allow clearance for the brass tube hard-soldered to the actuation port. The large radial hole in the stand is very important as it allows hot air to escape (which would otherwise lift up the bellows during soldering - leading to misalignment).

8.2.4 Assembling the plunger/guide/bellows subsystem

This assembly should be done on a hot plate (again, no point in doing this on a clean bench as it will only mess up the bench). All soldering should be done with rosin-free solder (a brand new spool of 60/40 Sn/Pb sold core solder has been purchased and should be in the appropriate place) and blue Superior No. 30 liquid soft solder flux. To be sure, just before soldering, gently clean the solder wire with Scotchbrite pads or steel wool. Note: Do not use 50/50 solder as it has very low mobility (melts into a pasty lump covered with a skin) and is not appropriate to this application. Several tests with different solders and fluxes have shown that 60/40 (Sn/Pb) solid core solder (no flux) with blue Superior #30 liquid flux¹⁰ works very well. This has far superior wetting and mobility compared to the usual multi-core solder wire (embedded rosin flux) used in simple electronics applications. Also, while the blue flux leaves corrosive residues (chlorides), they are water soluble and can easily be cleaned. Residues have never been a problem in the past. We have noticed that rosin-based fluxes are much more prone to blackening and probably not suitable in this application, especially for joints that may need re-flowing to fix leaks (it is not practical to simply throw away entire parts in such cases).

We have recently discovered the wonders of solder paste (a suspension of microscopic globules of the usual Sn/Pb alloy in a liquid flux base) that can be directly injected in small quantities¹¹. We would expect this to work quite well since it **might** remove the necessity of most of the procedures outlined below by being able to simply “glue” together the parts with this paste and heat the joints locally with a hot air gun to activate the paste and finish the joints. However, this has not been tested and is not guaranteed to work (but is certainly worth the effort).

1. First determine the melting point of solder by placing it on the hotplate and gently raising the power (primitive¹² hotplate dials are usually calibrated in terms of power used rather than final temperature which depends on the thermal mass placed on it). The setting that will be used for this assembly should be slightly higher than this setting to allow for the thermal mass of the components. Ultimately, all we can suggest is to use hotplate settings that work well.
2. The guide is the piece with the actuator port and it forms one of the end caps for the housing (can) while the seat forms the other. The guide is mounted on a stand, which is a copper block with an axial hole to allow clearance for the brass tube hard-soldered to the actuation port. The large radial hole in the stand is very important as it allows hot air to escape (which would otherwise lift up the bellows during soldering

¹⁰Manufactured by Superior Flux & Mfg. Co. and distributed by several suppliers, such as Amtech, INC., Deep River, CT, USA.

¹¹We have used this solder paste for soldering leads onto sub-mm sized chip resistors used within the cell. Model# KE1507-ND, Kester solder paste 63/37 with no-clean flux from Digi-Key Corp., Thief River Falls, MN, USA

¹²Primitive = those without temperature feedback circuits.

- leading to misalignment). The copper block allows for good thermal contact with the hotplate and increases the thermal mass to avoid large temperature fluctuations during soldering.
3. The bellows should fit snugly onto the first radially outward step on the guide and onto the matching ridge on the plunger. It may be easier to start with the heat off and assemble this tower first. Dab some flux on either side of the bellows so as to coat it evenly. Ensure that flux doesn't coat any other part of the bellows as any solder flowing onto the ribs of the bellows would be catastrophic.
 4. After placing the bellows securely on the guide and the plunger on top (everything should be snug), dab a little more flux over each joint. Again, it should wet the joint only. It might be helpful to use a housing can piece as a support while doing this. Use screws to hold it in place. Using the Puritan cotton-tipped applicators has been problematic for me in the past, especially for cramped spaces so one may wish to build one's own applicators somehow. The Puritan company (among others) makes lint free applicators that are thin and flat - perfect for soldering applications. These are primarily used as cleanroom swabs (good keywords for an internet search). This is (at least for me) a very nontrivial tool. Ultimately, we used lint-free, polyester, cleanroom swabs from Texwipe¹³ with very good results. In fact, we continue to use these swabs exclusively in every application (gluing, fluxing, cleaning, etc.) that previously required cotton swabs.
 5. Since the joint is a thin circle, squash the solder wire between the jaws of a clean pair of pliers or something similar to make it a flat ribbon like wire. **Update:** I found that this doesn't make solder wire thin enough, plus it might be the reason for contamination in my first piece - most pliers lying around do not have very clean jaws and crimping might get dirt/debris wedged onto the wire that is not easy to later clean out. Instead, one should do clamp the wire between two *clean* steel slabs and hammer it out into a ribbon (or use a manual or hydraulic press for this purpose). The cleaning part is important as getting dust on the solder and squashing it embeds the dust in it and can never be cleaned away.
 6. Start up the hotplate and allow the flux to sizzle. The places where the flux has been applied should become a nice, clean, shiny copper-like surface. This is the moment when you'll see how well you applied the flux. Any flux on the vertical walls of the first step leads to grief later on as the solder tends to leak over them and fall down onto the indium seal surface. Note that the joint must be superleak-tight so bubbling solder is a strict no-no¹⁴. If the flux has been applied properly and if the temperature is correct, the gently applied solder should uniformly coat the joints. Yuki says that you might have to lead the solder along (like that strange winter olympics sport with

¹³Model: Micro Absorbond TX759B from ITW Texwipe, Kernersville, NC, USA.

¹⁴I verified this personally - to my utter dismay.

the brooms on ice - called “Curling”, I think) with gentle brushes of flux so it flows evenly.

7. Be especially careful that no solder gets on the flat surface of the guide with the groove, which will hold the indium seal. Flux will inevitably get on there but may be cleaned thoroughly with isopropanol on a cotton applicator. Desperately avoid re-polishing this surface (although the only real danger here is if solder gets on - tiny scratches don't matter as the indium gets squashed and forms a good seal despite them)¹⁵.
8. This is tricky as both joints have to be done fast and nearly simultaneously (once solder coats one joint, it is not wise to keep heating the joint for too long so the other joint must be finished reasonably rapidly).
9. Keep a wooden base nearby. Once the soldering is done, turn off the hotplate and leave the assembly to cool slowly on the plate. Do not move it at all until it has visibly hardened. If it suffers any shock before hardening, it may shock-cool and crystallize, ruining the joints and opening pinholes. A bit later, lift the assembly (holding the copper base block we used to elevate it) and place it on the sheet to cool down some more. To speed up cooling at this point, one may put the assembly on large brass or copper stock to bleed the heat away. The slower the cool time, the stronger the joint and the less the probability of leaks due to cracking.
10. ***Very important note:*** Do **not** store the assembled piece sideways (especially with the bellows exposed). This was (I believe) the cause for misalignment in one of the valves constructed. Always store the piece in a housing can, even when cleaning. This is extremely important. I cannot stress this enough.
11. Once this is done, clean the assembly in the ultrasonic cleaner if it looks dirty but otherwise alright. For blue No. 30 flux, the assembly must be cleaned first in DI water, with a little bit of surfactant (mild soap) thrown in. Without a surfactant, the ultrasonic cleaner will not work (surface tension of water is too high to allow bubble cavitation). Once the soap water wash is done (10 minutes or less), it is very important that the assembly be thoroughly rinsed with water (DI preferred but tap water is alright as long as it is not allowed to dry on the joints. Immediately begin the solvent sequence (methanol-acetone-IPA) and blow dry when done. The soap must be removed quickly as otherwise alkaline cleaners tend to induce pitting in the solder (only if concentrated or left too long but remnants might do this over time as well - no sense in taking chances). Note that getting anything inside the bellows assembly can

¹⁵Actually, the error bar is more forgiving here due to the fact that we have an indium seal on this surface. If you do get a tiny bit of solder on this surface during soldering, quickly (while it's still on the hot plate) spread it around so it forms a VERY thin layer. Then later, one can use finest grit sandpaper and gently polish the surface to near-smoothness. The important thing is to have uniformly flush contact (a few mils off should be taken care of by the indium) with the outer can. It is therefore MORE dangerous to have a solder layer that juts out radially from the assembly.

make it very difficult to get it back out. In general high-purity acetone or IPA are fine as slightly warming the piece with a hair-dryer or low power heat gun can evaporate these solvents. Dangling the assembly with the actuation port tube sticking out of the water is a safe way to clean it in soap water.

Practice this soldering step on bad pieces or dummy parts first to avoid ruining good ones.

Leak-testing:

- Check that bellows inflate upon gently blowing air into the actuation port.
- A quick leak-test may be done at this point by blowing nitrogen or compressed air into the actuation port while the assembly is immersed in IPA. Check for bubbles to see if the joints are roughly leak-tight. Be sure to start gas flow **before** dunking it in IPA, else the IPA can form sealant films on tiny pinholes and make it appear as if it is leak-tight when in fact it isn't.
- A more thorough leak-test should also be done at this point (be careful though, the bellows are vulnerable) with a mass spectrometer leak-checker. Be sure not to bang it around while doing this. In fact, mounting it upside down on a housing can and then clamping the can might make a huge difference in safe handling. While doing this test, note carefully whether bellows compress at the initial pumpdown. This shows whether or not a line is blocked. If bellows do not react to pumpdown/venting of the leak-checker test port, the lines are blocked and the bellows assembly is not being tested at all.

Important note:

The final leak-check above is very important before proceeding to the next tests because if the bellows leak, the actuation line cannot be guaranteed to be evacuated and the bellows retracted from the seat. In that case, the bellows can fill up with air and when you pump out the seat ports for leak-checking those lines, you can actuate the valve with ~ 1 atm pressure, thus damaging the tip at 300 K. It is therefore advisable to ensure that the bellows assembly is leak-free and can be pumped out to high vacuum (and bellows show motion) before moving on to final assembly and leak-checks.

8.2.5 Final assembly

From this point on, everything must be done on the clean bench.

1. With indium wire of the right length, snip the ends at a slant so they match up upon laying down in the groove. A tiny bit of Apiezon vacuum grease (M or N) can be dabbed onto a clean piece of indium wire as long as it is gently wiped off with lint-free paper (lens paper works well), prior to sealing the can. If this is done, inspect

the indium wire to ensure no dust has gotten stuck to it. The grease is essential for seamlessly removing these seals. Without it, some indium residue will nearly always remain on surfaces after breaking seals. Ideally, these valves should never be opened again, but the bellows assembly being what it is, it would be nice to be able to reuse it in case the tip gets smashed at some point. It is also possible that things leak the first time and one may have to open it up, fix any leaks and try it again.

2. Screw on the Torlon tip so that the raised ridge on the plunger fits into the depression on the bottom of the tip piece. A rod may be inserted into the hole in the side of the tip piece and tightened as much as possible (remember that it is plastic and adjust your strength accordingly). Shine a light to ensure that the tip sits flush against the surface, else it may be misaligned (bent).
3. Note that the housing, guide and seat pieces all have a flat chord across their cross-sections for the purpose of alignment. This may be used as a reference while assembling the pieces.
4. Seal the guide onto the housing with a torque-driver to ensure consistent pressure from all screws. The torque-driver should be dialed up in stages to the maximum value of about 80-100 in-oz, making sure that diametrically opposite screws get attention in turn.
5. Then seal the seat onto this assembly in the same manner.

8.2.6 Inline filters

As seen in Fig. 8.13, simple inline filters can be used at the entrance to each of the three ports to prevent contamination of the sensitive surfaces inside the cryovalve. These filters ensure years of reliable operation without worrying about particulate contamination that can mess up the tip-seat mating surfaces and cause leaks in the cryovalve.

Two telescoping sizes of copper-nickel tubing are obtained. Two lengths of the smaller size tube (several inches long) and one length of larger tube (at least $\sim 2''$) are cut, deburred and cleaned. A small piece of cigarette paper (makes for a nice lint-free filter) is rolled into a cone and forced into the larger tube using one of the smaller tubes. The two smaller tubes meet in the middle of the larger tube with the filter paper between them. Now, the two ends of the larger tube are soft-soldered shut. Care must be taken during this step to not burn the paper. A simple way to ensure this is by wrapping a small wet wipe around the center of the large tube (where the paper is) to conduct heat away from it.

The impedance of these filters should not be too high. A simple rule of thumb is that gas passed into one end of this filter at ~ 1 atm of differential pressure should bubble vigorously when the other end is placed in alcohol (not water - to avoid contamination).

One of these filters should definitely be placed just prior to the actuation port and the off-axis port (since these carry gas from the outside in). We have not used a filter on the




Figure 8.13: An inline filter. See text for details.

outlet of the valve (axial port) that leads into the cell since that side is always clean. The filter not only guards against particulate impurities, which could disrupt the seal, but also prevents oil fumes from backstreaming into the cell and plugging up the aperture arrays. As such, this is a very important component for such experiments.

8.3 Testing the cryovalve

These tests should be performed in order and the valve should pass each test before moving on. Failure at any step may be fatal (to the valve, not the tester). Refer to Fig. 8.14 for a general plumbing layout that may be useful while testing. A simple jig (see Fig. 8.15) may be constructed so that the cryovalve may be safely lowered into a small liquid nitrogen dewar for the cold tests describe below.

8.3.1 First test - 300K Leak checks

1. Just after assembly, a quick leak-test should be performed on the actuator side of the valve to check the seals. In fact, this can (and should) be done after assembling just the guide onto the housing. This will be repeated after the full assembly is complete.
2. After valve is assembled and mounted on the test jig, the order of leak-checks must be followed (for very important reasons) as outlined below:
 - a) First leak-check the actuation line, including buffer can. The helium should be sent into the (off-axis) fill port while the seat port is open to atmosphere. If the actuation line leaks, do *not* proceed with further leak-checks unless and until it is

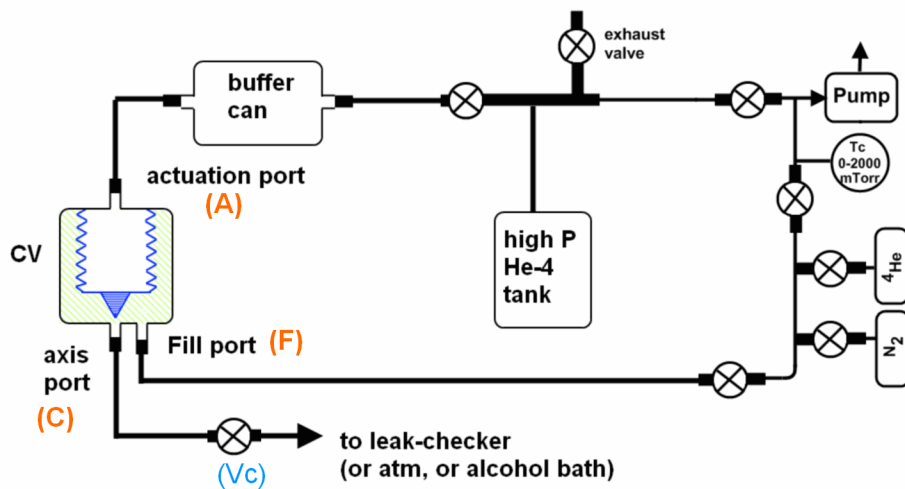


Figure 8.14: Plumbing circuit for testing cryovalve under various conditions (see text).

fixed. If actuation line leaks, it cannot be pumped out and (depending on the size of the leak) *could* reach atmospheric pressure while the seat side is being pumped out by the leak-checker. This would put a differential pressure of ~ 1 atm (15 psi) across the tip, actuating the valve and deforming the tip.

- b) If the actuation line is leak-free, first pump it out to high vacuum and plug it off. This makes the tip side pressure ~ 0 psi absolute so that when the seat is pumped out by the leak-checker, the tip never sees positive differential pressure trying to actuate it. Then, the seat side lines and the valve housing may be leak-checked by spraying helium from the outside.

Both tests must be passed to move on.

8.3.2 Second test - 300K - plunger motion

Then, a rough operation check should be performed to check for plunger motion. This is described below:

1. Hook up a pump and Tc millitorr gauge to the inlet port (off-axis) while leaving the outlet open to atmosphere. The pressure should reach some on-scale equilibrium (~ 700 microns of Hg depending on the impedance of the lines).
2. Then actuate the valve with a few psi (less than 5psi) of clean Helium from the high-pressure line. There should be no ambiguity about pressures here. The actuation pressure is always defined as the differential pressure across the seat. In this case, the pressure beyond the seat is atmospheric (~ 15 psi absolute¹⁶) so the actuation pressure

¹⁶Absolute pressure is measured above a perfect vacuum, which is also ~ 30 inches of Hg below atmospheric.




Figure 8.15: Photo of testing jig with cryo valve mounted. Brass tubes permanently hard-soldered to the 3 cryo valve ports are temporarily connected via copper-nickel capillaries to larger tubes on the jig with KF ends for easy connections. The entire jig is lowered into a small liquid nitrogen dewar for cold tests. Buffer can shown is necessary to avoid over-pressuring bellows under various conditions (leaks, temperature rises, etc.) and is a required safety component to protect the valve (also during actual operation).

will be 2-3psi above atmospheric (the zero on the tank regulator gauges is usually atmospheric so in this case, the proper (differential) actuation pressure is the same as the regulator pressure. Be sure to send a little helium into the buffer can and equalize with the inner pressure a few times and then control it directly from the tank.

3. At some point, (very quickly) the inlet line pressure should drop suddenly to base pumping pressure (~ 30 -50 microns). This denotes motion of the plunger - a test that must be passed before moving on.
4. Warning: do NOT actuate the valve beyond ~ 5 psi here as the tip can be squashed against the seat (Torlon is soft at this temperature).

8.3.3 Third test - 77K - the ink test

The entire valve should be pumped out (both sides) before cooling down in liquid nitrogen to avoid freezing air inside. A dab of a permanent marker on the seat prior to this should leave a nice circular mark on the tip after actuation to about 20-30 psi. It should be noted

here that the ink test is a VERY strong test and has worked well in the past in identifying successful valves¹⁷.

Note that this test is done to protect the bellows and tip from distorting completely if the parts are slightly misaligned. It may be possible to recover from this failure by using a different seat or by re-polishing the seat (in case a microscopic examination shows that the seat hole was off-balance). Proceeding directly to the final test (where larger pressures are imposed) may end up completely distorting the bellows.

8.3.4 Fourth test - 77K - final test to determine helium leak tightness and closing pressure

In this test (see Fig. 8.14), the leak-tightness of the valve is finally checked and the actuation pressure required for sealing the valve is experimentally determined. Note that the valve marked “Vc” in the figure is very useful (as will be seen in what follows). The buffer can **must** be used during this test for protecting the valve from mishaps.

300 K steps

1. The actuation line (A) is completely pumped out - this opens the cryovalve path from the off-axis fill line (F) to the axial cell line (C).
2. The actuation line and bellows subsystem is leak-checked again by squirting some helium into either (F) or (C). Do not proceed if a leak is found here or you risk actuating the valve at 300 K and damaging the tip.
3. Now, with valve (Vc) closed, the cryovalve chamber (hatched volume in the figure) is evacuated through (F). It may speed things up to first open (Vc) and flush out the lines and chamber with nitrogen gas several times before evacuating the chamber. This dramatically reduces the helium background in the lines, making it easier to perform further leak-checks.
4. With (Vc) now closed and all lines evacuated, the chamber is leak-checked through (F). All solder joints as well as housing can indium seals are leak-checked.

77 K steps

Assuming no leaks so far, the (now completely evacuated valve) is immersed in liquid nitrogen and allowed to reach thermal equilibrium (~ 1 hour is sufficient for all internal parts to thermalize - since we will be sending in a constant helium flow through (F), thermalization should occur quite rapidly).

¹⁷Yuki posits that a successful ink test has always led to a successful valve for him

1. Start by actuating the valve a little bit by pressurizing (A) to ~ 10 psig¹⁸. Note that the chamber is evacuated at this point so that the differential pressure across the (C)-hole will become $\sim 10 + 15 = 25$ psia. This is done so that the impedance for the (F) \rightarrow (C) path will be sizeable enough that the helium we send in next won't overwhelm the leak-checker.
2. Now, pressurize (F) with about an atmosphere of helium while (C) is continuously exposed to the leak-checker. The leak-checker will register a leak rate (possibly in gross leak mode). Do not expose too long as it may saturate the background (just set leak-checker to pump on (C) with the diffusion pump without being in leak mode).
3. Now, we slowly (increments of 5 psi at first, slowing to ~ 2 psi once leak rate drops significantly) increase the actuation pressure at (A). The leak rate should slowly go down to stable values (may have to wait several minutes for the background to get pumped out enough). Record three values at each actuation pressure step: P(A) (actuation pressure), P(F) (helium sent in) and leak rate in std. cc/s.
4. At some point, the leak rate should sink into the leak-checker background. This is the first closing pressure. Go another 20-30 psi in actuation pressure (to "train" the valve) but not beyond ~ 70 psig.
5. At this point, slowly start relieving the actuation pressure. At some point, the leak rate will rise out of the background. Be careful here as this could be very sudden and could overwhelm the leak-checker. Also be careful while relieving the actuation pressure to atmosphere, not to go below 5 psig (to prevent air from going in and freezing).
6. Once the valve is obviously open (finite leak-rate), repeat the above process several times. For a freshly made cryovalve, such "training" is needed to get the hysteresis out of the system and make sure that the tip is seating properly. Done properly, a lowering of the actuation pressure should be observed (see the result plots that follow) with each closing cycle. The final closing pressure should be noted as *the* pressure to be used.

Important note

To ensure that the seal is superleak-tight in actual experiments, an actuation pressure of $\gtrsim 20$ psi higher than the final closing pressure above is chosen. This has worked well in the past.

Closing pressure results

Results for three valves following the procedure described in Section 8.3.4. Cryovalve # 3 (internal numbering) had an outer leak problem due to a bad hard-solder joint on the seat.

¹⁸Psig = gauge pressure (relative to atmospheric pressure) and psia is absolute pressure (relative to a perfect vacuum).

A new seat that matches axially to the otherwise assembled parts needs to be made. The bellows subsystem has been checked and is fine.




Figure 8.16: Cryovalve # 1: With ~ 5 psig He-4 at off-axis seat port. Note that the point $(5, 10^{-7})$ for cycle 1 is put in by hand. The valve suddenly opened while reducing the actuation pressure and overwhelmed the leak-checker (leak rate $> 10^{-3}$).




Figure 8.17: Cryovalve # 2: With ~ 1 atm abs. (0 psig) He-4 at off-axis seat port. During cycle 2, the leak-checker background was saturated. After flushing out the lines and allowing the background to subside, the later two cycles show normal behavior. Note that cycle 2 is still consistent with the trend discussed in the text (of steadily lowered actuation pressures over cycling).




Figure 8.18: Cryovalve # 4: With ~ 1 atm abs. (0 psig) He-4 at off-axis seat port.

Chapter 9

Laboratory infrastructure

In this chapter, we discuss some important infrastructure issues and improvements related to general low-noise operation of the SHeQUID as well as quieter and automated reorientation of the entire experimental platform for the specific purpose of Sagnac interferometry.

9.1 Acoustic and vibration isolation

The quantum whistle is a low level acoustic signal that can easily get buried in the ambient noise of a common research lab. This includes noise sources related to generic lab machinery (pumps, instrument and computer fans) and the more specialized mechanical infrastructure related to reorienting the cryostat for experiments focusing on Sagnac interferometry using the SHeQUID. Beyond these (which are at least within our control) lie the ubiquitous and continuously present ambient hum of a modern-day building, which includes noise sources like air blowing out of ventilation vents, sound modes within air ducts, the main air handling systems embedded within the building and of course the sounds arising from human activity. Beyond this, there are tilts and wobbles of the cryostat in response to external noise stimuli, which couple specifically to a SHeQUID and add a rotational noise component to the interferogram. This section focuses on identifying the paths by which this vibrational and acoustic noise makes its way to the experimental cell and providing solutions and guidelines that we have found of some use in ameliorating the worst effects of this noise. We cannot over-emphasize the importance of noise isolation in our experiments. It is the one factor over which the experimenter has only limited control over and as such can be the deciding factor in whether such experiments succeed.

Refer to Fig. 10.1 for a block diagram of the experimental setup, including the layouts of the three separate rooms used. All three rooms are two stories underground, for added isolation from university noise.

9.1.1 Ambient sound

The cryostat is placed in the sound room and shielded (to a large degree, but not entirely) from ambient sound coming from outside the room by lining the walls and ceiling with sound-absorbing foam. This can be readily improved upon by using some of the many commercial products presently available for acoustic shielding. These typically have foam backed by heavy vinyl or lead sheeting (or even sheet rock). The idea is that the foam faces the sound source and the backing faces the object to be shielded (the cryostat in this case). Any incident sound is absorbed by the foam (essentially converted to heat) and any residue is reflected by the heavy backing and absorbed on the second pass. This kind of system has been in commercial use for a while now and a lot of literature exists on the subject. Having a second sound box just around the cryostat can also be a good way to attenuate the residual sound that makes it into the sound room.

We put most of the noisy instruments (with cooling fans) and the bath pump in an adjoining room that we call the pump room in the layout diagram. The pump room and sound room have a thick concrete block between them and cables are transported through pipes embedded in the concrete (one of the pipes is used as the pump line with connecting flanges on either side). Any unused pipes are stuffed with foam and plugged up with rubber stoppers or wood. The concrete block damps out a significant fraction of the sound transmitted along the pipes. The bath pump is placed in its own enclosure, which is a double-walled box with foam on the inside and thick plywood walls. The two sets of walls should not touch because the air gap is an effective barrier in its own right. We note that the room walls (and concrete block) by themselves seems to be sufficient to shield the sound room from sound carried through the air, to the extent that we see no measurable difference in noise appearing at the displacement sensor with the sound box lid open or closed. However, we find that placing the pump (and any other instruments) on a platform of wood–foam–wood *does* cause a measurable reduction in picked-up noise. Floor vibrations are therefore quite important and must be addressed. Yuki Sato (personal communications) reports good results with a pump that is suspended from the ceiling.

Any instruments (such as room lighting, preamps, etc.) still in the sound room must be thoroughly inspected to ensure there is no residual hum or any other audible noise. Fluorescent room lights are extraordinarily loud in this context and we invariably keep them off during data-taking (the humble incandescent light-bulb comes in handy in this case). Locking oneself in the sound room and just listening very carefully (with lights off) is probably the most useful diagnostic tool one can have (sound level meters are typically not very useful). Using a mechanic's stethoscope (just a doctor's stethoscope with a long, thin metal tube attached) to test various walls and instruments can make noise sources very clear. If sources cannot be removed from the room, enclosing them in small sound boxes can help. The worst offenders (once all the things under our control are taken care of) are invariably air ducts embedded in the walls, ventilators and building machinery. Human noise is **not** negligible – a small scrape of a chair or a door banging can show up as a large perturbation on the highly sensitive displacement sensor or kick up whistles or cause battery

state transitions quite easily, even with all the precautions described in this chapter¹.

All computers (and researchers) go in the control room, which is separated from the sound room by a double wall with an air gap. All cables coming out of the sound room are snugly ensconced in their breakout holes in industrial putty. We find that the experiment is largely oblivious to control room activity once all the doors are properly closed (it is in fact the best shielded wall in the sound room). This leads us to suspect that adding false walls and ceilings with another layer of (vinyl or lead-backed) foam would be a promising next step if better acoustic isolation were desired (it always is).

9.1.2 Ground vibrations

Three nitrogen-filled pneumatic springs (covered by green half-boxes in Fig. 9.4) lift the entire experimental platform $\sim 1/8'' - 1/4''$ above the ground. The (hollow) platform is filled with around a thousand pounds of lead ballast. The springs are kept as “soft” (low pressure) as possible to keep the coupling between the platform and the ground low.




Figure 9.1: Pneumatic spring suspension system.

This suspension system essentially acts as a low-pass filter against ground vibrations with an effective cut-off frequency related to the small displacement resonance of the spring. This frequency can be estimated by a simple linearized analysis of the suspension (see Fig. 9.1). The system can be modeled as one large piston with an effective area A (approximately equal to the sum of all three piston areas) supporting a total mass M with a pressure P in the spring canister². Newton’s second law for the mass M gives: $M\ddot{y} = PA - Mg$, where y is the height of the piston air column. Further, $y = y_0 + z$, where h is the steady state piston position and z is a small perturbation in the position. The ideal gas law (for temperature T , number of moles n and gas constant R) gives: $P(y) = \frac{nRT}{V} = \frac{nRT}{Ay}$, which lets us determine the steady state position from $\ddot{y} = 0$, so that $P(y_0)A = Mg$ and finally, $Mg = nRT/y_0$. We can rewrite the equation of motion using these expressions in terms of the perturbation z :

¹This can make for some interesting work hours forced upon hapless experimenters.

² P is the gauge pressure so we can ignore the atmospheric pressure on the piston.

$M\ddot{z} = \frac{nRT}{y_0+z} - Mg = \frac{nRT}{y_0} (1 + z/y_0)^{-1} - Mg$. Expanding the first term in the small quantity z/y_0 and using the steady condition obtained above, we finally obtain: $\ddot{z} = -\frac{g}{y_0}z$. This is an interesting equation because it also describes a small angle pendulum with length y_0 and fundamental resonant frequency $\omega = \sqrt{g/y_0} \sim 0.7 \text{ Hz}$ (for a piston height of $\sim 20''$). None of the other quantities come into play (at least where the resonance is concerned).

The idea is that frequencies greater than this will be attenuated by the spring as they pass to the cryostat. As we saw in the previous chapter (Section 7.5), we should try to push the normal modes of the cryostat frame as high as possible (for, for instance, stiffening it) so that they are far from this spring frequency and will therefore couple less to the (attenuated) vibrations that do leak through this filter.

The lead ballast serves both to stabilize the platform and suppress force noise. At this point, we should clarify what we mean by force noise and distinguish it from displacement noise. Displacement noise is something that (approximately) does *not* scale down by the inertial mass that it drives, while force noise is something that does. An example of the former would be building motion or seismic motion. The latter could include things like sound waves in air or through a medium (like pump hose skin). Force noise can be attenuated simply by putting a large mass in the way, so that the entire system doesn't accelerate as much in response to the noise. This is what we do at several places in this chapter, so anytime we talk about tying down something to a large mass, this is the implied goal.

9.1.3 Pumping lines

A double-gimbal bellows system [81] couples the pump inlet to the main pumping line to reduce the transmission of vibrations from the pump to the experiment. Past the concrete barrier, a flexible plastic bellows line connects the pump to a valve block with 3 different sized valves used to control the pumping rate for rough temperature regulation. This line is secured to the wall at several places. Past the valve block, we use (large bore) soft gum rubber tubing that is weighted down to a fixed structure in the room with cloth bags full of lead shot (to attenuate force noise, as discussed in the previous section). Note that up to this point, we carefully avoid increasing the impedance seen by the pump, because we are saving that for the final step before we get to the cryostat. The gum rubber tubing connects to a relatively short length (just enough to make it to the cryostat and still be able to reorient the Dewar) of 3/8" ID, 1/2" OD³ clear PVC tubing (simple hardware store variety). This tubing is very light and is secured by elastic cords to a lead brick suspended from the ceiling by a bungee cord (a small piece of reinforced tubing slit lengthwise serves as a protective sheath to prevent crushing the tubing at the tie-down point). Some heavy industrial putty helps damp out vibrations (in the tubing skin) at this and other places.

Of course, none of this does much to prevent sound from leaking in via the rarefied gas in the pumping lines. The gas pressure is around 5 kPa at T_λ and we typically work a few tens of mK within this temperature. We can use Eq. (6.18) to estimate the mean free path

³Wall thickness should be enough to support a rough vacuum and not collapse.

for helium (hard sphere radius $d = 2.2 \text{ \AA}$ and temperature in the external lines $\sim 300 \text{ K}$) as $\lambda \sim 4 \text{ }\mu\text{m}$. The gas ceases to look like a continuum (and to support sound waves) when the sound wavelength gets to be on the order of λ . However, $\lambda \sim 4 \text{ }\mu\text{m}$ still corresponds to a maximum sound frequency of $\sim 55 \text{ MHz}$, which means that the gas in the pump lines will not attenuate any of the incoming sound just because it is low pressure. We do not presently do anything about this issue, but we would suggest putting in some kind of sound baffles (or cotton balls) in the lines to reflect and/or absorb the pump noise coming in through the continuously evacuated gas. A baffle could be created simply by inserting a small pipe cleaner into a pump line. Alternatively, it could be a thin support rod with small, thin plastic washers strung along it with space in between and small, staggered cuts made in each washer to let the gas pass (but reflect most of the sound back).

A final issue here is a cryostat wobble arising from the slight, unbalanced force acting on the Dewar neck ring where the bath pumping port resides (a torque is applied to the Dewar, which is pivoted at its base). Refer to Section 7.2.3 for more details about this issue. Essentially, the pumping port is moved to the top plate and made vertical to remove this torque.

9.1.4 Cables

The cables that interface to the cryostat can be a significant source of noise and acoustic spikes if not properly secured (especially during reorientation as they bump into each other and the Dewar). We bunch up all the cables using elastic cords and sink the entire cluster to the same suspended lead brick mentioned in the previous section. This large mass and liberal amounts of industrial putty help damp out vibrations traveling along the cables. The cables are also sunk (prior to the brick) to some large fixed structure in the sound room (to attenuate force noise).

9.2 An improved, quiet rotation system

In all our previous work (prior to early 2012), the SHeQUID loop vector \mathbf{A} was reoriented point-by-point to change the Sagnac phase-shift (Section 1.3.2) $\Delta\phi_{rot}$ seen by the loop. This was done because our previous reorientable cryostat included homemade hydraulic or pneumatic bearings that injected considerable noise into the system. This noise increased the base noise level of the displacement sensor (reducing sensor resolution) and also interfered with temperature regulation during rotation. Consequently, the bearings had to be turned off during a measurement.

The situation was exacerbated while using the chemical potential battery due to its hysteresis (Section 2.4.4). The bearing noise and the large temperature fluctuations caused by it would often knock the battery state to nearby stable or (worse) metastable states. Because of the frequency dependent Fiske gains associated with each battery state (Section 2.3), we

might now be tracing out a different interference curve. An example⁴ of such an occurrence is shown in Fig. 9.2. Such amplitude gain jumps are intolerable since all we can do in practice during the feedback is keep the amplitude at a fixed value. Whether a fixed amplitude means a fixed phase depends on our ability to stay confined to one interference curve out of the family of curves corresponding to the nearby battery states. The SHeQUID feedback calibration is thus frequency dependent (at least in a system with the many cell resonances that we observe) and any battery fluctuations or jumps can render it useless and the feedback unusable. A future improvement to the SHeQUID might be an additional feedback system that kept the battery state fixed by small, occasional adjustments to the inner cell heater power.




Figure 9.2: Example of an interferogram where the battery state switched from one Fiske resonant mode (large gain) to another (small gain) and back again while taking continuous whistle data for incrementally stepped up heat-pipe heater powers. The data is divided into two sets that correspond to be two different battery states (distinguished by whistle frequency) and the two-slit model is fit to the major dataset. The minor dataset is displayed with the major dataset and its fit.

For these reasons, we built a quieter rotation stage that employed a simple commercial ball bearing rather than the previously used fluid bearings (see Fig. 9.3 for a schematic and Figs. 9.4 and 9.5 for photographs).

⁴Data for this particular curve was collected by Yashwant Gowda and it is shown here because it is the clearest such example of back and forth gain switching we have heretofore observed (usually, the time scales of such jumps are longer than the typical time needed for a single interferogram so that only a single jump is observed unless there is sufficient noise to trigger more jumps).




Figure 9.3: Dewar on bearing bolted to air springs. Symmetric drive system with kevlar belt. All platforms are rigidly bolted together for stability. Bottom wood platform rests on 3 screw jacks (not shown) for stability when rotation system not in use.

The Dewar is mounted on a 2 ft. diameter wooden base (7/8" thick plywood) bolted to a 3/8" thick aluminum platform on a 1" thick wooden base. This assembly is bolted to a 3/4" wide, 6" diameter pulley with a short shaft that is press-fit into a simple mechanical ball-bearing⁵, which itself is rigidly bolted to a 3/4" thick wood platform. The wooden platform is part of a vibration isolation stage using air springs and ~ 1000 lbs of lead ballast. It is placed on a set of vibration damping rubber pads that sit on top of the lead bricks inside the platform cavity.

The bolt heads used to mount the various pieces together should be counter-sunk or counter-bored appropriately to avoid snagging with the drive belt. Since this can be a bit tricky, we describe the mounting sequences here. Referring to Fig. 9.3, the center Al and bottom wood stages are bolted to the central pulley and shaft by 4 long hex bolts (with heads that sit in counter-bores in the top surface of the Al) that screw into taps in the pulley. The top wood stage is fixed to the Dewar bottom (which has threaded blind holes) by 4 hex cap screws counter-bored up into the bottom of the wood stage.

Two mu-metal shells shield the cryostat components from the Earth's magnetic field. Using a commercial magnetometer at room temperature, we have verified that the volume enclosed by both shields is free of DC magnetic fields down to the sensor resolution of ~ 1 nT (recall that the Earth's field is on the order of tens of thousands of nT). The inner mu-metal

⁵BearingsOn UCF215-48 flange-mounted 3 inch diameter bearing.

shield is an open mouth can with clearance holes for the hex cap screws holding the top wood stage to the Dewar bottom (the mu-metal base is sandwiched between these two and held in place securely). The outer mu-metal shield just sits on the top wood stage (which also has a cloth measuring tape tacked onto its curved side - see Fig. 9.5). The top wood stage is bolted to the bottom assembly (Al+bottom wood stage) using long hex bolts inserted upwards whose heads sit in wells counter-bored on the bottom surface of the bottom wood stage. This last set of counter-bores is to prevent the bolt heads from hitting the drive belt since the pulley is almost flush with this surface. The bolts are fastened to nuts on the top which can be seen in Fig. 9.5.

A drive system consisting of a dc motor, idler pulleys and a tensioner pulley is arranged symmetrically so that torque is applied evenly during rotation. A 1/2" wide Kevlar⁶ belt minimizes speed variation, jerks and slips due to belt stretch.

This new rotation stage injects no observable noise into the temperature regulation system⁷, thus allowing the maximum stability of ~ 20 nK. This is an almost 1000-fold improvement in thermal stability during rotation over prior experiments in our lab.

9.3 Automating rotation sweeps (with redundant safety features)

This section describes the newly developed automatic reorientation system used for obtaining Sagnac interference curves by sweeping the cryostat clockwise and counterclockwise for about 2 days. We can now obtain Sagnac curves in under a half hour in a fully automated way – a significant improvement over the entirely manual procedure used up to this point (which could take several hours of painstaking work to see a full interferogram). Most of the system is self-explanatory from the figures provided and from the rotation stage descriptions in the previous section.

A real-time measure of the cryostat angular position is obtained via a tachometer used as a digital angle meter (see p. 223) whose pulsed output is read by the digital input channels of a data acquisition card (DAQ) connected to a computer. The tachometer has a rotor over which we fix a rubber O-ring and this assembly is pushed up against the central pulley so that it rotates with the bearing (this way, we do not have to worry about belt slippage, as we would if we instead monitored the drive, tensioner or idler pulleys).

A backup to this angle meter is a webcam aimed along a diameter of rotation stage platform at a vertical needle, which acts as a reference crosshairs for capturing the reading

⁶Gates-Mectrol belt model F8: welded Kevlar embedded in flat polyurethane matrix. Other materials are available for the matrix that are rated for higher friction if needed. We found this belt to be adequate for our (slow rotation) needs. Since Kevlar is used specifically for its low-stretch properties, flat (not crowned) pulleys should be used, since crowned pulleys essentially rely on the belt stretching to a certain extent to center itself on the pulley. Refer to manufacturer website or manuals for details.

⁷up to some maximum cryostat rotation rate on the order of tens of deg/min, which is determined by the specifics of the experiment (cabling, weight distribution, etc.)

on a tailor’s tape measure tacked onto the rotating platform (see Fig. 9.5). This webcam image is saved automatically every minute or so to back up or verify the tachometer reading if needed. This webcam also serves to calibrate the tachometer, whose Boolean transitions are interpreted as discrete angle steps. The stage is turned through a known angle (as measured by the webcam) and the number of Boolean transitions (pulses) from the tachometer are recorded. This is repeated a few times for different angles and the slope of a fit to this data (steps vs. angle in deg) yields the tachometer resolution (step size in deg), which is then programmed into the Labview VI responsible for angle measurement. Thereafter, the angle is directly computed from the Boolean transitions and added to or subtracted from a given start angle (set by the user – this can be completely arbitrary) depending on the direction of motion. The direction of motion is known because it is the computer that controls it via a Boolean output from the DAQ. Note that in terms of the notation used in Eq. (1.18), this angle is simply θ . The reference angle θ_0 defined for that equation is the angle read by the angle meter described above at the point along the reorientation when the sense loop area vector is pointing due East or West. As long as θ and θ_0 are measured relative to the same origin angle (arbitrary), we can stay consistent with our definitions.

The output of a home-made optical switch (Figs. 9.9 and 9.10) is digitized by the DAQ into a Boolean signal, which is high when the switch is unblocked and low when it gets blocked by a piece of cardboard screen attached to the rotating stage. The output is tuned to standard Boolean levels by choosing $R_s \sim 1k\Omega$ and power supply $V_{cc} \sim 5 - 6V$ for a commercial CdS photoresistor (nominally $\sim 130\Omega$ in the dark). With the photoresistor embedded in a wood dowel as shown in the figure, putting a piece of 1/4” thick black foam over it with only a tiny hole in it (so that it takes a laser aimed directly into the hole to actuate the switch) makes the switch output $V_s \approx V_{cc}$ with the laser shining on it and $\approx V_{cc}/10$ with the laser blocked (even in a lighted room, because of the foam). In digital input mode, the DAQ sifts these into Boolean High and Low respectively. Two such screens can be placed on the stage to define the “soft” limits of the rotation. A Labview VI detects the screens by detecting the transition between low and high Boolean levels from the switch. A motor, which is controlled by the computer-interfaced circuit shown in Figs. 9.7 and the photographs in this section, is turned on and off, and its rotation direction reversed, when the VI detects that the soft limits have been reached. The speed is set manually by the user via a knob on the circuit box.

This way, the VI can automatically reorient the cryostat in one direction, detect the soft limit, stop and reverse direction and continue on to the other soft limit, and repeat this indefinitely; all the while taking whistle data and analyzing it immediately to produce interferograms on the fly. All this with no human intervention needed.

For safety, the Booleans have been set up in such a way that any computer problems would merely halt the rotation (we make it difficult to rotate, with multiple switches needing to be on at the same time and very easy to stop everything). A final layer of safety (which can be seen in the photographs in this section) is a pair of micro-switches mounted on the floor near the cryostat. These are normally closed and open when pushed (even gently) by wooden blocks fixed on the rotating stage. These can be thought of as the “hard limits”

of the sweeps. In the event that something goes wrong and the stage keeps rotating past the soft (optical switch) limits, these micro-switches actuate and turn off the power to the motor. Power stays off because the stage is now stopped. This requires human intervention to unstick the switch and get it going again. Three screw jacks (see Fig. 9.3) are arranged symmetrically around the Dewar (under the rotating stage). These can be raised to rigidly support the Dewar when rotation is not needed (to prevent wobbles and general instability). They are obviously kept lowered during rotation but not too much (just about 1/4" or less), so that if the Dewar sways or starts tipping for any reason, the screw jacks are there as a backup set of supports.

Finally, we note that this system is not shown in the main layout block diagram of Fig. 10.1 (in the next chapter), because of its complexity. It is quite easy to see how it is wired up in the larger scheme of things by noting that the motor power supply has its own cable (which has been hacked into in order to incorporate the safety switches and digital control). The laser for the optical switch (a cheap pointer) has been hacked to run from a DC power supply, while another DC supply provides the power to the LDR circuit for the optical switch. These remain inside the sound room. The optical switch output, as well as the tachometer power and output are sent to the control room via the several extra coax cables strung between the sound and control rooms. The main data PC (#1 in the layout diagram) controls this entire system via the digital I/O channels on its attached DAQ-1. These DIO channels are broken out of the main BNC-2110 accessory into its own sub-breakout (a custom-built box with BNC connectors) via a ribbon cable. Refer to the next chapter for details about the bigger picture.




Figure 9.4: photo of working setup for auto-rot (annotated).




Figure 9.5: Tachometer close-up and wood sandwich closeups, etc. (annotated).




Figure 9.6: Photo of control room side rot ckt. (annotated)

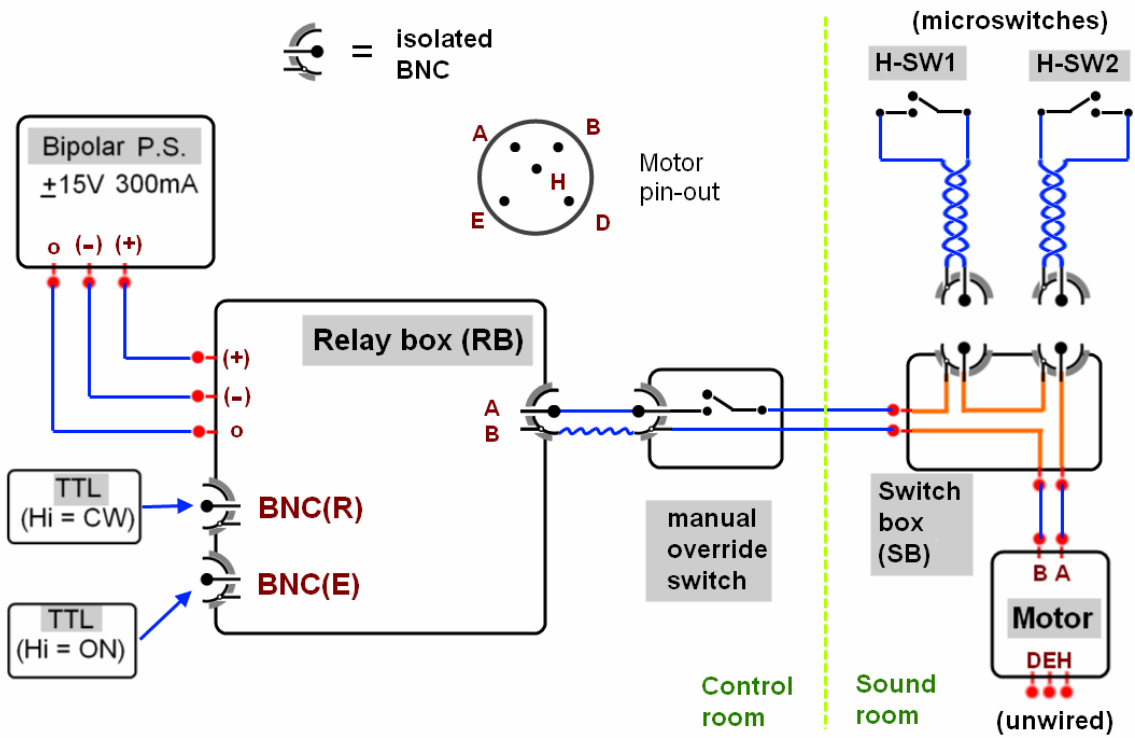


Figure 9.7: Full circuit used for auto-rot. PC not shown. See Fig. 9.8 for a detail view of the relay box (RB) module. Micro-switches are normally closed.

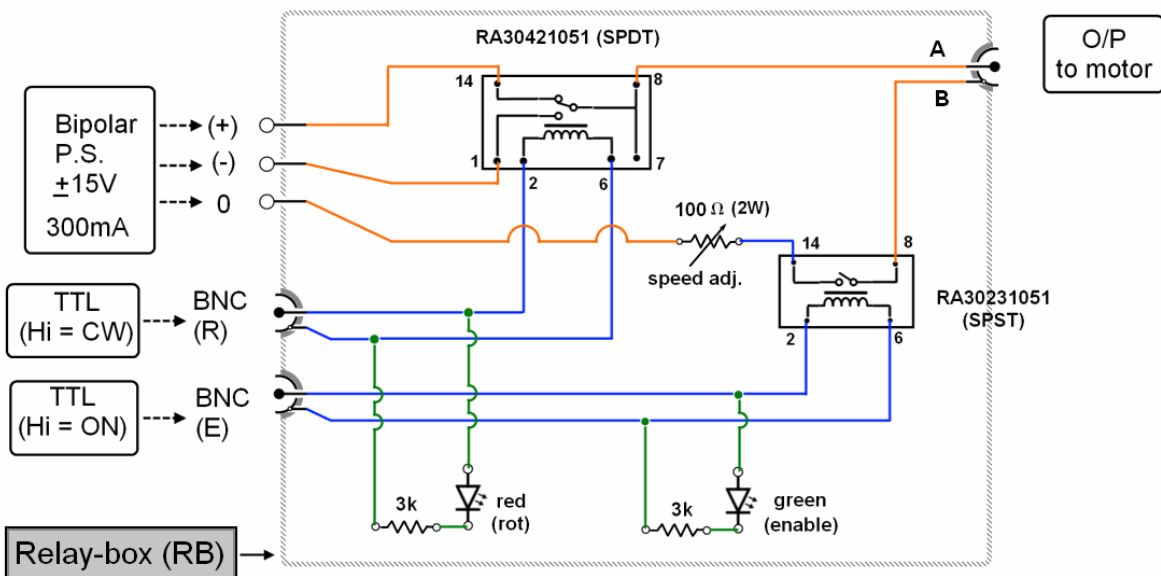


Figure 9.8: Detail view of the relay box (RB) module shown in Fig. 9.7.




Figure 9.9: Home-made optical switch initially used as Dewar angular position sensor and most recently (after switching to a commercial digital angle meter) used for sensing soft sweep limits in auto-reorientation circuit.




Figure 9.10: A simple LDR (light-dependent resistor) circuit used in the optical switch shown in Fig. 9.9. The LDR used here is a common Cadmium Sulfide (CdS) photo-resistor from Radioshack with $R \sim 130\Omega$ in the dark and $R = \infty$ with a commercial red laser shining directly at it. In a well-lighted room with black foam tape over it and just a small hole cut in the foam, the resistance is about a hundred $k\Omega$.

Chapter 10

Operation

This chapter describes in some detail the procedures involved in running a full experiment dealing with the kind of work that we have talked about on a somewhat more abstract level till now. It is important to note that since different experiments require differing levels of detail in the operating procedures, we have chosen a particular *single weak-link cell* run performed in 2010¹ to demonstrate the superset of procedures required. However, everything in this chapter is relevant to running a SHeQUID, except² perhaps for the section on current-phase relations (which is mostly of diagnostic importance in this context, so we mention it only briefly).

10.1 Layout and overview

Referring to Fig. 10.1, we present here an overview of the entire measurement apparatus used for experiments. Note that we will be referring frequently to material already included in the chapters describing the cryostat (7), the experimental cells (3) and the laboratory infrastructure (9) quite frequently in what follows. Please refer to those chapters for more details on particular components or concepts (if not otherwise referenced here).

Our interface with the cryostat can be conceptually divided into two broad sections: thermometry and everything else, where the latter includes such things as the various cell heaters, displacement sensor and bias electrodes. This division is made because the temperature regulation is typically always running in the background and is best kept independent of the rest of the experiment (which is usually extremely volatile in terms of rapidly changing setups).

All instruments that are controllable over a GPIB interface (this includes the Lakeshore 340 temperature controller, the SR 830 lock-in analyzer and the DS 345 function generator

¹Weak-link cell 6 run 1 (Feb - Apr 2010.)

²Of course, a single weak-link experiment can be quite useful as a gentler starting point if one is just beginning to work with SHeQUIDs. It is also wiser to start with this in order to validate the non-trivial technology behind fabricating the aperture arrays.

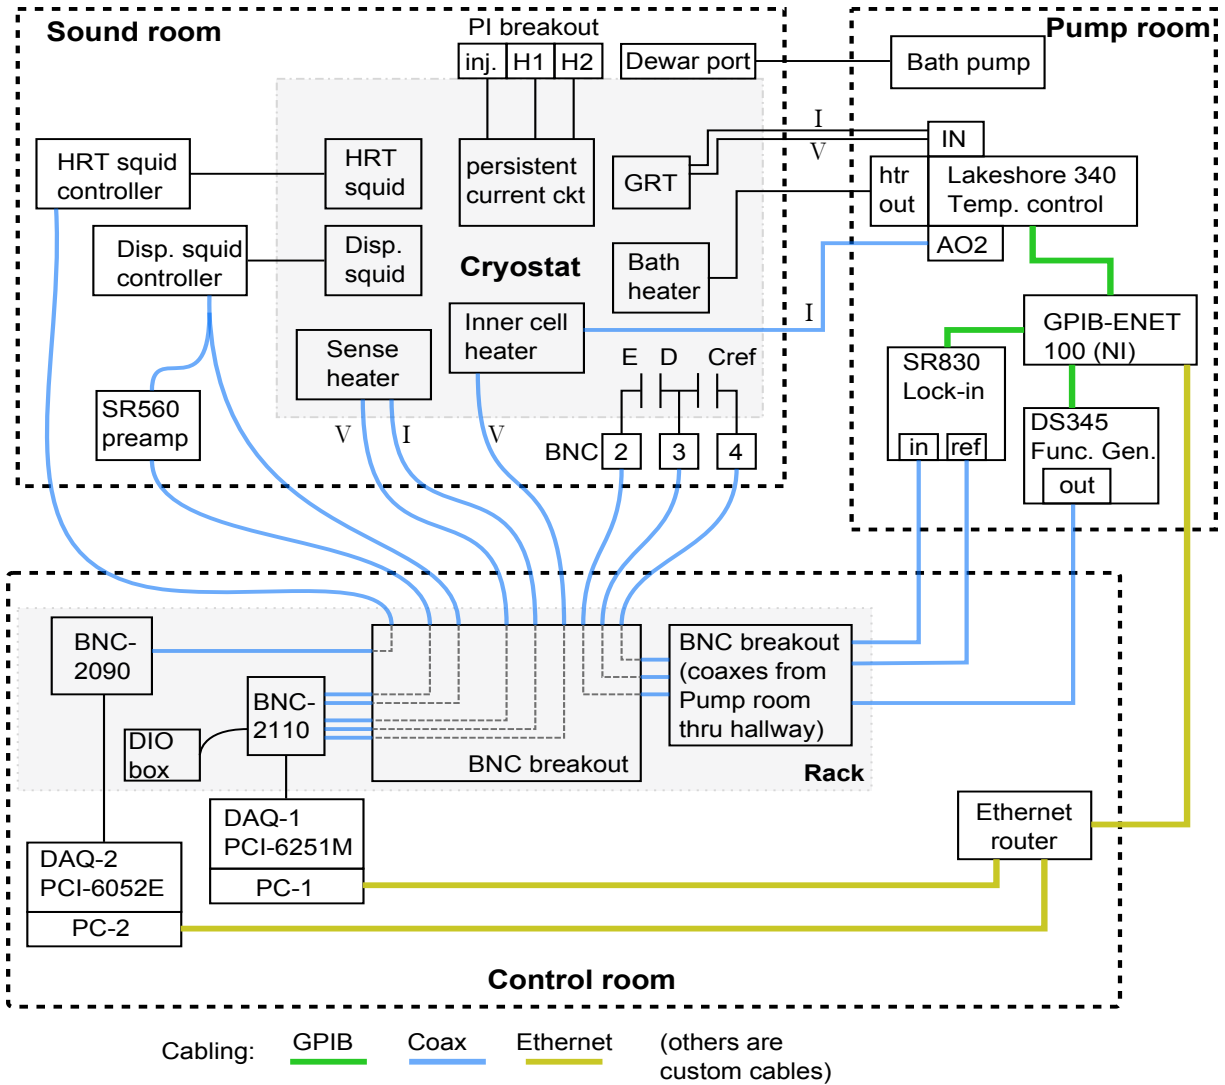


Figure 10.1: Measurement block diagram. This does not include the special automatic reorientation setup described in the previous chapter. The DAQ digital channels (from DAQ-1) used in that chapter are all accessed via the recently built breakout box labeled “DIO box”, which is connected to the BNC-2110 accessory by a ribbon cable. The configuration shown here for the capacitance leads (2,3,4) is for the capacitance bridge setup. For imposing a voltage bias across D-E, we ground the electrode (2) and bias the diaphragm (3) using a bipolar amplifier (BOP - not shown) whose input from DAQ-1 is controlled by a VI.

are connected to a single GPIB-ENET 100 hub, which is itself connected to an Ethernet router so that any computer can access these instruments over the lab network. All of this is located in the pump room to keep instrument fan noise away from the experiment. Two separate computers (PC-1 and PC-2) in the control room are both connected to this lab network and they can access any GPIB instrument as well as talk to each other (usually through Labview [82] global variables).

All analog signals are transported to and from the sound room and the control room (and between the control room and pump room) via low noise coax cables. Any multi-wire cables are always individually shielded, twisted pairs to minimize noise pickup.

10.1.1 Thermometry

As described in Chapter 7, our primary thermometer (for absolute temperature) is a Germanium Resistance Thermometer (GRT) whose resistance changes with temperature, with an additional secondary High Resolution Thermometer (HRT) whose magnetic susceptibility (read by a commercial dc SQUID) changes with temperature.

The GRT resistance is monitored by a commercial AC resistance bridge (the Lakeshore 340 temperature controller). The Lakeshore is capable of temperature regulation using PID³ feedback, but we do not use it for that purpose. This is because we can use the GRT only for crude temperature regulation with a stability of $\sim 20\mu K$ (limited by the resolution of the GRT). We need a customized two-stage feedback control⁴ that can switch between the GRT and HRT as needed for regulating the temperature. This temperature controller is implemented as a Labview VI (virtual instrument) running on a dedicated personal computer (“PC-2” in the figure). The temperature controller VI is henceforth simply called the TC-VI for brevity.

The GRT resistance measured by the Lakeshore 340 is digitized and sent to PC-2 over the network, where the TC-VI converts it to absolute temperature using an empirical calibration (Section 10.5). PC-2 also has within it a data acquisition card⁵ with a BNC breakout box (model BNC-2090) – labeled DAQ-2 in the figure – to read and digitize the output of the HRT-SQUID voltage (changes in this voltage are proportional to temperature changes - this conversion also done by the TC-VI using empirical calibrations).

During temperature regulation, changes in the bath temperature are detected and a PID feedback performed by the TC-VI to nullify them. The feedback output value is sent back to the Lakeshore 340 (as a manual output), which puts out heater power via a (manganin wire) bath heater installed on the cryostat (BNC-5,6 in the figure). Clearly, the bath must always be on a cooling trend for this to work (controlled by varying the rate at which the bath is being pumped on). Multiple levels of PID tuned parameters are available for crude and fine GRT and HRT feedback stages. Ultimately, we can stabilize the bath to better than $\sim 20nK$ when things are properly tuned. The TC-VI is computationally intensive as it

³Proportional-Integral-Differential

⁴developed by Emile Hoskinson

⁵National Instruments E-series PCI-6052 DAQ with 16-bit input resolution

has to work in real time. This is the main reason for splitting this function off from the rest of the experiment. Note that the TC-VI is capable of several different inputs and outputs (beyond just the Lakeshore) and being a VI, can easily be expanded for other custom setups.

Most recently, we have used one of the two (additional) digitally controlled outputs on the Lakeshore 340 for driving the inner cell heater (to generate stable battery states) as these outputs are quieter and more stable and are independent of the circuitry used to take whistle data (discussed next). The biggest issue there is the communication delays in the GPIB interface. Any additional communication with the Lakeshore can significantly interfere with the temperature regulation. This brings us to the second main reason why the regulation function is split off from the rest of the experiment and provided its own dedicated computer - Labview (at least the versions we have seen thus far) can only have one GPIB route running at a time. So, if everything was on PC-1, anytime we wished to use the DS345 function generator or SR830 lock-in analyzer, it would mess with the Lakeshore's operation. It is possible that there are workarounds to this, but all around, keeping things separated like this seems to be an optimal scenario.

Finally, important data (such as the bath temperature and any Lakeshore output values) are available to the networked PCs via Labview global variables for use during data collection and continuous analysis. It is also possible for other networked PCs to set the temperature setpoint remotely through these global variables and wait for stability before taking data automatically, and so on. This setup allows for a remarkable degree of automation, which is necessary given the limited window between bath helium transfers for actual data collection.

10.1.2 Cell

The main controller for the experiment is a separate computer (PC-1), which can run various other Labview VIs designed for specific experiments. PC-1 has within it a newer data acquisition card⁶ with a BNC breakout (model BNC-2110) – labeled DAQ-1 in the figure. As stated in the previous section, PC-1 can both read and set temperature setpoints from the PC-2 TC-VI, as well as control GPIB instruments over the network.

Given this background, we can consider the various input and output interfaces for the cryostat. The main output from the cryostat is the displacement sensor SQUID voltage (proportional to the diaphragm displacement in the cell). This is recorded by DAQ-1 in both its raw form (for DC values) as well as its AC-coupled and amplified (by the SR 560 preamp in the sound room) form (for spectral analysis of displacement oscillations). The cryostat SQUID sensor⁷ connects to an external SQUID controller, which outputs the final voltage.

There are two other heaters (besides the bath heater): the inner cell heater R_{in} and the sense arm (heat-pipe) heater R_{sense} . Each of these has a 4-wire (2 twisted pairs) cable interface, where one twisted pair is used for running current through the heater and the other

⁶National Instruments M-series PCI-6251 DAQ with 16-bit input resolution

⁷Quantum Design model 550 SQUID with model 50 controller

pair used for monitoring the voltage across it for more accurate measurement and control of the heater power. R_{sense} is run from DAQ-1, while R_{in} can be run either from DAQ-1 or from the Lakeshore 340 (as discussed in the previous section). The former is more appropriate for finer real-time control of a changing inner cell power (such as during transient experiments or for the kind of ramp rate analysis seen in Section 2.4.4). The latter is ideally suited for long-term stable battery states that only require minor, infrequent adjustments in cell power.

Finally, we have the cell capacitance C_{DE} defined by the fixed electrode (E) and flexible diaphragm (D). A mica capacitor (used as a reference capacitor C_{ref}) of similar size as C_{DE} is placed in a shielded box on the cryostat and one end connected to the diaphragm lead coming out of the cell – this joint (used for the capacitance bridge) is BNC-3 in the figure. The other end of C_{ref} is BNC-4 and the electrode lead is BNC-2 (BNCs 0 and 1 are extra, unused leads). These can either be hooked up to a capacitance bridge for monitoring C_{DE} (used as a diagnostic during cell-filling and for calibrations) or to DAQ-1 for putting a electrostatic force (DC steps or AC drives) on the diaphragm for various experiments. AC drives are also applied at times using a function generator (DS 345) in the pump room.

At this point, the various connection possibilities for the various instruments and the cryostat can be deduced from the layout figure.

10.1.3 Other

A superconducting liquid helium level meter module⁸ installed on the cryostat connects directly to its controller in the sound room and is used only rarely as it severely disrupts temperature regulation (a capacitive level meter would be a much better choice given the kind of thermal stability needed here).

A completely separate and very well-shielded cable and connector interface to the persistent current circuit used for the displacement sensor (see Section 5). This connector on the cryostat is kept inside a metal enclosure during normal operation to lessen electrical noise injection down into the displacement sensor. Only when the persisted current needs changing do we connect power supplies to the current injection leads and the two heaters. In the past, researchers have used mechanical switches to completely decouple the cryogenic circuit from the room temperature leads and connector. However, we have found that with judicious shielding and filtering, that added complexity is unnecessary.

Finally, analog and digital pressure gauges monitor cell and cryovalve pressures, while a higher resolution solid state pressure gauge⁹ can be used to monitor bath pressure while calibrating the GRT against the 4He vapor pressure curve.

All wires are shielded (as described in Section 7.4) and the cryostat is shielded from DC magnetic fields by 2 mu-metal shells (see Section 9.2).

⁸Lakeshore model 241 level meter

⁹Sensym/Honeywell model ASCX15AN

10.1.4 Chronology

A brief chronology of a typical experiment (with details in subsequent sections):

- Cell internal assembly, plumbing assembly with only capacitor leads hooked up
- Evacuate and leak-check cell and cryovalve
- SQUID (and others) final wiring, full wiring check
- Insert cryostat into Dewar, repeat wiring check
- Evacuate bath, leak-check cryostat top plate
- Cool to 77K with liquid nitrogen (LN2), wiring check - wait to equilibrate
- Cool to 4K with liquid helium (LHe), wiring check, check SQUIDS (displacement sensor and HRT)
- Go below T_λ to get quieter, measure vacuum resonance
- Calibrate GRT and HRT - can also be done later
- Fill cell
- Close cryovalve
- Test transients near T_λ
- Pump on fill line (usually about a day) through a cold trap to remove the helium above the cryovalve since it is a primed bomb during subsequent transfers.
- find approximate $T_\lambda(\text{cell})$
- Main experiments (and calibrations, if required)
- End of experiment - open cryovalve and empty out cell
- Warmup

10.2 Preliminary tests

10.2.1 Capacitance bridge

Our primary displacement sensor is the superconducting system described in Chapter 5. However, that only works below 4 K and its dynamic range is quite low so that it's better suited for low amplitude AC oscillations rather than large amplitude DC excursions of the flexible diaphragm. Having a second (independent) measure of the diaphragm position via

measuring the capacitance C_{DE} with a capacitance bridge comes in very handy for diagnostics at 300 K and 77 K and is useful in cell calibrations below 4 K. Our capacitance bridge is home-built and easily reconfigurable with a design that has been used for many years in our group. The electrode and diaphragm form C_{DE} (labeled C_x in general). C_{ref} is usually the reference capacitor installed on the cryostat. Note however, that this setup can be used to measure any capacitance with respect to any reference. The bridge input is a $\sim 1kHz$ reference oscillation from a lock-in analyzer (SR 830) passed through a 1:3 isolation transformer. A ratio transformer (split 1 : α where α is some fraction) is used to balance the bridge. The bridge output (sent back to the lock-in and nulled for balance) is the voltage between the $C_x - C_{ref}$ joint and the transformer tap (which is usually grounded). For the bridge setup shown in Fig. 10.18 (where the bridge output (3) is sent directly to the SR 830 lock-in), the unknown capacitance C_x is given by:

$$C_x = C_{ref} \frac{\alpha}{1 - \alpha} \quad (10.1)$$

Phasing and bridge balance

Since the output voltage (AC) is a two-vector, the bridge needs to be phased while balancing. This is done by changing the relative phase between the lock-in reference and the lock-in input (i.e. bridge output) so that the reactive and dissipative parts of the bridge output are separated as much as possible into the two orthogonal lock-in channels. Practically, this should be reflected in the dissipative channel not changing at all when the bridge is pushed off balance (by, e.g. the motion of the diaphragm or an off-balance ratio transformer setting). In practice of course, the dissipative channel will change a bit, but phasing merely involves minimizing the change in any *one* channel for a chosen extreme excursion in bridge stimulus (effective α change), which is also the chosen bridge dynamic range for which the bridge will behave linearly and predictably.

Henceforth, we will call the reactive channel X and the dissipative channel Y. At optimal phasing and balance, X will be within the electronic noise level and Y will be some constant (not necessarily small), where the change in Y will be much smaller than the change in X if the bridge is subsequently unbalanced. Note that the lock-in analyzer and the ratio transformer are in different rooms. Performing the actions in this and subsequent sections requires using a VI (provided by Stanford Research instruments) to control and view the front panel of the SR 830 on the computer in the control room.

Bridge calibration

Once the bridge is balanced and phased, the ratio transformer α is step-wise dialed between the endpoints of the chosen bridge dynamic range. The off-balance bridge output δX is recorded (so is δY , to ensure that we are phased correctly) vs. the ratio change $\delta\alpha$. It can be shown by an elementary analysis of the bridge circuit that $\delta X \propto \delta\alpha$. A linear fit ($V_x = S_b \cdot (\alpha - \alpha_0) = S_b\alpha - S_b\alpha_0$) to this data gives us the bridge sensitivity $S_b = \text{slope}$ (to

changes in α) and the true alpha balance $\alpha_0 = -\text{intercept}/S_b$ (our balancing before was a first approximation).

The bridge α is then set to α_0 and the bridge has been calibrated. Now, if the diaphragm position changes, it changes $C_x = C_{DE}$ and throws the bridge off-balance so that $V_x \neq 0$. Note that if the bridge sensitivity S_b is negative (just a phasing choice of 0 or 180°), the off-balance voltage V_x is negative for positive changes in α .

The following argument is a bit subtle. Now, imagine that we had adjusted α to null the bridge again. By how much ($\delta\alpha$) would we have needed to adjust it? The answer lies in the calibration we just did. The change in bridge voltage with α is $dV_x/d\alpha = S_b$, so a change δV_x in the bridge voltage comes from an equivalent change $\delta\alpha = \delta V_x/S_b$ and this change can be nullified by *an equal and opposite change* in α of $\delta\alpha_{needed} = -\delta V_x/S_b$. After nulling, the new α is now $\alpha_{new} = \alpha_0 - \delta V_x/S_b$. The capacitance in this hypothetical nulling scenario can now be found as usual from Eq. (10.1).

The intuitive leap comes when we realize that the changed capacitance calculated in our hypothetical scenario is identical to that in the scenario where the bridge remains unbalanced with the bridge set to our original α_0 . The diaphragm capacitance is therefore given by:

$$C_x = C_{ref} \frac{\alpha_0 - \delta V_x/S_b}{1 - \alpha_0 + \delta V_x/S_b} \quad (10.2)$$

where the calibration is strictly valid only within our calibration dynamic range (for which the bridge was properly phased). In practice, once we get outside this range, we will start seeing larger and larger changes in the dissipative channel Y. The deviations from the bridge formula will therefore keep increasing until we hit some kind of artificial catastrophe in the equation and we see absurd values for the capacitance. We state these things here because it is important to know why that happens and how the dynamic range can be extended. The catastrophe is not a rare occurrence - it happens nearly every time we evacuate the cell and C_{DE} increases past calibration and shows exactly this behavior (of diverging to ∞ and coming back up from $-\infty$).

The Labview VI “The Ultimate strip chart” can be used to monitor the capacitance over time given calibration parameters, which can be obtained by running a calibration sequence using the “Bridgesense” VI¹⁰.

We note in passing that the reference capacitor on the cryostat is first measured against a standard reference (this is repeated for each new temperature that the bridge is used at) and then subsequently used as a reference to measure C_{DE} during experiments.

10.2.2 Cell evacuation

Since it is possible at this point that there are severe issues with the aperture arrays or the cell has leaks, it is pointless to waste time connecting anything beyond the bare minimum of cryostat umbilicals to the cell. So, we just solder the cryostat fill line to the cell (using a

¹⁰Here and in other places, we omit the version numbers for all these VIs.

tiny amount of blue liquid flux and being careful not to contaminate the cell) with the top plate fill valve open so that all gas volumes are at atmospheric pressure. We further connect the cell wires from the diaphragm and electrode to the leads terminating in the 3 BNCs for the capacitance bridge as described previously. Now, we can monitor the inner cell pressure by its effect on the diaphragm position. The capacitance signal will be very noisy at this point because of pressure and temperature fluctuations and the fact that the diaphragm is floppy at 300 K and is very susceptible to static cling. Still, it should be good enough to monitor gross changes in cell pressure.

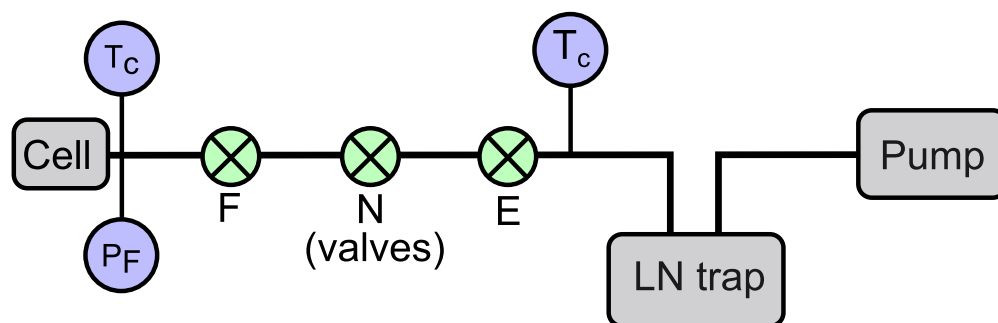


Figure 10.2: Plumbing setup for cell evacuation. F is the top plate fill valve, N is a Hoke metering valve that allows for very small flows and E is an extra diaphragm valve before the mechanical pump.

The cryovalve actuation line (bellows) should already be at high vacuum at this point (which it will be unless this is its first use) so that the valve is open and the fill line exposed to the cell. The cryovalve and its actuation line should be leak-checked first to ensure that any leaks don't make the valve actuate when we start pumping on the cell.

Referring to the plumbing setup in Fig. 10.2 (and its caption for valve labels), the setup is assembled with all valves open and the pump off. The lines are all evacuated with (F) and (N) closed (with the cell and the F-N space both at 1 atm). If (N) closes all the way, opening (F) at this point should only transiently perturb the capacitance C_{DE} . Now (N) is slowly opened, and the capacitance should rise (as the outer cell is evacuated before the inner cell – the diaphragm bulges out). (N) is opened to a point where C_{DE} remains within calibration range and is not too high. Depending on how sensitive (N) is, and how large the apertures are, one can throttle (N) so that C_{DE} is only slowly drifting up. This means that the pumping rate for the outer cell is approximately equal to the inner cell emptying rate through the holes. The pumping rate is kept low enough not to put too much differential pressure across the aperture array (it should easily withstand an atm but it's safer not to push this) and high enough that the cell is (rough) evacuated in a few hours. It is wise to use an inline filter¹¹ for the pumping line to prevent oil fumes from going in and contaminating

¹¹a commercial $2\mu\text{m}$ inline filter from Swagelok has worked well for us.

the apertures. The filters in the cryovalve ports should also help with this but we don't want to clog them either. Remember that the cryostat lines are typically difficult to replace and nearly impossible to clean.

Over this time, C_{DE} will eventually start coming back down as the cell gets emptied and the diaphragm relaxes. All this was simply the baseline process. In extreme cases, the flow conductance can be very small, so that the diaphragm shorts to the electrode the moment we start pumping. There is nothing to be done but to keep pumping at a reasonable (as defined by the experimenter) rate and watch for the diaphragm to unshort. Hooking up an ohmmeter and monitoring the resistance over time can help here. Once the cell is roughly pumped out (all this was with a simple mechanical pump), we can proceed to a leak-check.

10.2.3 Leak-checks and final prep

All cryostat plumbing is now thoroughly leak-checked, including the cryovalve chamber and all cell seals. After fixing any leaks, we finish the cell installation by cleaning and etching all superconducting wires in nitric acid and cleaning all superconducting joint pads with sandpaper and solvents (details in Chapter 5) and connect up the displacement sensor leads. The cell heater, sense arm heater and bath heater are likewise wired up and secured in place. All loose wires and components are tied down or taped to reduce microphonics. Radiation baffles and other important surfaces should be cleaned prior to cooldown.

After this, the cell should now be pumped out through a nitrogen trap using a diffusion pump (or other high vacuum pump) for several hours (typically overnight). For this, the metering valve is removed and the cell is pumped full on. We can do this because there is now negligible differential pressure across the aperture array and we can pump with impunity. We cannot over-stress the importance of the nitrogen trap and any inline filters mentioned before while pumping on the cell¹².

10.2.4 Electrical checks

All electrical components are checked after final prep. This includes checking for shorts of all leads (especially (D) and (E)) to the cryostat body, resistances of all heaters and the persistent current circuit leads (whose 300 K resistances should be known to a good degree in advance for diagnostics). All 4-wire resistances (this includes the two cell heaters, the GRT and the level meter) should be known in detail (i.e. pairwise resistances) so that any breaks or internal shorts immediately stand out during electrical checks¹³.

All wiring checks should be repeated after the cryostat is inserted into the Dewar to ensure that bumps and movement did not hurt anything. The relative fragility of the aperture arrays must *always* be kept in mind while moving the cryostat.

¹²We keep two separate sets of plumbing lines, with one set *specifically* set aside only for clean applications far from a pump - these are the lines that we use for cell filling, etc.

¹³We saved a considerable amount of time this way by locating an internally broken wire on a GRT purely through these pairwise checks.

10.3 Cooldown

We describe briefly a typical cooldown since there are special precautions and techniques needed for these kinds of cells and components.

10.3.1 To 77 K

Once the cryostat is sealed in place and the cell is at high vacuum (typically a few tens of mTorr with no noticeable outgassing after an overnight pumping), the bath space is evacuated. Since it is a large space that cannot be baked out, it must be purged with nitrogen a few times between evacuations to get to a reasonable vacuum (a hundred mTorr or so is sufficient as long as most of it is nitrogen). The cryostat top plate is leak-checked if desired but this is very difficult without hours of pumping and purging (there are innumerable air pockets in the cryostat and cell). In most cases, the leak background is very high and we resign ourselves to checking for large deviations when spraying helium at the obvious top plate breakouts and seals. As long as the Dewar can hold a rough vacuum for a few hours with only a very slow rise in pressure (due to outgassing), it is likely to be just fine (based on our experience). The Dewar jacket should also be evacuated to a high vacuum (purging helps, especially for a jacket with Mylar super-insulation).

Now, the bath is back-filled with N₂ gas and the liquid nitrogen (LN₂) transfer tube is inserted all the way down with a continuous N₂ stream going into the bath to keep air out. With the GRT temperature and C_{DE} being monitored, the bath is filled with LN₂ (over about an hour for a 20 L Dewar to prevent thermal shocks to components). The boiloff is huge so a large bath port is opened to atmosphere to prevent overpressuring the bath space.

Once the LN₂ fill is done, the top plate is thawed carefully with a heat gun, the LN₂ tube plugged up and the bath venting slowly to atmosphere through a one way (check) valve (at around a psi in the bath).

10.3.2 77 K tests

Electrical checks should be repeated here and 77 K values recorded for future use. The capacitance bridge will need to be re-calibrated due to the reduction in the (resistive) impedance of the lines. Most importantly, C_{DE} should now be much more stable and quiet since the diaphragm is stretched taut and there is no gas to randomly jostle it.

At this point (especially after a few hours at 77 K), the cell pressure should plateau at less than 10 mTorr and never rise after this (as any outgassing is suppressed and gaseous impurities freeze on tubing walls).

10.3.3 To 4 K

The cryostat is kept at 77 K for at least a few hours to let it thermalize. However, we prefer to time things so that we can leave it at 77 K overnight. To prep for the first liquid helium

(LHe) transfer, we start by pressurizing the bath space (from a top plate port at about 5-10 psig) with He gas with the LN2 tube almost all the way down. This forces the LN2 up the tube and it is collected outside the Dewar. The LN2 tube is lowered completely to get the last dregs out and we wait a few minutes after this to allow any liquid to evaporate.

Now, the LN2 tube is removed and the bath space evacuated and back-filled with He. LHe is transferred slowly with the transfer tube as low as possible so as to use the enthalpy of the escaping gas to the fullest for cooling the cryostat. Like LN2, the first LHe transfer is done over an hour or so at a reasonably low filling pressure. While blowing out the LN2, we can perform another leak-check on the cell and cryovalve lines to see if any cold leaks have opened up. If one has an LN2 level meter inside the Dewar, this test can help find the location of the leak (if any exists), which helps considerably given the insidious nature of cold leaks.

In contrast to the first LHe transfer, any subsequent transfers are performed quickly, typically over ~ 15 min using a drive pressure of about 3 – 4 psig in the storage Dewar. Further, the transfer tube is held higher in the bath so as to prevent bubbled warm LHe in through the bottom where the bulk of the cryostat thermal mass resides (and which is typically colder than the helium higher up). Pumping the system down to low temperatures prior to the transfer and immediately resuming pumping after the transfer keeps the bulk of the cryostat mass from warming up too much (which in turn increases the hold time as less of the helium needs to be used up to re-cool everything). Evaporated helium is recovered through a large diameter (at least 1/2") hose capable of withstanding cryogenic temperatures to prevent pressure build-up in the bath. When the bath is backfilled with helium gas at the beginning of the transfer (to bring it up to atmospheric pressure prior to inserting the transfer line), we observe the fill and actuation line pressures closely. The fill line pressure should not increase by more than a 5 – 10 mTorr, while the actuation line pressure (due to the presence of the buffer) should be limited to a less than ~ 5 psig increase¹⁴. These pressures should of course go back down to their nominal values when the transfer is finished.

For a Dewar with a vacuum jacket and aluminized Mylar (super-insulation), we observe that the top plate can easily get very cold and frost up during the transfer. We keep it defrosted by heating it with a common heat gun, being careful not to heat any wiring or plumbing breakouts or other sensitive components (especially the SQUID breakouts). This is where a long neck-ring comes in handy as a convenient surface to heat to indirectly keep the top plate warm.

10.3.4 4 K tests

Immediately after the first transfer, we check both the displacement sensor and HRT SQUIDS to make sure they tune and behave properly. There should be some sensitivity in the PI-style sensor even without any injected current because of trapped vorticity [26, p. 251] so we can

¹⁴Obviously, this statement applies only to a full cell with cryovalve closed and fill line above the valve evacuated (see later sections on cell-filling).

even test the sensor by listening to the SQUID output on a pair of headphones and making noise at the Dewar. As always, electrical checks should be repeated.

10.3.5 Persistent current injection*

Note that this was not relevant in the case of the single weak-link cell run referenced in this chapter since we had used the magnet style displacement sensor for this run (which attains its fixed sensitivity by the time we reach helium temperatures (4 K)).

In any case, the proper way to inject (or change) the persistent current was explained in Chapter 5. At the end of this step, we assume that (one way or another), the displacement sensor is active and has some (useful) sensitivity to diaphragm displacements.

10.4 The vacuum resonance (simple and mass-loaded)

A function generator (DS345) was used to drive the diaphragm electrostatically. The drive frequency was swept over smaller and smaller ranges to narrow down the resonance once it was found. The raw displacement sensor SQUID signal was FFT'd and the peak amplitude at the drive frequency measured. The results are as shown in Fig. 10.3. The fit function is the response of a forced damped simple harmonic oscillator. Note that this form is only qualitatively correct but since the fit is so good, we can at least be confident that the *location* of the resonance as determined from the fit is correct (if not the other parameters related to the inertial, restoring and damping terms).

A better way to obtain the vacuum resonance is simply by exciting the diaphragm into free resonant vibrations with a sharp electrostatic kick (a square voltage pulse applied between the diaphragm and electrode). The time-series data when spectrally analyzed immediately yields the fundamental resonant frequency (and any higher modes). Checking the ratios of the amplitudes of these modes relative to the fundamental mode can verify whether the diaphragm is truly in the membrane regime since we have concrete predictions for these mode numbers.

The resonant frequencies of simple and mass-loaded diaphragms vibrating in the membrane regime are discussed in Section F.2. The lowest eigenvalues x_0 for Eqs. F.9 and F.10 (for physical parameters used in this run¹⁵) are 0.218223 and 2.40483 respectively. Note that there's nearly an order of magnitude difference between them (as also between the fundamental frequencies)! This is not a "minor correction".

The mass-loaded value was used in this run to calculate the effective spring constant from the observed (vacuum) resonant frequency (assuming the design value for the diaphragm and magnet radii) and it matched well with previously observed values of the spring constant (2000-3000 N/m), which would have been off by nearly two orders of magnitude if we had ignored this issue. This option has now been incorporated into the "CxvsVb" VI that does this part of the calibration.

¹⁵(Magnet radius: $a \sim 1/32''$ and Diaphragm radius: $b \sim 3/16''$)




Figure 10.3: Vacuum resonance

10.5 Thermometry calibrations

Refer to Section 7.3 for descriptions of the various thermometers used. We discuss calibrations of each of the thermometers in this section.

10.5.1 Vapor pressure: Sensym

The Sensym output is a voltage signal proportional to the pressure and to the input power supply voltage applied to run it. Therefore, it must be calibrated for a given supply voltage. We do this by exposing the bath space (or any other large volume) to both the Sensym and a mercury manometer and tabulating the Sensym output (V) against the absolute pressure (kPa) measured by the manometer. This gives us an absolute calibration that is subsequently used during the run to conveniently obtain the bath pressure as a voltage that can be digitized and saved by a data acquisition card (DAQ). An example of such a calibration is shown in Fig. 10.4.

10.5.2 Primary: GRT

The bath is pumped down to low temperatures ($\sim 1.5 \text{ K}$) and the pumping valve throttled until the temperature is approximately stable. The GRT resistance R_{GRT} (measured by the




Figure 10.4: Calibrating the Sensym pressure gauge using a mercury manometer. Data from SHeQUID#2 [2/18/2011].

Lakeshore 340 or other AC bridge) and the absolute bath pressure (measured by the now calibrated Sensym) are now continuously measured as we allow the bath temperature to drift up slowly (by throttling the valve some more). The Sensym pressure data is converted to absolute temperature T using published [72] P vs. T data for ^4He . A suitable model ($T = [\sum_{n=0}^{10} a_n (\ln R)^n]^{-1}$) is fit to the R_{GRT} vs. T data and the fit model thereafter used as a calibration in the TC-VI. An example of such a calibration is shown in Fig. 10.5.

10.5.3 Secondary: HRT

The HRT SQUID voltage V_{HRT} and the bath temperature T (from the calibrated GRT) are measured using the TC-VI. An example plot of V_{HRT} vs. T is shown in Fig. 10.6. As the temperature rises, V_{HRT} rises with it¹⁶ till it hits the dynamic range limit of the SQUID, whereupon it resets to 0 and the process repeats. Continuous data between two consecutive resets is fit to a straight line and the slope (V/K) is the HRT sensitivity for a temperature

¹⁶The sign of V_{HRT} is purely arbitrary and depends on how the SQUID leads are wired up. We have a ± 1 multiplier in the TC-VI to always make the HRT temperature rise and fall with the GRT temperature.




Figure 10.5: Calibrating the GRT against vapor pressure data and standard curves. Data from SHeQUID#2 [2/18/2011].

T_{bin} , which is the average GRT temperature within that data bin. This sensitivity data is shown in Figs. 10.7 and 10.8 for an older HRT (using a paramagnetic salt as in Ref. [74]) and in Fig. 10.9 for the one presently in use (using a PdMn alloy as in Ref. [73]).

10.5.4 Bath temperature regulation

This HRT sensitivity data is (one way or another) programmed into a subVI that the TC-VI can use to obtain the sensitivity dV_{HRT}/dT for any given temperature. As we briefly discussed at the start of this chapter, the TC-VI is used to regulate the bath temperature by performing a PID feedback using a distributed manganin heater in the bath. Initially, the GRT temperature is used to calculate an “error signal” (away from the setpoint) and this quantity used in a feedback routine to calculate the heater power output needed to nullify the error¹⁷.

¹⁷clearly, the bath must be drifting down in temperature for this to work – we find for our setup that a drift rate of $\sim 10 \mu K/s$ is optimal so that only a few hundred mW of bath heater power is needed to maintain stability. Too much power means that we’re wasting helium while too little power can lead to instability and over-sensitivity to fluctuations. Also, for long-term stability, the lowering of the necessary bath power with




Figure 10.6: (Paramagnetic salt) HRT: SQUID raw voltage (V) vs. temperature (K). Data from (Single) Weak-link cell #2 [4/21/2010]. Each reset section is fit to a straight line and the slope is the HRT sensitivity in V/K (shown in subsequent figures).

Once the bath is roughly stabilized (within tens of μK) using the GRT, the TC-VI automatically switches to the HRT, using the last GRT temperature T_0 as a starting point. The sensitivity at temperature T_0 is obtained from the subVI. The new HRT temperature (in the next iteration of the feedback loop) is computed as $T_{HRT} = T_0 + \frac{\delta V_{HRT}}{dV_{HRT}/dT}$. This new temperature is now set as the starting point, the sensitivity obtained for this temperature and the process repeated to find the next temperature value. We are essentially “navigating inertially” around the GRT starting temperature. The absolute temperature will therefore be only as accurate as the GRT resolution allows, even though temperature *excursions* can be measured nearly around orders of magnitude more accurately! This approach only works within a small temperature range around the starting point as errors will build up, especially if there is drift in the HRT or the SQUID. However, if the feedback is quickly switched over to using the HRT as input, the temperature can now be quickly stabilized to within 20 nK.

Sometimes, acoustic noise spikes or temperature spikes can inject a fixed offset between the setpoint and the actual stable temperature because the spike confuses the initial value used for the “inertial navigation” described above. In such cases, there is an option in the TC-VI to simply reset the HRT temperature to the present value of the GRT temperature. After the feedback oscillations die down (over a few minutes), the temperature stabilizes to the correct value.

decreasing bath level should be taken into account or the feedback can run out of (lower) dynamic range.




Figure 10.7: (Paramagnetic salt) HRT sensitivity (V/K) vs. temperature (K). Data from (Single) Weak-link cell #2 [4/21/2010] for $T > 1.685\text{ K}$. The sensitivity decreases linearly for increasing temperature. This linear fit is programmed into a subVI and used to obtain the sensitivity at the present bath temperature, which is further used to calculate temperature deviations using SQUID voltage deviations during bath T-regulation.

10.6 Locating the lambda point

10.6.1 Bath

This value is obtained by observing the vapor pressure decrease as we cool the bath by pumping on it. When we reach $T_{\lambda, (bath)}$, the vapor pressure levels off for a short duration due to the specific heat singularity at the lambda point, after which it resumes its decreasing trend (see Fig. 10.10). We can roughly associate this plateau with the lambda point pressure P_{lambda} and read off the temperature from standard, published [72] vapor pressure curves for ^4He . In practice, as we saw in a previous section, the GRT is already calibrated against the standard vapor pressure curve, so the temperature can be directly read off the GRT monitor when we see it plateau. There is some ambiguity in this temperature because T_{λ} depends on the ambient pressure, which is the saturated vapor pressure only at the bath surface.




Figure 10.8: (Paramagnetic salt) HRT sensitivity (V/K) vs. temperature (K). Data from (Single) Weak-link cell #2 [4/21/2010] for $T < 1.685$ K. See caption for Fig. 10.7 – the data shown here is for the temperature range where the sensitivity starts changing non-linearly. A large degree polynomial is fit to it and programmed into the subVI mentioned in the previous figure.

Hydrostatic pressure increases as we go below the surface so that $T_{\lambda, bath}$ would decrease towards the bottom of the bath space (where the cell sits). Still, this value is not used for any calculations so we do not have to worry about this ambiguity.

10.6.2 Cell

$T_{\lambda, cell}$ is the important quantity that must be known for many cell calculations and calibrations. The easiest way to do this is in a working weak-link cell (or SHeQUID). As we will see in later sections of this chapter, putting a step voltage bias across D-E (or the inner cell heater) can excite transient quantum whistles. As we go above $T_{\lambda, cell}$, the whistles disappear and purely normal flow begins. Somewhere between these regimes must be a point at which the whistles just disappear. We can slowly go up stepwise in temperature and determine this point manually. We can usually determine $T_{\lambda, cell}$ this way to within 0.5 mK, which is




Figure 10.9: (PdMn) HRT sensitivity (V/K) vs. temperature (K). Data from SHeQUID#2 [2/18/2011]. This HRT is (at 2 K, for example) almost 50 times larger than the one with the paramagnetic salt from previous figures. This is the model we are presently using.

sufficient for any computational needs perceived thus far.

Another way of doing this (a bit more rigorously) is to measure the Helmholtz frequency f_h vs. temperature¹⁸, especially close to T_λ . Since this frequency decreases as $\sqrt{\rho_s}$ (see Eq. (F.12)), we can fit that $\sqrt{\rho_s}$ dependence (or the full model in the f_h equation) to the data and extrapolate to $f_h = 0$ to find $T_{\lambda, cell}$.

10.7 Filling and emptying the cell

10.7.1 Cell-filling

A note on cell-filling pressure

Since the cell is filled and closed at a pressure much greater than the vapor pressure (at ~ 1.5 K), the T_λ of the bath is higher than that of the cell (see Fig. 10.11). $T_{\lambda, (bath)}$ was observed to be ~ 2.177 K while pumping down (see Section 10.6, but note that the

¹⁸We will do this anyway in Section 10.8.7 for a different purpose, so it isn't a huge inconvenience.

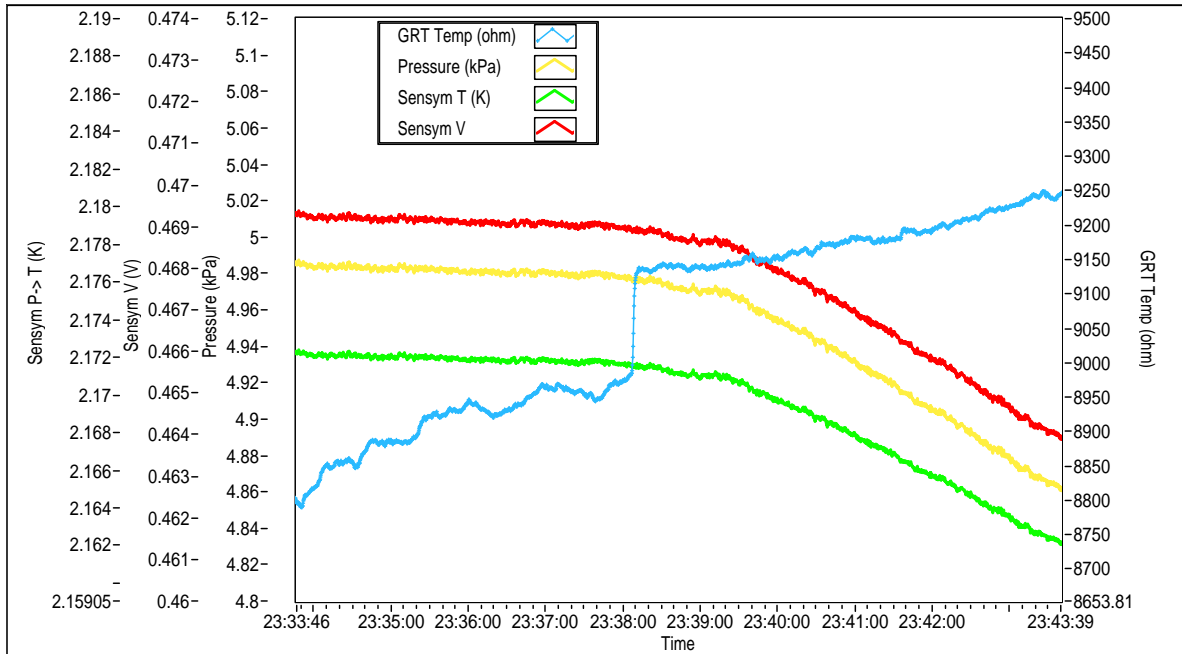


Figure 10.10: Locating T_λ of the bath. The raw data here is the Sensym voltage and GRT resistance. The Sensym calibration is used to convert the voltage to pressure and the published vapor pressure curve [72] used to convert this pressure to temperature. All 4 of these quantities are plotted vs. time. The plateau referred to in the text can be clearly seen.

determination there was for a different experiment). From the approximate observed value of $T_{\lambda, (cell)} \sim 2.166 K$, P_{cell} can be read off the lambda line as $\sim 111 kPa$ (a little more than an atmosphere). During filling, the pressure at the top plate was estimated as $\sim 24'' Hg = 80 kPa$. Right neighborhood, but clearly the real pressure is a larger. The cryovalve actuation is where the ambiguity comes in since during actuation, the fill line pressure is observed to rise by nearly $2/3 atm$! We see that controlling the cell pressure during filling and actuation can only be very approximate. We have observed that actuating the cryovalve very slowly (say over half an hour) keeps the fill line pressure from rising too much. Carefully studying the temperature of the bath (initially stabilized by throttling the pumping rate) can serve as a very sensitive indicator of how much helium is going in (both while filling and actuating). The warm gas going in also contributes to raising the pressure in the fill line so going slowly (letting the bath soak up the extra heat) seems to help.

We do not require a very accurate measure of the closing cell pressure at the moment (since there are much more accurate ways (see Section 10.6) to measure $T_{\lambda, (cell)}$, which instantly gives us the cell pressure via the lambda line). If ever needed, I suspect that a buffer volume for the fill line (like the one we use for the actuation line) would promptly remove this ambiguity, since the main problem is that tiny volume and temperature changes near the cryovalve cause large pressure changes. Artificially increasing the fill line volume at




Figure 10.11: Lambda line (from data in Ref. [83] converted to the ITS-90 scale)

the top plate should do the trick (and keep the fill line pressure from opening the cryovalve in case of too fast a temperature rise).

Calculations

There are signatures for when the cell gets full, but these have rapid onsets. It is wise therefore, to know how many moles of gas we need to fill the cell at a given pressure and carefully keep track of the gas used up so that we can at least know when to get ready for stopping the filling process. As we will see in the next section, we first fill a large keg with helium that has been cleaned through a nitrogen trap (with activated charcoal). The pressure of this keg can be read continuously. We then take this clean helium and (after passing it through yet more traps) fill it in the cell through a system of capillaries and metal hoses. Our final goal is to fill the cell until it is full of liquid helium, typically up to the bath helium level in the cryostat fill line, above which it is gas. We want some final pressure P_G in the system. The total volume of liquid V_L in the filling system can be easily estimated (volume up to the bath level), and the rest of the filling system volume is full of gas V_G . The liquid can be assumed to be at the bath temperature T_L and the gas (since most of the volume is outside the cryostat with only a tiny capillary volume inside, where it is still warm near the Dewar neck) is mostly at room temperature T_G . Given this final desired state, we can easily write down an estimate for the total number of moles of helium needed to make this happen:

$$n_{tot} = \frac{\rho V_L}{M_4} + \frac{P_G V_G}{RT_G} \quad (10.3)$$

where ρ is the total density of liquid helium at temperature T_L of the bath and M_4 is the helium molar mass.

Since we don't have a "mole-meter" in the lab, we can simply convert this to keg pressure at a constant temperature to see how much pressure we should lose from the keg. In this narrow context, the keg pressure can be taken as a direct measure of moles filled in the cell. Using the ideal gas law for the keg pressure:

$$P_{keg} = \frac{n_{tot} R T_{keg}}{V_{keg}} \quad (10.4)$$

There are some important caveats here. For the filling setup shown in Fig. 10.12 and typical SHeQUID volume and closing pressures used, this keg pressure can be on the order of 500 kPa (comes out to around 3-4 moles). Depending on how much control we would like during filling and the structural limits of the filling manifold, we can muster keg pressures of around 200 kPag¹⁹. The setup shown is engineered to be able to isolate the filling end from the keg end so that the keg can be refilled with clean helium before resuming the fill. We have found these estimates extremely robust (to within a small fraction of a mole) and it takes a lot of the guesswork out of filling the cell. The final caveat is that the keg needs more and more pressure relative to the cell, once the cell starts getting filled, so we must add another 1-2 bar to the P_{keg} estimate above to know how much keg pressure is needed in practice.

Procedure and signatures

- Referring to Fig. 10.12 for the plumbing setup, the main precaution to be followed here is that the aperture arrays should not see any rapid changes in pressure. All change should be gradual and as slow as possible. Given this, the kind of pressures (on the order of a bar) we use are not particularly dangerous to aperture arrays of even average strength.
- The capacitance bridge is recalibrated and the capacitance C_{DE} monitored throughout the filling as a diagnostic tool. The SQUID displacement sensor should also be kept on to listen to what's going on in the cell using headphones. With all the pumps running, the vacuum resonance should be kicked up all the time. The moment some gas gets into the cell, this should disappear immediately – a good sign that things have begun. Remember that the bridge excitation will be ringing the diaphragm – this frequency should be kept far from the vacuum resonance to be able to see this signature. Note the equilibrium capacitance at the beginning – multiplying it by the dielectric constant of liquid helium should be what we see at the very end of cell-filling as the space between the diaphragm and electrode gets filled with liquid.
- The bath temperature must be as low as possible (1.6 K or so) to increase the superfluid fraction and thereby minimize any differential pressures across the aperture array.

¹⁹A suffix of -g to any pressure unit signifies gauge pressure (above atmosphere).

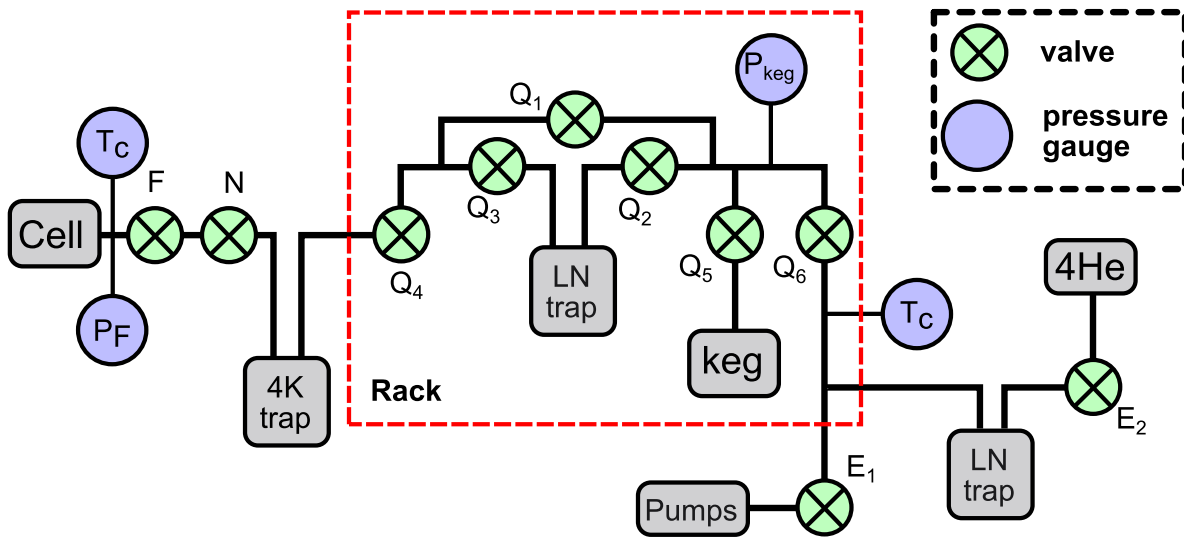


Figure 10.12: Plumbing setup for filling cell. Q_1 is a bypass valve used for cleaning out and leak-checking the system. The purpose of the other components should be self-explanatory.

- The entire apparatus is leak-checked and cleaned out (not every time it is used, but at least on first use and sporadically between run) by purging with nitrogen, leak-checking and then purging (and back-filling) with helium – it should typically be stored with a positive helium pressure so that air doesn't bleed in through leaks or diffusion.
- We clean the helium from a standard lab grade gas bottle (marked “ ^4He ” in the figure) through an activated charcoal nitrogen trap before storing it in a keg. The keg pressure can be monitored on a standard analog gauge (capable of several hundred kPag) and as previously mentioned, serves as a direct measure of the number of moles filled. There is another nitrogen trap in series with a 4K trap (which is just a long-double-walled tube submerged in an LHe storage dewar such that the entering helium is forced to travel all the way down before exiting the trap – this freezes out any residual impurities on the inner wall of the trap). This duo of traps is used to clean the (already quite clean) keg helium before sending it on to the cell. These precautions are followed to avoid closing any of the apertures via impurities freezing in them.
- Once the fill line is hooked up to the filling system, the entire setup is pumped out with a mechanical pump and then diffusion pump prior to cooling all the traps. The keg is filled through the LN trap and the gas bottle is thereafter isolated. Common sense precautions should be followed to avoid air intake (starting gas flow *before* hooking up lines, etc.) in the system. To begin the cell filling, valve (F) is slowly opened (into vacuum). A small shot of helium is let into the traps from the keg and the keg closed off. The metering valve (N) is now opened slowly (into a very low pressure clean helium environment) and the gas gets sucked into the cell fill line.

- This injection of gas into the outer cell initially bulges the diaphragm inward, reducing the capacitance (refer to Fig. 1.21). This should relax quickly as the gas enters the inner cell and the pressures equilibrate. Continue to put in such single shots until the capacitance goes down one last time and comes back and sticks to its equilibrium value. If things look good, one can even throttle all valves including the keg and start sending in a very slow but steady flow of gas. This capacitance behavior is a sign of the gas in the cell reaching the saturated vapor pressure at that temperature. After this point, we can speed up the gas input rate quite a bit (carefully) without fear of damaging the aperture arrays because the helium maintains the outer (and inner) cell pressures at saturated vapor pressure via condensation on the cell walls and superfluid film flow. Re-fill the keg as needed. Typically, we keep a pressure differential of around a bar between the cell and filling system by throttling valves to achieve a comfortable filling rate. The metering valve (N) rate is pretty consistent across runs so once a good valve setting is found, it can be quite reproducible in terms of filling rate.
- Occasionally, we check the situation by closing (F) momentarily – if the cell is not full, the pressure at the top plate will decay over time. A sign of a cell close to filling is a long decay time for this. Once the liquid level reaches the diaphragm-electrode in the cell, the SQUID usually starts going haywire and C_{DE} responds appropriately to the pressure differentials. Around this time, C_{DE} might start showing step-like behavior as different electrical parts get submerged in liquid (changing the dielectric constant) and it may also go off-scale for a while as the entire outer cell can get filled but the sense arm and inner cell is still relatively empty. This is ok, as long as it's not permanent.
- During filing, we can hear strange (highly reproducible) melodies on the SQUID sensor whose character and periodicity are constant over hours (essentially sounding like sharp, liquidy tinkles with xylophone flourishes at the end – we suspect them to be literally drops of condensed helium on the cell walls falling on the diaphragm or in the already collected liquid pool). Another, very useful signature is related to the bath temperature. If we maintain the bath temperature at a constant level by only throttling the bath pump rate, the entering warm gas heats up the bath and causes the temperature to rise slowly. The rate of temperature rise is a very sensitive indicator of the flow rate into the cell! Using all these diagnostic tools, we can make the cell filling extremely predictable and safe.
- Once we get close to our estimated keg (end) pressure, we begin performing the test with closing (F) and watching the pressure decays much more often. Once the cell pressure no longer decays (with enough drive pressure from the keg), we pressurize the cell to the desired value (see the lambda point discussion at the beginning of the cell-filling section) and close (F), stopping the filling process. The filling system is carefully secured by pumping it out (with the keg closed) and cleaning out the traps.

If desired, the system can be left pressurized with helium once the traps are warm. The keg should only be opened to the system with traps cold to maintain its cleanliness²⁰.

- At the end of this, once (F) is closed, C_{DE} should relax to a new equilibrium (enhanced by a factor of the LHe dielectric constant over its vacuum value). The SQUID sensor can be turned back on (it should be turned off during filling if it gets too noisy, especially if it goes into constant resets, because it is an active device and it's not worth stressing out its feedback system over junk signals). If everything works, tapping the Dewar should now produce very liquidy reverbs instead of the sharp twangs from the vacuum state). Essentially, listening to it should feel like listening to cymbals underwater.

10.7.2 Cryovalve closing

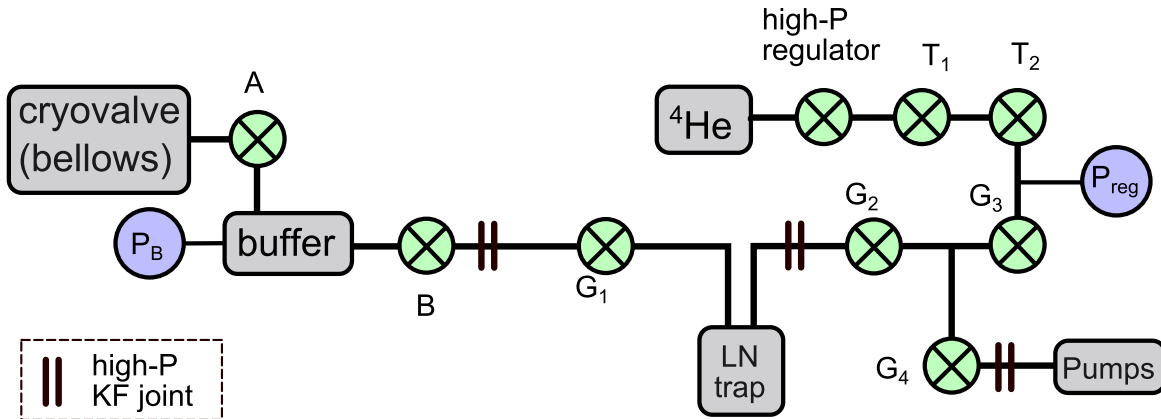


Figure 10.13: Plumbing setup for closing the cryovalve.

- While closing the cryovalve, the fill line should have a trap-protected line connected to a pump, in case the cell pressure rises too high.
- Referring to Fig. 10.13 for the plumbing setup, the plumbing is pumped out with a diffusion pump to the low mTorr level and the trap is cooled (as we did prior to cell-filling). These joints are all high-pressure (capable of withstanding at least a 100 psig, preferably more in case of an accidentally warmed trap) and all KF joints are made with high-pressure retainer rings. With no aperture arrays to protect, things are less dangerous in this case. We start off by bleeding small shots of helium into the buffer can through the trap with the actuation valve (A) closed and then exposing just the buffer volume to the actuation line with the helium source isolated. After a while, we can set up a steady flow directly from the helium bottle (still through the trap).

²⁰Being excessively paranoid about contamination is a good thing here, because in this case Nature really is out to get you.

- We need to watch the fill line pressure and make sure it doesn't go out of control (it will rise by several inches of mercury during actuation, which is ok). Once the valve has been actuated²¹ to sufficiently high pressures (typically 70-80 psig), we close off the buffer valve (B) but leave the actuation valve open to the buffer can, and then safely bleed and pump out the plumbing.
- At this point, the cell is full, the cryovalve is closed so that the cell is isolated from the fill line and the cryovalve bellows volume is full of helium and is connected to the large buffer can volume (all the time, so that if the bath temperature rises - during a transfer for instance - the resulting expansion of the bellows gas doesn't raise the actuation pressure by more than a few psi. This protects the valve and prevents it from leaking due to large swings in actuation pressure).
- At this point, the fill line should be pumped out to some extent and then some tests can be done to see if the whistle can be heard, etc. If one is planning to do some serious experiments past this point, we strongly advise taking a day to completely pump out and empty the fill line (above the cryovalve). Staying above T_λ helps considerably. One might even have an extra heater wire wrapped around the appropriate fill line length to speed things along (this has not been tried). If this is not done, one has to hook up a pump to the fill line each time one transfers and pump out any excess pressure. The bigger problem is that the upper fill line volume is a primed bomb and a sharp spike in pressure can slap the cryovalve momentarily open (it doesn't close quite right after the spike, causing the superleak)²² or even bounce the valve tip sideways very slightly without us finding out about it (except over a long time as the cell T_λ drifts). The biggest advantage in pumping out the line is that we have a thermocouple pressure gauge on the fill line that typically reads 1-2 mTorr (at 4 K) after pumping out the line. If there is a small super-leak, we should start seeing it on the gauge at some point. We have always found (over several runs) that there is no measurable change in this pressure over periods of up to 2 months of run-time.

10.7.3 Cryovalve opening and cell emptying (at run's end)

We won't discuss these aspects in great detail because the main precautions and ideas have been already conveyed. If the aperture arrays will be reused (typically doesn't happen), one should take just as much care opening the cryovalve as the initial pressure spike when the opens can be dangerous. To protect the aperture arrays, we can re-fill the fill line up to

²¹We actuate it to pressures typically 15-20 psi higher than the pressures needed to make it leak-tight in the cryovalve tests described in Chapter 8. This is to allow for super-leaks – a rule of thumb that has worked consistently for a long time now.

²²We revamped the entire procedure after the first (and only) time this happened, and never looked back. Luckily, the cryovalve suffered no permanent misalignment damage as a result. Chapter 8 should have convinced the astute reader that properly made cryovalves are worth their weight in gold (and a bit more when measured in blood, sweat and tears, but we digress).

the closing pressure before opening the cryovalve so that the chips don't see a huge pressure differential once the cryovalve opens. Aside from that, we open the valve in shots as usual – closing (A) and pumping out the buffer, closing (B) and bleeding the valve into the buffer, etc. The valve should open quickly (as signaled by the fill-line pressure beginning to rise up from ~ 0), but if it sticks (should not happen if properly made and used), we can pressurize the fill line side to push it open (the precise eventuality that we tried to prevent before).

Once a capacitance spike related to the valve opening is observed and the actuation pressure is low enough (so that the valve is now nearly fully open), we can start pumping on the fill line to empty the cell. Once both pressures are similarly low (takes several hours at bath temperatures above T_λ), we can continue pumping on both together through an LN trap and raise the bath temperature to help speed things along. Staying at 4K is fine too and the bath should have enough liquid helium in it to cover the cryovalve and cell for as long as there is still some liquid left in the cell²³. It can take up to a day or so for emptying completely.

There will be another capacitance spike as the sense arm empties out and the level falls below the diaphragm, accompanied by capacitance steps (for the same reasons as during the cell-filling).

Once the cell is empty (reasonably stable capacitance signatures as well as a fill line pressure that doesn't rise much when the pumping is halted), one can burn off the remaining bath helium fairly quickly using the bath heater²⁴ and/or warming and cooling back past T_λ a few times (we use up a lot of enthalpy in passing through the specific heat singularity at the lambda point) and then soften the vacuum in the Dewar jacket with a small infusion of N2 once we are above 77 K (this will obviously need to be evacuated before the next cooldown – labeling things really helps in a lab!). This enables fast turnaround times but we do not recommend the jacket-softening unless one is in a blinding hurry to cool back down afterward (the time difference is about a week). We can replace the He in the bath with N2 and let the cryostat warm in the Dewar, where water won't condense on it. This replacement with N2 should be done especially when things will be left as they are for a long time (order of weeks) as He can diffuse through the fiberglass (G10) neck of the Dewar and soften the vacuum jacket. Residual He is much more difficult to get out as compared to N2, because the latter would just freeze harmlessly in the jacket below 77 K.

²³We have not had the dubious pleasure of watching a cell blow up this way, but we imagine that it must be at least mildly fascinating.

²⁴We have also used a home-made dipstick with a high power resistor to enable faster heating but we do not recommend it as there are sensitive components on the cryostat. Around 5 W is probably as high as one should go, as a good long-term practice.

10.8 Cell calibrations

10.8.1 Need for calibrations

There are many different experiments that can be performed with the apparatus in this state. What calibration are needed will depend on what exactly one wants to measure. Refer to Table 1.1 for symbol definitions.

- At the simplest level, if one is interested merely in running the SHeQUID as a phase-meter (gyroscopy or monitoring other phase influences), there are literally no calibrations necessary past this point. This is because all we need measure in that case is the whistle amplitude in arbitrary units (say, raw SQUID voltage). Whether we are building interferograms or performing a flux-locked feedback (as in the main results of this dissertation – Chapter 11), nothing more is needed.
- For diagnostic purposes (such as checking the SEM imaging results of the aperture arrays to see how effective our fabrication process is), especially if a quantum whistle is not observed or is very faint, we might wish to have an independent measure of the average aperture size. This entails either performing the normal flow transient calibration to obtain the normal conductance β_n or the Helmholtz frequency series calibration to obtain the superflow conductance β_s and calculate the aperture size from either one or both.
- For detailed investigations into aperture arrays characteristics, such as phase-slip and Josephson dynamics, current-phase relations in different coupling regimes and some of the experimental tests of the chemical potential battery model (described in Section 2.4.4) – essentially, anything that delves into the detailed dynamics of the superfluid weak-links – we need a plethora of calibrations. Chiefly, we need to convert raw SQUID voltage into mass current in kg/s and diaphragm displacement into flow velocity (in the apertures) and into cell pressure. γ_1 (the pressure calibration) is the main conversion factor needed for this (see Eq. (G.10)). Also needed is the ratio A^2/k . Finally, the inner cell volume V_{cell} is needed in order to compute the inner cell heat capacity and the effective thermal boundary resistance of the inner cell R is needed to model heat loss from the cell. V_{cell} can be known from design values, but because of the use of epoxy to glue and seal things in place, this can be difficult to obtain, so it can be obtained as fit parameters to a model that involves the heat capacity.

There are several ways to obtain these calibrations under ideal circumstances and some options even under non-ideal conditions (when certain components are crippled and useless). We will discuss these possibilities in subsequent sections. Note that the dependencies of the various parameters on each other can be quite variable depending on what exactly we are trying to do (e.g. some applications can tolerate uncertainties arising from using design values while some cannot). We have attempted to order the following calibrations in a way

that respects these dependencies. However, it is best that the reader think of each section as separate and refer to the *Summary* section 10.9 to see the actual orders for various calibration sequences available to us.

10.8.2 Normal flow: β_n (using A^2/k)

Since $T_\lambda(\text{cell}) < T_\lambda(\text{bath})$, we stabilize the temperature between these 2 temperatures so that the cell helium is in the normal regime and take some flow transients using pressure steps (i.e. applying a step bias voltage, and hence a step electrostatic force between the electrode and diaphragm). We can easily derive the behavior of the fluid under these conditions by setting $\rho_s = 0$, $\rho = \rho_n$ in Eq. (1.23) to get $I_n = \rho_n \frac{\beta_n}{\eta} \Delta P$ and $I_t = I_n$ in Eq. (1.24) to get $I_t = \rho \frac{A^2}{k} \dot{\Delta P} = I_n$. Finally, we obtain a differential equation for the pressure:

$$\frac{A^2}{k} \dot{\Delta P} = -\frac{\beta_n}{\eta} \Delta P \quad (10.5)$$

For a pressure step excitation, this equation describes a diaphragm that relaxes exponentially to its new equilibrium position with a characteristic time constant τ_n described by

$$\beta_n = \frac{\eta_n A^2}{\tau_n k} \quad (10.6)$$


where β_n is the normal flow conductance and η_n is the viscosity of the normal fluid.

The time constant τ_n is obtained by an exponential fit to the data shown in Fig. 10.14 and Eq. 10.6 gives us $\beta_n = 1.73489 \cdot 10^{-19} \text{m}^3$. *This is an important parameter as it can be used to calculate the average hole diameter of the apertures.* Using Eq. (G.2)) with this β_n gives an average hole size of $\sim 83 \text{nm}$ in this run (for comparison, the size obtained from SEM viewing was $\sim 75 - 80 \text{nm}$).

10.8.3 Fountain calibration (series): V_{cell} , A^2/k (and R , using β_n)

The fountain calibration consists of a series of *subcritical* fountain transients taken over a range of temperatures from 1.5 to 1.9K (or even closer to T_λ , so long as the flow remains subcritical – i.e. doesn't start whistling) with a set of heater power steps $\{W_h\}$ for each temperature. W_h is chosen to be small enough so that the flow stays subcritical throughout the transient. This is to enable us to use the vastly simplified expressions that result since we can now take the chemical potential difference $\Delta\mu$ to be approximately 0 during the transient. Since we have already derived a model for subcritical flow in Section 2.4, we can immediately solve the ordinary differential in Eq. (2.3) for a heater power function that is suddenly stepped up to (a thereafter constant value of) W_h at $t = 0$ and obtain the solution for $\Delta P(t)$:

$$\Delta P(t) = \frac{W_h}{b} + \left(\Delta P_0 - \frac{W_h}{b} \right) e^{-t/\tau_f} \quad (10.7)$$

Figure 10.14: Normal flow ($T = 2.1658K$)

where $\tau_f \equiv a/b$, $a = \left(\frac{c_p V_{cell}}{sM_4} + sT\rho \frac{A^2}{k} \right)$ and $b = \left(\frac{1}{s\rho R} + sT \frac{\rho\beta_n}{\eta} \right)$ (a and b were defined back in Section 2.4 and we have used $C_p = \frac{c_p V_{cell}}{\rho M_4}$). Here, c_p is the molar heat capacity and M_4 is the molar mass of ^4He .

We can see from this solution that the heater step W_h causes a fountain pressure in the cell that drives the diaphragm to a new equilibrium position (corresponding to a new, final equilibrium pressure difference

$$\Delta P_f = \Delta P(t \rightarrow \infty) = W_h/b \quad (10.8)$$

between the inner and outer cells) with a characteristic time constant τ_f . These 2 parameters are measured by fitting an exponential form to the raw displacement sensor transient data.

Note that ΔP_f is known from the difference in SQUID voltages ΔV_f at the start and end of the transients *only if* the pressure calibration constant $\gamma_1 \equiv \Delta P_f/\Delta V_f$ is known (from the whistle calibration in Section. 10.8.4 or from the initial Cx/SQUID calibrations in Section. 10.8.5).

By fitting the modeled behavior of the time constant, or more precisely (merely for computational convenience), the quantity

$$\frac{\tau_f W_h}{\Delta P_f} = \frac{1}{s\rho} \left(\frac{c_p V_{cell} \rho}{M_4} + (s\rho)^2 T \frac{A^2}{k} \right) \quad (10.9)$$




Figure 10.15: Fountain series fit.

vs. the temperature T (see Fig. 10.15) for an entire series of such fountain transients, we can obtain the inner cell volume V_{cell} and A^2/k as fit parameters (where A is the movable diaphragm area and k is the diaphragm spring constant). Note that the fit model (the right hand side of the above equation is a highly non-trivial function of temperature since most of the helium parameters are temperature dependent. This can be handled in Matlab (and presumably in other computational software) by using a script file with either tabulated data or cubic splines for the helium parameters that are called by a function *file* (instead of an analytical fit function). A non-linear least-squares fit using this “function” does the job.

In extreme cases, one can even leave γ_1 in as a fit parameter and do a 3-parameter fit using the raw ΔV_f data (with a corresponding increase in fit uncertainty).

We will describe a much better calibration method for A^2/k and γ_1 (these are used to obtain the mass current directly from the SQUID voltage signal) in Section 10.9.

Thermal boundary resistance

Once we have β_n (from the normal flow transients), we can go back to the Fountain series data and calculate the thermal boundary resistance R from an expression obtained by manipulating Eq. (10.8):

$$R = \left(\frac{s\rho W_h}{\Delta P_f} - \frac{(s\rho)^2 T \beta_n}{\eta} \right)^{-1} \quad (10.10)$$

A convenient function (see Fig. 10.16) is fit to it so that R may be obtained easily

during further analysis. This quantity is needed for integrating the temperature model while calculating the current-phase relation or any analysis involving the chemical potential battery.




Figure 10.16: Thermal boundary resistance.

10.8.4 Whistle calibration: γ_1

A pressure step is electrostatically applied to the diaphragm at $t = 0$ near T_λ . The fluid oscillates at the Josephson frequency

$$f_J = \frac{\Delta\mu}{h} = \frac{m_4}{h} \left(\frac{\Delta P}{\rho} - s\Delta T \right) \quad (10.11)$$

At $t = 0$, $\Delta\mu$ only has the pressure term since it is the resultant superflow that changes the temperature in the cell, thus making ΔT non-zero. However, the whistle frequency is very difficult to measure right at $t = 0$ so we take a series of FFT windows along the time series, get the whistle frequency and extrapolate back through time using a parabolic fit (see Fig. 10.17) to get the whistle frequency $f_{J,0}$ at $t = 0$. The pressure difference at this time can be written as $\Delta P_0 = \gamma_1 \Delta V_0$ where ΔV_0 is the difference between the final equilibrium SQUID voltage and the initial SQUID voltage just after the initial compression step.

γ_1 is then found from a fit to $f_{J,0} = \gamma_1 \frac{m_A \Delta V_0}{h\rho}$ ideally for a series of different transients (with different steps)²⁵. Here, we just use one transient and merely use that expression to calculate γ_1 . Note that we are simply assuming here that the Josephson frequency relation is exactly true and using it to calibrate our cell. This gives a more robust value of γ_1 than that obtained from the Cx/SQUID calibration steps.




Figure 10.17: γ_1 calibration using whistle.

10.8.5 Cx vs. bias voltage: β_2

A series of voltages step $\{V_b\}$ are applied across the capacitor formed by the diaphragm and electrode and the capacitance change ΔC_x vs. V_b is measured once the diaphragm reaches its new equilibrium position. Referring to the circuit shown in Fig. 10.18, most of the components should be self-explanatory (filters to suppress noise, the capacitance bridge as already discussed and a high voltage supply (BOP) to impose the large forces needed to produce a measurable displacement and capacitance change). A preamp is used to amplify the off-balance bridge signal before sending it to the lock-in because it can be very small (tiny diaphragm displacements). The DC blocking capacitor protects the lock-in input from the large DC step change applied to D-E. Regardless, this initially overloads the lock-in, so we wait a few seconds before taking data from the lock-in. For a different purpose, we tested a variant of this circuit where the DC blocker was replaced by a computer-controlled relay circuit, which is initially open (while the step is being applied) and is closed a few ms after the step. This prevents the initial overload and improves the recovery time of both

²⁵For an example of a full calibration like this, see Ref. [44].

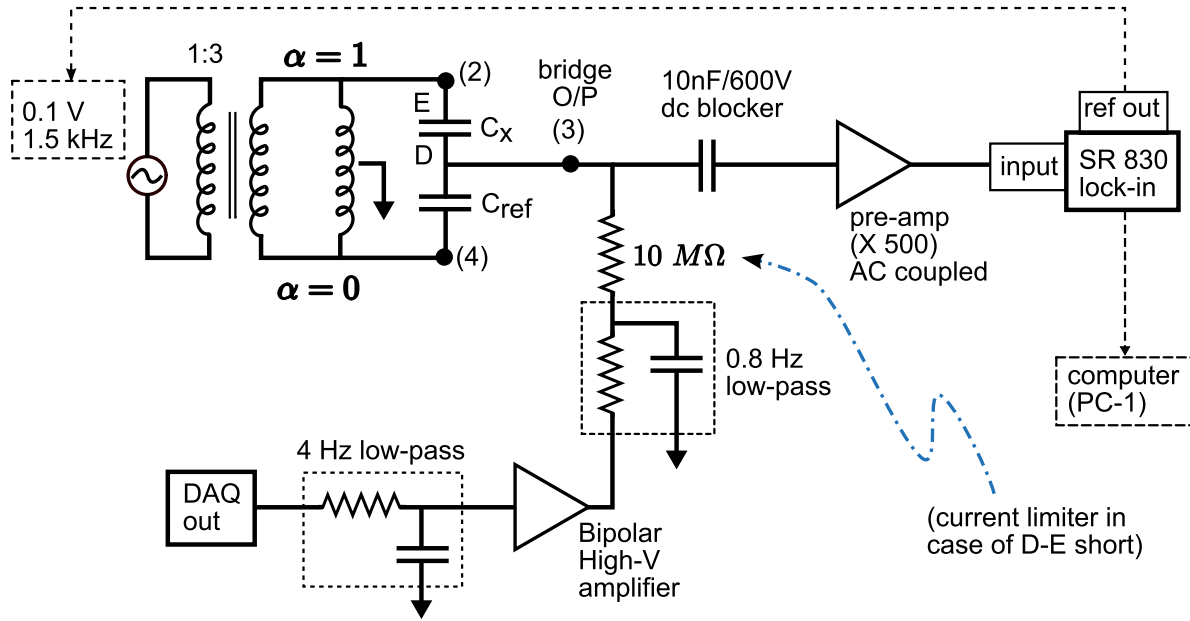


Figure 10.18: (D-E capacitance) C_x vs. V_b circuit setup.

amplifiers. However, that would be overkill for the present purpose and the circuit shown works just fine. Note that we can sample (and average) for a long time at each step.

A parabolic form²⁶

$$\Delta C_x(V_b) = \beta_2 V_b^2 - 2\beta_2 V_0 V_b + x_0 \sqrt{2k\beta_2} + \beta_2 V_0^2 \quad (10.12)$$

is fit to this data (Fig. 10.20), where V_0 and x_0 are the bias and spacing at equilibrium; and k and A are the diaphragm spring constant and area. This fit yields the parameter β_2 , whose theoretical value is: $\beta_2 = \frac{\epsilon^2 A^2}{2kd^4}$, where ϵ is the medium permittivity and d is the average (equilibrium) spacing between the electrode and diaphragm. Knowing k (from the vacuum resonance) and A (by design), we can find d from β_2 . Alternatively, the ratio A^2/k found from fountain calibration can be used directly.

Note however, that β_2 is merely the raw coefficient of the quadratic term in the fit and contains no other input. It is a purely empirical fit parameter.

10.8.6 SQUID vs. bias voltage: β_1

Similarly to the previous section, the displacement sensor SQUID response (ΔV_{sq}) is measured vs. a series of step bias voltages $\{V_b\}$, using the circuit shown in Fig. 10.21 and a parabolic form²⁷

$$\Delta V_{sq}(V_b) = -\beta_1 V_b^2 + 2\beta_1 V_b V_0 - \beta_1 V_0^2 - \alpha x_0 \quad (10.13)$$

²⁶See Ref. [44] or Ref. [46, p. 44] for a derivation.

²⁷Again, see Ref. [44] or Ref. [46, p. 44] for a derivation.




Figure 10.19: C_x vs. V_b raw data from bridge output. Note the fairly constant Y channel, showing a properly phased bridge.




Figure 10.20: C_x vs. V_b data and quadratic fit.

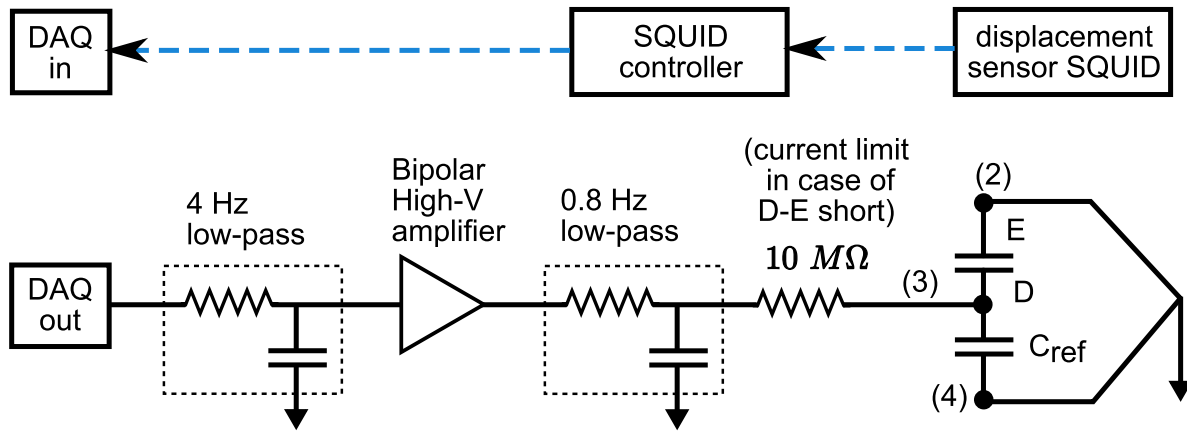


Figure 10.21: SQUID voltage vs. V_b circuit setup.

fit to the data (Fig. 10.22), where $\alpha \equiv \Delta V_{sq}/\Delta x$ is the SQUID sensitivity (V/m). This gives $\beta_1 = \frac{\epsilon\alpha A}{2kd^2}$ and using the values of A, k and d from the capacitance calibration, we obtain the SQUID displacement sensitivity α . With $\gamma_1 = \frac{\Delta P}{\Delta V_{sq}}$, we can easily show that $\gamma_1 = \frac{\epsilon}{2d^2\beta_1}$.

As in the case of β_2 of the preceding section, we carefully note that β_1 is merely the raw coefficient of the quadratic term in the fit and contains no other input. It is a purely empirical fit parameter.




Figure 10.22: SQUID voltage vs. V_{bias} data.

10.8.7 Helmholtz frequency (series): β_s

A series of pressure step transients (so-called Helmholtz transients) are taken over the $\sim 1.5K - 2.17K$ temperature range. When flow at the critical velocity ends (either phase-slips or Josephson oscillations), subcritical superfluid oscillations (with frequency $f_H = \omega_H/2\pi$) begin and eventually damp out with a characteristic time constant τ_H . This Helmholtz frequency is measured by spectral analysis of a section of the subcritical data. The decay time τ_H is measured by fitting a damped sinusoidal form to the actual time trace (the frequency may also be obtained from this fit). The quality factor Q of the oscillations is simply $Q = \omega_H\tau_H$.

The measured frequency vs. T is fit to the theoretical expression from Eq. (F.12) where the calibration constants V_{cell} and A^2/k are used from the fountain calibration calibrations. β_s is the superflow conductance for the aperture array and is the sole fit parameter for this fit (see Fig. 10.23). *This is an important parameter as it can be used to calculate the average hole diameter of the apertures using Eq. (G.22).* This gives $\sim 72nm$ for the average aperture size (compare this to nominal, SEM and normal flow values of 65, 75 – 80 and 83nm respectively).




Figure 10.23: Helmholtz series: frequency f_H .

Theoretical values of τ_H and Q are calculated using Eqs. (F.13) and (F.14) respectively and the result of the frequency fit above (β_s) is used as a “known parameter” in these calculations. The actual data for τ_H and Q (the latter calculated using experimental values for ω_H and τ_H) is also plotted for comparison (since there is no independent information to be obtained from further fitting).

Our results (including the discrepancies) are very similar to the ones obtained in Ref. [46, pp. 75-76], showing a systematic problem with the models used.




Figure 10.24: Helmholtz series: decay time τ_H .




Figure 10.25: Helmholtz series: quality factor Q .

The actual observed Q 's are smaller than the expected ones. Clearly, the hydrodynamic model used is missing some key dissipation term. It is interesting to note that as we approach

T_λ , this mysterious dissipation term becomes *less important* and the data approaches the prediction! It would be safe to say that normal fluid damping is probably not the missing ingredient (since it would show the exact opposite trend in that case). The tiny upward kink in Q very close to T_λ is most likely an artifact of the lessening accuracy of the frequency formula (which is derived for strong coupling).

10.9 Summary: calibration sequences

We have presented several methods in preceding sections for empirically obtaining various cell parameters. We summarize here the various possible calibration sequences that can be followed, including a new method that has several advantages over the others.

On a purely empirical level, the calibration procedures described above give us the following information (without using any other parameters found in this list - so this is a dependence-free list presented in no particular order):

- Normal flow: decay time τ_n
- Whistle calibration: γ_1
- Cx vs. Vbias: β_2
- SQUID vs. Vbias: β_1
- Helmholtz frequency series: β_s
- Vacuum resonance: k/A
- Fountain calibration: by itself, it gives us nothing but raw data of final SQUID voltages ΔV_f and fountain transient decay times τ_f vs. temperature T .

The relation used to convert a raw SQUID voltage signal [V] into mass current [kg/s] is $I_t = \rho \left(\frac{A^2}{k} \right) \gamma_1 \dot{V}_{squid}$ (from Eq. (1.24)). So, the mass current calibration constants are $\frac{A^2}{k}$ and γ_1 . Note that for a sinusoidal oscillation at frequency ω , $V_{squid} = V_0 \sin \omega t$ so that the mass current amplitude $I_{t,0}$ corresponding to a SQUID signal amplitude V_0 will be $I_{t,0} = \rho \left(\frac{A^2}{k} \right) \gamma_1 \omega V_0$.

In practice, we measure the SQUID voltage vs. time (we call a chunk of this data a “time-series”) and obtain its power spectral density numerically using a Fast Fourier Transform (see Figs. 10.26 and 10.27). The integrated SQUID voltage amplitude is converted to mass current amplitude as explained previously. The plots show a slightly different method, where the entire displacement PSD (V^2/Hz) is converted²⁸ to a PSD for velocity amplitude

²⁸We do this during data collection because it is the velocity spectrum that is more meaningful in these cells since it is the mass current, proportional to flow velocity that drives the diaphragm. It is therefore the velocity spectrum that we should be looking at when attempting to locate the whistle.

of the diaphragm $((V/s)^2/Hz)$ using the ideas on p. 345 (essentially involves a factor of ω^2). For such an integrated velocity peak $V_{vel,0} = \omega V_0$ (as opposed to displacement peak), the calibration will be (as expected): $I_{t,0} = \rho \left(\frac{A^2}{k} \right) \gamma_1 V_{vel,0}$.




Figure 10.26: Sample time-series data from SHeQUID 3, run 1 (with a stable battery state running)




Figure 10.27: Power spectral density (PSD) of Fig. 10.26 showing whistle peak and integrated power under peak

1. Use k/A from vacuum resonance and design value of A with β_2 from Cx vs. V_{bias} to get D-E spacing d . Use d , k , A and helium permittivity ϵ with β_1 from SQUID vs. V_{bias} to get γ_1 (and SQUID sensitivity α [V/m], if desired). This gives us γ_1 , A and

- k*. Note that since A and k were used to find γ_1 , there is no point in performing the fountain calibration (which uses γ_1) to obtain a more precise A^2/k (though it needs to be done to get V_{cell} and R vs. T).
2. Find γ_1 directly from whistle calibration. Use with fountain calibration to get A^2/k . This is one of the most precise methods. There are two caveats. In a SHeQUID, with cell resonances, whistle calibration is difficult and error-prone as the whistle sticks to resonances during a transient, making it hard to extrapolate accurately to $f_{J,t=0}$. Also, this needs to be repeated each time the SQUID current is changed (and it is more difficult, manually done and tedious than a simple SQUID vs. V_{bias} run, which is automated). The upside is that Cx vs. V_{bias} and vacuum resonance calibration (can be problematic for magnet-loaded membranes) need not be done.
 3. Our new method involves the following observation. We know that $\beta_1 = \frac{\epsilon_1}{2d^2} \gamma_1$ and $\beta_2 = 2 \left(\frac{\epsilon_2}{2d^2} \right)^2 \left(\frac{A^2}{k} \right)$. Define a new parameter $\beta_3 \equiv \frac{\beta_2}{(\beta_1)^2}$ and note that $\beta_3 = 2 \left(\frac{A^2}{k} \right) \gamma_1^2 \left(\frac{\epsilon_2}{\epsilon_1} \right)^2$. If Cx vs. V_{bias} and SQUID vs. V_{bias} are done at the same temperature, the permittivity terms cancel. However, even if they are not, this ratio (worst-case) can be estimated using the values of the helium permittivity at the extreme ends of our working temperature range: $\left(\frac{\epsilon[1.6K]}{\epsilon[2.17K]} \right)^2 \sim 0.998955$, which is essentially 1 within the errors of the calibrations. So, the main conclusion here is:

$$\beta_3 = \frac{\beta_2}{(\beta_1)^2} = 2 \left(\frac{A^2}{k} \right) \gamma_1^2 \quad (10.14)$$

This means that given the two V_{bias} calibrations β_1 and β_2 , γ_1 and A^2/k are no longer independent and they are related by a now *known* parameter β_3 (calculated solely from β_1 and β_2 , which are themselves purely empirical fit parameters to data with no real theoretical input).

Using $\Delta P_f = \gamma_1 \Delta V_f$ and the β_3 equation above, we can rewrite the fountain fit function from Eq. (10.9) as:

$$\frac{\tau_f W_h}{\Delta V_f} = \sqrt{\frac{\beta_3}{2}} \frac{1}{s\rho} \left[\frac{c_p V_{cell}}{M_4} + \left(\frac{A^2}{k} \right) (s\rho)^2 T \right] \left(\frac{A^2}{k} \right)^{-1/2} \quad (10.15)$$

The procedure is now clear. Starting with no information, do the Cx vs. V_{bias} and SQUID vs. V_{bias} calibrations (preferably at similar temperatures, but this is not necessary) and get β_1 and β_2 . Calculate $\beta_3 = \beta_2/(\beta_1)^2$. Do a fountain calibration series and fit the above model to the raw data (using ΔV_f and τ_f vs. T) to get fit parameters A^2/k and V_{cell} . Then, β_3 and A^2/k gives γ_1 from Eq. (10.14).

Note that these are cell parameters and only need to be found once (regardless of SQUID sensitivity). β_2 also needs to be found just once. Now, if the SQUID sensitivity

is changed (different current injected), one need only repeat SQUID vs. V_{bias} (easy and automated) to get a new β_1 . Then, $\gamma_1 = (1/\beta_1)\sqrt{\frac{\beta_2/2}{(A^2/k)}}$ for this new β_1 .

If needed, the other parameters mentioned can be easily found from this point:

- Normal flow τ_n and A^2/k : β_n
 - Fountain series data and β_n : R vs. T
 - A^2/k and β_2 : d
 - Vacuum resonance + model: k/A
 - k/A and γ_1 : α
 - A^2/k and k/A : A and k separately.
4. We end this list with a scenario we recently faced, wherein the deposited metal from the diaphragm had flaked off and shorted to the fixed electrode. This did not affect anything else in the run, but it did make it impossible to apply any bias voltages on the diaphragm. Since we only had the fountain series calibration available, we obtained k/A from the vacuum resonance, and with the design value of A (not too uncertain if glued carefully), obtained A^2/k . We fit the fountain data to the original model in Eq. (10.9) (with γ_1 and V_{cell} as unknown (fit) parameters. This lets us calibrate the current but β_n , $R(T)$ and β_s remained unobtainable, except by calculation using the nominal aperture sizes (not too bad).

10.10 Transient analysis

Procedures in this section are based on the seminal experiments in ^4He weak-link physics performed by Hoskinson, *et al.* [1, 2, 38]. We describe our re-enactment of some of these experiments to fit them into the theoretical and experimental framework built in this dissertation. The essence of these experiments is to impose a step voltage across the D-E capacitance to pull on the diaphragm (called a Helmholtz transient) or put a step voltage across the inner cell heater (called a Fountain transient). In either case, this imposes a momentary chemical potential difference (driving Josephson or phase-slip oscillations), which subsequently decays until a new equilibrium is reached. We can record both the DC value of the SQUID displacement sensor voltage as well as its amplified, AC-coupled form.

10.10.1 Whistle frequency vs. chemical potential

The chemical potential difference $\Delta\mu$ is calculated by numerically integrating the temperature equation (1.27) over time from the start of the transient to the point where subcritical

Helmholtz oscillations begin, in order to obtain $\Delta T(t)$. This step requires most of the calibration constants obtained thus far²⁹. ΔP is obtained directly from the raw displacement SQUID voltage over time using the pressure calibration $\Delta P = \gamma_1 \Delta V_{SQUID}$. A plot of the whistle frequency vs. $\Delta\mu/h$ (see Fig. 10.28) is linear to first approximation with a slope of ~ 1 , thus verifying the generalized Josephson frequency relation (Eq. 10.11).




Figure 10.28: Josephson frequency relation.

The deviation from linearity towards the end of the graph (left to right also shows increasing time) is most likely due to temperature drifts leading to errors building up in the integration to get $\Delta\mu$ (insufficient temperature stability). This is why we can't take the extra step to just go ahead and integrate $\Delta\mu$ to get the phase (via the Anderson phase-evolution equation $\frac{\partial\Delta\phi}{\partial t} = -\Delta\mu/\hbar$) and obtain the current-phase relation³⁰. The phase calculation if started from $t = 0$ turns into a hopeless mess by the time Josephson oscillations actually begin. Section. 10.10.2 describes what we do instead – a method developed by Simmonds [26, p. 280] for ^3He and refined for the more complicated calculations in ^4He by Hoskinson [39, p. 66].

²⁹This is simply a cruder form of the method used to obtain the current-phase relation. The references in Section 10.10.2 have further details on the integration of the temperature equation.

³⁰We tried.

10.10.2 Current-phase relation

The current is calculated from the SQUID signal during a transient using the calibration constants already found. The phase difference is calculated by integrating the Anderson equation. To avoid phase drift, we use the Helmholtz oscillation peaks as waypoints to keep the phase integral on track. Specifically, several Helmholtz peaks and valleys are located and used to chop the data into sections for integration. The temperature model (for ΔT) is integrated numerically starting from a Helmholtz peak, where the current and fluid velocities are 0 so the phase difference is also 0. This constraint on the phase is used to recalculate the initial ΔT at the beginning of each section and this prevents any errors from accumulating.

Two representative current-phase relations for the strong and weak coupling regimes are shown in Figs. 10.29 and 10.30 respectively. These were taken at $1.75mK$ and $0.85mK$ away from T_λ respectively. To be precise, the latter is in the cross-over regime between the strong and weak regimes and is beginning to look sinusoidal.




Figure 10.29: Strong coupling - linear $I(\phi)$ (trans 15, 3/13/10)

10.10.3 Critical currents

For a transient in the strongly coupled regime ($T_\lambda - T = 1.75mK$), the data is sectioned and the amplitude (I_1) of the first harmonic and the frequency of the whistle are extracted for each small section (chosen to be as small as possible while still providing sufficient cycles for the FFT). For a perfect sawtooth (synchronous phase-slip oscillations), $I_1 = I_{full}/\pi$. Also, the expected amplitude of phase-slip oscillations is just the slip size for N apertures




Figure 10.30: \sim Weak coupling - $I(\phi)$ starting to look sinusoidal (trans 23, 3/13/10)

$\Delta I_s = \kappa/L_l$ (where κ is the ^4He quantum of circulation and L_l the hydrodynamic inductance of the aperture array³¹). The quantity $\pi I_1/\Delta I_s$ should therefore equal unity if the phase-slip oscillations are synchronous. Figs. 10.31 and 10.32 show this to some extent with the added complication that there are two large (presumably) cell resonances that kick up the amplitude in their vicinity. The background level is clearly a flat ‘1’ though.

10.11 Frequency response of cell

To investigate the resonant features of the cell, we can set up a typical sweep run as shown in Fig. 10.33 (also see the main block diagram – Fig. 10.1 – at the start of this chapter). The fixed electrode should be grounded (here and anytime a bias voltage is applied between it and the diaphragm³²) and an AC drive applied to the diaphragm using a function generator. The response of the diaphragm (via the displacement sensor) is sent to the input of a lock-in analyzer whose reference signal is the SYNC (clock pulse train) of the function generator. The lock-in determines the amplitude of the displacement signal at the drive frequency for the first N harmonics of the drive frequency (N is usually chosen to be 2). The response is recorded for each drive frequency and a plot such as the one shown in Fig. 2.1 is obtained.

³¹I like to think of κ as the circulation flux, which instantly gives us $\kappa = \Delta I_s L_l$

³²The pancake coil is right next to the electrode, so grounding the electrode helps considerably lower the noise pickup by the pancake coil. We have tried it both ways and putting the bias on the electrode is just not optimal in terms of noise.




Figure 10.31: Whistle amplitude vs. frequency (strong coupling).




Figure 10.32: Whistle amplitude vs. frequency (strong coupling) - zoomed in.

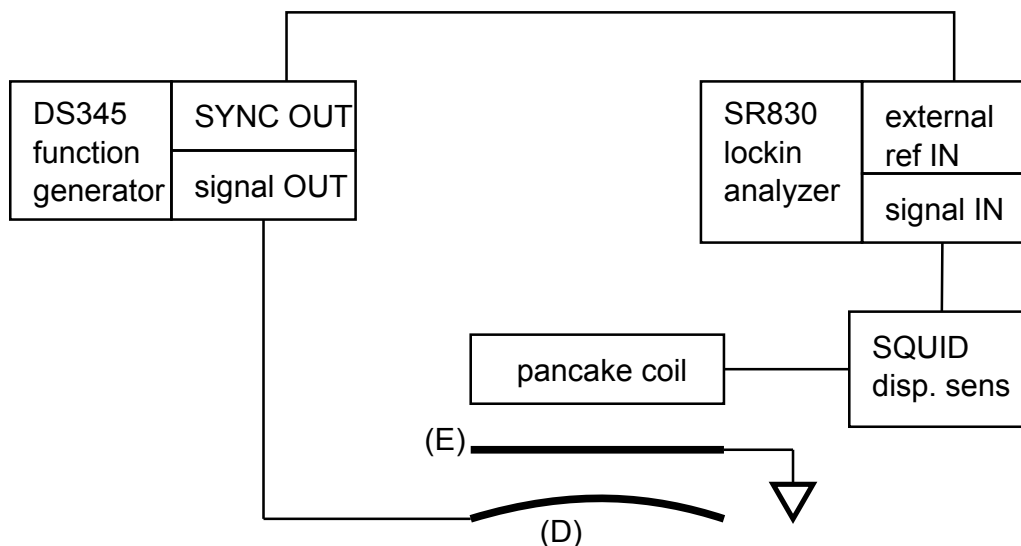


Figure 10.33: Setup to perform resonant frequency sweeps with capacitive drive.

We have set up an automated system to control both these instruments over a GPIB interface through Labview, step through a set of drive frequencies, wait for the cell to settle, average many measurements of the amplitude and save everything to file. We can even input a set of ranges with frequency steps to be completed in order. We have found it very useful to begin with a broad sweep (say, 0 - 6 kHz by 100 Hz) and then perform successive sweeps from 50 to 6050 by 100, 25 to 6025 by 100, etc. to fill in the gaps. This way, a more and detailed picture slowly builds up over time while still preserving a birds-eye view³³ (rather than just getting one small range with very high resolution).

10.12 Chemical potential battery

Battery states are obtained using an automated VI. The main data vi has several “co-running” VIs (that are not subvis). One of these is a cell heater ramper VI, which ramps the cell heater power between two values at a specified rate. This VI can use either the onboard DAQ-1 (for greater control over the ramps) or the Lakeshore extra analog output (AO2 as shown in the layout diagram Fig. 10.1) to inject current into the cell heater. The virtue of the co-running VI is that these ramps can be performed independently of the data-taking VI so that we can closely observe the behavior of the whistle during the ramp. The data shown in Section 2.4.4 was obtained using such a technique. A similar VI performs the same function for the sense arm heater. A practical issue here is communicating with the Lakeshore 340 without disturbing the bath temperature regulation. We get around this

³³Eliminates some of the nail-biting suspense involved in these slow sweeps.

(severe) problem by sending the necessary GPIB messages about demand values of AO2 to the TC-VI via global variables over the LAN. These messages are concatenated to the fairly regular messages between the TC-VI and the Lakeshore as this way, we don't add any time lag to this communication.

The DAQ-1 option is more noisy and is best suited for scenarios where the ramp timing is more important (during investigations of the battery for instance) while the Lakeshore option is ideally suited for interferometry where all we want is to obtain a very stable battery state and the ramp timing is not crucial (beyond ensuring a slow enough and fairly gradual ramp). Of course, one can easily get the best of both worlds by using a more sophisticated instrument as the current source - a low noise, high resolution current source that can be controlled via an ascii-programmable interface. This last feature is important because merely using an amplifier will just amplify the input noise of whatever voltage source (typically a DAQ-output) one uses to program it. An ascii-programmable interface (such as GPIB) isolates the programming from the output. Such sources are readily available (as of this writing). Using the Lakeshore in this way is merely a patchwork solution that nonetheless works quite well within the stated constraints.

No matter how the ramp is done, we have provided some practical advice for obtaining battery states quickly and reliably in Section 2.4.4.

The raw displacement sensor signal along with its AC-coupled and amplified version are recorded. The power spectrum ($|\mathcal{F}(V_{SQ})|^2$) of the amplified signal is scanned for the whistle peak. With multiple peaks present, the whistle peak can be identified as the moving peak and is typically most easily distinguishable if one looks at a live spectrum while listening to the amplified signal on headphones. The spectral power under the peak is integrated and the square root of this result is the diaphragm displacement amplitude (in volts). This can be converted, if desired, into a current amplitude in kg/s using procedures explained in Section 10.9.

10.13 Interferometry

We explain how interferometry measurements are performed in the next chapter, which also details the main new results of this dissertation.

Chapter 11

New results

This chapter reproduces Sections 5,6 and 7 of our published work in Ref. [11].

11.1 Gyroscopy with continuous cryostat reorientation

As we saw in Chapter 2, the Fiske amplified chemical potential battery enables continuous operation of the SHeQUID with the added bonus of improving the signal as well as the phase sensitivity [10]. In Fig. 11.1, we show the measurement of a continuously varying rotational flux using such an enhanced device – the first of our two main results in this paper. The data in this figure were obtained by a continuous (as opposed to point-by-point) reorientation¹ of the SHeQUID in the rotating reference frame of the Earth. To obtain the data in Fig. 11.1, we capture short (~ 80 ms long) timeseries of the position of the diaphragm. For each time segment, we perform a Fast Fourier Transform (FFT) to obtain the frequency spectrum and then integrate the whistle peak (at about 1080 Hz in this case²) to find the diaphragm displacement amplitude, which is converted to the whistle amplitude I_t using calibrations described in Section 10. Automated Labview [82] programs do this continuously so that we have an almost real-time measure of I_t .

Referring to Section 1.3.2 for a detailed analysis of the Sagnac effect, we can directly use Eq. (1.18) for the Sagnac phase-shift in terms of the cryostat angular position θ :

$$\Delta\varphi_{\text{rot}} = 2 \left(\frac{m_A}{\hbar} \Omega_E A \cos \lambda \right) \sin(\theta - \theta_0) \equiv 2 c_{\text{rot}} \sin(\theta - \theta_0) \quad (11.1)$$

where Ω_E is the Earth's angular velocity³ ($\sim 7.29 \times 10^{-5}$ rad/s) and A is the magnitude of the area vector \mathbf{A} of the interferometer sense loop. In our apparatus \mathbf{A} is oriented horizontally so that when the SHeQUID is reoriented, the vector sweeps out a circle in a plane parallel

¹See Section 9.2 for a discussion on issues related to such continuous reorientation.

²The duration of a timeseries chunk is chosen to be long enough to include several tens of whistle cycles for accurate FFT results and can be made smaller when the battery state is dialed higher in frequency.

³Here we use the length of a sidereal day (= 23.9344696 hours) for computing Ω , since it is the Earth's absolute rotation relative to the fixed stars that the SHeQUID senses.

to the ground. λ is the latitude of the location of the experiment (37.9° North in this case) and θ_0 is the cryostat position for which \mathbf{A} points due West (or East)⁴. θ is measured from an arbitrary reference zero using a commercial digital angle indicator⁵ whose resolution is $\sim 0.1^\circ$.

The expression for current (Eq. (1.13) with $\Delta\varphi = \Delta\varphi_{rot}$ from Eq. (11.1)) is fit to the data in Fig. 11.1 to obtain c_{rot} as a fit parameter. Using Eq. (11.1) and the known values of λ and Ω , we can obtain the effective area A of the sense loop ($\sim 10.9\text{cm}^2$), which differs from the design value by $\sim 2\%$. The design value ($\sim 10.7\text{cm}^2$) is the area of the loop defined by the axial paths along all tubular elements in Fig. 1.21 (which should be close to the averaged path “seen” by the superflow [84]). We see that the axial path is a good approximation to the effective boundary of the sense loop.




Figure 11.1: Interference from continuously reorienting the SHeQUID. This changes the rotation flux incident on the sense loop (and thus the relative phase difference between the aperture arrays). Data taken as described in the text is bin-averaged for every 1° of Dewar rotation before plotting.

⁴Refer to the discussion about using the sine instead of the cosine for the dot product on p.17.

⁵Renco Encoders Inc. model E-series optical encoder, which was originally part of a tachometer. We tapped into the displacement signal for our purposes here.

11.2 Flux-locked and linearized gyroscope for measuring continuously changing rotation fields

11.2.1 Flux locking and linearization using a heat current

The main idea with flux locking is to use a heat-induced phase shift to counteract the rotation-induced (or otherwise externally influenced) phase shift, thereby keeping the phase (and therefore the whistle amplitude) constant. The element used for flux locking is the heat-pipe shown in Fig. 1.21 and described in more detail in Section 1.3.2. A power injected into the heat-pipe heater R_{sense} creates a counterflow in the heat-pipe, with the superfluid flowing towards the heater and normal fluid carrying heat away from it, towards a thermal sink S (a thin, roughened copper sheet). The phase difference $\Delta\varphi_{heat}$ induced between the interferometer arms due to this counterflow is given by Eq. (1.21):

$$\Delta\varphi_{heat} = 2 \left[\frac{l}{\sigma} \frac{\pi m_A}{h} \frac{\rho_n}{\rho_s \rho T_s} \right] \dot{Q} \equiv 2 c_h \dot{Q} \quad (11.2)$$

where ρ_n and ρ_s are the normal and superfluid densities respectively, T is the temperature in the cell, l is the spacing between the arms and σ is the cross-sectional area of the heat-pipe.

Fig. 11.2 is an example of an interference pattern obtained by sweeping the heat-pipe power. We have fit the same functional form to this data as we did for Fig. 11.1, except with $\Delta\varphi = \Delta\varphi_{heat}$ from Eq. (11.2) instead of $\Delta\varphi_{rot}$. This fit yields the fit parameters a , b and c_h , where c_h is related to the change in heat-pipe power ($\dot{Q}_{2\pi} = \pi/c_h$) needed to make a 2π phase-shift across the sense arm. If we compare the fit value of c_h to the theoretical value from Eq. (11.2), we find that they agree within $\sim 1.4\%$ (the main source of systematic error being the effective length l , which is affected by flow details at the tube intersections [85]).

Now, the total phase-shift in the sense loop will be $\Delta\varphi_{total} = \Delta\varphi_{heat} + \Delta\varphi_{rot} + \varphi_{offset}$, where the offset term includes any constant phase-bias or drifts. Therefore, the heat-induced phase can be adjusted to cancel any changes in the rotation-induced phase. From this and Eq. (11.2), we see that the amount of heat-pipe power required to hold the SHeQUID at a point of constant phase is a linear measure of the rotationally induced phase-shift. The parameter c_h provides a calibration for this device by translating heater powers to phase-shifts, while a and b are used to optimize the automated flux-locking routine described in the following section.

11.2.2 Feedback

Previously [9], we demonstrated a static, point-by-point feedback as proof of principle. In that work, the cryostat was initially oriented so that the SHeQUID was biased at one of the points of steepest slope in Fig. 11.2 and the average amplitude recorded. The cryostat was then reoriented to a new position, thus changing the Sagnac phase. The heater power in the heat-pipe was then manually changed to bring back the amplitude to its original value.




Figure 11.2: Interference pattern obtained by sweeping the heat-pipe power (proportional to phase), as described in the text. The two types of biasing points of steepest slope are marked U (upslope) and D (downslope).

This process was repeated for many angular positions. The feedback input (heat-pipe power) was shown to depend linearly on the Sagnac phase. Such a method is of limited utility for monitoring time-varying rotation signals unless such variations are extremely slow. One of the main products of the present work is an automated computer system that can generate a battery state, perform a calibration, assist in optimal biasing and finally, adjust the feedback output power in response to the external phase changes that we wish to track. The process is as follows:

1. Once a stable battery state (with acceptable Fiske gain) is obtained, the computer program calibrates the device (as described in the previous section) and obtains the parameters a , b and c_h .
2. For the device to have sufficient dynamic range to track (+) and (-) phase shifts, the system must be able to add or remove heater power from the heat-pipe. Therefore, we impose an initial power offset of several 2π cycles (a few μW of power here).
3. The operating point (of maximum phase sensitivity) is found by adjusting the power offset until the amplitude is equal to the value at the steepest points in Fig. 11.2 ($I_{S,max} = b^{1/4}a$ from a simple analysis of the derivatives of Eq. (1.13)).

4. The whistle amplitude is now continuously measured and the change in whistle amplitude δI between consecutive measurements constitutes the “error signal” in a proportional feedback scheme⁶. This is used to compute the heat-pipe feedback power adjustment $\delta\dot{Q}$ required to nullify the error. At the steepest point, this is

$$\delta\dot{Q} \approx \delta I \cdot \left(dI_t/d\delta\dot{Q} \right)_{max}^{-1} = \delta I \cdot \left[a \left(1 - \sqrt{b} \right) c_h \right]^{-1} \equiv \delta I \cdot G_0$$

with G_0 being a good initial value for the proportional gain⁷. We record the total heat-pipe power in each iteration of the feedback loop and this is our feedback output signal \dot{Q}_{FB} required to keep the measured whistle amplitude at a fixed value. See Fig. 11.3 for an example of raw data from such an experimental run.

5. As mentioned previously, the total phase seen by the SHeQUID given both rotation and heater power is $\Delta\varphi_{total} = \Delta\varphi_{heat} + \Delta\varphi_{rot} + \varphi_{offset}$. If this total phase is kept constant using feedback, we should observe that $\Delta\varphi_{heat} = -\Delta\varphi_{rot} + \text{constant}$. We use the feedback power \dot{Q}_{FB} , calibration c_h and Eq. (11.2) to compute $\Delta\varphi_{heat}$. Together, they constitute the *measured* phase-shift (i.e. measured by the flux-locked SHeQUID). An observed drift rate⁸ in this phase is independently measured prior to this run and the drift is subtracted from $\Delta\varphi_{heat}$.
6. The cryostat angular position θ (independently recorded during the feedback run using the digital angle meter mentioned earlier) and c_{rot} obtained from the fit to Fig. 11.1, are used with Eq. (11.1) to compute the “actual” Sagnac phase-shift ($\Delta\varphi_{rot}$).
7. We plot $\Delta\varphi_{heat}$ (with drift subtracted) vs. $\Delta\varphi_{rot}$ in Fig. 11.5 and see that our SHeQUID is indeed, continuously tracking the Sagnac phase correctly (with a systematic error of $\sim 1\%$ in the expected slope of -1). This is the second main result for this paper.

⁶An additional integral term and gain are kept available but have thus far been found unnecessary for the work shown here.

⁷In practice, about half this value seems to work well to avoid feedback oscillations.

⁸See Section 12.1 for more details. As discussed on p. 229, the total drift (fairly constant) during this drift run is $\sim 30nW$. Compare this to around $1200nW$ of total change in \dot{Q}_{FB} .




Figure 11.3: Raw data: Whistle amplitude (kept constant) and feedback output power \dot{Q}_{FB} needed to do so.




Figure 11.4: Raw data: Independently recorded cryostat angle θ and feedback output power \dot{Q}_{FB} .




Figure 11.5: Example of dynamic feedback demonstration in a continuously operating SHe-QUID. The expected slope is -1 (see feedback process step 5 in Section 11.2.2). The slope of a linear fit to the data is -0.98 . The vertical intercept corresponds to ~ 7 cycles of initial (time goes from left to right here) heat-pipe offset, plus the phase required to bias at the steepest point (see feedback process steps 2 and 3 in Section 11.2.2). Phase drift has been subtracted from the data (see Section 12.1).

Chapter 12

Noise and drift

Sections 8 and 10 of our published work [11] have been reproduced in this chapter, with additional details included.

12.1 Phase drifts

To be most useful as a continuous monitor of long-term changes in some parameter (e.g. small changes in the Earth's rotation rate), it would be best to have no intrinsic long-term drifts in the SHeQUID. At temperatures further below T_λ , ever greater heater power \dot{Q}_{IN} is required in the inner cell to reach the critical velocity and generate a continuously whistling battery state. We find that when \dot{Q}_{IN} is on the order of microwatts, the SHeQUID phase drifts noticeably over time. We have measured this drift as a function of time and temperature during several different runs, each run consisting of the ~ 2 day period after transferring liquid helium into our cryostat. A sample drift run is shown in Fig. 12.1(a) and the results of the entire series of drift runs are summarized in Fig. 12.1(b).

From these drift measurements, we note two significant features: (1) there is a clear decrease in drift rates with decreasing inner cell powers (correspondingly being closer to T_λ); and (2) the drift rate in each run decreases smoothly in time, dropping by about a factor of two over a period of 48 hours. This drift (measured independently of the feedback) has been subtracted from the data shown in Fig. 11.5. The feedback test in Fig. 11.5 was performed at 9 mK below T_λ at ~ 7.5 inches of bath level where the drift rate (of ~ 30 nW/hr) was fairly constant over the duration of the test of ~ 1 hr (compare that to over 1200 nW of feedback output change during the same period of time).

We verified that this drift does not come from drifts in the heat-pipe power by shorting the resistive heater in the heat-pipe during some drift runs; instead, obtaining families of plots like Fig. 11.1 by sweeping the cryostat angular position back and forth. An automated system (with several safeguards) was developed to perform these reorientation sweeps continuously over periods of days in order to verify these drift measurements and eliminate the heat-pipe resistor as a cause. This automated system is described in Section 9.3.

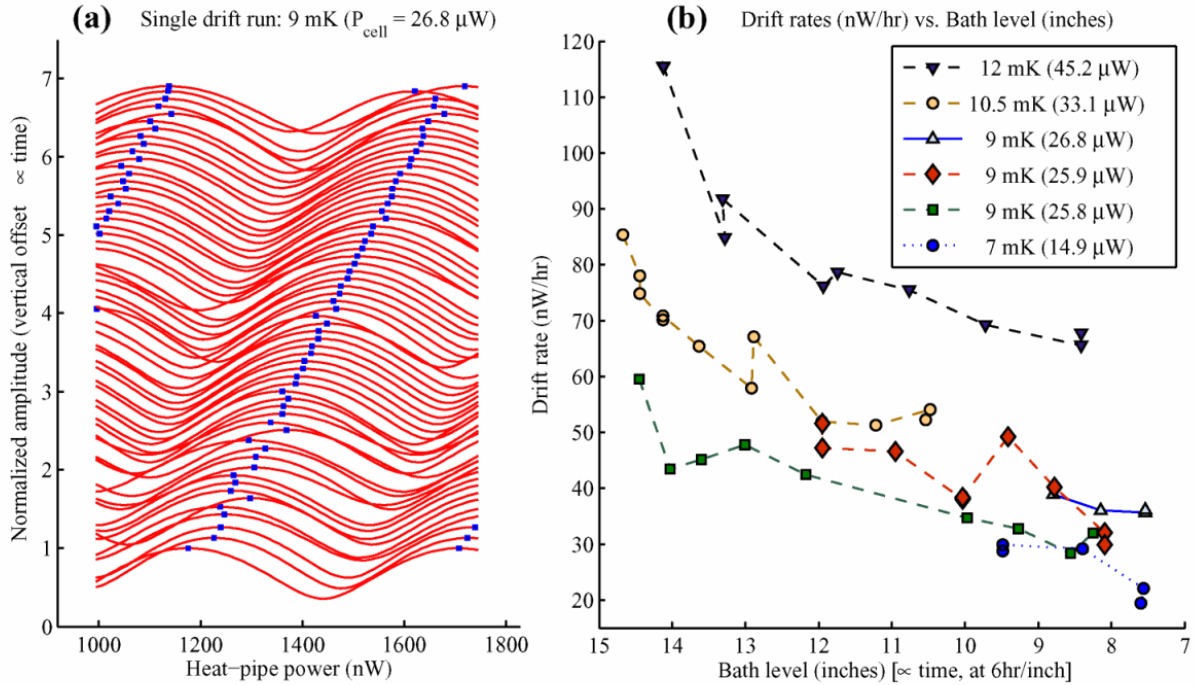


Figure 12.1: (a) Plots from a representative drift run described in Section 12.1. Each plot is a fit of the two-slit amplitude function (Eq. (1.13) with $\Delta\varphi = \Delta\varphi_{heat}$ from Eq. (1.21)) to an interferogram created by sweeping the heat-pipe power back and forth between two (fixed) extreme values. The fit result amplitude is normalized (so that we can focus solely on the phase information) and the curves are plotted with approximately equal vertical offsets (proportional to the actual time of each sweep) with an average duration of ~ 12.8 min between sweeps. The locations of the maxima (shown as solid squares) can then be used to obtain the rate at which the phase is drifting. This is done by binning the maxima location data over time and finding slopes in each bin. (b) This drift rate is observed to vary over time and with temperature (in mK below T_λ), which also dictates the inner cell heater power needed for the battery (see Section 2.4). We have plotted drift rate data from several drift runs against the helium bath level (which is proportional to time at ~ 6 hr/inch). We can see that the drift rates decrease with time (and bath level) as well as with inner cell heater power (included parenthetically in the figure legend).

A null test was also performed on the whistle measurement system (which includes the SQUID, any amplifiers, the data acquisition cards and the computer that performs the spectral analysis). This test consisted of driving the diaphragm capacitively (to ground) at similar frequencies (~ 1080 kHz) as the whistle with the chemical potential battery turned off and measuring the diaphragm response (designed to be approximately the same size and with the same amplifier gain and SQUID scale as usually set for the whistle) over ~ 16 hours. We found no drift within the noise level (which was similar to previous noise levels) in this null test, showing that the measurement systems were not the source of the drift.

It must be noted that there was a way to narrow down the source to the chemical potential battery more conclusively by obtaining interference curves with a whistle feedback (transient) system instead of the battery (as described in Section 2.1.1). However, we had a problem with the electrodes used to apply electrostatic forces on the diaphragm and could not make the feedback work during this (final) run.

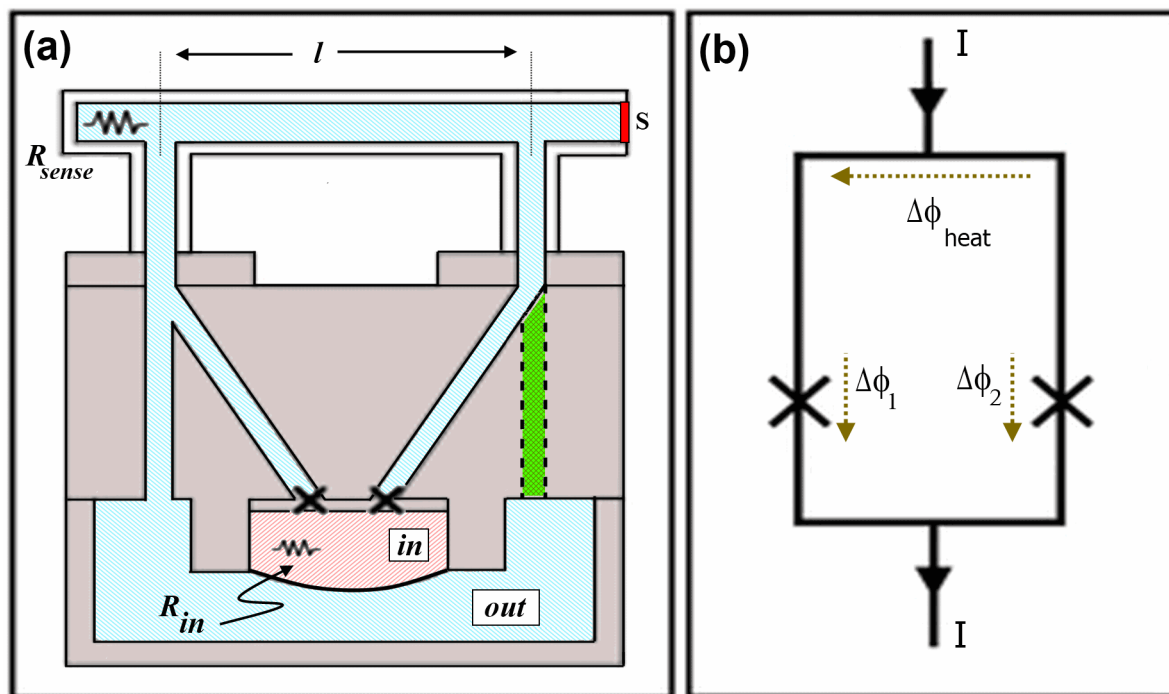


Figure 12.2: (a) Cell schematic and (b) equivalent circuit (dotted arrows denote the direction of increasing phase for the phase-drops shown). Resistive heater (R_{sense}) and roughened copper-foil sink (S) in the top “heat-pipe” produce a superfluid counterflow from S to R_{sense} . This is equivalent to a phase-difference $\Delta\varphi_{heat}$ between the two vertical side arms (and thus between the two aperture arrays marked X). The darker, dashed (green) flow path shown does not exist yet – it is a feature considered for future versions of the cell as a possible way to ameliorate the heat-pipe drift as described in Section 12.1.

Thus far, we have not discovered the mechanism of the cell heater dependent drift in the

present apparatus. However, we suspect that some of the inner cell heater power leaks into the heat-pipe, the leakage amount drifting as the Dewar's liquid helium level falls. We do not have a detailed model of the heat flow out of the inner cell but some of the heat must be transferred to the surrounding bath and some may leak into the heat-pipe. If the ratio of these two conductance paths changes in time, it would cause a drift in the interferogram.

We hope to make changes in a future apparatus that would eliminate this possibility (or at least reduce its impact). Specifically, adding a second, symmetric return path to the SHeQUID loop (the darker, dashed, vertical flow path (colored green) shown in Fig. 12.2(a)) or converting both return paths into superleaks might remove the unbalanced conductance path leading from the inner cell to the heat-pipe. An additional helium-filled, temperature regulated, sealed enclosure around the cell might prevent a time variation in heater power flowing from the inner cell to the heat-pipe.

The smallest drift rate seen in an independent device [10], which also used a Fiske-enhanced battery state (albeit at lower cell powers of only a few hundred nW compared to the tens of μW we use in this work), is ~ 0.01 rad/day [86]. This is consistent with the trend (in drift rate vs. cell power) that we observe and suggests that lowering the battery power (by, for instance, using smaller size apertures) might lower the drift rate to previously observed values.

12.2 Performance figures of the flux-locked SHeQUID

This section reproduces Section 8 of our published work in Ref. [11].

12.2.1 Phase noise

The current noise δI_n is calculated as the standard deviation of the mean of the whistle amplitude measurements in a single iteration of the feedback loop. By biasing the device at the steepest part of the interferogram, we can measure the noise in the system with the greatest sensitivity. We obtain the phase noise $\delta\varphi_n$ (in rad) from the current noise δI_n by using the phase sensitivity at the steepest point S_m (measured from a sample interferogram):

$$\delta\varphi_n \approx \delta I_n \cdot (dI_t/\Delta\varphi)_{max}^{-1} \equiv \delta I_n/S_m = \delta I_n \cdot \left[a \left(1 - \sqrt{b} \right) / 2 \right]^{-1}$$

which is then normalized to one second of measurement time to obtain the phase noise density (PND).

For the flux locked device reported here, we find a PND of $9 \times 10^{-2} rad/\sqrt{Hz}$ at our operating temperature of 9 mK below T_λ . The noise limitation is due to fluctuations in the SQUID-based displacement sensor that are an order of magnitude greater than the intrinsic noise in the SQUID. By comparison, the best-quoted phase noise in a Fiske-enhanced SHeQUID [10] is $3 \times 10^{-3} rad/\sqrt{Hz}$ and in a multi-turn device [12] is $3 \times 10^{-2} rad/\sqrt{Hz}$.

A rather comprehensive discussion on sources of noise (besides the ones we have already addressed in the course of this dissertation) in superfluid interferometers may be found in [26, pp. 224-233].

12.2.2 Response time and slew rate

This is ultimately determined by the time required to determine the amplitude of the quantum whistle. It takes us on the order of a second to acquire sufficient data for this purpose in each iteration of the feedback loop.

We can define the slew rate of this device (in the same manner as that for SQUIDS) as the maximum rate of phase change that the feedback can track without losing lock. While the intrinsic slew rate for this device is presently unknown, we observed that the feedback could maintain lock for phase change rates up to a maximum¹ value of 20 mrad/s. The phase change rate is $\Delta\dot{\varphi}_{rot} = (4\pi m_A/h) A \cos \lambda \left[\Omega \dot{\theta} \cos(\theta - \theta_0) \right]$, which is just the time derivative of Eq. (11.1). This formulation allows us to see more explicitly that a *constant* cryostat reorientation rate $\dot{\theta}$ imposes time-varying phase change rates upon the sense loop during a continuous reorientation. If we interpret the quantity in square brackets as an effective angular acceleration (as seen by the sense loop), we can restate the limit stated above as a maximal angular acceleration of $\sim 180 \text{ nrad/s}^2$ that the feedback has been shown to successfully track.

12.2.3 Dynamic range

This is set by the usable range of heat-pipe power values. As observed in Ref. [87], when one increases the heat-pipe power above a critical value, the superflow in the pipe becomes turbulent and vortices crossing the pipe cause the SHeQUID phase to oscillate through 2π . For the heat-pipe used in Ref. [87], the dynamic range actually observed corresponded to a phase-shift of $\sim 250 \times 2\pi$. We note that the heat-pipe dimensions and heater powers used in this work are nearly identical to the ones used in Ref. [87] and also to those in a further experiment where this dynamic range was subsequently verified.

Since the sensing points (the two vertical side-arms in the sense loop) are designed to be far from the ends of the heat-pipe, we have observed only a linear relationship between heater power and the phase response (up to the turbulent limit). Non-linearity considerations come into play only for multi-path interferometers (such as the 4-path “grating” interferometer described in Ref. [13]) where small differences in distances between the sensing points get amplified with increasing heater power and distort the interferogram. For the single-loop SHeQUID in the present work, linearity is not a problem.

¹We can call this a lower bound for the effective slew rate for this device at the present time, since this maximum value occurs during the fastest cryostat reorientation rate ($\sim 9 \text{ deg/min}$) that we have been able to sustain under the temperature regulation and battery stability constraints noted in Section 9.2.

12.3 The resonant landscape and frequency-dependent Fiske gains

We have already discussed in Chapter 2, the effects of the cell resonant modes (known and unknown) on the battery states. We have also seen (from the Fiske-state switching example in Fig. 9.2) that acoustic noise can inject enough energy into the system to cause transitions between neighboring resonant states. We finish that discussion here by considering the effect of a finite frequency width in a resonant state. Referring to Fig. 12.3, we can see (intuitively) that the whistle interaction with the resonance should amplify the whistle by a frequency-dependent gain that should roughly scale with the peak profile. If we call this frequency-dependent gain $G(f)$, the amplitude of the observed (Fiske-enhanced) whistle should be related to that of the bare whistle by $A_{obs} = A_{whis}G(f)$. In that case, if the whistle frequency fluctuates by an amount δf to the fluctuations in the chemical potential difference (which, we recall, can come from acoustic or thermal fluctuations injecting energy into the system), the corresponding fluctuations in the observed whistle amplitude will be:

$$\delta A_{obs} = A_{whis}G'(f)\delta f \quad (12.1)$$

We conclude that the walls of the broad peaks seen in the cell resonant landscape can lead to such “resonant noise” in the whistle.




Figure 12.3: A cartoon depiction of a broad resonant peak. See Fig. 2.1 for some actual broad peaks observed in our cell.

Chapter 13

The superfluid diffraction grating

We are especially grateful for helpful discussions on this subject with Ty Volkoff, Alexander Fetter and Yuki Sato.

13.1 Introduction

13.1.1 History

Interferometers employing neutral particles such photons, neutrons, cold atoms as well as neutral superfluids like helium-3 and helium-4 have been used to explore fascinating physical phenomena. They have shown great potential to be ultra-sensitive gyroscopes and find applications in fields as diverse as seismology, geodesy and inertial navigation. Their use as gyroscopes to measure extremely small rotational signals has been amply demonstrated for interferometers in the classic double-path configuration (after the fashion of Michelson and Morley [88]).

For superfluid helium-4, the natural progression to a multiple-path interferometer to achieve greater sensitivity to phase-shifts has already occurred [13] [14] but has so far not been used to measure rotations (rather, it has been used to detect phase-shifts due to a heat current). We describe here a feasibility study of an experiment to design and build what we shall henceforth call a SHeQUIG (Superfluid Helium 4 Quantum Interference Grating) in a configuration to measure the rotation of the Earth with a greater *rotation* sensitivity than previous attempts. Optimization of these devices for sensitivity to small rotations is a matter of considerable practical importance in these fields but leads to some rather counter-intuitive results when closely scrutinized. Specifically, it has been suggested [13] that using the analogue of a “diffraction grating” in superfluid helium gyroscopes might improve their sensitivity to rotation signals. It appears that what works for optical interferometry does not in this case carry over to superfluid interferometry.

In this chapter, we present an analysis of this issue as well as numerical flow simulations and experimental evidence supporting our claim that if we hold the grating length and the

total flow (which scales with the total exposed aperture area) fixed, there is no rotation sensitivity gain in increasing the number of slits in a grating. In other words, we can show that the purported N^2 gain in sensitivity is limited to the sensitivity with respect to the nearest-neighbor phase difference in a grating but is only an N -fold gain for the actual rotation sensitivity.

13.1.2 SHeQUIG implementation

The SHeQUIG implementations proposed thus far can all be shown to be equivalent to one archetypal design, reproduced here in Fig. 13.1 from Fig. 5 of Ref. [13]. Specifically, the linear grating we discuss in the experimental and numerical simulation sections here is simply the toroidal structure in the figure cut at the septum and unfolded out into a straight pipe. Topologically (and where phase gradients are concerned), there is no meaningful distinction. However, a direct sensitivity analysis of the toroidal SHeQUIG is provided in passing, in the final section (13.5) of this chapter.

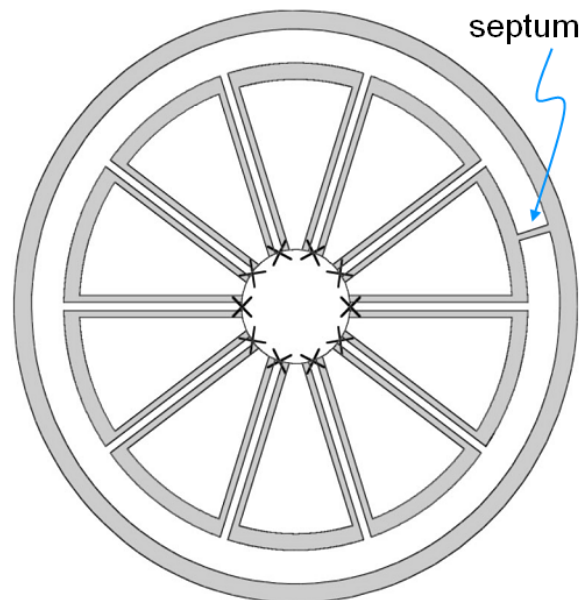


Figure 13.1: Reproduced from Fig. 5 of Ref. [13]. The septum is just a blocking wall in the circumferential flow path. It is needed for encoding information about the rotating reference frame into the superfluid phase (see Section 13.4.2). The X's are weak-links and the modulated currents interfere in the central chamber where they can be measured.

13.2 Linear grating theory

For the purposes of this chapter, we will simply refer to an aperture array as a “slit” and assume it behaves as one point source of superfluid oscillations (the oscillation being the so-called quantum whistle [ref to original whistle paper]). Subtleties arise when the internal structure of the slit is considered [ref. to Yuki diffraction paper], but that is simply a matter of judiciously demarcating the boundaries of an “aperture array” – specifically, we can think of single rows or columns or other groups of apertures as independent “slits” *if the quantum phase at each aperture in the group is identical within a given threshold*. The multi-slit expressions used here are discussed and derived in some detail in Appendix D.

13.2.1 Total amplitude

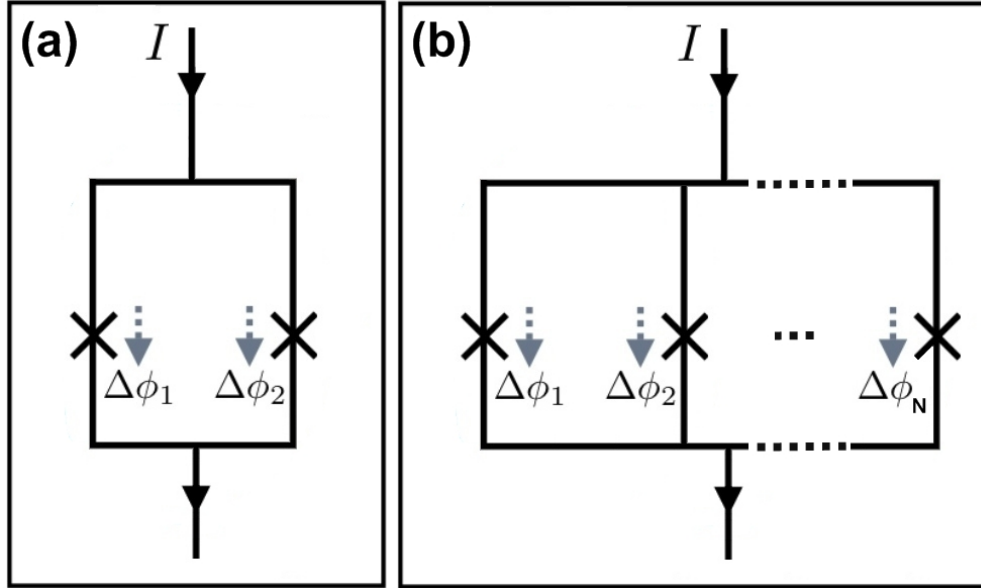


Figure 13.2: (a) Two slit interferometer and (b) N-slit grating interferometer. The $\Delta\phi_k$'s are phase-differences across each weak-link, and (for a constant chemical potential difference across the weak links) $\Delta\phi_k = \omega t + \varphi_k$, where φ_k is a time-independent phase-offset at the k'th weak-link.

We have already seen the implementation of a two-slit SHeQUID in the preceding chapters. That idea can be easily extended to placing more than two parallel paths with a “slit” in each path (see Fig. 13.2 b). The total mass current I_{tot} for N slits in parallel can then be written as:

$$I_{tot} = \sum_{k=0}^{N-1} I_{0,k} \sin(\omega t + \varphi_k) \quad (13.1)$$

where the $\{I_{0,k}\}$ and $\{\varphi_k\}$ are the amplitudes and phases of the oscillating currents at each individual slit.

Assuming identical slits¹, the $\{I_{0,k}\}$ are all the same (and equal to, say, I_0) and Eq. (13.1) can be rewritten as $I_{tot} = \left(I_0 \sqrt{\sum_{n=0}^{N-1} \sum_{k=0}^{N-1} \cos(\varphi_k - \varphi_n)} \right) \sin(\omega t + \lambda)$ where λ is a constant phase offset (unimportant here). The observed oscillation amplitude is then just the prefactor:

$$I_t = I_0 \sqrt{\sum_{n=0}^{N-1} \sum_{k=0}^{N-1} \cos(\varphi_k - \varphi_n)} \quad (13.2)$$

For the special case where the phase changes linearly between the parallel paths, the phase at the k^{th} slit is: $\varphi_k = k\theta$, **where θ is defined as the phase difference between adjacent paths (identical for all slits)**². Physically, this arises when there is a uniform phase gradient and therefore a uniform velocity superflow flowing along a channel transverse to all the paths. Then, the sum in Eq. (13.2) can be performed exactly to obtain the amplitude:

$$I_t = I_0 \left| \frac{\sin(N\theta/2)}{\sin(\theta/2)} \right| \quad (13.3)$$

This is the *linear grating formula* and is largely independent of physical system. The same expression crops up in any system with coherent oscillations in parallel with a linear phase gradient. We can see that it is modulated by the phase gradient along the grating. Note once again that θ is the nearest neighbor phase-difference (assumed constant along the grating).

13.2.2 A superflow sensor

A linear grating as described above is essentially a *superflow velocity sensor* when used in a superfluid system. The physical picture is as follows:

1. A physical influence sets up a stable superflow near the grating. *The nature of this flow is independent of the grating itself.* For any sensible analysis, the grating itself must be thought of as an external “phase-meter” brought in to measure the superflow.
2. This superflow, in the wavefunction picture of a superfluid is equivalent to a phase gradient along the slits in the grating so that adjacent slits see a relative phase difference. A uniform velocity corresponds to a linear phase gradient since $\mathbf{v}_s = (\hbar/m_4) \nabla\varphi$ (from

¹Corresponding expressions for non-identical slits are more difficult to obtain but are available in Appendix D.

²Note that this is an important definition and the entire chapter hinges on the distinctions between variously defined phase differences. For this reason, we will **always** consider θ to be defined in this way throughout this chapter.

Eq. (1.2)). The phase difference $\Delta\varphi$ between two fixed points separated by a distance x can then be written as:

$$\Delta\varphi_x = x (m_4/\hbar) v_s \quad (13.4)$$

3. By the formalism derived previously, the total whistle amplitude observed is modulated by this phase difference. Or, physically, an “interference pattern” is observed by changing the flow velocity.

The superflow can be set up by (among other things), a rotating reference frame such as the Earth (Section 1.3.2) and a heater-sink combo that generates a superflow towards the heater (Section 1.3.2).

The superflow velocity generated by a rotating rectangular cell will scale with the rotation rate Ω and the linear dimension of the cell (say R_{eff}). For a fixed size cell, we can therefore say that $v_s \propto \Omega$. Similarly, for a heater power \dot{Q} , $v_s \propto \dot{Q}$. The proportionality constants in these relations are independent of the interferometer slits and have to do solely with the cell geometry used for generating the flow. Putting in different gratings in identical cells will not affect their values.

13.2.3 Figure of merit

To determine an appropriate figure of merit related to sensitivity in a grating, we need to decide what it is that we’re trying to measure, *fix it*, and bring in the various instruments to try to measure it so that we can see which one is better. In all this, it is (as should be obvious once stated this way) essential that the physical effect being measured be kept the same when comparing different instruments. The way to do that in the context of the linear superfluid grating is to focus on the superfluid velocity. It is the sensitivity of the grating to changes in this induced superflow velocity (dI/dv_s) that ultimately determines the figure of merit for this grating. Given this criterion, and given a superflow that exists in some region of space, all that one can ask is: is an N-slit grating superior in superflow sensitivity to a 2-slit grating, and if so, by how much? It would be rather arbitrary to put any further restrictions on the spacing between the slits because all we would be interested in is whether the *best*³ such N-slit grating is superior to the *best* such 2-slit grating, given some ultimate physical size limit on the apparatus. It is this (eminently practical) consideration that underlies our eventual assumption of a fixed total grating length with slits added in between. We note for clarity that the superflow sensitivity does scale with total length, which is where the idea of maximizing the length comes from. Put another way, if one has only a certain maximum size limitation, why precisely should one use an N-slit grating over a 2-slit grating?

If one objects to this and demands that we choose a phase-difference instead, the relevant quantity must be the phase-difference between *two fixed points* along a tube (for instance) and comparing the performance of various gratings in measuring changes in that phase-difference. In the toroidal scheme mentioned previously, this merely implies keeping the

³See the discussion revolving around Eq. (13.9) for a visual aid to this argument.

total toroidal length fixed and adding varying numbers of slits within that constraint. It is strange that this essential starting point does not generally seem to be automatically accepted in the literature on the grating versions of SQUIDs, which is why we needed to belabor this point here⁴.

The essence of this entire discussion is that the phase-variable that always appears in the literature is the nearest neighbor phase-difference θ and that it is not an “instrument-agnostic” variable. The superflow velocity (or, equivalently, any phase-difference between two *fixed points in space*) on the other hand, *is* instrument-agnostic and must therefore be the relevant variable for comparing sensitivities.

13.2.4 Sensitivity with respect to . . .

For a generic phase pattern characterized by a set of phase values $\{\varphi_k\}$, Eq. (13.2) tells us that there is no unique phase to characterize the flow. Only in the case of a linear grating can the flow be easily characterized by a single phase parameter (θ) as seen in Eq. (13.3). This also is not unique because one might just as well use the phase difference between the ends of the grating array as a parameter (which is technically the correct thing to do, as we argued in Section 13.2.3). However, we will continue to use θ to demonstrate its inadequacies.

We can consider either the change in whistle amplitude with θ or with rotation rate (or equivalently, with superflow velocity). *These sensitivities scale differently and this difference is the central issue that this paper seeks to address.*

Nearest neighbor phase-difference (θ)

Fig. 13.3 shows the normalized amplitude from Eq. (13.3) plotted vs. θ for different numbers of slits N . The patterns shown always repeat in a 2π interval so we concern ourselves only with the interval from 0 to 2π .

We notice that there are always N lobes for an N slit grating and that the largest maxima occur at $\theta = 0$ and 2π , where the amplitude is I_0N (found by taking the appropriate limits of Eq. (13.3)). The zero crossings are precisely those points for which the numerator of Eq. (13.3) is 0 but the denominator is finite and this occurs at $\theta = 2\pi s/N$ (for $s = 1, 2, \dots, N - 1$). All N lobes are therefore seen to be of equal width ($= 2\pi/N$).

The maximum phase sensitivity (steepest slope) should therefore scale as the height of the largest lobe divided by its width:

$$S_\theta \equiv \frac{dI}{d\theta} \sim \frac{\text{peak height}}{\text{peak width}} \sim \frac{I_0N}{2\pi/N} \sim 0.16 (I_0N^2) \quad (13.5)$$

We can numerically compute the location of the steepest point (θ_c) in Eq. (13.3) and find that $\theta_c = 4.1633/N$ for N larger than about 20. Using this location, a more precise value of

⁴Lest the reader think we are being overly pedantic here, this issue of which phase-difference to use is still generally considered a matter of personal choice, to our deep bewilderment.




Figure 13.3: Simulated modulation amplitude for different numbers (N) of identical slits plotted versus the nearest neighbor phase difference θ . This has been normalized by dividing by N .

maximal sensitivity for large N is:

$$\boxed{S_\theta = 0.2181 (I_0 N^2)} \quad (13.6)$$

This constant is independent of physical quantities but depends on the linearity of the grating. This is a standard result that is quoted almost universally in the grating literature.

Phase-difference between two fixed points

As we explained at some length in Section 13.2.3, this is the (uniquely) meaningful phase one must consider for comparing two phase monitoring devices. For the special case: $x = D$ where D is the nearest neighbor spacing, Eq. (13.4) describes the nearest neighbor phase-difference:

$$\theta = D (m_4/\hbar) v_s = \frac{L}{N-1} (m_4/\hbar) v_s \quad (13.7)$$

where $L = (N-1)D$ is the total length of an N -slit grating. Used with Eq. (13.4) for $\Delta\varphi_x$, this equation yields:

$$\theta = \Delta\varphi_x \frac{L}{x} \frac{1}{(N-1)} \quad (13.8)$$

where we emphasize that x is independent of any grating dimensions (it is simply the distance between the two fixed test points). Therefore, the sensitivity with respect to this fixed $\Delta\varphi_x$ is:

$$S_f \equiv \frac{dI}{d\Delta\varphi_x} = \frac{dI}{d\theta} \frac{d\theta}{d\Delta\varphi_x} = S_\theta \frac{L}{x} \frac{1}{(N-1)}$$

Finally, using Eq. (13.6), this becomes:

$$\boxed{S_f = 0.2181 I_0 \left(\frac{N^2}{N-1} \right) \frac{L}{x}} \quad (13.9)$$

It is worthwhile to pause and consider this equation in light of our figure of merit discussion in Section 13.2.3. This tells us that for a given N -slit grating, the phase-sensitivity is maximized by maximizing the total grating length L . Now, the critical argument is that this L can be maximized equally well for any N ! In other words, if we wish to gauge the performance of the grating, it makes no sense to choose arbitrarily different L 's for gratings with different N 's. This means that the fraction L/x in the boxed equation above is irrelevant and the (true) phase-sensitivity scales only as the ratio $N^2/(N-1)$, which is approximately equal to N to within 5% for $N > 20$.

This is a rather straightforward result, and it tells us that for the most meaningful sensitivity test (and definition) we can devise, a linear superfluid grating gains in sensitivity only by a factor of N over its 2-slit counterpart (more precisely, by a factor of $N/4$ for $S_f(N)/S_f(2)$).

Superflow velocity

Using Eq. (13.7) and a similar process as the previous section, we can obtain the sensitivity to superflow for an N -slit grating:

$$S_{v_s} \equiv \frac{dI}{dv_s} = \frac{dI}{d\theta} \frac{d\theta}{dv_s} = S_\theta D \frac{m_4}{\hbar} = 0.2181 \frac{m_4}{\hbar} I_0 (N^2 D)$$

As before, with the total grating length $L = (N-1)D$,

$$\boxed{S_{v_s} = 0.2181 \frac{m_4}{\hbar} I_0 \left(\frac{N^2}{N-1} \right) L} \quad (13.10)$$

Comparing this to Eq. (13.9) and the discussion following it, this sensitivity is also seen to scale up only as $\sim N$.

From Section 13.2.2, this will also be the scaling for sensitivity to rotation or to heat-induced superflow (again, keeping the total length of the linear grating the same while changing the number of slits used).

Rotation rate and heat-pipe power

Rotation rate Ω is the crucial quantity for a practical rotation sensor. Following the same reasoning as before while noting that the superflow along the grating will scale as $v_s = \Omega R_{eff}$ (for a characteristic length R_{eff} that depends only on the box dimensions and not on the grating characteristics), we find that the sensitivity to rotation is $S_\Omega \equiv (dI/d\Omega) = (dI/dv_s)(dv_s/d\Omega) = S_{v_s} R_{eff}$. Using Eq. (13.10), we have

$$S_\Omega = 0.2181 \frac{m_A}{\hbar} I_0 R_{eff} \left(\frac{N^2}{N-1} \right) L \quad (13.11)$$

For the flow induced in a heat-pipe (such as the one described in Section 1.3.2), the relationship between heater power \dot{Q} and superflow velocity v_s is derived from Eqs. (1.19) and (1.20):

$$\dot{Q} = (\sigma \rho T s \rho_s / \rho_n) v_s \quad (13.12)$$

This allows us to write the sensitivity with respect to heater power: $S_{\dot{Q}} \equiv (dI/d\dot{Q}) = (dI/dv_s)(dv_s/d\dot{Q})$. And finally,

$$S_{\dot{Q}} = 0.2181 \frac{m_A I_0}{\hbar \sigma T s \rho_s / \rho_n} \left(\frac{N^2}{N-1} \right) L \quad (13.13)$$

As before, both these sensitivities are also seen to scale up only as $\sim N$ over the corresponding 2-slit sensitivity.

Summary

Any meaningful definition of grating sensitivity with respect to a well-defined physical quantity that is *independent of the grating dimensions* is therefore seen to scale proportionally with the number of slits N . The only definition (S_θ) that does scale up as N^2 has been shown to be deeply flawed as a figure of merit for gratings because the nearest neighbor phase-difference θ contains within it grating parameters whose scaling is thereby hidden from view. Put simply, by using S_θ , we mask the fact that we are automatically changing the size of the physical influence being measured when we change N - not a fair or realistic way of comparing two different gratings.

The linear ($\propto N$) sensitivity increase can be attributed simply to the N -fold increase in current from having N -slits. Thinking back to our intuitive estimation of the sensitivity as the lobe height divided by the lobe width (see p. 240), this means that the lobe height increases with N , but the lobe width (in an interference curve plotted against the meaningful variables above) does not change appreciably. In fact, we will see from simulations and from superfluid grating experiments already performed that the lobe width (i.e. the horizontal compression in the interference pattern) essentially contributes nothing of importance to

enhancing the sensitivity. Detailed simulations show that the lobe width actually starts low (!) and converges quickly to a higher stable value with increasing N. This means that after allowing for the lobe height increase (solely due to higher flow through the increased number of slits), the 2-slit sensitivity is actually marginally *greater* than that for larger N!

We can see this immediately from our analytical results. Using $v_s = \Omega R_{eff}$ from the previous section in Eq. (13.7), we have

$$\theta = \frac{L}{N-1} (m_4/\hbar) \Omega R_{eff} \quad (13.14)$$

This enables us to substitute for θ in the total current expression of Eq. (13.3) and obtain the current as a function of Ω . Using this new current function, with a fixed grating length $L = 2$ inches and $R_{eff} \sim 4$ inches (on the order of the cell dimensions), we can plot the current as a function of Ω for various values of N. These currents (normalized by N to equalize the maximum amplitudes) are shown in Fig. 13.4 for several values of N. Note that box and grating parameters were chosen so that several diffraction lobes could be seen upon changing Ω from 0 to Ω_{earth} . Since we know that the first lobe width in θ should be $2\pi/N$, the width in Ω can be found from Eq. (13.14): $\delta\Omega = \frac{(h/m_4) \frac{N-1}{N}}{R_{eff}L}$, which approaches $\delta\Omega_{N \rightarrow \infty} = \frac{(h/m_4)}{R_{eff}L} = 1.93 \times 10^{-5} rad/s$ for large N (and this is exactly the limiting value seen in the inset of Fig. 13.4).

A deeper puzzle

It is instructive to ask where the grating analogy referenced to optical interferometry goes astray. SQUIDS and SHeQUIDS (despite their physics) do not exhibit, in the strictest sense of the word, *quantum interference*. The interference is quantum mechanical only in the limited sense that the very idea of the phase in a superfluid context has its origin in the quantum mechanical order parameter that describes the superfluid. The main feature of “quantum interference” is that the amplitudes are added and *then* squared to obtain probabilities. This is true of photons and electrons and atomic beams. It is (ironically) even true of sound interference and the classical electromagnetic picture of optical interference, where the amplitudes (or electric fields) are coherently superposed and this coherent sum is *then squared* to obtain the acoustic or optical intensities that are the measured quantities. This adding in quadrature is what gives us the extra factor of N in the increased sensitivity (merely the square of Eq. (13.3)).

Recall that in SQUIDS and SHeQUIDS, the weak-links produce coherent current oscillations, which are then coherently superposed *but not subsequently squared*. It is therefore a bit unreasonable to expect a similar sensitivity enhancement from superfluid gratings.

Epilogue: so what?

What does all this mean as far as the practical question of making and using superfluid gratings is concerned? We have seen from the various analyses provided that a grating




Figure 13.4: Current (normalized by N) as a function of Ω for various values of N . Inset shows closeup of the first zero-crossings (or approaches) of the interference curves on the horizontal axis, which define the first lobe widths. Note that the first lobe width does approach the predicted limit for large N ($\sim 1.93 \times 10^{-5} \text{ rad/s}$). Also note that after allowing for the N -fold increase in lobe height, the lobe width actually *increases* with N as predicted. This makes a two-slit SHeQUID marginally more sensitive than an ($N > 2$)-slit grating (with the same total current).

can only increase rotation sensitivity linearly with the number of slits N . Since a superfluid grating would be constructed either as a true two-dimensional array (with rows along the flow and columns transverse to the flow) or a set of individual aperture arrays (as was done with 4 aperture arrays in Ref. [13]) laid out linearly, the sensitivity increase for measuring superflow (or rotation or heat currents) will scale only as the number of slits N , which is fully attributable to the increased flow due to the overall larger exposed area from the aperture arrays.

Therefore, one can merely increase the number of holes in individual aperture array chips (by increasing the linear dimension of the square arrays by a factor of \sqrt{N}) or in the case of the true 2D arrays, increase the number of holes in the columns of the array by a factor of N to increase the maximum sensitivity by the same magnitude as having an N -slit grating.

Considering the technical difficulties involved in fabricating very long aperture array

gratings (of the order of an inch or so), it is therefore more sensible to build two-slit (two-array) interferometers with the largest possible slit extent (L) and pack as many apertures into each slit as possible. There are limits to increasing the total flow for technical reasons, but that is a problem that will afflict both 2-slit and N -slit SHeQUIDs alike so that is not a point of comparison between them.

It is important to note that if it *were* possible to enhance one of the meaningful sensitivities by N^2 instead of N , we would have to increase the number of holes in each column by N^2 to gain the same enhancement as having N slits (this would become impractical very quickly from the standpoint of fabrication and other distortions will also start setting in once the column extent starts becoming too large for the phase to stay uniform along it). In that case, an N -slit grating *would* yield a significant advantage because the increase would be N times larger than could be explained by just the linear increase in the number of holes.

Optical gratings work for precisely this reason. The oscillating current in our base equations 1,2 and 3 is directly analogous to the oscillating electric field in an optical diffraction grating. The measured quantity there is intensity, which is the *square of the electric field*. That is where the extra factor of N comes from in optical gratings and why the analogous gratings are not as powerful for SQUIDs and SHeQUIDs (where it is phase-coherent particle fluxes, i.e. mass currents, that interfere with each other).

13.3 Experimental evidence

The only experiments (to the best of our knowledge) done with superfluid helium-4 gratings ($N > 2$) are for superflows created by heat currents (and not for Sagnac phase-shifts). Since the critical currents per aperture can be quite different for different temperatures and different hole sizes, we cannot directly compare the actual sensitivities to check the above analysis (because we cannot fully account for the differences in total current between different apertures arrays across cooldowns).

However, we can check the data to see whether there is any horizontal compression in the interference pattern (for when the whistle amplitude is plotted against an actual physical quantity like the heat-pipe power or superfluid velocity in a channel). Since the alleged sensitivity enhancement comes from one factor of N owing to the increased current (increased first lobe height) and another factor of N owing to the horizontal narrowing of the pattern and hence a decreased first lobe width, verifying that one of these factors is absent would be a good way to support our hypothesis.

In what follows (as before), we define the first lobe of the pattern as the tallest half lobe as seen in Fig. 13.3. We can refer to that figure and the experimental plots to read off the first lobe width. Now, we need a prediction within our analytical framework for the first lobe width in heat-pipe power \dot{Q} as a function of the grating parameters. From Eqs. (13.7) and (13.12), we obtain a relationship between θ and \dot{Q} :

$$\dot{Q} = \left(\frac{\hbar}{m_4} \frac{\sigma \rho T_s \rho_s}{\rho_n} \right) \frac{N-1}{L} \theta \quad (13.15)$$

Since the first lobe width in θ (see the previous sub-section on “Nearest neighbor difference”) is $\delta\theta = 2\pi/N$, the above equation gives us a prediction for the first lobe width (on the heat-pipe power axis):

$$\boxed{\delta\dot{Q} = \left(\frac{h}{m_4} \frac{\sigma T_s \rho_s}{\rho_n} \right) \frac{N-1}{N L}} \quad (13.16)$$

Clearly, we should expect no decrease in the first lobe width! On the contrary, the function $(N-1)/N$ starts low at $1/2$ for $N = 2$ and converges to its maximum value of 1 as N becomes arbitrarily large.

13.3.1 $N = 2$

There are several experiments for this case, including runs described in this dissertation.

Fig. 13.5 shows data reproduced from work by Sato, *et al.* [8]. The first lobe width for this interference curve is observed (directly from the figure) to be $\sim 319nW$ for a temperature of $T_\lambda - T \approx 16mK$. For the grating parameters used in that paper ($L = 1$ inch and $\sigma = 3.78 \times 10^{-6}m^2$), Eq. (13.16) predicts $\sim 341nW$, which is within the systematic uncertainty quoted in that paper. We repeated this comparison for the lobe widths (half the power needed for 1 cycle) summarized in Fig. 3 of that paper and found similar agreements with the predictions.

A clearer example is the 2-slit SHeQUID described in this dissertation (and in Ref. [11]). Fig. 3 of this reference (reproduced here as Fig. 11.2) shows a first lobe width of $\sim 275nW$ at $T_\lambda - T \approx 9mK$ and Eq. (13.16) predicts a width of $\sim 274nW$ at that temperature with grating parameters: $L = 1$ inch and $\sigma = 4.45 \times 10^{-6}m^2$.

13.3.2 $N = 4$

Fig. 13.6 shows data reproduced from Ref. [13]. In this case, the horizontal axis has already been converted to the nearest neighbor phase difference θ (the reference calls this $\Delta\phi$). So, Eq. (13.16) is equivalent to simply a prediction (for $N = 4$) of a lobe-width of $\sim 2\pi/4 = 0.25 \times 2\pi$. As seen in the figure, the actual first lobe width at $T_\lambda - T \approx 4mK$ is $\sim 0.24 \times 2\pi$.

13.3.3 $N = 75$

Fig. 13.7 shows data reproduced from work by Narayana, *et al.* [14]. This is an interference pattern produced by a single aperture array, with one slit corresponding to one column of 75 apertures. There are 75 such rows in the aperture array and so we have 75 slits with a $2\mu m$ separation between each slit.




Figure 13.5: 2 slit interferogram reproduced from Fig. 2 of Ref. [8].




Figure 13.6: 4 slit interferogram reproduced from Fig. 2 of work by Sato, *et al.* [13]. Note that the nearest neighbor phase difference θ is denoted as $\Delta\phi$ on the graph horizontal axis.

For design values of $L = (75 - 1) \cdot 2\mu\text{m}$ and $\sigma = 1.22 \times 10^{-5}\text{m}^2$, the predictions from Eq. (13.16) are (for temperatures of $T_\lambda - T = 5, 8$ and 12 mK) respectively, 168 , 234 and $311 \mu\text{W}$. The first lobe widths read directly from the figure are 165 , 226 and $286 \mu\text{W}$ for the respective temperatures; all within the systematic uncertainty of $\sim 10\%$ seen in the referenced paper.




Figure 13.7: 75 slit interferogram reproduced from Fig. 3 of Ref. [14].

13.3.4 Summary

We see that the horizontal compression in the interference pattern closely matches the predictions of the model developed here. Therefore, we should expect only an N -fold (and not an N^2) enhancement in either of the physically meaningful sensitivities defined in this chapter. The data bear this out.

13.4 Simulations

As a further test of our claims in this chapter, we simulate⁵ the rotation-induced flow in a box full of superfluid and find that the flow is approximately linear (implying a linear phase gradient) at the box walls, away from the edges. In accordance with the picture presented previously in this chapter, we then place a linear grating with N slits along one of the walls so that the phase at each slit in the grating is set by the superflow (and varies approximately linearly along the grating).




Figure 13.8: Arrangement for simulation (schematic only).

13.4.1 Procedure

We simulate the flow using a numerical PDE solver to solve Laplace's equation for the superflow phase texture in the inertial frame of the fixed stars (based in part on Ref. [60]). The dimensions of the box are kept fixed at 4 inches square. Fig. 13.8 shows this arrangement. The flow is solved for different rotational speeds Ω for a slice normal to the rotation axis and passing through the grating (the blue vertical plane shown in Fig. 13.8). An aperture array grating (of total length L and number of slits N) is then placed along one face of the box and the phase texture along the grating (due to the flow) causes the transverse currents through the aperture columns (the "slits") to interfere with differing phases. The transverse currents (vertical in the figure) are the quantum whistles, which are generated and the interfered oscillations detected by techniques described previously in this dissertation.

⁵Details and Matlab code for these simulations are provided in Appendix H.

13.4.2 Phase textures and invariance under boosts

It is important to note that the phase texture is invariant under a rotation boost [30] [89] so that the only influence of the rotation on the phase texture comes from the direct action of the walls and not from the boost to the rotating frame attached to the grating⁶. The phase changes non-trivially only due to the force exerted on the superfluid by the walls (mathematically, rotation information enters via the boundary condition of the moving walls). Incidentally, this is why a perfect cylinder would not work and why a rectangular (or elliptical: see Ref. [90]) box is needed to encode information about rotation into the phase.

13.4.3 Grating uniformity: equispaced and “optimal” gratings

The interference pattern is distorted because the superflow is not exactly uniform along the grating (due to edge effects from the finite sized boundary). We show here the results of using an equispaced (linear) grating and also a so-called “optimal” grating (in the sense that the slits are spaced so as to make the phase gradient linear and the nearest neighbor phase differences all equal for one value of the rotational speed Ω and use those fixed positions for subsequent values of Ω). We will see that the label “optimal” is ill-considered as far as maximum sensitivity goes.

13.4.4 Results

Phase textures and velocities

A representative phase texture is shown in Fig. 13.9 for one value of the rotation rate $\Omega = \Omega_{earth}$. The flow velocity (boosted to the non-inertial rotating frame attached to the box walls/grating) is shown as a superimposed vector field. We calculate such phase textures for several values of Ω and this generates our master data set.

Given the procedure so far, if all physical quantities are input in SI units, the velocity will be obtained in real units as well. For rotation rates similar to that of the Earth and box dimensions of the order of an inch, we obtain maximum flow speeds (in the rotating frame R , i.e. relative to the box walls) on the order of $\mu m/s$. Fig. 13.10 shows the speed in m/s as an intensity plot, with the velocity in the rotating frame superimposed as a vector field plot.

Interference curves and sensitivities

Now, different gratings can be placed on the box (here we place gratings centered on the top boundary of the box) and the phases $\{\varphi_k\}$ at each slit position k in the grating are collected from the simulated data.

The total current amplitude is directly computed according to Eq. (13.2). This constitutes one point in our interference pattern of amplitude vs. rotation rate. This is repeated for

⁶See Section H.2.4 for a more detailed explanation of this issue.




Figure 13.9: Sample phase texture in the inertial frame for a 4×4 inch square box and $\Omega = \Omega_{earth}$. The vector field shown is the fluid velocity referred to the rotating frame of the box.

each value of Ω (1000 steps between 0 and $\sim 10 \cdot \Omega_{earth}$)⁷ and the amplitude is plotted vs. Ω in Fig. 13.11 (for $N = 10$ slits). The equispaced and “optimal” slit positions are shown in the inset figure. Note that the equispaced pattern is slightly steeper (higher sensitivity) than the “optimal” grating. However, this doesn’t mean the optimization is useless - it does make a given grating more linear in the phase-gradients set up near it.

The maximum sensitivity with respect to rotation rate (S_{Ω}^{max}) occurs at the steepest part of this curve. We compute this interference curve as described above for a fixed length grating (2 inches long) for different numbers of slits placed between the (fixed) endpoints ranging from $N = 2$ to $N = 200$ and directly measure the maximum sensitivity for each curve (by numerically differentiating the curve). (S_{Ω}^{max}) normalized by the current amplitude

⁷This large value of Ω is used in the simulation to increase the number of lobes visible. Note that because these simulations are all to scale, this fact tells us much about what sort of box sizes would be needed to detect the Earth’s rotation with a device like this.




Figure 13.10: The vector field shown is the fluid velocity referred to the rotating frame of the box. Intensity plot is the magnitude of the velocity field in m/s. Speeds are on the order of $\sim \mu\text{m/s}$.

of one slit is plotted against the number of slits N in Fig. 13.12.

For comparison, the maximum sensitivity with respect to the nearest neighbor phase (S_{θ}^{max}) (normalized by the current amplitude in one slit I_0) is plotted against N in Fig. 13.13. We note that a parabolic fit to the “optimized” curve in (a) recovers our result in Eq. (13.6) (plus small numerical errors due to the inclusion of small N data). Similarly, the essentially linear rise in S_{Ω}^{max} with N seen in Fig. 13.12 confirms our prediction of Eq. (13.11).

We note in passing that the distorted curve (equispaced grating) actually has a $\sim 16\%$ larger sensitivity than the “optimized” one (the constant factor in Eq. (13.6) becomes 0.2536 instead of 0.2181).




Figure 13.11: Sample interference pattern vs. Ω for a 10 slit grating and slit positions for equispaced and optimal gratings (inset).

Horizontal compression in interferograms

Now, we can have a final confirmation of our analytical and experimental results on the lack of horizontal compression for higher order gratings as compared to 2-slit gratings. Fig. 13.14 shows six gratings with different numbers (N) of slits (but the same total extent) placed on our simulated box to compute interference patterns. The vertical axes show that the first lobe height does indeed scale up linearly with N but we can see (yet again) that the first lobe width actually increases and plateaus at a limiting value for very large N .

13.5 The dc SQUID: misconceptions and clarifications

There are instances of grating analogues in the SQUID literature [91] [92] (typically called SQUIGs or SQIGs) and they seem to fall prey to the same fallacy that we analyzed in such




Figure 13.12: Maximum sensitivity with respect to rotation rate (S_{Ω}^{max}) vs. number of slits N in a fixed length grating 2 inches long.




Figure 13.13: Maximum sensitivity with respect to nearest neighbor phase (S_{θ}^{max}) vs. number of slits N in a fixed length grating 2 inches long.

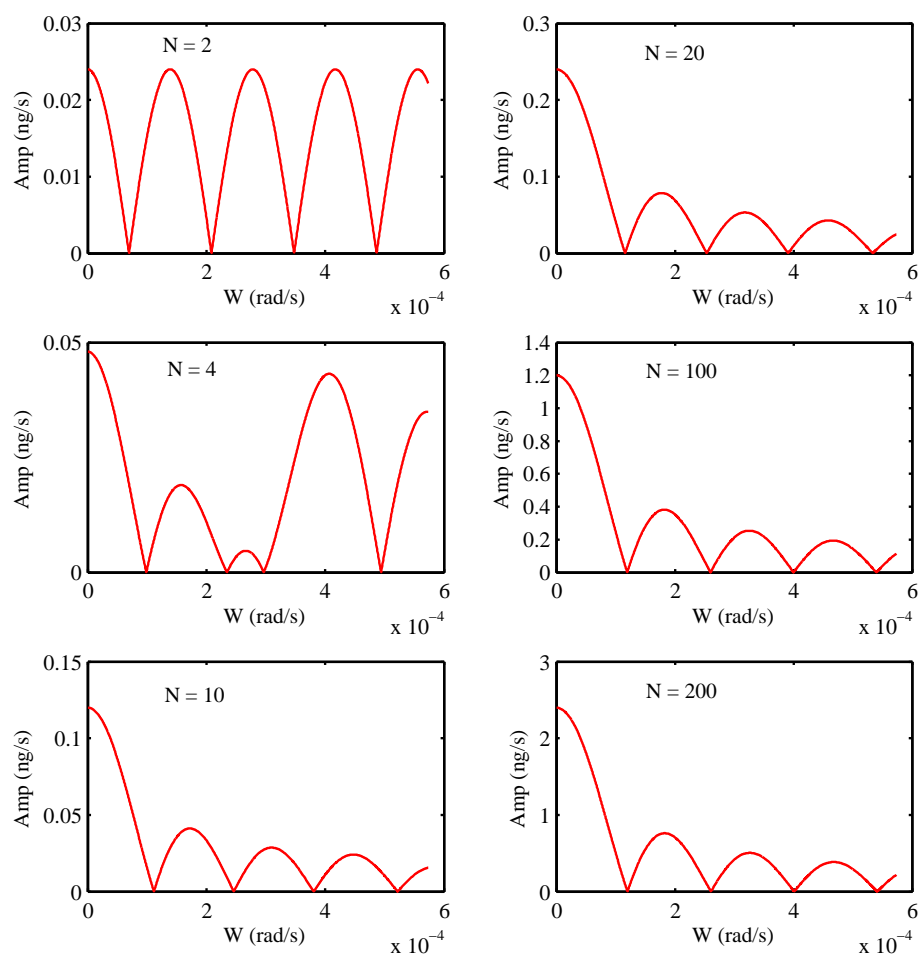


Figure 13.14: Interference curves for 6 different gratings with varying numbers of slits (N) as obtained from flow simulations. These are all equispaced gratings. Note that once we get above $N = 10$, the curves look essentially identical, with only a vertical scaling factor of $\sim N$ to distinguish between gratings. This is the main message of this chapter and we have now presented analytical, experimental and simulation-based evidence to support it.

detail in Section 13.2.3. Fortunately, that issue is much easier to clarify in the context of an external magnetic flux being measured by a SQUID than it was in the preceding sections. Therefore, while addressing the issues with the SQUIG literature, we will construct a parallel framework for SHeQUIGs (used as gyroscopes) by discussing it in terms of external rotation flux. In this framework, the archetypal SQUIG (or SHeQUIG) looks like Fig. 13.15.

As before, we should define a figure of merit such that the physical influence (magnetic field, rotation field) being measured is confined to a fixed region (say, some total area A_T) and we ask how well a given N -slit grating performs as compared to a simple 2-slit grating when it comes to detecting changes in this total flux.




Figure 13.15: A schematic SHeQUIG (or SQUIG) with (time-independent components of the) phase-drops across the weak-links shown. Each of the hatched loops has area A .

Eqs. 3 and 5 of Ref. [92] or Eq. 1 and Fig. 1 of Ref. [91] show that their variable of interest (with respect to which they define their sensitivity) is the flux through *a single loop* (i.e. through just one of the hatched regions shown in Fig. 13.15), which the reader will recall, is just our nearest neighbor phase difference θ . So, their quotations of an N^2 gain in sensitivity apply only to S_θ , and we have already seen (ad nauseum perhaps) the problems with using that as a figure of merit. The choice of θ as the sensitivity variable is simply not instrument-agnostic.

Further, the gratings in these references (with total area $A_T = A \cdot (N - 1)$) are compared to a 2-slit grating with area A , which seems entirely unreasonable given our discussions on the matter! This choice is cause for further perplexity when one considers the fact that the sense-loop area has already been a factor to be maximized for maximal sensitivity in most of the Sagnac interferometry literature (regardless of the working substance). Note that the toroidal implementation of the SHeQUIG shown in Fig. 13.1 makes the necessity of the

fixed total area constraint (and adding more radial paths with a weak-link in each one) more immediately obvious.

So, in the simplest terms possible, these references are obtaining an additional factor of N in sensitivity gain by taking advantage of the increased area of exposure to the magnetic field being measured and incorrectly ascribing this gain to the grating.

In closing, we can quickly derive a rotation sensitivity for the SHeQUIG pictured in Fig. 13.15. The phase difference θ between neighboring weak-links is found by quantizing the circulation integral in the loop (of area A) defined by those two weak-links. To wit, $\theta \equiv \varphi_{k+1} - \varphi_k = 2\pi\Omega A/\kappa_4$, where $\kappa_4 = h/m_4$ is the ^4He quantum of circulation. So, the rotation sensitivity is $S_\Omega \equiv \frac{dI}{d\theta} \frac{d\theta}{d\Omega} = S_\theta 2\pi A/\kappa_4$. Using Eq. (13.6), we obtain: $S_\Omega = (0.2181 I_0 N^2) \frac{2\pi A}{\kappa_4}$. Finally, defining the (fixed) total grating area to be $A_T \equiv A \cdot (N - 1)$ we obtain:

$$S_\Omega = \left(\frac{0.4362\pi}{\kappa_4} A_T \right) (I_0 N) \frac{N}{N - 1} \quad (13.17)$$

As before (Eq. (13.11)), we see that the sensitivity to rotation scales up as N, solely due to the increased current due to multiple slits, with the grating contribution (the last term involving N) actually getting slightly *worse* with increasing N (but quickly plateauing to a limiting value).

The same general equations and trends hold true if we were to find the SQUIG sensitivity to changes in a magnetic field that is incident on a fixed area A_T .

Bibliography

- [1] Emile Hoskinson, Richard Packard, and Thomas Haard. Oscillatory motion: Quantum whistling in superfluid helium-4. *Nature*, 433(7024):376, 1 2005. doi: 10.1038/433376a. URL <http://dx.doi.org/10.1038/433376a>.
- [2] Emile Hoskinson, Yuki Sato, Inseob Hahn, and Richard Packard. Transition from phase slips to the josephson effect in a superfluid ^4He weak link. *Nat Phys*, 2(1):23, 1 2006. doi: 10.1038/nphys190. URL <http://dx.doi.org/10.1038/nphys190>.
- [3] Yuki Sato, Emile Hoskinson, and Richard Packard. Transition from synchronous to asynchronous superfluid phase slippage in an aperture array. *Physical Review B (Condensed Matter and Materials Physics)*, 74(14):144502, 2006. doi: 10.1103/PhysRevB.74.144502. URL <http://link.aps.org/abstract/PRB/v74/e144502>.
- [4] Emile Hoskinson, Yuki Sato, and Richard Packard. Superfluid ^4He interferometer operating near 2 K. *Phys. Rev. B*, 74:100509, Sep 2006. doi: 10.1103/PhysRevB.74.100509. URL <http://link.aps.org/doi/10.1103/PhysRevB.74.100509>.
- [5] Lorne A. Page. Effect of earth’s rotation in neutron interferometry. *Phys. Rev. Lett.*, 35:543–543, Aug 1975. doi: 10.1103/PhysRevLett.35.543. URL <http://link.aps.org/doi/10.1103/PhysRevLett.35.543>.
- [6] S. A. Werner, J. L. Staudenmann, and R. Colella. Effect of earth’s rotation on the quantum mechanical phase of the neutron. *Phys. Rev. Lett.*, 42:1103–1106, Apr 1979. doi: 10.1103/PhysRevLett.42.1103. URL <http://link.aps.org/doi/10.1103/PhysRevLett.42.1103>.
- [7] Emile Hoskinson, Yuki Sato, Konstantin Penanen, and Richard Packard. A chemical potential “battery” for superfluid ^4He weak links. *AIP Conference Proceedings*, 850(1):117–118, 2006. doi: 10.1063/1.2354630. URL <http://link.aip.org/link/?APC/850/117/1>.
- [8] Yuki Sato, Aditya Joshi, and Richard Packard. Direct measurement of quantum phase gradients in superfluid ^4He flow. *Phys. Rev. Lett.*, 98:195302, May 2007. doi: 10.1103/PhysRevLett.98.195302. URL <http://link.aps.org/doi/10.1103/PhysRevLett.98.195302>.

- [9] Yuki Sato, Aditya Joshi, and Richard Packard. Flux locking a superfluid interferometer. *Applied Physics Letters*, 91(7):074107, 2007. doi: 10.1063/1.2772659. URL <http://link.aip.org/link/?APL/91/074107/1>.
- [10] Yuki Sato. Fiske-amplified superfluid interferometry. *Phys. Rev. B*, 81:172502, May 2010. doi: 10.1103/PhysRevB.81.172502. URL <http://link.aps.org/doi/10.1103/PhysRevB.81.172502>.
- [11] Aditya Joshi and Richard Packard. A continuously operating, flux locked, superfluid interferometer. *Journal of Low Temperature Physics*, pages 1–13, 2013. ISSN 0022-2291. doi: 10.1007/s10909-013-0857-y. URL <http://dx.doi.org/10.1007/s10909-013-0857-y>.
- [12] Supradeep Narayana and Yuki Sato. Superfluid quantum interference in multiple-turn reciprocal geometry. *Phys. Rev. Lett.*, 106:255301, Jun 2011. doi: 10.1103/PhysRevLett.106.255301. URL <http://link.aps.org/doi/10.1103/PhysRevLett.106.255301>.
- [13] Yuki Sato, Aditya Joshi, and Richard Packard. Superfluid ^4He quantum interference grating. *Phys. Rev. Lett.*, 101:085302, Aug 2008. doi: 10.1103/PhysRevLett.101.085302. URL <http://link.aps.org/doi/10.1103/PhysRevLett.101.085302>.
- [14] Supradeep Narayana and Yuki Sato. Quantum coherence in a superfluid josephson junction. *Phys. Rev. Lett.*, 106:055302, Feb 2011. doi: 10.1103/PhysRevLett.106.055302. URL <http://link.aps.org/doi/10.1103/PhysRevLett.106.055302>.
- [15] Packard group research website. <http://www.physics.berkeley.edu/research/packard/>. Accessed: 2/10/2013.
- [16] D.R. Tilley and J Tilley. *Superfluidity and superconductivity*. Graduate student series in physics. IOP publishing Ltd, Bristol, Philadelphia, 3rd edition, 1990.
- [17] Christian Enss and Siegfried Hunklinger. *Low-temperature physics*. Springer-Verlag, Berlin, Heidelberg, 2005.
- [18] Tony Guenault. *Basic Superfluids*. CRC press, 2002.
- [19] L. Tisza. Transport phenomena in helium ii. *Nature*, 141:913, May 1938. doi: 10.1038/141913a0.
- [20] L. D. Landau and E. M. Lifshitz. *Fluid Mechanics*, volume 6 of *Course of Theoretical Physics*. Pergamon Press, New York, 2nd edition, 1987.
- [21] E L Andronikashvilli. A direct observation of 2 kinds of motion in Helium-II. *Zhurnal Eksperimentalnoi I Teoreticheskoi Fiziki*, 16(9):780–785, 1946. ISSN 0044-4510.

- [22] Evgeni Burovski, Jonathan Machta, Nikolay Prokof'ev, and Boris Svistunov. High-precision measurement of the thermal exponent for the three-dimensional xy universality class. *Phys. Rev. B*, 74:132502, Oct 2006. doi: 10.1103/PhysRevB.74.132502. URL <http://link.aps.org/doi/10.1103/PhysRevB.74.132502>.
- [23] B.D. Josephson. Possible new effects in superconductive tunnelling. *Physics Letters*, 1(7):251 – 253, 1962. ISSN 0031-9163. doi: 10.1016/0031-9163(62)91369-0. URL <http://www.sciencedirect.com/science/article/pii/0031916362913690>.
- [24] P. W. ANDERSON. Considerations on the flow of superfluid helium. *Rev. Mod. Phys.*, 38:298–310, Apr 1966. doi: 10.1103/RevModPhys.38.298. URL <http://link.aps.org/doi/10.1103/RevModPhys.38.298>.
- [25] B.S Deaver Jr. and J.M Pierce. Relaxation oscillator model for superconducting bridges. *Physics Letters A*, 38(2):81 – 82, 1972. ISSN 0375-9601. doi: 10.1016/0375-9601(72)90496-3. URL <http://www.sciencedirect.com/science/article/pii/0375960172904963>.
- [26] Raymond Wiley Simmonds. *Josephson Weak Links and Quantum Interference in Superfluid ^3He* . PhD thesis, University of California at Berkeley, 2002. http://www.physics.berkeley.edu/research/packard/past_research/theses/RSimmondsThesis.pdf.
- [27] Talso Chui, Warren Holmes, and Konstantin Penanen. Fluctuations of the phase difference across an array of josephson junctions in superfluid ^4He near the lambda transition. *Phys. Rev. Lett.*, 90:085301, Feb 2003. doi: 10.1103/PhysRevLett.90.085301. URL <http://link.aps.org/doi/10.1103/PhysRevLett.90.085301>.
- [28] David Pekker, Roman Barankov, and Paul M. Goldbart. Phase-slip avalanches in the superflow of ^4He through arrays of nanosize apertures. *Phys. Rev. Lett.*, 98:175301, Apr 2007. doi: 10.1103/PhysRevLett.98.175301. URL <http://link.aps.org/doi/10.1103/PhysRevLett.98.175301>.
- [29] Emile Hoskinson. personal communication.
- [30] E Varoquaux and G Varoquaux. Sagnac effect in superfluid liquids. *Physics-Uspekhi*, 51(2):205, 2008. URL <http://stacks.iop.org/1063-7869/51/i=2/a=L09>.
- [31] R. Anderson, H. R. Bilger, and G. E. Stedman. “Sagnac” effect: A century of Earth-rotated interferometers. *American Journal of Physics*, 62:975–985, November 1994. doi: 10.1119/1.17656.
- [32] B H W Hendricks and G Nienhuis. Sagnac effect as viewed by a co-rotating observer. *Quantum Optics: Journal of the European Optical Society Part B*, 2(1):13, 1990. URL <http://stacks.iop.org/0954-8998/2/i=1/a=002>.

- [33] Y. Aharonov and D. Bohm. Significance of electromagnetic potentials in the quantum theory. *Phys. Rev.*, 115:485–491, Aug 1959. doi: 10.1103/PhysRev.115.485. URL <http://link.aps.org/doi/10.1103/PhysRev.115.485>.
- [34] Frederick G. Werner and Dieter R. Brill. Significance of electromagnetic potentials in the quantum theory in the interpretation of electron interferometer fringe observations. *Phys. Rev. Lett.*, 4:344–347, Apr 1960. doi: 10.1103/PhysRevLett.4.344. URL <http://link.aps.org/doi/10.1103/PhysRevLett.4.344>.
- [35] Martin Wilkens. Quantum phase of a moving dipole. *Phys. Rev. Lett.*, 72:5–8, Jan 1994. doi: 10.1103/PhysRevLett.72.5. URL <http://link.aps.org/doi/10.1103/PhysRevLett.72.5>.
- [36] Haiqing Wei, Rushan Han, and Xiuqing Wei. Quantum phase of induced dipoles moving in a magnetic field. *Phys. Rev. Lett.*, 75:2071–2073, Sep 1995. doi: 10.1103/PhysRevLett.75.2071. URL <http://link.aps.org/doi/10.1103/PhysRevLett.75.2071>.
- [37] Yuki Sato and Richard Packard. On the feasibility of detecting an aharonov-bohm phase shift in neutral matter. *Journal of Physics: Conference Series*, 150(3):032093, 2009. URL <http://stacks.iop.org/1742-6596/150/i=3/a=032093>.
- [38] Emile Hoskinson and Richard Packard. Thermally driven josephson oscillations in superfluid ^4he . *Phys. Rev. Lett.*, 94(15):155303, 4 2005. doi: 10.1103/PhysRevLett.94.155303.
- [39] Emile Hoskinson. *Superfluid ^4He weak links*. PhD thesis, University of California at Berkeley, 2005. http://www.physics.berkeley.edu/research/packard/past_research/theses/Emile_thesis.pdf.
- [40] D. D. Coon and M. D. Fiske. Josephson ac and step structure in the supercurrent tunneling characteristic. *Phys. Rev.*, 138:A744–A746, May 1965. doi: 10.1103/PhysRev.138.A744. URL <http://link.aps.org/doi/10.1103/PhysRev.138.A744>.
- [41] R. W. Simmonds, A. Loshak, A. Marchenkov, S. Backhaus, S. Pereversev, S. Vitale, J. C. Davis, and R. E. Packard. dc supercurrents from resonant mixing of josephson oscillations in a ^3he weak link. *Phys. Rev. Lett.*, 81:1247–1250, Aug 1998. doi: 10.1103/PhysRevLett.81.1247. URL <http://link.aps.org/doi/10.1103/PhysRevLett.81.1247>.
- [42] Ho Jung Paik. Superconducting tunable-diaphragm transducer for sensitive acceleration measurements. *Journal of Applied Physics*, 47(3):1168–1178, 1976. doi: 10.1063/1.322699. URL <http://link.aip.org/link/?JAP/47/1168/1>.
- [43] Ho Jung Paik. *Analysis and Development of a Very Sensitive Low Temperature Gravitational Radiation Detector*. PhD thesis, STANFORD UNIVERSITY., 1975.

- [44] Yuki Sato and Richard Packard. dc superconducting quantum interference device based neodymium magnet displacement sensor for superfluid experiments. *Review of Scientific Instruments*, 80(5):055102, 2009. doi: 10.1063/1.3129942. URL <http://link.aip.org/link/?RSI/80/055102/1>.
- [45] Yuki Sato and Richard Packard. On the feasibility of detecting an aharonov-bohm phase shift in neutral matter. *Journal of Physics: Conference Series*, 150(3):032093, 2009. URL <http://stacks.iop.org/1742-6596/150/i=3/a=032093>.
- [46] Yuki Sato. *Experiments using ^4He weak links*. PhD thesis, University of California at Berkeley, 2007. http://www.physics.berkeley.edu/research/packard/past_research/theses/Yuki_Sato_thesis.pdf.
- [47] Aditya Joshi, Yuki Sato, and Richard Packard. Fabricating nanoscale aperture arrays for superfluid helium-4 weak link experiments. *Journal of Physics: Conference Series*, 150(1):012018, 2009. URL <http://stacks.iop.org/1742-6596/150/i=1/a=012018>.
- [48] B.P. Beecken and W. Zimmermann Jr. The fabrication of 100 to 200-nm-diameter orifices in free-standing 100 and 200-nm-thick foils. *J. Vac. Sci. Technol.*, 3:1839, 1985.
- [49] P. Sudraud, P. Ballongue, E. Varoquaux, and O. Avenel. Focused ion-beam milling of a submicrometer aperture for a hydrodynamic josephson-effect experiment. *J. Appl. Phys.*, 62:2163, 1987.
- [50] K. S. Ralls, R. A. Buhrman, and R. C. Tiberio. Fabrication of thin-film metal nanobridges. *Appl. Phys. Lett.*, 55:2459, 1989.
- [51] A. Amar, R.L. Lozes, Y. Sasaki, J.C. Davis, and R.E. Packard. Fabrication of sub-micron apertures in thin membranes of silicon nitride. *J. Vac. Sci. Technol.*, 11:259, 1993.
- [52] Alexander Loshak. *Development of and experiments with superfluid ^3He weak links*. PhD thesis, University of California, Berkeley, 1999.
- [53] Gary H. Bernstein and Davide A. Hill. On the attainment of optimum developer parameters for pmma resist. *Superlattices and Microstructures*, 11(2):237 – 240, 1992. ISSN 0749-6036. doi: 10.1016/0749-6036(92)90260-C. URL <http://www.sciencedirect.com/science/article/pii/074960369290260C>.
- [54] John R. Vig. Uv/ozone cleaning of surfaces. *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, 3(3):1027–1034, 1985. doi: 10.1116/1.573115. URL <http://link.aip.org/link/?JVA/3/1027/1>.
- [55] D. S. Campbell. *Handbook of thin film technology*, page 12. McGraw-Hill, New York, 1970.

- [56] R. W. Hoffman. *Physics of nonmetallic thin films*, page 273. Plenum Press, New York, 1976.
- [57] C. W. Li, C. C. Toh, J. K. Fu, and K. Lahiri. Effect of thermal cycling on stresses in thin silicon nitride and silicon carbide films. *J. Mater. Sci. Lett.*, 15:2177, 1996.
- [58] Quentin Diduck. personal communication, 2008. Cornell University.
- [59] Neils Bruckner. *Development of a superfluid ^4He phase slip gyroscope*. PhD thesis, University of California at Berkeley, 2002.
- [60] Keith Charles Schwab. *Experiments with superfluid oscillators : design and microfabrication of a superfluid gyroscope ; modulation of a ^4He rf SQUID Analog by the Earth's rotation ; vortex nucleation in superfluid ^4He* . PhD thesis, University of California, Berkeley, 1996.
- [61] Robert Richardson and Eric Smith, editors. *Experimental techniques in condensed matter physics at low temperatures*. Addison-Wesley, Redwood City, California, 1998.
- [62] Keith Schwab. personal communication, 2012.
- [63] Emile Hoskinson and Richard Packard. Calibration technique for superfluid ^4He weak-link cells based on the fountain effect. *AIP Conference Proceedings*, 850(1):119–120, 2006. doi: 10.1063/1.2354631. URL <http://link.aip.org/link/?APC/850/119/1>.
- [64] M. Savard, C. Tremblay-Darveau, and G. Gervais. Flow conductance of a single nanohole. *Phys. Rev. Lett.*, 103:104502, Sep 2009. doi: 10.1103/PhysRevLett.103.104502. URL <http://link.aps.org/doi/10.1103/PhysRevLett.103.104502>.
- [65] Geoffrey Nunes Jr. Sec. 2.4: Pumps and plumbing. In Robert Richardson and Eric Smith, editors, *Experimental techniques in condensed matter physics at low temperatures*, pages 55–60. Addison-Wesley, Redwood City, California, 1998.
- [66] P. Clausing. The flow of highly rarefied gases through tubes of arbitrary length. *Journal of Vacuum Science and Technology*, 8(5):636–646, 1971. doi: 10.1116/1.1316379. URL <http://link.aip.org/link/?JVS/8/636/1>. Translated from the German [Ann. Physik (5) 12, 961 (1932)].
- [67] Saul Dushman. Recent advances in the production and measurement of high vacua. *Journal of the Franklin Institute*, 211(6):689–750, 1931. ISSN 0016-0032. doi: 10.1016/S0016-0032(31)90550-4. URL <http://www.sciencedirect.com/science/article/B6V04-49WPXK1-1BJ/2/666a952377f36ad0123eff34edcacedc>.
- [68] L. Marino. Experiments on rarefied gas flows through tubes. *Microfluidics and Nanofluidics*, 6(1):109–119, 2008. doi: 10.1007/s10404-008-0311-7.

- [69] A. K. Sreekanth. Transition flow through short circular tubes. *Physics of Fluids*, 8(11):1951–1956, 1965. doi: 10.1063/1.1761142. URL <http://link.aip.org/link/?PFL/8/1951/1>.
- [70] Mathematica. *Version 4.1*. Wolfram Research, Inc., Champaign, Illinois, 2001.
- [71] Frank Pobell. *Matter and Methods at Low Temperatures*. Springer, 1991.
- [72] James S. Brooks and Russell J. Donnelly. The calculated thermodynamic properties of superfluid helium-4. *Journal of Physical and Chemical Reference Data*, 6(1):51–104, 1977. doi: 10.1063/1.555549. URL <http://link.aip.org/link/?JPR/6/51/1>.
- [73] S. Narayana and Y. Sato. Compact miniature high-resolution thermometer. *Applied Superconductivity, IEEE Transactions on*, 20(6):2402–2405, dec. 2010. ISSN 1051-8223. doi: 10.1109/TASC.2010.2087021.
- [74] Paul B. Welander and Inseob Hahn. Miniature high-resolution thermometer for low-temperature applications. *Review of Scientific Instruments*, 72(9):3600–3604, sep 2001. ISSN 0034-6748. doi: 10.1063/1.1396663.
- [75] I. R. Walker. Removal of enamel from ultrafine monofilamentary wires. *Review of Scientific Instruments*, 75(5):1169–1174, 2004. doi: 10.1063/1.1666992. URL <http://link.aip.org/link/?RSI/75/1169/1>.
- [76] Tom Irvine. Bending frequencies of beams, rods, and pipes. URL: <http://www.vibrationdata.com/beams.htm>.
- [77] Simulation Multiphysics. *version 2012*. Autodesk, Inc., San Rafael, California, 2012.
- [78] Pat R. Roach. Kevlar support for thermal isolation at low temperatures. *Review of Scientific Instruments*, 63(5):3216–3217, 1992. doi: 10.1063/1.1142582. URL <http://link.aip.org/link/?RSI/63/3216/1>.
- [79] L. Duband, L. Hui, and A. Lange. Thermal isolation of large loads at low temperature using kevlar rope. *Cryogenics*, 33(6):643–647, 1993. ISSN 0011-2275. doi: 10.1016/0011-2275(93)90124-7. URL <http://www.sciencedirect.com/science/article/pii/0011227593901247>.
- [80] Neils Bruckner, Scott Backhaus, and Richard Packard. An improved low temperature valve. *Czechoslovak Journal of Physics*, 46(5):2741–2742, 1996. ISSN 0011-4626. doi: 10.1007/BF02570357. 21st International Conference on Low Temperature Physics (LT 21), Prague, Czech Republic, Aug 08-14, 1996.
- [81] W. P. Kirk and M. Twerdochlib. Improved method for minimizing vibrational motion transmitted by pumping lines. *Review of Scientific Instruments*, 49(6):765–769, 1978. doi: 10.1063/1.1135609. URL <http://link.aip.org/link/?RSI/49/765/1>.

- [82] LABVIEW. *version 8.2.1*. National Instruments Corporation, Austin, Texas, 2006.
- [83] Robert D. Mc Carty. Thermodynamic properties of helium 4 from 2 to 1500 k at pressures to 10⁸ pa. *Journal of Physical and Chemical Reference Data*, 2(4): 923–1042, 1973. doi: 10.1063/1.3253133. URL <http://link.aip.org/link/?JPR/2/923/1>.
- [84] O. Avenel, P. Hakonen, and E. Varoquaux. Detection of the rotation of the earth with a superfluid gyrometer. *Phys. Rev. Lett.*, 78:3602–3605, May 1997. doi: 10.1103/PhysRevLett.78.3602. URL <http://link.aps.org/doi/10.1103/PhysRevLett.78.3602>.
- [85] Yuki Sato and Richard Packard. On detecting absolute quantum mechanical phase differences with superfluid interferometers. *Physica E: Low-dimensional Systems and Nanostructures*, 43(3):702 – 706, 2011. ISSN 1386-9477. doi: 10.1016/j.physe.2010.07.033. URL <http://www.sciencedirect.com/science/article/pii/S1386947710004236>.
- [86] Yuki Sato. personal communication.
- [87] Yuki Sato, Aditya Joshi, and Richard Packard. Observation of an extrinsic critical velocity using matter wave interferometry. *Phys. Rev. B*, 76:052505, Aug 2007. doi: 10.1103/PhysRevB.76.052505. URL <http://link.aps.org/doi/10.1103/PhysRevB.76.052505>.
- [88] Albert A. Michelson and Edward W. Morley. On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, 34:333–345, 1887.
- [89] Andrew Moulthrop. *Tests of Macroscopic Wave Function Theories of Helium II*. PhD thesis, UNIVERSITY OF CALIFORNIA, BERKELEY., 1984.
- [90] Alexander L. Fetter. Vortex nucleation in deformed rotating cylinders. *Journal of Low Temperature Physics*, 16:533–555, 1974. ISSN 0022-2291. URL <http://dx.doi.org/10.1007/BF00654901>. 10.1007/BF00654901.
- [91] A. Th. A. M. De Waele, W. H. Kraan, and R. De Bruyn Ouboter. A superconducting interference grating. *Physica*, 40(2):302 – 308, 1968. ISSN 0031-8914. doi: DOI:10.1016/0031-8914(68)90027-X. URL <http://www.sciencedirect.com/science/article/pii/003189146890027X>.
- [92] J.T. Jeng, K.H. Huang, C.H. Wu, K.L. Chen, J.C. Chen, and H.C. Yang. Characteristics of the superconducting-quantum- interference-grating magnetometer consisting of grain-boundary josephson junctions. *Applied Superconductivity, IEEE Transactions on*, 17(2):691 –694, june 2007. ISSN 1051-8223. doi: 10.1109/TASC.2007.898270.

- [93] E. T. Kornhauser and D. Mintzer. On the vibration of mass-loaded membranes. *The Journal of the Acoustical Society of America*, 25(5):903–906, 1953. doi: 10.1121/1.1907216. URL <http://link.aip.org/link/?JAS/25/903/1>.
- [94] Philip M. Morse. *Vibration and sound*. American Institute of Physics, 1976.
- [95] Harold L. Weissberg. End correction for slow viscous flow through long tubes. *Physics of Fluids*, 5(9):1033–1036, 1962. doi: 10.1063/1.1724469. URL <http://link.aip.org/link/?PFL/5/1033/1>.
- [96] Richard Feynman. *The Feynman Lectures on Physics: Volume 3*. Addison-Wesley, Boston, 1963.
- [97] Yu. B. Rumer and M. Sh. Ryvkin. *Thermodynamics, Statistical Physics and Kinetics*. Mir Publishers, Moscow, 1980.
- [98] Alexandre J. Chorin and Jerrold E. Marsden. *A mathematical introduction to fluid mechanics*. New York : Springer-Verlag, 1993.
- [99] MATLAB. *version 7.9 (R2009b)*. The MathWorks Inc., Natick, Massachusetts, 2009.
- [100] G. B. Hess and W. M. Fairbank. Measurements of angular momentum in superfluid helium. *Phys. Rev. Lett.*, 19:216–218, Jul 1967. doi: 10.1103/PhysRevLett.19.216. URL <http://link.aps.org/doi/10.1103/PhysRevLett.19.216>.

Appendix A

Engineering drawings, fabrication and assembly

This chapter includes engineering drawings for several parts that do not appear in the main body of the dissertation. The first section also details the fabrication of the some of the parts that are especially difficult or tricky. Some custom tools and new techniques are also provided to assist in proper fab and assembly of the parts.

A.1 Fabrication: SHeQUID structural parts

The design philosophy, constraints and essential parts for the modular SHeQUID were described in Chapter 3. We addressed the two most complicated components of the SHeQUID – the aperture arrays and the superconducting displacement sensor – in chapters 4 and 5, respectively. We now turn to the fabrication¹ and finishing of the remaining (largely mechanical) components. Note that since most researchers get parts fabricated from professional machine shops, we will only provide such details if they are not deducible from the drawings. Finishing (polishing, etc.) is described in its own section. We will end this sequence with a description of the final assembly of the SHeQUID.

Note that many of the parts and techniques (including assembly steps) discussed here are also relevant to single weak-link cells. Exploded views of scale drawings of the assembled SHeQUID were provided in Figs. 3.10 and 3.11. Engineering drawings for SHeQUID parts and drawings for single weak-link cell parts can also be found in this chapter.

Note also that all dimensions used are Imperial² (inches) and all screw thread systems are Unified Thread Standard (UTS)³.

¹We are grateful to Yuki Sato and the lab notebooks of Emile Hoskinson for the design basis on which we have developed some of the newer techniques and tricks described here.

²A consequence of the machine shop conventions in the United States

³Commonly used in Canada and the United States, typically UNC or UNF. Fortunately, conversions and information are readily available on the web.

Finally, the engineering drawings included are all vector graphics and preserve all structure so that they can be viewed (on screen) in detail at arbitrary zoom levels⁴. For this reason (and due to the large amounts of information contained in the drawings without adding appreciably to file size), they are likely to be much more useful in electronic form (but still largely legible in print). This was a practical choice as optimizing this chapter for print would have increased the length of this dissertation to an intolerable extent.

A.1.1 The inner cell piece

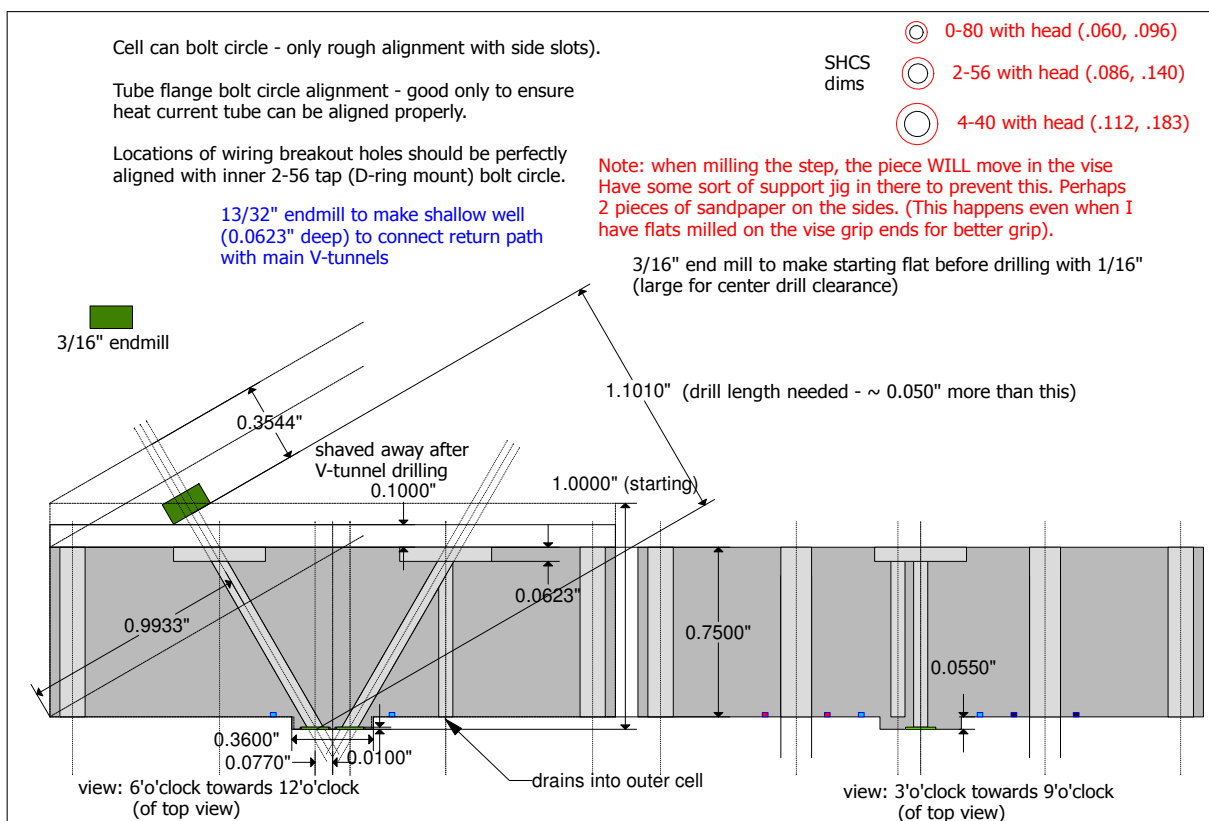


Figure A.1: Inner cell piece side view

This is machined from aluminum. There was a time when a very specific alloy of aluminum was used to match the thermal contraction of black Stycast and silicon chips, but we have found over the years that this is not necessary. See Figs. A.1, A.2, A.10, A.11 and A.12 for reference. This is probably the most complicated part to make. Note that the dimensions used in the side-view assume a certain method of fabrication (discussed here) and may need changes if a different path is followed.

⁴Having said that, all CAD files are freely available upon request from the author (because reinventing the wheel is *not* conducive to progress).




Figure A.2: Inner cell piece (ICP) top and bottom views

We start with an aluminum disk of the maximum size shown and turn the bottom view features on a lathe (chip-holder island, indium seal groove and rough polish). Proper alignment is done with a dial indicator at all times. Two precise flats are then milled on the curved side and the piece is then held in a mill vise at a 60° angle from the vertical using an angle block with the bottom view side down (the flats help hold the piece in place – equally thick pieces of sandpaper can be used to hold it more securely). A small ledge is cut on the top side with an end-mill to provide a flat surface for drilling the V-tunnel. The process is repeated for the other tunnel. Obviously, there are multiple ways to do the V-tunnel – this is merely the easiest way we could find, given the quality and precision of the machines available to us. After this point, the side flats can be used to properly align the piece with our coordinate system and the rest of the features cut straightforwardly.

The astute reader/machinist will notice that the bolt holes for the indium seals are not arranged in a nice, equispaced circle. We are fortunate to have freed ourselves early on from the tyranny of the even bolt circle. Since the precision of a dividing head is usually worse or no better than that of a nice, digital XY indicator, such constraints are largely a historical convention and unnecessary at present. Removing this constraint provides great flexibility in creating compact cell designs. As long as the bolt spacing is not wildly variable

and the arrangement somewhat symmetric around the seal pattern, indium seals tend to be extremely robust and forgiving. A happy consequence of this design choice is that we can choose bolt locations that are nice, round figures for rapid dialing on a digital indicator.

A.1.2 D-ring




Figure A.3: D-ring

See Figs. A.3 and A.13 for reference. The critical cut here is the angled tunnel for the inner cell heater leads. A smaller hole is drilled for the (4.2 mil SC) twisted pair leads and a larger one drilled on the outside to fit a lead tube for shielding as far down as possible. This angled cut is made in the same way that we did the V-tunnels for the ICP. The extra starting thickness here is designed for the same purpose. Different machining methods would need some reworking of the design in just this respect.

Another important feature we introduced in the most recent incarnation of the cell is the annular ledge in the main circular cavity. This ledge has two functions: (i) to serve as a gluing platform for the inner cell heater (on the side facing the chips), and (ii) to shield

the pancake coil (at least to some extent) from direct EM coupling to the cell heater and its leads.

Brass tubes are soft-soldered into the breakout holes as shown. These should be done right after most of the machining work is done. Polishing can be done after soldering. We have found that making the step for the tubes slightly undersized and then press-fitting the tubes in prior to soldering significantly simplifies soldering (which can be subsequently done on a hotplate using the blue flux and flux-less solder from Section 8.2.4).

After this, we come to another new feature that requires some explanation. Making the screw joints between a wire and the diaphragm (or electrode) tab can sometimes rip through the Kapton and short the tabs to the cell body. To guard against this, we have been using a rather simple trick of making an embedded washer out of potted black Stycast (2850FT). In the hollow step indicated on the drawing, Stycast is poured (scotch tape can be used to mold open wells) and after setting, machined or sanded flush with the surface and screw threads tapped in (for the tab joint). Now, diaphragms can be glued flush over this and joints can be made without fear of shorts.

A.1.3 E-ring

See Fig. A.4 for reference. The E-ring (made of brass) is the simplest piece and requires no special instructions. Note the epoxy well on the side-wall for making a screw joint between the electrode tab and its wire.

A.1.4 Cell can

See Fig. A.5 for reference. This is yet another straightforward piece (made of brass). The only thing to be noted here is the depth of the cell can, which can be made much shorter given the other parts described here. In our most recent run, we reduced the depth of this can by 0.875" using a special aluminum insert. One should probably start with this modified depth in future cells.

The cell fill line is another special feature. A small "island" is made on the can's wall (see drawing) by milling a "moat" around the fill line hole. A Cu-Ni capillary is inserted through the hole and soft-soldered in place – the moat reducing the heat loss to the large thermal mass of the can. It still requires a small butane torch to make this joint. The cell can might be lead plated for shielding the electrical components. In that case, the soldering should be done with a clean Cu-Ni capillary *after* plating and thoroughly cleaning out the can. Although we invariably lead-plate the can, we still use lead sheets to fashion a shield around the can, so the plating can be optional.

A.1.5 Wiring breakouts

The design of our wiring breakouts has changed considerably over the years. Several of the early experiments in our group used essentially the same basic idea of sliding a cap machined




Figure A.4: E-ring

out of Stycast 1266 onto a tube emerging from the cell can that carried the wires. The wires would go through small holes in the cap, that were subsequently sealed with fresh Stycast. Teflon capillaries could be used as strain relief outside the cell can and these wires would subsequently be connected to a shielded station with wiring terminals loosely held onto the cap with set screws. These terminals would provide an interface to the outside world. This worked just fine in the past, but we decided to modify the design to attempt to reduce cross-talk between the wires by reducing the length during which the wires would “see” each other unshielded (by lead tubes). The other disadvantage we perceived with the old design was that superconducting joints to the pancake coil leads were made by spark-welding a very short length of leads *after sealing the wires into the breakout cap* just outside the cell. At that point, there is no room for error and considering the violence of a spark-weld so close to the assembled cell, it seemed unnecessarily dangerous.

Since then, we have been using breakout seals that work the same way but have additional layered steps to enable sealing in lead tubes on both the inside and outside of the cell (see, for instance, our breakout designs for the modular single-weak link cell in Figs. A.32 and A.33). Since our modular cells do not require dismantling these breakouts each time we change a chip or sense arm, the wires from the cell are now kept continuous (and continuously shielded




Figure A.5: Cell can




Figure A.6: Wiring breakout

with lead tubes) all the way to their final destination (the capacitance joint box, PI box and cell and sense arm heater joint boxes).

Fig. A.6 represents the most recent design, with a slight increase in complexity (for conventional machining only) and the added benefit of improved strain relief on both sides of the breakout as well as a dramatically reduced unshielded gap (where the wires just have to be sealed with Stycast) whose impact is further minimized by making these gaps staggered for breakouts with multiple wires. Since Stycast 1266 is soft and very easy to cut, we could make the deep, inside cut (to slip the cap over a tube) with a rather long, custom-ground groove tool (that would never survive in metal-work). However, these pieces are most easily formed by potting aluminum molds with Stycast and subsequently etching the aluminum away using NaOH (see the next section for details). The piece is designed to be as short as possible (to reduce cell can depth), which is the main reason for the extra length inside the hollow cap (rather than just raising the upper length to compensate and keeping the cap that fits over the tube just a simple hollow). For breakouts that are outside the cell, this need not be done, and the single weak-link cell breakout shown here would be a much better choice (it can be easily machined from stock).

A.1.6 The sense arm




Figure A.7: Sense arm tube flange and side-arm




Figure A.8: Sense arm

Tube flange

The reusable part of the sense arm is the tube flange shown in Fig. A.7. The two brass disks are straightforward to machine. Indium seal surfaces should be polished on a lathe to fine grit (1500+). The stainless steel tube sections that make up the side-arms should be cut as equal in length as possible and their ends smoothed and flush. In line with our discussion in Section A.1.2, making the soldering step in the brass flange slightly undersized and press-fitting the steel tube into it can help keep things properly aligned. Note here that alignment is particularly important here, since there are multiple pieces that are glued together to make up the sense arm. We have come up with a simple trick to do this alignment. We mount the brass flange in a lathe with its face properly aligned orthogonal to the lathe axis. The steel tube to be press-fit is mounted gently in a drill chuck facing the flange. The drill chuck is carefully guided into the flange hole and forced inside with gentle, manual rotations of the lathe. This should keep the tube perfectly normal to the flange face. Alternatively, the brass flange can be gently heated on a hotplate and the steel tube immediately pushed into the central hole after being dipped in LN2. A small chamfer on the hole edge should be available for the solder to flow freely.




Figure A.9: Sense arm mold

These pieces should be hard-soldered so that they can be indefinitely reused with new and different Stycast pieces. We note in passing that the raised lip on the bottom of the flange fits loosely into the matching wells on the top side of the ICP (see Fig. 3.10). This feature was included to ensure that the squeezed indium from the flange seals did not block the return path on one of the flanges (this would be – in technical terms – “a very bad thing”). The Stycast parts that are glued to these flanges can be easily removed by light heating with a small butane torch, after which they can be used for other sense arms.

Heat-pipe

We have made this heat-pipe (main design shown in Fig. A.8) using two different methods in the past. One is simply a brute force approach, drilling the heat-pipe out with a long drill (readily available). However, Stycast 1266 is not very heat-resistant and must be thoroughly lubricated whilst cutting into. The long tunnel implies a significant problem with getting lubricant to the cutting tip of the drill. Cutting dry at any point means that the Stycast can (and does) form micro-bubbles or even larger cracks and defects on the surface of the tunnel. While we are given to understand that the resulting roughness should be much larger than

the dimensions over which vortices are nucleated, such defects might make the flow through the heat-pipe turbulent or at least distorted (compared to modeled flow patterns) at lower flow speed limits than we would otherwise expect.

We have therefore taken the path of making the main heat-pipe in an aluminum mold. We find that this process is much simpler and its results far more optimal than conventional machining. Since we recommend using this method at several places in the text, we will describe this in some detail. Fig. A.9 shows the aluminum mold that we use to pot the Stycast in. A commercially available aluminum tube (polished and cleaned) is slid inside the main mold body and the entry points sealed shut with simple super-glue. Stycast 1266 is poured into the mold and allowed to set. The Stycast-in-mold is then trimmed to a length of 2.5" and the top, uneven surface trimmed by 50 mils (see drawing). The metal mold helps us locate our coordinate system precisely and the side-arm flow holes are now drilled to connect to the main heat-pipe tunnel. At this point, the Stycast piece can be easily removed (via differential contraction) from the mold by dipping it in LN2. The heat-pipe end cuts can be made now (one side houses the copper sink plate and the other houses the module containing the sense arm heater).

Once all the cuts are done, we just have to etch out the aluminum tube from within the heat-pipe⁵. We have been informed that a 1 molar NaOH solution is optimal for aluminum etching in tight spaces (for lowered viscosity, we presume). We have done this successfully by etching the piece in a test tube of 1 M NaOH, placed inside a water bath that is heated on a hot plate, and it takes a few hours for the entire tube to be etched. The Stycast is impervious and doesn't get etched at all. However, it *is* sensitive to heat, so the temperature should only be warm to touch, not hot. Since the greatest impediment to the reaction seems to be precipitation on the aluminum, we recommend that this etching be done by immersing the test tube in an ultrasonic cleaner (which typically also has a heater in the bath) to prevent deposition on the metal surface. Note that this is why we used a tube and not a rod – to allow the etchant free access to the metal. With a rod, we would be etching at just the tiny end faces of the rod. The etchant may need to be replaced once in a while. Once etched, the tunnel should be mechanically cleaned (in addition to the usual cleaning steps) by dragging some cotton thread or a wooden applicator through. This concludes the heat-pipe fabrication.

Side-arms

The side-arms are just machined from stock. As with the tube flanges, the depths of the holes and steps are properly matched between the two copies. Steps are drilled with normal drills but finished with flat drills. Flat drills are not that easy to find but can be manually ground from regular drills.

⁵A good mold release agent can also be sprayed on the tube and dried prior to pouring in the Stycast. Depending on how the release is activated, it may be easy or difficult to actually remove the tube this way.

Heater module

The heater module shown in the drawing is a new design meant for the thin-film chip resistors we have recently begun to use. A #68 hole is drilled along its axis for the 4.2 mil SC twisted pair heater leads to exit the module (later sealed off with Stycast). A #55 tunnel is drilled as far down along this hole as possible to fit a lead tube for shielding the wires (this is sealed in along with the wire). The end that sits in the heat-pipe is made snug-fitting and a small notch is carved on the small end and filed by hand to seat the chip resistor without blocking (or gluing over) its two large faces.

A.1.7 More drawings

Additional (detail) drawings for the parts described previously are collected in this section.




Figure A.10: ICP detail view: chip slots

A.2 SHeQUID assembly

A.2.1 Part finishing and cleaning

All parts should be deburred (smoothed), especially any place that will be touched by wires. We use small reamers, files and sandpaper (sanding twigs and cord are also available) for this purpose. Parts to be lead-plated can be sandblasted for greater exposed area. All surfaces to be glued to should be sanded to around 1000 grit in a figure-8 pattern to prevent unevenness. For the D-ring, the two breakout tubes will prevent face-down sanding of the




Figure A.11: ICP detail view: coordinates




Figure A.12: ICP detail view: more coordinates




Figure A.13: D-ring detail view

diaphragm gluing surface, so we compromise in this case by using a smoothed plastic rod with a piece of sandpaper stuck on top with double-stick tape that is narrow enough to fit between the breakout tubes and rubbing the D-ring surface over this sandpaper strip. We then go in with finer sandpaper and scuff up the lines gently.

All parts, prior to assembly, are cleaned ultrasonically in our standard sequence of soap-water \rightarrow acetone \rightarrow isopropanol (IPA) baths and blow-dried with compressed air. The pieces are rinsed under tap water after the soap-water bath to remove any traces of soap and dipped in the successive baths without letting water dry on the surface. Water stains can be bad spots for gluing and could lead to leaks. Cleanliness of parts is crucial, as any contamination can easily get transferred to the aperture array chips during the experiment and block the holes. Small conductive debris (like metal dust) can get between the electrode




Figure A.14: Lead shield cutouts

and diaphragm and short the two together. There are all sorts of exciting possibilities of this flavor that one can cheerily contemplate for hours on end. Alternately, one can just clean the parts thoroughly.

Note that Stycast pieces should not be left in IPA for more than 10 minutes and they should not be cleaned in methanol or acetone at all.

A.2.2 Stycast considerations

As already mentioned, we use two flavors of Stycast extensively in our cells – 1266 (clear, parts A and B) and 2850FT (black, with catalyst 24LV)⁶. These are both 2-part epoxies that are used largely because they have been comprehensively tested in cryogenic applications over decades. We try not to buy too much so that we can afford to not use these (at least for joints that need to be superleak-tight) too far beyond their shelf lives. We mix them in small, clean, aluminum weighing dishes⁷ using sterile wooden applicators⁸. All applicators should be disposable and one must be on guard against contamination of the main stock. We degas the mixed epoxy in a polycarbonate bell jar hooked up to a vacuum pump. Mixing and degassing are done (rigorously) for at least 10 minutes (timing helps maintain discipline in this respect).

⁶Manufactured by Emerson & Cuming and distributed within the USA by Ellsworth Adhesives (<http://www.ellsworth.com>)

⁷Clean lab supplies for such chemistry applications can be obtained from places like Fisher Scientific.

⁸Puritan company

Stycast 1266: Part A is very clear and thin when new and gets cloudier and thicker over time. It can be rejuvenated by taking a small amount in an aluminum dish and heating it from below with a small heat gun. If this is done, one must wait till it cools down before mixing in part B as otherwise, the setting reaction will be accelerated.

Stycast 2850FT: This is thick and black and must be stirred each time so that the quartz particles (that are the main reason for its flexibility and strength at cryogenic temperatures) are evenly distributed instead of sinking to the bottom. It tends to crystallize over time, especially in colder weather, and can even form one big, hard lump over time. Refreshing this is not as easy as the 1266 because if we simply broke off a chunk and melted it, the quartz concentration would be unpredictable and probably substandard. Best practice is to not let this happen in the first place by storing it at the recommended ambient temperatures. Failing that, the entire can must be melted on a hotplate. We have read anecdotal reports of $\sim 60^{\circ}\text{C}$ being a good temperature, but in practice, we have found it useful to just heat it slowly at low heat while stirring with a thick wooden stick or a very clean metal rod all the time. We stop heating immediately when the entire volume has become liquid.

Potting: We make machinable stock from these epoxies by potting them in small plastic tubes (or other forms). Small quantities (around 10 gms) are usually fine, but for large quantities, the heat generated in the (exothermic) setting reaction can cause air bubbles and other intrinsic defects. To prevent this, we simply insert the potting tubes in putty or other stands and immerse them as much as possible in a water bath to conduct the generated heat away. Already mixed and degassed epoxy can be poured very gently (we have noticed that the process is not unlike that used for pouring soda in a glass without generating too much foam) into the molds and even injected from the bottom up using disposable syringes.

Viscosity: Both epoxies are pretty thin when freshly mixed and grow steadily more viscous over time. Any wire seals in small holes should be done with fresh Stycast to allow it to flow into tiny crevasses and make a robust seal. During assembly, we refer to at least three stages of viscosity – fresh, viscous and very viscous – where the latter two stages correspond to around 1.5 hrs and 3 hrs respectively.

A.2.3 Indium seals

We have multiple indium seals in our cell. The same techniques apply to all of them.

Surface prep: Both seal surfaces should be flush and all bolt holes and screw threads should be as precisely normal to the surface as possible to prevent relative tilt (and hence uneven pressure leading to leaks). All seal surfaces should be finely polished such that the polish striations are closed loops around the sealed off volume. This, more than the polishing grit, is what determines a robust seal as the indium conforms to these striations in the surface.

The surfaces should be inspected closely for any scratches that cross the closed loops, thus providing a leak path through the seal. For the ICP and D-ring off-axis seals, we first polish the entire surfaces on-axis (on a lathe). Then, we use custom-built mandrels (nylon threaded rod with a rubber or cork stopper impaled on one end) with small sandpaper disks of the right size (these can also be custom made using standard or home-made punches to cut them out of sheets) stuck on with double-stick tape, to polish the off-axis seals and create well-defined closed loop striations on them (see Fig. A.15). We use a low speed drill press for this purpose. This can take a while to set up the first time, but it's easy for subsequent uses. For a lathe capable of off-axis mounting, this need not be done, however, polishing the off-axis seals without disturbing the on-axis seal might be very difficult on a lathe. Note that it may take several passes with smaller and smaller grit sizes to attain a good finish. It is unlikely that an aesthetically pleasing polish is absolutely necessary for a robust indium seal – however, the high polish makes it much easier to see any radial scratches⁹.

Indium prep: Note that the indium wire used dictates the groove size in the drawings. To minimize waste, we can measure the wire lengths needed directly from the drawings. We also make circular or square metal forms with sizes such that the indium wire wrapped around them fits perfectly in the center of the groove (this dummy form should have a cross-section composed of segments running along the centers of the grooves that are pulled inward by half the wire width). We clean the wire by wiping gently with an IPA-wetted piece of lint-free paper (lens paper works well) and then wiping again on some lens paper with a bit of Apiezon (M or N) vacuum grease dabbed on. Most of the grease should be gently wiped off, leaving just a very thin surface layer to improve the indium flow during compression and (more to the point) enabling easy removal of the indium later. The wire is cut (preferably directly on the dummy forms) at a shallow angle using a clean blade such that the two ends mate together at an angle (see Fig. A.16). This prevents any leaks at the mating point. Two freshly cut ends joined together immediately can almost cold-weld together and form a nice, clean joint. We use stainless steel tweezers to manipulate the wire and press the ends together (and place the closed loops on their respective grooves).

Making the seal: With the indium wire sandwiched by the two seal surfaces (and sitting securely in its groove), the bolts are inserted and hand-tightened gently in a star pattern to avoid distortion and asymmetric indium extrusion. They are tightened the rest of the way in a similar manner using an Allen wrench. A torque wrench can also be used to ensure that we are tightening each screw to the same extent each time and stepping up the torque gradually by the same amount for each screw. Once this done a few times, we tend to develop a feel for it and no longer need the torque wrench. For designing purposes, we have found that #2-56 is the smallest screw size that can provide enough pressure for an indium seal without stripping the threads. We use this size for most of the seals used here. #0-80 is just too

⁹And heck, it just looks nice, so that's that.




Figure A.15: Three different custom-made mandrels used for off-axis seal polishing. From left to right: brass threaded rod with rubber stopper (sanding disk stuck on with double-side tape); brass threaded rod with cork stopper; nylon threaded rod with rubber stopper. The protruding rod lengths help stabilize the polishers in their respective holes. The nuts are added as stops against which the drill chuck jaws can rest. This is needed more for the nylon rods since these cannot be tightened in the drill chuck so much and can therefore slide in after a while.




Figure A.16: Cutting indium wire just right.

fine of a screw thread and we have had the threads get destroyed before a good seal could be made.

The most important thing here is to apply uniform pressure across the seal so that the indium gets extruded equally everywhere.

A.2.4 Cell and sense arm heaters

We discuss this separately because soldering leads to the new thin-film chip resistors¹⁰ that we have recently switched to (from the old metal film resistors) is highly non-trivial since these chips can be as small as a mm wide! The one used for the sense arm heater is slightly bigger (2 mm), but not by much.

We use 4.2 mil SC wire (NbTi with Cu-Ni cladding covered with formvar) with only the formvar stripped off. The 2 parallel wires are taped to a glass slide under slight tension with the strip boundaries aligned as shown in Fig. A.17(a). The wire spacing is kept slightly less than the resistor width so that the wires can gently hold the resistor in place. This is all done under a low power microscope. A tiny drop of solder paste (a suspension of microscopic globules of the usual Sn/Pb alloy in a liquid flux base)¹¹ is laid down on each joint using a piece of small gauge wire as an applicator. As long there's a small puddle around the joint, we should be fine (precise positioning of the paste is not required). A small butane torch with an attachment that converts it into a heat gun with a very narrow air stream is now aimed at the resistor. Just a few seconds of heating melts the solder globules and the activated flux forces it to flow to just the exposed Cu-Ni of the wire and the tinned terminals of the resistor. Once the joints are covered in nice beads, we stop heating immediately and trim off the excess leads. The result is shown in Fig. A.17(b). This may take some practice to get right. The wires are now carefully twisted together. We have tried making a twisted pair prior to soldering. However, since we need at least an inch of lead length for soldering, twisting the final inch once everything's done is surprisingly difficult owing to the considerable elasticity of this SC wire. It also puts tremendous strain on the delicate solder joints and can easily break them. So, we recommend twisting the wires only after the soldering step and that too, gently by hand. Once a few inches have been twisted, one can tape it down and twist the rest by motorized screwdriver as usual. In the case of the inner cell heater, one lead is bent around the resistor body so that both leads come out on the same side, in a direction parallel to its body length.

¹⁰The inner cell heater is a Model PFC-W0402LF-03-1001-B: $1\text{ k}\Omega \pm 0.1\%$, rated at 62.5 mW and 25 V (we have found it to be fine even at 225 mW if immersed in LHe). The package is coded as 0402, which is $1 \times 0.5 \times 0.35\text{ mm}$ in size.

The sense arm heater is a Model PFC-W0805LF-03-1: $1\text{ k}\Omega \pm 0.1\%$, rated at 250 mW and 100 V . The package is coded as 0805, which is $2 \times 1.25 \times 0.4\text{ mm}$ in size. Both are manufactured by International Resistive.

¹¹Model# KE1507-ND, Kester solder paste 63/37 with no-clean flux from Digi-Key Corp., Thief River Falls, Minnesota, USA.

We have tested this brand of resistors and find that it does indeed remain very stable under temperature changes (temperature coefficient spec of ~ 25 ppm/ $^{\circ}$ C). Still, any new resistor should always be checked using the 4-wire arrangement provided over the experimental temperature range to verify its stability. Small changes are not fatal, provided one tabulates the resistance vs. temperature (and it is a stable function) and uses the appropriate value for any heater power calculations.




Figure A.17: Viewed under a low-power microscope: (a) Unsoldered 4.2 mil SC wires taped onto glass slide with slight tension holding a chip resistor between them. (b) Resistor after soldering and trimming leads.

A.2.5 Assembly steps

Assembly steps for the SHeQUID are presented in order. See Figs. 3.10 and 3.11 for reference during assembly. Annotated photographs of several intermediate states may be found in Figs. A.18, A.19, A.20, A.21, A.22 and A.23 – we will refrain from referencing each photo separately during assembly steps. We start with all parts finished and cleaned as already described. All work is now done on a clean bench (laminar flow hood to keep dust out). Nitrile gloves are used to prevent part contamination and all tools and surfaces within the bench are wiped clean with IPA. All paper, cloth and anything else that can produce lint are kept away from the bench. Cleanroom supplies are recommended (wipes, applicators, swabs, etc.) This is all for the purpose of protecting the aperture array chips from contamination. See the section on Stycast considerations for definitions of its various stages of viscosity. Many of the gluing steps can be done in parallel and must be planned out properly. Each

gluing step should be allowed at least 8 hours to set strong enough to handle (overnight works well).

1. The pancake coil (PC) is glued into the E-ring by holding both pieces flush, face down on top of a flat piece of Teflon stock. Fresh black Stycast is painted on the former sides before inserting it into the E-ring cavity. The E-ring top surface must be kept very clean and the former should be very snug so that Stycast does not leak out and wet this surface. A small amount of viscous Stycast is used to secure the former from the back side and very viscous Stycast filled in to finish the job. A lead tube enclosing the PC leads is cut to the right length and embedded in this Stycast.
2. The (normal metal coated) electrode is now glued onto the E-ring top surface (metal side up). We do this by allowing the metal side to (static) cling to a clean glass slide and placing it gently onto the surface, which is wetted with (fresh) clear Stycast. The wetting can be done by dabbing small drops all over with a toothpick and spreading it with either a folded Kimwipe or a lint-free swab¹². The electrode tab is placed slightly to askew of the screw hole so that a joint can be made to a wire.

The electrode almost never falls perfectly in place, so we adjust its position gently with a pair of tweezers. We also lift up the electrode and lay it down gently to let the glue push out any air bubbles and enable flush gluing. This is important (not having air bubbles¹³) as the spacing between the diaphragm and electrode (and pancake coil) critically determines the displacement sensor sensitivity – recall that we use spacers that are just 76 μm thick and the electrode and diaphragm are themselves just 7.6 μm thick. The amount of Stycast used should also be minimized so that it does not add appreciably to the thickness. We would strongly recommend calibrating this technique on blank brass or aluminum pieces first to see how much of thin layer one can get away with and still form superleak-tight and mechanically robust seals.

3. Once dry, a *new* bladed X-acto knife (multiple blade sizes are available – small ones work best) is used to trim the excess Kapton from the electrode and cut away clearances for screw holes, etc. 4.2 mil SC wire with just the formvar stripped off, is inserted into a small piece of Cu-Ni capillary, which is then bent into a half circle before being crimped with a cross-hatched pair of pliers. This makeshift washer is carefully smoothed to avoid tearing the Kapton. A tiny piece of Teflon tubing is slid over the free end of the wire prior to sliding a piece of lead tube over it (this isolates the lead tube from the washer). A nylon screw is tightened over another stainless steel washer, which goes over the Cu-Ni washer we just made and the joint is complete. The half-circle washer should be placed in such a way that tightening the screw does not move the washer away from the tab. This completes the E-ring piece.

¹²Lint-free, polyester, cleanroom swabs (Micro Absorbond TX759B) can be obtained from ITW Texwipe, Kernersville, NC, USA.

¹³Smartphone users in the early 00's attempting to install screen protectors will empathize with this frustrating issue (although it is more of a functional, rather than aesthetic issue in this case).

4. The cell heater (H) previously made is glued into the D-ring cavity (from the ICP side). A small piece of thin Kapton is placed under the heater to isolate it electrically from the body. This can be done in two steps, once to just secure the heater in place with a few tacks of quick-setting superglue, followed by viscous black Stycast to paint entirely over the leads and around the heater. It is better to leave some heater surface exposed but black Stycast is not too much of a thermal insulator, so as long as it is not drowned in Stycast, we should be fine. The heater leads escape tunnel is also sealed off at this point. First seal it off from the inside of the cavity (while securing the heater). Later, during one of the other gluing steps, we can pour in more black Stycast from the outside end and then insert a lead tube over the wires to embed the lead tube securely in the tunnel.
5. The diaphragm is now glued on similarly to the electrode. We have to be much more careful about the initial placement of the diaphragm onto the Stycast wetted surface, as its central part is meant to be flexible. Too much adjustment will result in an annular ring being wetted by the Stycast and being stiff after drying – this reduces the effective (usable) area of the diaphragm and will inject errors into certain calibration methods (in addition to reducing displacement sensor sensitivity). It is even more important to practice this step on blank pieces and inspect the resulting joints to see how bad things get. The breakout tubes (at their present spacing) will not allow clearance for a standard glass slide. We therefore made a thick, clear plastic plate that does fit in and that is sanded as flat as possible. Cutting a glass piece would be preferable. A screw joint is made in a similar way as with the E-ring.
6. The E-ring is now lowered over the D-ring, with the diaphragm facing the electrode. 4 Kapton (300HN - 76 μm thick) spacers, which are just small rectangles with square holes cut in them using a blade, are placed between the 2 rings and the pieces secured together using 4 #0-80 screws.
7. All lead resistances should be checked at this point, including shorts to the body or to the lead tubes. Everything should be nominal.
8. All 4 leads are now fed through their respective breakouts – PC and E in one and H and D in the other – and sealed in place. A few drops of (fresh) clear Stycast is dribbled along each wire so it goes into the tiny gaps in the breakout caps. After this, viscous (2 – 3 hrs) Stycast is used to seal off each breakout with lead tubes embedded on either side of the breakout. One must be careful here to wait long enough for the right viscosity, else the Stycast will just get sucked into the lead tubes and not suffice to make a proper seal. The lead tubes should be cleaned and lightly sanded to provide a good mechanical grip for gluing (otherwise, lead doesn't stick well). The leads are made long enough to reach their final destination on the cryostat (no more intermediate breakout boxes).

9. The sense arm gluing (all done with Stycast 1266) can take multiple steps (which can be done in parallel with other things). A dummy ICP (just the threaded bolt holes for the 2 tube flanges) is made of aluminum to assemble the sense arm for the first time.
 - a) The tube flanges are screwed into the dummy ICP with washers cut out of plastic transparencies to protect the seal surfaces on the flange bottoms.
 - b) The side-arms are glued onto the tube flange steel tubes and the main heat-pipe glued onto the side-arms. This can be done in the same step, just by painting a tiny bit of fresh Stycast onto each surface to be glued. Then, viscous Stycast is used to make external seals around each joints. The dummy ICP (rectangular ingot) makes it easy to hold the sense arm in any orientation for most effective gluing. Be careful not to let Stycast seep inside the flow paths (could block or distort the flow).
 - c) In parallel, the sense arm heater is glued into the heater module (taking similar care as the inner cell heater to keep the resistor surfaces free of Stycast). The wires are taken out and the breakout hole sealed and lead tube embedded in Stycast.
 - d) Once this is all dry, the heater module can be sealed onto the appropriate end of the heat-pipe.
 - e) A ~ 5 mil thick copper disk is made by punching it out of a sheet. The sheet is one that has been sand-blasted to increase the effective area (it is a heat-sink after all) and then dipped in some dilute hydrochloric acid to clean it of oxide layers. The disk is cleaned in the usual series of ultrasonic baths.
 - f) This disk is very carefully glued in to the other end of the heat-pipe with viscous Stycast. Be very careful here that Stycast doesn't seep on to the central part of the disk as we want that part to be a good heat conductor. Practicing with such disks on a dummy piece helps one figure out the right way to do this.
10. The aperture array chips are cleaned (see Step. 14 on p. 69) only when everything else is done up to this point and the cryostat is ready to go otherwise. Once the chips are cleaned, the clock starts ticking – we should evacuate the cell and cool down within a few days (a week at most), else the chips could get contaminated again by virtue of being out in air (even on a clean bench). If one can do this entire assembly within a clean room, perhaps this time could be extended.
11. The chips are now glued into the ICP slots made for this purpose. First, 4 thin dots of (fresh) black Stycast are applied on the slot corners with a very sharp-pointed applicator and only this quantity (no more) is spread evenly around the perimeter of the slots. A chip is held in tweezers and gently placed in the slot, being careful to avoid the sensitive nitride membrane in the center. The chip is gently massaged in place to ensure that the bottom glue layer adheres to it. Now comes the hard part – we wait for around 2.5 hours (or until the Stycast is quite viscous, but still forms drops

if held in an applicator). This viscous Stycast is now carefully applied around the chip to completely seal it in. We have to be very careful here because the Stycast can be sucked up behind the chip (or creep along the top surface) and block the apertures on the chip.

12. The chip-gluing can be checked immediately (and should be checked after drying as well) by shining a light on the chip and using a small mirror under the ICP V-tunnel to see that a square window of light shines through.
13. Once everything is dry, we make the triple indium seal between the ICP and the D-ring. Appropriate indium wire pieces are cut and formed around dummy forms (as explained earlier) and placed carefully on the ICP (which remains with chips facing up). The (now-sealed) wires inside lead tubes from the D-ring breakouts are fed into the breakout clearance holes in the ICP (around which we have the two off-axis seals) – carefully lest they damage the chips. 4 hands can be a big help at this point. The triple seal is made using #2-56 screws, going slowly to let the inner cell pressure equilibrate.
14. The ICP is now inverted and indium-sealed onto the cell can.
15. The completed sense arm is now sealed onto the ICP. For practical reasons, our indium grooves are on the sense arm, so we massage the circle-formed indium wires onto the grooves – that;s sufficient to hold them in place even when we turn the flanges over. Be careful not to put any pressure on the many Stycast joints of the sense arm (could open a superleak). Effective use of lab-stands helps here. The indium seal is made, pretending that the two separate flanges are actually one flange and alternating accordingly so that both flanges are lowered together.

This completes the cell assembly. The cell is mounted on the cryostat and the wires hooked up to the cryostat interfaces. See the chapters on the cryostat (7) and operation (10) for chronology and details beyond this point.




Figure A.18: Pancake coil glued into E-ring with lead tube covered wires in background. Electrode will be glued over this.




Figure A.19: Bottom side of D-ring. Blue chip resistor (cell heater) can be seen in the cavity.




Figure A.20: Top side of D-ring ready for diaphragm gluing. Cell heater has been glued on other side with leads exiting out this side.




Figure A.21: Bottom side of ICP showing triple seal (central seal defines inner cell and off-axis seals provide wiring breakouts). Aperture array chips have been glued in their slots. White dots at their centers are the actual suspended silicon nitride membranes with aperture arrays shot through them.




Figure A.22: E-ring and D-ring (with electrode and diaphragm glued on) have been screwed together with Kapton spacer. This assembly has been sealed onto the ICP.




Figure A.23: The assembly in Fig. A.22 has been sealed onto the cell can. The sense arm has been sealed onto the top side of the ICP. Inset shows blow-up of fully assembled heater module.

A.3 Displacement sensor drawings

Included here are additional engineering drawings of displacement sensor parts referred to in Chapter 5.




Figure A.24: Tank inductor former for PI circuit.




Figure A.25: Inductor former for PI circuit injection line chokes (need 2).




Figure A.26: Shielded box for injection line chokes.




Figure A.27: Plastic (Lucite or polycarbonate) base for pancake coil winder.




Figure A.28: Evaporation mask for diaphragms and electrodes.

A.4 Single weak-link cell drawings

These are engineering drawings for single weak-link cell parts used in the partially modular design described in Section 3.4.1. As described there (and in the one following it – Section 3.4.2), the inner cell piece shown here can be modified to make the cell fully modular for rapid investigations of the physics of different kinds of aperture arrays. This cell was used to obtain the basic weak-link physics results discussed in Chapter 10.

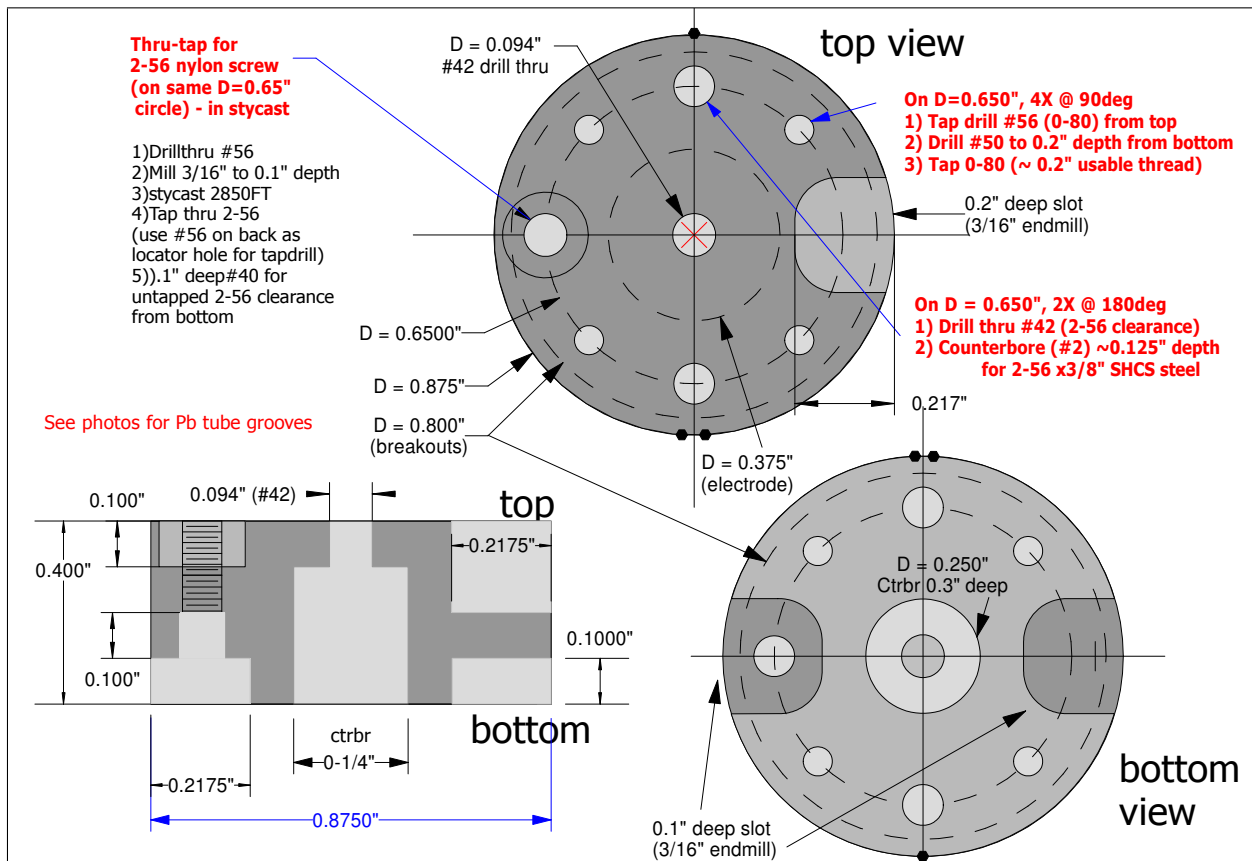


Figure A.29: Single weak-link cell: pancake coil holder




Figure A.30: Single weak-link cell: inner cell ring piece




Figure A.31: Single weak-link cell: Modular cell can




Figure A.32: Single weak-link cell: wiring breakout (single set of leads)




Figure A.33: Single weak-link cell: wiring breakout (double set of leads)

Appendix B

Matlab scripts

In this chapter, we include some scripts that do not fit in elsewhere in this dissertation. More scripts may be found in their respective topical sections.

B.1 Interference curves

This script is used to generate interference curves and model fits from saved data. It is capable of batch processing multiple data sets within a master directory. The phase guess in the fit is also iterated to give the best goodness of fit (r^2). The only external functions used are the ones called in the user params section, which return standard ^4He parameters as a function of temperature from tabulated data. Data is assumed to be saved as plain text with one variable to a column – column numbers clearly labeled in the script for easy modification. The extensive documentation within the script should cover the rest.

Note that this code handles heat-pipe power interference curves. Small changes are needed for analyzing Sagnac reorientation data (uncommenting some constants in the user params section, a redefinition of MP and cguess in the fit section), but all necessary code is provided within. In that case, the “MP” data column is assumed to contain angular positions in degrees instead of heat-pipe heater power. We leave the plot label recoding as an exercise for the reader since it is purely aesthetic.

Listing B.1: make_interferograms.m

```
% Analysis date: 2012/6/4
%This script is used for analyzing
%Data properly saved by trackfftpeak_2012_06_02.vi (including sense arm
%heater power and voltage data)

%This file is used for data on 2012/06/04
%T = 2.166K

%% Clear variables and command window
clear all;
close all;
clc;
```

```

%% User input

%Main Data Folder
datafol = 'C:\DATA\SHeQUID5_run1\analysis\multi_Fiske_mod_example\data_2012_08_09\';

folder_filter_str = 'single*'; %Only folders inside the main data folder
%matching this pattern will be assumed to contain modlists (one text file
%for each separate mod power)
%To process one folder only, put exact folder name above with no wildcards

folder_num = 0;%If user wants only a number from modlist folder to label the result
%files, else user wants full folder name appended to result files
%The first option is good for an automated series of mod curves such as
%the ones in a drift run (folders are labeled by incrementing counters
%in labview in that case). For random single or multiple folders being
%analyzed, set this flag to 0

savpathstr = 'C:\DATA\SHeQUID5_run1\analysis\multi_Fiske_mod_example\results\';
%All data will be saved here

appendstr = '2012_08_09_';
pausing = 0; %Pause between files to view plots
minrsq = 2; %Make this >1 if you want to find the best rsquare possible
    % (by iterating the phase offset guess: dguess) instead of just stopping
    % at an arbitrary minimum.

% Rsense = 99.946; %in ohms
G = 10; %Pre-amp gain (saved amplitude is G times larger than actual).
gamma = 0.042368;
A2byk = 7.831524e-13;
% Ic1 = v*gamma*A2byk*rho/G;

%cell data
L = 1*0.0254; %spacing between sense arm probe points (1 inch) in (m)
dsens = 0.09375*0.0254; %Sense arm tube bore diameter (#44 drill) in (m)
sig = pi*(dsens^2)/4; %Sense arm tube bore area in (m^2)
% AD = pi*((0.375*0.0254)^2)/4; %Diaphragm area

%Constants for Sagnac curves
%Dewar angle (in proper units depending on what units the data was saved in
%- tape reading in inches or dewar angle in deg) for which loop area vector
%is along true north-south direction.
%theta0 = 11;

%Asens = 10.7e-4; %Area of the sense loop in m^2
%lat = 37.87167; %37° 52' 18" north latitude (in deg) for Berkeley, CA
%Wp = 2*pi*(1/(3600*23.9344696))*cos(lat*pi/180); %Using sidereal day length
%Wp = component of Earth's rotation field (angular velocity parallel to the
%Earth's surface)

%Fundamental constants
hbar = 1.05457172647e-34; %hbar in J.s
amu = 1.66053886e-27; %atomic mass unit in kg
m4 = 4.002602*amu; %Mass of helium 4 atom in kg

T = 2.166; %Temperature for data in K
% Tlambda2 = Location of Tlambda according to Rbath calibration
% This is the actual cell Tlambda accurate to 0.5mK (source: where normal
% flow begins).
Tlambda2 = 2.175; %in K

```

```

%Helium 4 params from Donnelly data
Tlambda = 2.1767999; % Value for Tlambda used by Donnelly ref data. (K)
T0 = Tlambda - Tlambda2; %Offset (K)
Tp = T + T0; %This is the temperature to use for accessing any of the
%Donnelly data curves for helium 4 (K)

s = fpentropy(Tp); %Entropy in J/kg.K
rhos = rhosf(Tp); %Superfluid density (kg/m^3)
rho = rhohe4f(Tp); %Total density (kg/m^3)
rhon = rho - rhos; %Normal fluid density (kg/m^3)

%NOTE THAT Tp is used ONLY for accessing Donnelly data.
Tlambda_minus_T = 1000*(Tlambda2 - T); %in mK, Using cell Tlambda and data temperature
lstring = sprintf('T_\\lambda-T_=%.2f_mK', Tlambda_minus_T); %for use in graphs

% END USER ENTERED PARAMS
%
%
%
%% START ANALYSIS

% If results folder doesn't exist, create it
if (exist(savpathstr, 'dir') ~ 7)
    mkdir(savpathstr);
end;

cd(datafol); %Make current folder

%List of directories. Each directry has a set of textfiles. Each textfile
%has data for ONE sense arm power. Each directory will give ONE mod curve.
dfols = dir(folder_filter_str);
folnames = cell(size(dfols));

for j = 1:length(dfols)
    folnames{j} = dfols(j).name; %Get bare name
end;

%Initialize data arrays
%Cell arrays that will hold num arrays from each list in each cell

v1 = cell(1,length(folnames));
f1 = v1; %frequency for file (just for reference to see if whistle jumps to new state)
DF1 = v1; %peak width
P1 = v1;
fits = v1; %each cell is a cell array for each file -{fit, gof, output}

fnum = zeros(1, length(folnames)); %folder number for each list
phase = fnum; %Phase offset from fits
time = fnum; %Relative time in seconds at midpoint of a listfile

for i = 1:length(folnames)
    %Get file number
    fn = folnames{i};

    %If user wants only a number from modlist folder to label the result
    %files, else user wants full folder name appended to result files
    %The first option is good for an automated series of mod curves such as
    %the ones in a drift run (folders are labeled by incrementing counters
    %in labview in that case). For random single or multiple folders being

```

```

%analyzed, set this flag to 0
if folder_num
    [sn1 sn2] = regexp(fn, '([\d]+)'); %Finds boundaries of any number
    %in the foldername
    fn = fn(sn1:sn2); %Get only the number as string
    fnum(i) = str2num(fn); %#ok<ST2NM>
else
    fnum(i) = i;
end;
%This will be used for later labelling of stuff if needed

%Go to subdir containing lists (1 subdir = 1 mod curve; 1 listfile
%in each subdir is for 1 sense power to be averaged into 1
%modpoint)
cd(folnames{i});

%Get list of files (struct array)
tempdir = dir('*.txt');
fnames = cell(size(tempdir));

for k = 1:length(tempdir)
    fnames{k} = tempdir(k).name; %Get bare name
end;
fnames = sort(fnames); %Sorting also maintains chronological order
clear tempdir;

%Data collectors for average values of each listfile
Mv = zeros(length(fnames),1);
Msdom = Mv;
Mf = Mv;
MDF = Mv;
MP = Mv;
Mtime = Mv;
%For clarity -Mv1 vs. MP1 is 1 mod curve -these arrays get added to
%container cells in the master data cell array initialized before this
%loop started (v1, P1)

for j = 1:length(fnames)

    %Load data
    fdata = load(fullfile(datafol, folnames{i}, fnames{j}));

    %Get only necessary data columns and add to the cell arrays at the
    %end of this loop
    f = fdata(:,2);
    DF = abs(fdata(:,5) -fdata(:,4)); %Peak width in Hz
    amp = fdata(:,3); %sqrt(Integrated PSD Peak)
    P = fdata(:,12); %Sense arm power
    fftflag = fdata(1,13); %Any element of the last column will do.
    %Surely, no one would switch FFT methods midway through a
    %single listfile. (If your name is Shirley, my profound apologies).

    %0 = amplitude peak 1= velocity peak
    if fftflag
        v = amp;
    else
        v = 2.*pi.*f.*amp; %Velocity amplitude in V/s (unscaled)
    end;

```



```

%% Find velocity amplitude scaled and for fluid
v = (1e12)*v*gamma*A2byk*rho/G; %ng/s

%% Averaged data points representing each listfile (one phase point
%%in mod curve)
Mv(j) = mean(v);
Msdom(j) = std(v)/sqrt(length(v)); %Stdev of the mean = stdev/sqrt(num pts)
MP(j) = mean(P);
Mf(j) = mean(f);
MDF(j) = mean(DF); %Mean peak width for this phase point
Mtime(j) = fdata(floor(length(f)/2),1); %Time at approx. midpoint of
%listfile (in seconds -Labview timestamp)

clear f v P DF amp fdata;
end;

%Done collecting data for ONE mod curve
%Now analyze and do fits, etc. for single mod curve

%% Sort data according to the sense arm power
[MP, IX] = sort(MP);
Mv = Mv(IX);
Mf = Mf(IX);
MDF = MDF(IX);
Msdom = Msdom(IX);
Mtime = Mtime(IX); %Not necessary since it will be averaged for
%the entire modcurve anyway but what the hey ...

%% Phase difference between the SheQUID arms in radians
%dphi = (m4*rhon*L)*Psense/(hbar*rho*rhos*Tp*s*sig);
%So, define c such that dphi = 2*c*Psense; (See fit later)
%So, c is theoretically (m4*rhon*L)/(2*hbar*rho*rhos*T*s*sig);
%% Rescale Psense to nW. Now, c is rescaled by 1e-9.
%So, [P] = nW and [c] = 1/nW

%%Note that our MP above is dewar angle in degrees from an arbitrary
%0. Find the
%Omega.A = Omega_p.A since Omega normal to earth's surface does not contribute to flux
%dphi = 4*pi*Omega[dot]A/kappa4 = 4*pi*Omega_p*A*cos(alpha -alpha0) /kappa4

%This means that dphi/2 = (2*pi*Omega_p*A/kappa4) *Psense
%So, define c such that dphi/2 = c*Psense; (See fit later)

%dphi/2 is the argument of the sin and cos in the fit function -call this x.
%So, c is theoretically (2*pi*Omega_p*A/kappa4) where A is the loop area
%and Omega_p is the rotation rate parallel to the Earth's surface

% Uncomment next line if doing Sagnac curve
%MP = sin((MP -theta0)*pi/180); %Here MP was dewar angle in deg and theta0 is user-specified

%% Fitting to double-slit model

% Guess value a = It = (Ic1 + Ic2). Should be (max Ic) seen.
aguess = max(Mv)-min(Mv); %Just to be safe

% Guess value b = gamma^2. gamma = Imin/Imax
bguess = (min(Mv)/max(Mv))^2;

% c is a constant (dependent on T only): dphi_arms = 2*c*Psense
% Since a dphi of order 1 is ~ 100nW, this would make c ~ 0.01 in [1/nW]

```

```

cguess = (1e-9)*(m4*rhon*L)/(2*hbar*rho*rhos*T*s*sig);
%This includes the factor of 2 scaling

%For reorientation curve:
    %cguess = 2*pi*Wp*Asens/kap4; %See explanation above

%% Iterate over dguess values
%The fitting is repeated with incremented phase offset guess until
%either the rsquare becomes high enough (say >0.85) or we run out of
%phase guesses. In the latter case, the phase guess with the best
%rsquare is kept as the result
%Setting minrsq to >1 will force the routine to search for the best
%possible r-square always.

dguessSTEP = pi/6;
bestrs = 0;
bestguess = 0;
dguess = -pi/2; %d is a phase offset

while dguess <= pi/2

    sfit = fitoptions('Method','NonlinearLeastSquares',...
        'Lower',[0,0,0,-pi/2],...
        'Upper',[2*max(Mv),1,1,pi/2],...
        'Startpoint',[aguess bguess cguess dguess],'MaxFunEvals', 2000,...
        'MaxIter', 1000);
    set(sfit,'Weight',Msdom); %Weights

    %From SheQUID analysis book 1 (AJ and MR)pg.20
    fitfun = fitype('a*sqrt(((cos(c*x_+_d))^2)+_b*((sin(c*x_+_d))^2)),...
        'options',sfit,'coefficients',{'a','b','c','d'}, 'independent','x');
    %Here, the dependent var x is the heater power Psense

    [~,gof,~] = fit(MP,Mv,fitfun, sfit)

    %debug

    dguess
    gof.rsquare

    if gof.rsquare >= minrsq
        dguess = 10*pi; %to make loop stop
        bestrs = gof.rsquare;
        bestguess = dguess;
    else
        if bestrs < gof.rsquare %update bestguess so far so that if
            %no rsquare is >0.85, we at least have the best one possible
            bestrs = gof.rsquare;
            bestguess = dguess;
        end;
        bestguess
        bestrs
        dguess = dguess + dguessSTEP;
    end;
end;

%% Redo fit with bestguess
dguess = bestguess;

```

```

sfit = fitoptions('Method','NonlinearLeastSquares',...
    'Lower',[0,0,0,-pi/2],...
    'Upper',[2*max(Mv),1,1,pi/2],...
    'Startpoint',[aguess bguess cguess dguess],'MaxFunEvals', 2000,...
    'MaxIter', 1000);
set(sfit,'Weight',Msdom); %Weights

%From SheQUID analysis book 1 (AJ and MR)pg.20
fitfun = fittype('a*sqrt(((cos(c*x_+d))^2)+_b*((sin(c*x_+d))^2))',...
    'options',sfit,'coefficients',{'a','b','c','d'}, 'independent','x');
%Here, the dependent var x is the heater power Psense

[cfit,gof,output] = fit(MP,Mv,fitfun, sfit)

%% Plot fit vs. Power
figure(1); clf; hold on;
errorbar(MP,Mv,Msdom,'o');
% plot(MP,Mv,'o');
plot(cfit,'-r'); %FIT
title(['Modulation_curve_' lstring '](file#_' fn ')]);
xlabel('Sense_arm_power_(nW)');
ylabel('Whistle_amplitude_(ng/s)');
legend('Data', 'fit')
hold off;
saveas(1,fullfile(savpathstr,['Ic_vs_P_' appendstr num2str(T,'%0.4f') 'K_' fn '.jpg']));
saveas(1,fullfile(savpathstr,['Ic_vs_P_' appendstr num2str(T,'%0.4f') 'K_' fn '.fig']));

%% Get phase
%get fit coeffs
temp3 = coeffvalues(cfit);
aF = temp3(1);
bF = temp3(2);
cF = temp3(3);
dF = temp3(4);
ItF = aF;
gammaF = sqrt(bF);

%Phase difference between the SheQUID arms in radians
dphi = 2*cF*MP; %c is in 1/nW and MP is in nW, so dphi is in real rads

%Shift data by multiple of 2*pi so it starts close to 0 without messing with phase-offset
temp1 = floor(dphi(1)/(2*pi));
dphi = dphi - 2*pi*temp1;

dphi2min = dphi(1) - pi/10; %min for fit function
dphi2max = dphi(length(dphi)) + pi/10;

%Create a dense array for the fit function
npts = 100;
dphi2 = dphi2min:(dphi2max-dphi2min)/npts:dphi2max;
dphi2=dphi2';

% %Create a dense array for the fit function
% npts = 100;
% dphi2 = -pi/2:(max(dphi)-min(dphi))/npts:3*pi;
%Eval fit fun at these points
v2 = aF.*sqrt(((cos(dphi2/2 + dF)).^2) + bF.*((sin(dphi2/2 + dF)).^2));

%Create axis ticks

```

```

maxt = ceil((max(dphi) + pi/2)/(pi/2));
mint = floor((min(dphi) - pi/2)/(pi/2));
ticklabs = cell(1,(maxt-mint+1));
tns = mint:1:maxt;

for ti = 1:(maxt-mint+1)
    if mod(tns(ti),2) == 0 %even
        if tns(ti) == 2 %x = pi
            ticklabs{ti} = 'p';
        elseif tns(ti) == -2 %x = -pi
            ticklabs{ti} = '-p';
        else
            ticklabs{ti} = [num2str(tns(ti)/2), 'p'];
        end;
    else %ODD
        if tns(ti) == 1 %x = pi/2
            ticklabs{ti} = 'p/2';
        elseif tns(ti) == -1 %x = -pi/2
            ticklabs{ti} = '-p';
        else
            ticklabs{ti} = [num2str(tns(ti)), 'p/2'];
        end;
    end;
end;

% ticklabs

%Plot
figure(2); clf; hold on;
errorbar(dphi,Mv,Msdm,'s'); %DATA
plot(dphi2, v2,'-r','LineWidth',2); %FIT
title(['Modulation_curve_' lstring ' '(file#_ ' fn ')]);
xlabel('Sense_arm_phase_difference_(rad)');
ylabel('Whistle_amplitude_(ng/s)');
set(gca,'XTick',(min(dphi) - pi/2):pi/2:(max(dphi) + pi/2));
set(gca,'XTickLabel',ticklabs,'fontname','symbol');
%legend1 = legend;
%legend('data', 'fit');
%set(legend1,'FontName','Tahoma');
hold off;

saveas(2,fullfile(savpathstr,['Ic_vs_dphi_' appendstr num2str(T,'%0.4f') 'K_' fn '.jpg']));
saveas(2,fullfile(savpathstr,['Ic_vs_dphi_' appendstr num2str(T,'%0.4f') 'K_' fn '.fig']));

%% Save ONE modcurve's data to the master cell array
f1{i} = Mf;
DF1{i} = MDF;
v1{i} = Mv;
P1{i} = MP;
fits{i} = {cfit, gof, output};
phase(i) = dF;
time(i) = mean(Mtime); %Average time for the mod curve

%Status update
sprintf('Finished_modcurve_#%03d_of_%03d', i, length(folnames))
% If pausing between folder is enabled
if pausing
    userentry = input('Press_Enter_to_continue,_Ctrl-C_to_abort_script...')
end;

% clc; %Clear screen

```

```

clear frames;
%Return to parent dir to await the next subdir (=next modcurve)
cd ..

end;
%% Save full data as mat file
cd(savpathstr);
save(['analyzed_data_' appendstr num2str(T,'%0.4f') 'K.mat']);

```

B.2 Fountain calibration

Only the non-trivial (fitting) part of the front-end script is shown. The method used here is the newest one (#3 on p. 214).

Listing B.2: Fitting excerpt from fountain calibration front-end script

```

options1 = fitoptions('Method','NonlinearLeastSquares');
options1.StartPoint = [a1guess, a2guess];
options1.Lower = [0, 0]; %No negative values (gives errors)
options1.Upper = [Inf, Inf]; %No negative values (gives errors)
options1.MaxFunEvals = 8000; options1.MaxIter = 8000;
options1.Robust = 'on';
%options1.Algorithm = 'Levenberg-Marquardt';
options1.TolFun = 1e-9;
options1.TolX = 1e-9;

%Fitting function object

%Fountain fit to get A^2/k and Vcell
%beta3 is a KNOWN parameter defined at the start of this file and not used for fitting.
ft1 = fitype('fountfun(x,a1,a2,B)',options1, options1, 'coefficients',{a1,a2}, 'problem','B');

%fountfun.m should be in any defined path
%Note that this function file is the key to the whole thing. The procedure in this block may be
%used with any function M-file, no matter how bizarrely defined (using other M-files, splines,
%etc.) as long as the function is defined over the entire domain of x. Since our data will always
%be in the physically meaningful region of x, it is trivial to add an if or switch statement to
%set the non-meaningful regions of x to output 0 or whatever since it will never be used. But this
%must be done or the fit returns errors.

%% Fit
[yfit, gof, output] = fit(T, Y, ft1, 'problem',beta3);

```

Recall that β_3 is previously known from other calibrations (see the earlier page reference) and is fed to the fit function as a parameter (the “problem” switch in the calling code). The data used for fitting includes the absolute temperature T and the left-hand side of Eq. (10.15) (which is the Y-data used in the script).

The fit function file is shown below:

Listing B.3: Fit function file (fountfun.m)

```

function y = fountfun(x, a1, a2, beta3)

%x = temperature; a1 = 1e12*(A^2 /k), a2 = (1e8)*Vcell
%a1 and a2 should then be of order unity.
% A is diaphragm area and k is spring constant

```

```

%beta3 is beta2/(beta1^2)

a1 = a1*1e-12;
a2 = a2*1e-8;

%Fitting function for fountain calibration. This expression is equal to (1e8)*tau*Wh/DeltaVf where
%tau and DeltaVf are obtained from fitting the fountain transient to an exponential and Wh is the
%heater power for the transient. The factor of 1e8 is simply to make the function values large enough
%to avoid numerical errors in fitting related to tolerances. The data is accordingly scaled up by 1e8
%before fitting.
%spheat.m returns sp heat in J/mol.K -extra constants are to convert to J/m^3.K in the below
%formula(since Cp = cp*V*rho/molarmass). he4m = 4.0026e-3; 4He molar mass (kg/mol);

Tlambda = 2.1767999; % Value for Tlambda used by Donnelly ref data.

%Tlambda2 = Location of Tlambda according to Rbath calibration.
Tlambda2 = 2.1795; %accurate to 0.5mK (source: where normal flow begins)
T0 = Tlambda -Tlambda2; %Temp. offset for getting T90 scale from my GRT

if isvector(x) %can be vector or scalar -gives true
    y=zeros(size(x));
    for i=1:length(x)
        xi = x(i);
        xi2 = xi + T0; %Use for accessing Donnelly data and everything else

        if xi>1.2
            %y(i) = (1e8)*(spheat(xi2).*a2.*(1e-8)./(4.0026e-3) + (fpentropy(xi2).^2).*(rhohe4f(xi2)).*xi.*a1.*(1e-12))./
            fpentropy(xi2);
            y(i) = sqrt(beta3)*(1e8)*(spheat(xi2).*a2./(4.0026e-3) + (fpentropy(xi2).^2).*(rhohe4f(xi2)).*xi2.*a1)./(sqrt(2).
            *fpentropy(xi2).*(a1^0.5));
            %See Analysis book # 1 (A. Joshi and M. Ray) pg. 33 for
            %details of this new model
        end; %if not, do nothing -y(i) is already 0
    end;
else
    y = [];
    'ERROR!_x_must_be_a_scalar_or_vector' %#ok<NOPRT>
end;

end

```

Appendix C

Flow tests: further analysis

This chapter includes further details of calculations used in Section 6.1 of Chapter 6. Mathematica [70] code is included for the final computations of hole size.

C.1 Clausing's factor: $K[u]$

Clausing's factor was described in the section on aperture size (p. 109) and is used while calculating hole sizes in gas flow tests of the weak link aperture arrays. All symbols retain their definitions from that chapter.

The full expression for $K[u]$ is of breathtaking ugliness and admits of no physical insight. Nevertheless, for the sake of completeness, we state the procedure to derive it and also the final form. A Mathematica file has been written to do all analysis and may be obtained from the authors. It must be noted that this is extremely useful in that it makes the Clausing conductance formula valid for all aspect ratios of the aperture (from 0 to infinity) as opposed to approximations that are valid at one extreme (orifices) or another (long tubes). In any case, one does not have to actually bleed to derive it, a good math package is sufficient to do the computations without ambiguity since they are all elementary (algebraic even!) expressions.

From Eq.(42) in [66] (recast in terms of $u = R/L$ rather than R and L),

$$\begin{aligned}
 K[u] = & \frac{1 + 2u^2 - \sqrt{1 + 4u^2}}{2u^2} \\
 & + \frac{-1 + 4u^3 + (1 - 2u^2)\sqrt{1 + 4u^2}}{3u^2} \\
 & + \alpha \left(1 - \frac{1 + 2u^2 - \sqrt{1 + 4u^2}}{2u^2} - \frac{2(-1 + 4u^3 + (1 - 2u^2)\sqrt{1 + 4u^2})}{3u^2} \right)
 \end{aligned} \tag{C.1}$$

where $\alpha \equiv \alpha[u, v]$ (u is as before and $v \equiv \frac{x}{L}$ represents a dimensionless coordinate along the channel length. Thus, $v \in [0, 1]$). The full form of α is as follows:

$$\alpha[u, v] = - \frac{\left(-1 - \sqrt{4u^2 + (-1+v)^2} (-1+v) + 2v - v\sqrt{4u^2 + v^2} \right)}{2 \left(\frac{(-1+v)^2 v + u^2 (-2+4v)}{\sqrt{4u^2 + (-1+v)^2}} + \frac{(-1+v)v^2 + u^2 (-2+4v)}{\sqrt{4u^2 + v^2}} \right)} \quad (\text{C.2})$$

The key idea here is that $\alpha[u, v]$ is nearly independent of v for $u \gtrsim 0.1$ (corresponding to apertures with diameter $\gtrsim 12\text{nm}$ for a nitride thickness of 60 nm) so that one may simply choose a convenient value of v and obtain α as a function of u only. This weak v -dependence can be seen by direct plots of Eq. (C.2) (see Figs. C.1 & C.2).




Figure C.1: α on the vertical axis, u running from left to right and v from front to back. The rippling of the surface starts fading around $u \sim 0.1$

Further, $\alpha[u, v]$ is chosen in such a way that the conductance reduces to the correct forms in the large and small u limits. Clausing's choice (Eq.(43) in [66]) is $v = \frac{2\sqrt{7}u}{3+2\sqrt{7}u}$, which transforms Eq. (C.2) into our final form for α :




Figure C.2: α on the vertical axis, v running from left to right.

$$\begin{aligned}
 \alpha \left[u, \frac{2\sqrt{7}u}{3+2\sqrt{7}u} \right] = & \frac{1}{u} \left(3 + 2\sqrt{7}u \right)^2 \left(-3 + 2\sqrt{7}u - 4\sqrt{7}u^2 \sqrt{1 + \frac{7}{(3+2\sqrt{7}u)^2}} \right. \\
 & \left. + 3 \sqrt{4u^2 + \frac{9}{(3+2\sqrt{7}u)^2}} \right) \times \\
 & \left(\frac{(-69 - 18\sqrt{7}u + 84u^2 + 56\sqrt{7}u^3)}{\sqrt{\frac{(4+3\sqrt{7}u+7u^2)}{(3+2\sqrt{7}u)^2}}} \right. \\
 & \left. + \frac{4(9\sqrt{7} - 27u - 18\sqrt{7}u^2 + 84u^3 + 56\sqrt{7}u^4)}{\sqrt{4u^2 + \frac{9}{(3+2\sqrt{7}u)^2}}} \right)^{-1}
 \end{aligned} \tag{C.3}$$

Eq. (C.3) and Eq. (C.1) together fully define the Clausing constant $K[u]$ needed in the section on aperture size (p. 109). Mathematica code for the final form of $K[u]$ and for an automatic script to calculate the hole size given experimental data from Section 6.1.3 is given in the following sections. If you don't want to use Mathematica, simply program $K[u]$ from Eq. (C.1) and substitute α from Eq. (C.3) to obtain the final function needed for the later calculation.

C.2 Mathematica code

These code snippets can be pasted directly into a Mathematica worksheet (where they are interpreted as working cells) to get the function $K[u]$. All three snippets pasted (in order) are sufficient to create a working program that accepts all the experimentally obtained parameters discussed at the end of this section (and described more fully in Section 6.1).

C.2.1 $K[u]$ code

Listing C.1: Clausing factor $K(u)$

```

Notebook[{
Cell[BoxData[
  \[K[u_] := (\(1/u^2\) \((\(-\((1.` u \((\((3.`)\(\[
\InvisibleSpace]\)) + 5.291502622129181` \ u)\)^2 \((\(-3.`) +
3.` \ @\((4.` u^2 + 9.` \ \((\((3.`)\(\[
\InvisibleSpace]\)) + 5.291502622129181` \ u)\)^2\) +
u \ \((\((5.291502622129181`)\)\(\[InvisibleSpace]\)) \
-10.583005244258363` \ @\((u^2 \ \((\((1.`)\)\(\[InvisibleSpace]\)) + \
7.` \ \((\((3.`)\)\(\[InvisibleSpace]\)) + 5.291502622129181` \
u)\)^2)\))\))\)\(\((148.1620734196171` \ \((\(-0.7343486479483514` \) +
u)\) \ u \ \((\((0.634175914905256`)\)\(\[
\InvisibleSpace]\)) + 1.3012953574621922` \ u +
u^2)\))\)\(\(@\((u^2 \ \((\((4.`)\(\[
\InvisibleSpace]\)) + 7.937253933193772` \ u + 7.` \ \
u^2)\)\)\(\((\((3.`)\)\(\[InvisibleSpace]\)) + 5.291502622129181` \ \
u)\)^2)\)) + \((592.6482936784683` \ \((\((0.25303002342193637`)\)\(\[
\InvisibleSpace]\)) -0.8522971742241046` \ u +
u^2)\)) \ \((\((0.6351589567941865`)\)\(\[
\InvisibleSpace]\)) + 1.4192438837379453` \ u +
u^2)\))\)\(\(@\((4.` u^2 + 9.` \ \((\((3.`)\)\
\(\[InvisibleSpace]\)) + 5.291502622129181` \ u)\)^2)\))\))\)) +
0.5` \ \((\((1.`)\)\(\[InvisibleSpace]\)) + 2.` \ u^2 -
1.` \ @\((\((1.`)\)\(\[InvisibleSpace]\)) + 4.` \ u^2)\))\ \
\(\((\((1.`)\)\(\[InvisibleSpace]\)) + \((\((\((3.`)\)\(\[InvisibleSpace]\)) \
+ 5.291502622129181` \ u)\)^2 \ \((\(-3.`) +
3.` \ @\((4.` u^2 + 9.` \ \((\((3.`)\)\(\[
\InvisibleSpace]\)) + 5.291502622129181` \ u)\)^2\) +
u \ \((\((5.291502622129181`)\)\(\[InvisibleSpace]\
)\)) -10.583005244258363` \ @\((u^2 \ \((\((1.`)\)\(\[InvisibleSpace]\)) + \
7.` \ \((\((3.`)\)\(\[InvisibleSpace]\)) + 5.291502622129181` \ \
u)\)^2)\))\))\))\)\(\((u \ \((\((148.1620734196171` \ \
\((\(-0.7343486479483514` \) +
u)\) \ u \ \((\((0.634175914905256`)\)\(\[
\InvisibleSpace]\)) + 1.3012953574621922` \ u +
u^2)\))\)\)\(\(@\((u^2 \ \
\((\((4.`)\)\(\[InvisibleSpace]\)) + 7.937253933193772` \ u + 7.` \ \

```



```

\((2 u*L /. Join[sol1, \ {L -> \ LSiN}])\); \[IndentingNewLine]If
  full == 1, \[IndentingNewLine]{Print["<Gfit = \>", \
    Gfit*10\^18, \ "\< in nm.ns\>"]; \[IndentingNewLine]Print["<n \
sol1 gives u = \>", \
  u /. sol1]; \[IndentingNewLine]Print["<n Gth with sol1 is \>", \
\ tempg1, \ "\< in nm.ns. Gfit was \>", \
  Gfit*10\^18, \ "\<n %Difference is \>", \ \((tempg1 - \
Gfit*10\^18)\)\(\(Gfit*10\^18)*100 //
  ScientificForm, \ "\<% compared to Gfit \n\>"]; \
\[IndentingNewLine]Plot[Gth[LSiN, T], \ {u, tempu1 -0.1, tempu1 + 0.1}, \
  PlotLabel -> \ "\<Gth [nm.ns] vs. u\>", \
  PlotRange -> \ All; \[IndentingNewLine]Plot[\((2 u \ L)\) /.
  FindRoot[Gth[L, T] == \ Gfit*10\^18, \ {u, tempu1}], \ {L, 50,
  80}, \ PlotLabel -> \ "\<Hole diameter (nm) vs. nitride \
thickness (nm)\>"]; \[IndentingNewLine]Print["<For LSiN = \>",
  LSiN -15, "\< nm, D = \>", \((2 u \ \((LSiN -15)\))\)\) /.
  FindRoot[
    Gth[LSiN -15, T] == \ Gfit*10\^18, \ {u,
    tempu1}], \ "\< nm\>"];
Print["<For LSiN = \>",
  LSiN + 15, "\< nm, D = \>", \((2 u \ \((LSiN + 15)\))\)\) /.
  FindRoot[
    Gth[LSiN + 15, T] == \ Gfit*10\^18, \ {u,
    tempu1}], \ "\< nm\>"]; \[IndentingNewLine] \
\[IndentingNewLine]; \[IndentingNewLine] \

```

C.2.2 Usage

Codes from the previous sections when executed in order will enable the automatic script. If code spans multiple pages, copy the code from separate pages completely into a plain text file and then paste the full code into Mathematica (Version 4 or higher). If everything works, the code should show up properly formatted as in Fig. C.3. We have tested this in Mathematica 4 using the code pasted from the PDF version of this dissertation and verified that it works properly.

The module inputs are (in order):

τ Time constant from experiment in seconds(usually averaged over transients)

Vin Inner cell volume in m^3

n Number of holes in the array

T Absolute temperature (K) during transient

full A boolean switch. full=1 shows all debug steps and graphs and hole sizes for perturbed values of the nitride thickness (channel length) in case this is experimentally uncertain by a certain amount. full=0 shows only final computed diameter

The output of the procedure is the hole diameter in nm. If the “full” switch is enabled, intermediate calculations and plots are shown. See example inputs and outputs in Fig. C.4.

■ constants

```
mHe = 6.64647617 * 10-27;
kB = 1.3806503 * 10-23;
LSiN = 60;

Gth[L_, T_] := K[u] u2 *  $\frac{3}{8} \sqrt{\frac{32 \pi mHe}{9 kB T}} L^2$ ;
```

■ procedure

```
analysis[τ_, Vin_, n_, T_, full_] := Module[{soll, Gfit, tempg1, tempul},
  Gfit =  $\frac{mHe Vin}{\tau kB T n}$ ;
  soll = FindRoot[Gth[LSiN, T] == Gfit * 1018, {u, .6}];
  tempg1 = Gth[LSiN, T] /. soll;
  tempul = u /. soll;
  Print["Hole diameter in nm = ", (2 u * L /. Join[soll, {L → LSiN}])];
  If[full == 1,
    {Print["Gfit = ", Gfit * 1018, " in nm.ns"];
     Print["\n soll gives u = ", u /. soll];
     Print["\n Gth with soll is ", tempg1,
           " in nm.ns. Gfit was ", Gfit * 1018, "\n %Difference is ",
            $\frac{(tempg1 - Gfit * 10^{18})}{Gfit * 10^{18}} * 100$  // ScientificForm, "% compared to Gfit \n"];
     Plot[Gth[LSiN, T], {u, tempul - 0.1, tempul + 0.1},
          PlotLabel → "Gth [nm.ns] vs. u", PlotRange → All];
     Plot[(2 u L) /. FindRoot[Gth[L, T] == Gfit * 1018, {u, tempul}], {L, 50, 80},
          PlotLabel → "Hole diameter (nm) vs. nitride thickness (nm)"];
     Print["For LSiN = ", LSiN - 15, " nm, D = ", (2 u (LSiN - 15)) /.
           FindRoot[Gth[LSiN - 15, T] == Gfit * 1018, {u, tempul}], " nm"];
     Print["For LSiN = ", LSiN + 15, " nm, D = ", (2 u (LSiN + 15)) /.
           FindRoot[Gth[LSiN + 15, T] == Gfit * 1018, {u, tempul}], " nm"];
    }
  ];
]
```

Figure C.3: Mathematica screengrab of code snippets C.2 and C.3.




Figure C.4: Example of Mathematica procedure usage and (full=1) output. This is the result shown in column 2 of Table 6.1.

Appendix D

On the superposition of phase-coherent oscillations

This chapter contains analytical work related to multiple-slit interference patterns. Results obtained here are used in fits to interference data throughout this dissertation as well as discussions in Chapter 13. This chapter was made necessary because we did not find these results anywhere in the literature and had to derive them for ourselves.

D.1 General problem

D.1.1 Problem statement

We wish to combine an arbitrary number of sinusoidal oscillations having the same frequency but different phase offsets and different amplitudes, into a single sinusoidal oscillation that is modulated by some *time-independent* function. Let the number of oscillations be $M+1$ and the resultant oscillation be represented by:

$$I_{tot} = \sum_{k=0}^M A_k \sin(\omega t + \phi_k) \quad (\text{D.1})$$

D.1.2 Formal Solution

We start by assuming that such a representation does exist so that we can write:

$$G \sin(\omega t + H) = \sum_{k=0}^M A_k \sin(\omega t + \phi_k) \quad (\text{D.2})$$

Here, we demand that G and H be time-independent. They will, in general depend on the phase-offsets (ϕ_k) and the amplitudes (A_k). Expanding the two sides trigonometrically, we obtain:

$$G \sin(\omega t) \cos H + G \cos(\omega t) \sin H = \sum_{k=0}^M (A_k \sin(\omega t) \cos \phi_k + A_k \cos(\omega t) \sin \phi_k) \quad (\text{D.3})$$

A judicious re-arrangement yields:

$$\left[G \cos H - \sum_{k=0}^M A_k \cos \phi_k \right] \sin(\omega t) + \left[G \sin H - \sum_{k=0}^M A_k \sin \phi_k \right] \cos(\omega t) = 0 \quad (\text{D.4})$$

This equation is of the form $X \sin(\omega t) + Y \cos(\omega t) = 0$. Since this must be true for all times t , we are faced with an over-determined system of an infinite number of equations (one for each value of t , which generates distinct values for the “parameters” $\sin(\omega t)$ and $\cos(\omega t)$) in the two unknowns X and Y that are time-independent by construction. Therefore, only the trivial solution exists so that X and Y must each be zero. This is equivalent to saying that $\sin(\omega t)$ and $\cos(\omega t)$ are linearly independent (when considered as functions allowed to vary in their arguments) and so their coefficients X and Y must be identically zero in the above homogeneous equation.

$$G \sin H = \sum_{k=0}^M A_k \sin \phi_k \quad (\text{D.5})$$

$$G \cos H = \sum_{k=0}^M A_k \cos \phi_k \quad (\text{D.6})$$

The total phase-offset H can be found by dividing Eq. (D.5) by Eq. (D.6) to get:

$$H = \arctan \left[\frac{\sum_{k=0}^M A_k \sin \phi_k}{\sum_{n=0}^M A_n \cos \phi_n} \right] \quad (\text{D.7})$$

The total amplitude G can be found by squaring and adding Eq. (D.5) and Eq. (D.6) and using $\sin^2 H + \cos^2 H = 1$ to get:

$$G^2 = \left[\sum_{k=0}^M A_k \cos \phi_k \right]^2 + \left[\sum_{k=0}^M A_k \sin \phi_k \right]^2 \quad (\text{D.8})$$

D.1.3 A suggestive expression

Eq. (D.7) and Eq. (D.8) constitute the formal solution to our problem. This explicit construction can be evaluated by mathematical software for the required number of oscillators. However, these expressions yield no physical insight to the dynamical problem of superposed oscillators. We can re-cast the amplitude expression (Eq. (D.8)) into a more suggestive form as follows:

$$G^2 = \left[\sum_{k=0}^M A_k \cos \phi_k \right]^2 + \left[\sum_{k=0}^M A_k \sin \phi_k \right]^2 \quad (\text{D.9a})$$

$$= \left[\sum_{k=0}^M A_k \cos \phi_k \right] \left[\sum_{n=0}^M A_n \cos \phi_n \right] + \left[\sum_{k=0}^M A_k \sin \phi_k \right] \left[\sum_{n=0}^M A_n \sin \phi_n \right] \quad (\text{D.9b})$$

$$= \left[\sum_{n=0}^M \sum_{k=0}^M A_k A_n \cos \phi_k \cos \phi_n \right] + \left[\sum_{n=0}^M \sum_{k=0}^M A_k A_n \sin \phi_k \sin \phi_n \right] \quad (\text{D.9c})$$

$$= \sum_{n=0}^M \sum_{k=0}^M A_k A_n [\cos \phi_k \cos \phi_n + \sin \phi_k \sin \phi_n] \quad (\text{D.9d})$$

The quantity in brackets can be combined to give:

$$\boxed{G^2 = \sum_{n=0}^M \sum_{k=0}^M A_k A_n [\cos (\phi_k - \phi_n)]} \quad (\text{D.10})$$

Eq. (D.10) has the simple interpretation (reminiscent of Feynman diagrams) of oscillators interacting pairwise with each other, modulated according to the phase offset between them. The result of the superposition is just such a pairwise sum of all possible (abstract) “interactions”. In the next section, we investigate a special case where the phase gradient along the line of oscillators is uniform so that the phase-offset ϕ_k is simply incremented by a constant amount as k increases. *Note: $M+1$ is the total number of oscillations being superposed (not M).*

D.2 In retrospect

At the risk of sounding anecdotal, prior to developing the rigorous treatment described above, we tried to derive Eq. (D.10) using two familiar techniques that one usually employs in such situations: (i) phasors and (ii) complex algebra. Obviously, this failed. However, that was before we knew the form of the final expression. So, the above method is not so much a “derivation” or “proof” as it was an analysis - a road to obtain a solution rather than prove an already-known result and is therefore doubly useful in a sense. Now that we have the

solution in hand, it is instructive to seek out significantly shorter and more elegant proofs of this result and see exactly why our earlier efforts met with such singular failure. This retrospective analysis led to some delightful realizations.

D.2.1 Complex algebra

Eq. (D.1) can be written as follows:

$$\begin{aligned}
 I_{tot} = \sum_{k=0}^M A_k \sin(\omega t + \phi_k) &= \Im \left\{ \sum_{k=0}^M A_k e^{i(\omega t + \phi_k)} \right\} \\
 &= \Im \left\{ \left(\sum_{k=0}^M A_k e^{i\phi_k} \right) e^{i\omega t} \right\} \\
 &\equiv \Im \{ \chi e^{i\omega t} \}
 \end{aligned} \tag{D.11}$$

where,

$$\chi \equiv \sum_{k=0}^M A_k e^{i\phi_k} \tag{D.12}$$

Since χ is a complex number, we can write it in polar form $\chi \equiv R e^{i\lambda}$ (where $R = |\chi|$ and $\lambda = \text{Arg}[\chi]$) in Eqs. (D.11) to get:

$$\begin{aligned}
 I_{tot} &= \Im \{ \chi e^{i\omega t} \} \\
 &= \Im \{ R e^{i\lambda} e^{i\omega t} \} \\
 &= R \sin(\omega t + \lambda)
 \end{aligned} \tag{D.13}$$

It is worthwhile to note at this point that the amplitude R and phase λ are respectively G and H from Section. D.1.2 (and the rest of the document). Obtaining R is as simple as calculating the magnitude of χ , starting from Eq. (D.12):

$$\begin{aligned}
 R^2 = |\chi|^2 &= \chi \chi^* \\
 &= \left(\sum_{k=0}^M A_k e^{i\phi_k} \right) \left(\sum_{n=0}^M A_n e^{-i\phi_n} \right) \\
 &= \sum_{k=0}^M \sum_{n=0}^M A_k A_n e^{i(\phi_k - \phi_n)}
 \end{aligned} \tag{D.14}$$

$$= \left(\sum_{k=0}^M \sum_{n=0}^M A_k A_n \cos(\phi_k - \phi_n) \right) + i \left(\sum_{k=0}^M \sum_{n=0}^M A_k A_n \sin(\phi_k - \phi_n) \right) \tag{D.15}$$

But R^2 is clearly real, and since the A_k 's are also real, the second term in the last line above must vanish¹!

We then have (just like we obtained in Eq. (D.10)):

$$R^2 = \sum_{k=0}^M \sum_{n=0}^M A_k A_n \cos(\phi_k - \phi_n) \quad (\text{D.16})$$

D.2.2 Phasors

In the phasor diagram description (see Fig. D.1), the $(M + 1)$ oscillations (all at frequency ω) can be depicted as vectors in the XY plane whose magnitudes are the corresponding oscillation amplitudes (A_k) and whose arguments (angles measured from the (arbitrary) X-axis) are the relative phases (ϕ_k). If the vectors are all added together, the magnitude of the resultant vector (\vec{A}_{tot}) is the superposed oscillation amplitude (previously called G - see Section. D.1.2) and the total phase-offset (previously called H) is the angle of \vec{A}_{tot} from the X-axis.




Figure D.1: Phasor diagram for superposition of oscillations with unequal amplitudes.

Formal solution

Eq. (D.8) for the formal solution previously derived emerges instantly from Fig. D.1. $|\vec{A}_{tot}|$ can be found by resolving each vector (\vec{A}_k) into its components ($A_k \cos \phi_k$ and $A_k \sin \phi_k$) and

¹We do not need to prove this - the mathematics tells us so. Even so, it is worth seeing why. The sine function is odd so that each term indexed by $(k = x, n = y)$ will be cancelled by the term indexed $(k = y, n = x)$ when $k \neq n$. When $k = n$, $\sin(\phi_k - \phi_n) = \sin(0) = 0$

adding separately as usual. This gives the X and Y components and thus the magnitude of the resultant vector:

$$\begin{aligned}
 A_{tot,X} &= \sum_{k=0}^M A_k \cos \phi_k \\
 A_{tot,Y} &= \sum_{k=0}^M A_k \sin \phi_k \\
 \left| \vec{A}_{tot} \right|^2 &= A_{tot,X}^2 + A_{tot,Y}^2 \\
 &= \left[\sum_{k=0}^M A_k \cos \phi_k \right]^2 + \left[\sum_{k=0}^M A_k \sin \phi_k \right]^2
 \end{aligned} \tag{D.17}$$

...which was the formal solution previously derived. Nevertheless, we find this unsatisfactory and wish to derive the more useful form (Eq. (D.10)) directly from the phasor diagram.

Useful solution

Since we now know what to look out for, we can find $G \equiv |A_{tot}|$ in a more direct manner as follows:

$$\begin{aligned}
 \left| \vec{A}_{tot} \right|^2 &= \vec{A}_{tot} \cdot \vec{A}_{tot} \\
 &= \left(\sum_{k=0}^M \vec{A}_k \right) \cdot \left(\sum_{n=0}^M \vec{A}_n \right) \\
 &= \sum_{k=0}^M \sum_{n=0}^M \vec{A}_k \cdot \vec{A}_n \\
 &= \sum_{k=0}^M \sum_{n=0}^M A_k A_n \cos [\phi_k - \phi_n]
 \end{aligned} \tag{D.18}$$

The last step follows directly from the geometric meaning of the scalar product and Fig. D.1. We must emphasize again that the proper ways to apply these approaches were apparent to us only once the solution was obtained rigorously and free of ambiguity as described in the initial sections.

D.3 Example: The heat current experiment with multiple weak-links

D.3.1 Problem statement

Consider the case of $M + 1$ weak-links (labeled $0, 1, \dots, M$) in the Josephson regime with sinusoidal current-phase relations. The weak-links are arranged equidistant in a straight line along a channel with a steady superfluid counterflow set up along the channel using a heat current (see Fig. D.2).




Figure D.2: Uniform phase gradient set up using superfluid counterflow due to heat current. R is a heater and S is a sink. The green rectangles are the weak-links. The phase increases uniformly from left to right so that the weak-links see an incremental increase in phase (the increment being δ).

D.3.2 General solution for $M+1$ weak-links

A steady flow implies a uniform phase-gradient ($v_s = \frac{\hbar}{m_A} \nabla \phi$) so that each weak-link differs from its neighbors by equal amounts (say δ). Without loss of generality, let the phase-offset ϕ_0 of the 0^{th} weak-link be 0. Then, the phase-offset of the k^{th} weak-link will be:

$$\phi_k = k \delta \quad (k = 0, 1, 2, \dots, M) \quad (\text{D.19})$$

Using Eq. (D.19) in Eq. (D.10) gives:

$$G^2 = \sum_{n=0}^M \sum_{k=0}^M A_k A_n [\cos(k - n) \delta] \quad (\text{D.20})$$

Again, taking a hint from the suggestive nature of Eq. (D.10), we can rewrite this to group terms by the harmonics $\cos \delta$, $\cos 2\delta$ and so on to reflect the mixing between different

neighbors. This can be accomplished by a change of variables in the sum above. Let $p \equiv k - n$ so that $k = p + n$. We use the fact that $\cos x = \cos(-x)$ to note that all terms (except when $k = n$ so that $p = 0$) appear twice in the double sum above. So far, we have:

$$G^2 = \sum_{n=0}^M A_n^2 \cos(0) + 2 \sum_{p=?}^? \sum_{n=?}^? A_{p+n} A_n \cos(p \delta) \quad (\text{D.21})$$

The limits in the above sum must be determined. Since $p = 0$ has been taken care of, $p_{min} = 1$. p_{max} occurs when one of k or n is M and the other is zero so that $p_{max} = M$. This just reflects the fact that we have all harmonics in from $\cos 0$ to $\cos M\delta$ in the sum, which is clear from Eq. (D.20). Eq. (D.21) now becomes:

$$G^2 = \sum_{n=0}^M A_n^2 + 2 \sum_{p=1}^M \left[\sum_{n=?}^? A_{p+n} A_n \right] \cos(p \delta) \quad (\text{D.22})$$

We can say right away that $n_{min} = 0$. Further, it is clear from Eq. (D.22) that the limits on n will depend on the current value of p during the outer sum, since A_{p+n} cannot remain valid for all possible values of n and p (e.g. if $p = M$ and $n = M$ the term becomes A_{2M} which does not exist!) This suggests that $p + n_{max} = \text{maximum possible index} = M$ so that $n_{max} = M - p$. Putting all this together finally yields:

$$\boxed{G^2 = \sum_{n=0}^M A_n^2 + 2 \sum_{p=1}^M \left[\sum_{n=0}^{M-p} A_{p+n} A_n \right] \cos(p \delta)} \quad (\text{D.23})$$

D.3.3 Clarifying the patterns: combinatorial diagrams

Eq. (D.23) displays a beautiful pattern that exhibits itself upon writing out a few of its terms:

$$G^2 = \sum_{n=0}^M A_n^2 + 2 \left[\sum_{n=0}^{M-1} A_{n+1} A_n \right] \cos \delta + 2 \left[\sum_{n=0}^{M-2} A_{n+2} A_n \right] \cos 2\delta \dots \quad (\text{D.24})$$

The first term is the “self-interaction” (0δ), the second term is the “nearest neighbor interaction” (neighbors differing by 1δ) and so on. This simple “Feynman diagram” visualization (see Fig. D.3 for a diagrammatic illustration of 4 weak-links) should help to quickly write down the asymmetric resultant amplitude for any number of weak-links.

D.3.4 2 weak-link case with comparisons to the usual fitting formula

The resultant amplitude for 2 weak-links can be found by using Eq. (D.23) with $M = 1$ giving:

$$\begin{aligned}
(G_{two})^2 &= \sum_{n=0}^1 A_n^2 + 2 \left[\sum_{n=0}^0 A_{n+1} A_n \right] \cos \delta \\
&= (A_0^2 + A_1^2) + 2A_0 A_1 \cos \delta \\
&= (A_0^2 + A_1^2) \left[\cos^2 \left(\frac{\delta}{2} \right) + \sin^2 \left(\frac{\delta}{2} \right) \right] + 2A_0 A_1 \left[\cos^2 \left(\frac{\delta}{2} \right) - \sin^2 \left(\frac{\delta}{2} \right) \right] \\
&= (A_0^2 + A_1^2 + 2A_0 A_1) \cos^2 \left(\frac{\delta}{2} \right) + (A_0^2 + A_1^2 - 2A_0 A_1) \sin^2 \left(\frac{\delta}{2} \right) \\
&= (A_0 + A_1)^2 \left[\cos^2 \left(\frac{\delta}{2} \right) + \left(\frac{A_0 - A_1}{A_0 + A_1} \right)^2 \sin^2 \left(\frac{\delta}{2} \right) \right] \\
G_{two} &= (A_0 + A_1) \left[\cos^2 \left(\frac{\delta}{2} \right) + \left(\frac{A_0 - A_1}{A_0 + A_1} \right)^2 \sin^2 \left(\frac{\delta}{2} \right) \right]^{\frac{1}{2}} \tag{D.25}
\end{aligned}$$

where we have used the trigonometric identities: $\cos^2 \theta + \sin^2 \theta = 1$ and $\cos^2 \theta - \sin^2 \theta = \cos 2\theta$ to re-write the two-slit expression (easily obtained from our master equation in the second line above) in a form commonly used in past publications. The combined asymmetric amplitude expression from Equation 4 of [4] (re-written using the symbols introduced here) is:

$$G'_{two} = (A_0 + A_1) \left[\cos^2 \left(\frac{\delta}{2} \right) + \gamma^2 \sin^2 \left(\frac{\delta}{2} \right) \right]^{\frac{1}{2}} \tag{D.26}$$

where $\gamma \equiv \frac{A_0 - A_1}{A_0 + A_1}$ is known as the (2-slit) asymmetry factor. Eq. (D.25) and Eq. (D.26) are manifestly identical, which bodes well for the accuracy of our master equation (Eq. (D.23)).

Further, it is clear from Eq. (D.24) that an N-slit interferometer will have (N-1) unique asymmetry factors, though in general they will have to be defined as the coefficients of each of the $\cos(p\delta)$ terms (for non-zero p) normalized by the first (constant) term. For example, in the 2-slit case, the generalized asymmetry factor would be $2A_0 A_1 / (A_0^2 + A_1^2)$ rather than the γ conventionally used.

0	1	2	3	modulation	coefficient
				$\cos 0$	$(A_0^2 + A_1^2 + A_2^2 + A_3^2)$
				$\cos \delta$	$2(A_0A_1 + A_1A_2 + A_2A_3)$
				$\cos 2\delta$	$2(A_0A_2 + A_1A_3)$
				$\cos 3\delta$	$2(A_0A_3)$

Figure D.3: Illustration of the “Feynman diagram” approach to writing the asymmetric amplitude for the case of 4 chips. Note that the first term has only single loops (only one way to connect a weak-link with itself). The rest of the terms have 2 ways each (shown in red and green) and so have the overall factor of 2 in their coefficients. The entire expression for the 4 weak-link resultant amplitude is given by Eq. (D.28).

D.3.5 Fully worked out 4 weak-link example

The resultant amplitude for 4 weak-links can be found by using Eq. (D.23) with $M = 3$ giving:

$$(G_{four})^2 = \sum_{n=0}^3 A_n^2 + 2 \left[\sum_{n=0}^2 A_{n+1}A_n \right] \cos \delta + 2 \left[\sum_{n=0}^1 A_{n+2}A_n \right] \cos 2\delta + 2 \left[\sum_{n=0}^0 A_{n+3}A_n \right] \cos 3\delta \quad (\text{D.27})$$

Writing it out explicitly (also see Fig. D.3), we obtain:

$$(G_{four})^2 = (A_0^2 + A_1^2 + A_2^2 + A_3^2) + 2(A_0A_1 + A_1A_2 + A_2A_3) \cos(\delta) + 2(A_0A_2 + A_1A_3) \cos(2\delta) + 2(A_0A_3) \cos(3\delta) \quad (\text{D.28})$$

D.4 Heat Current Experiment: closed form solution for equal amplitudes

If the amplitudes A_k are all equal, the resulting series can be summed. However, we can use a trick to find a closed-form expression for this case in a much easier fashion if we start from Eq. (D.1) with $A_k = A$ (note that $\Im\{x\}$ denotes the imaginary part of x):

$$I_{tot} = \sum_{k=0}^M A \sin(\omega t + \phi_k) \quad (\text{D.29a})$$

$$= \Im \left\{ \sum_{k=0}^M A e^{i(\omega t + \phi_k)} \right\} \quad (\text{D.29b})$$

$$= \Im \left\{ A e^{i\omega t} \sum_{k=0}^M e^{i\phi_k} \right\} \quad (\text{D.29c})$$

$$= \Im \left\{ A e^{i\omega t} \sum_{k=0}^M e^{ik\delta} \right\} \quad \dots \text{from Eq. (D.19)} \quad (\text{D.29d})$$

$$= \Im \left\{ A e^{i\omega t} \sum_{k=0}^M (e^{i\delta})^k \right\} \quad \dots \text{(see note below)} \quad (\text{D.29e})$$

$$= \Im \left\{ A e^{i\omega t} \left(\frac{(e^{i\delta})^{M+1} - 1}{e^{i\delta} - 1} \right) \right\} = \Im \left\{ A e^{i\omega t} \left(\frac{e^{i\delta(M+1)} - 1}{e^{i\delta} - 1} \right) \right\} \quad (\text{D.29f})$$

Note: A geometric sum is performed in Eq. (D.29e) - $\sum_{k=0}^M r^k = \frac{r^{M+1} - 1}{r - 1}$, where $r = e^{i\delta}$.

Applying the identity $e^{ix} - 1 = e^{i\frac{x}{2}} (e^{i\frac{x}{2}} - e^{-i\frac{x}{2}}) = e^{i\frac{x}{2}} 2i \sin \frac{x}{2}$ separately to the numerator and denominator of the fraction in Eq. (D.29f) and rearranging, we get:

$$I_{tot} = \Im \left\{ A e^{i\omega t} \left(\frac{e^{i\frac{\delta(M+1)}{2}} 2i \sin \frac{\delta(M+1)}{2}}{e^{i\frac{\delta}{2}} 2i \sin \frac{\delta}{2}} \right) \right\} \quad (\text{D.30a})$$

$$= \Im \left\{ A e^{i\omega t} \left(\frac{e^{i\frac{\delta(M+1)}{2}} e^{-i\frac{\delta}{2}} \sin \frac{\delta(M+1)}{2}}{\sin \frac{\delta}{2}} \right) \right\} \quad (\text{D.30b})$$

$$= \Im \left\{ A e^{i\omega t} e^{i\frac{M\delta}{2}} \left(\frac{\sin \frac{\delta(M+1)}{2}}{\sin \frac{\delta}{2}} \right) \right\} \quad (\text{D.30c})$$

$$= \Im \left\{ A e^{i(\omega t + \frac{M\delta}{2})} \left(\frac{\sin \frac{\delta(M+1)}{2}}{\sin \frac{\delta}{2}} \right) \right\} \quad (\text{D.30d})$$

Finally,

$$I_{tot} = \left[A \frac{\sin \frac{(M+1)\delta}{2}}{\sin \frac{\delta}{2}} \right] \sin \left(\omega t + \frac{M\delta}{2} \right) \quad (\text{D.31})$$

where the quantity in brackets is the modulated amplitude. The superposed oscillation, as expected, has the same frequency as its constituent oscillations.

D.5 Plots and code listings

The formulae derived above are plotted in this section for various numbers of aperture array chips in the SHeQUID. Note that we have already displayed the plot for a two chip-SHeQUID




Figure D.4: SHeQUID with 4 chips. Critical current in each chip is 1 (or less). When all critical currents are equal to 1, we get a symmetric SHeQUID. Total maximum amplitude is therefore 4 in this case. The plot legend shows a 4-vector for each plot, denoting the values of the 4 critical currents (randomly generated in the range $[0, 1]$).




Figure D.5: SHeQUID with 10 chips. Critical current in each chip is 1 (or less). When all critical currents are equal to 1, we get a symmetric SHeQUID. Total maximum amplitude is therefore 10 in this case. See Fig. D.6 for plot legend, which shows a 10-vector for each plot, denoting the values of the 10 critical currents (randomly generated in the range $[0, 1]$).




Figure D.6: Zoomed in version of Fig. D.5.

■ **The total amplitude by adding currents from "numchips" number of chips is:**

`In[1]:= VF[A_, i_] := A[[i + 1]]`

$$G[\text{numchips_}, \mathbf{A}_] := \sqrt{\left(\left(\sum_{n=0}^{\text{numchips-1}} \text{VF}[\mathbf{A}, n]^2 \right) + 2 \left(\sum_{p=1}^{\text{numchips-1}} \left(\sum_{n=0}^{\text{numchips-1-p}} \text{VF}[\mathbf{A}, n] \text{VF}[\mathbf{A}, p + n] \text{Cos}[p \delta] \right) \right) \right)}$$

- Here, $\delta \equiv \frac{\Delta\phi}{\text{numchips}-1}$ is the phase-offset between two adjacent chips, where $\Delta\phi$ is the phase difference between the first and last chip. This of course assumes that the chips are placed uniformly in a line and further that the phase gradient along the line is also uniform (say, like the heater current phenomenon).

■ **Symmetric amplitude (for identical chips)**

`In[3]:= S[numchips_] := Abs[$\left[\text{IO} \frac{\text{Sin}\left[\frac{\text{numchips}}{2} \delta\right]}{\text{Sin}\left[\frac{\delta}{2}\right]} \right]$] /. {IO → 1};`

■ **Testing the formulae: M = numchips - 1 (in writeup)**

`In[4]:= G[3, Array[A, 3]]`

`Out[4]=` $\sqrt{(A[1]^2 + A[2]^2 + A[3]^2 + 2(A[1]A[2]\text{Cos}[\delta] + A[2]A[3]\text{Cos}[\delta] + A[1]A[3]\text{Cos}[2\delta]))}$

`In[5]:= S[n]`

`Out[5]=` $\text{Abs}\left[\text{Csc}\left[\frac{\delta}{2}\right] \text{Sin}\left[\frac{n\delta}{2}\right]\right]$

Figure D.7: Mathematica code for applying results in this chapter

Appendix E

Effect of finite fluid compressibility

This chapter is a detailed derivation of the equations used in Section 3.2.1 to set constraints on the inner cell volume in weak link experimental cells.

E.1 Introduction - the zero order approximation




Figure E.1: Single weak link cell. Number of weak links is irrelevant (i.e. the analysis holds) as long as the structure of the inner cell is unchanged. Note that a magnet-style displacement sensor is shown (blue rectangle hanging from the diaphragm is the magnet). The analysis holds regardless of the presence or absence of the magnet.

In a weak link cell (see Fig. E.1), we usually make the assumption that the fluid in the inner cell is incompressible so that some entering fluid (having mass δm) directly leads to a *final observed* increase in the inner cell volume (δV_{obs}) given by -

$$\delta V_{obs} = Ax_{obs} = \frac{\delta m}{\rho_0} \quad (\text{E.1})$$

where ρ_0 is the total density of the fluid inside the cell. The diaphragm area is A and the spring constant is k .

This assumption is incorrect in at least two aspects:

1. The entering mass of fluid increases the cell pressure, thereby compressing the fluid in the cell. This means that the actual observed volume change will be *less* than the simple approximation in Eq. (E.1).
2. The density of the fluid in the cell is different from the normal density (at that temperature) for the same reason - the fluid is compressed. These quantities are all tangled up together in potentially complicated ways.

We will try to untangle these relationships at increasing levels of sophistication.

E.2 Compression with constant density

As a first approximation, we continue to assume that the inner cell fluid density remains constant at ρ_0 with flow into and out of the cell (at whatever normal value it's supposed to be at a given temperature - but with no dependencies other than temperature). Obviously, this is not a consistent approximation because fluid compression cannot be decoupled from a density change - you must have both or neither in your model. Still, it helps get us started mathematically (and as we shall see in the next section, the error that results from this assumption is much smaller than the 1st order deviation we obtain in this section).

The story unfolds as follows:

1. Let the volume of the inner cell at perfect equilibrium be $V_0 \equiv AH$, so that the the position of the diaphragm is $x = 0$, where x is always measured relative to the relaxed position, the pressure (relative to the outer cell) and temperature in the inner cell is $P_0 = 0$ and T_0 respectively and the total density of the inner cell fluid is ρ_0 . Here, H is the height of the inner cell.
2. At some instant of time (say, during the cell expansion part of an oscillation), suppose some fluid (with mass δm) enters the cell over a time δt . To avoid any confusion, we imagine a gatekeeper who shovels the mass into the inner cell and closes the gate at the end of the duration δt , leaving the system to settle down. Makes it easier to visualize what's going on.
3. At the end of this time duration¹, suppose the entering fluid has pushed out the diaphragm by a final observed amount equal to x_{obs} , so that we can define an observed

¹ δt is assumed to be larger than the time required for any internal equilibration related to the compressibility. This should happen at timescales related to the speed of first sound ($c_1 \approx 220m/s$ near T_λ). The

volume change in the inner cell as $\delta V_{obs} \equiv Ax_{obs}$. Again, these volume changes are relative to the equilibrium (relaxed) inner cell volume V_0 .

4. The increase in the diaphragm restoring force $\delta F = kx_{obs}$ must be balanced by the increase in inner cell pressure:

$$\delta P = \frac{kx_{obs}}{A} \quad (\text{E.2})$$

5. This pressure change must have compressed the fluid by an amount defined by $\kappa = 1/B \equiv -\frac{1}{V_0} \frac{\partial V}{\partial P}$, where κ is the compressibility of the fluid in units of $1/Pa$ (and B is the bulk modulus). So, the volume change (a decrease [-] for pressure increase [+]) due to compression only is well approximated by:

$$\delta V_{comp} \approx -\kappa V_0 \delta P \quad (\text{E.3})$$

for sufficiently small changes in pressure so that the compression is small compared to V_0 .

6. For an ideal incompressible fluid, the observed volume change would have been:

$$\delta V_{obs,ideal} = \frac{\delta m}{\rho_0} \quad (\text{E.4})$$

7. A compressible fluid would therefore have this reduced by the amount in Eq. (E.3):

$$\delta V_{obs} = \frac{\delta m}{\rho_0} - \kappa V_0 \delta P \quad (\text{E.5})$$

8. Using Eq. (E.2) and solving for x_{obs} :

$$\begin{aligned} \delta V_{obs} &= \frac{\delta m}{\rho_0} - \kappa V_0 \delta P \\ Ax_{obs} &= \frac{\delta m}{\rho_0} - \kappa V_0 \frac{kx_{obs}}{A} \\ Ax_{obs} + \kappa V_0 \frac{kx_{obs}}{A} &= \frac{\delta m}{\rho_0} \\ Ax_{obs} \left(1 + \kappa V_0 \frac{k}{A^2} \right) &= \frac{\delta m}{\rho_0} \\ x_{obs} &= \frac{\delta m}{A\rho_0} \left(1 + \kappa V_0 \frac{k}{A^2} \right)^{-1} \end{aligned}$$

upper limit for the equilibration time will be for the largest dimension in the inner cell ($\sim 3/8''$) and is $\sim 50\mu s$. Events that happen over timescales shorter than this will be smeared together. This corresponds to an oscillation frequency of $\sim 20kHz$. So, longer times are good, i.e. frequencies shorter than this are observable without ambiguity of equilibration. Of course, these numbers are extreme because the vertical distances are probably more important than lateral ones.

Finally,

$$x_{obs} = \frac{\delta m}{A\rho_0} \left(1 + \frac{V_0}{V_c}\right)^{-1} \quad (\text{E.6})$$

where we have defined a characteristic volume $V_c \equiv \frac{A^2}{\kappa k}$.

The fractional deviation ϵ in the observed displacement from the ideal incompressible case is then:

$$\epsilon \equiv \frac{\frac{\delta m}{A\rho_0} \left(1 + \frac{V_0}{V_c}\right)^{-1} - \frac{\delta m}{A\rho_0}}{\frac{\delta m}{A\rho_0}} = \left(1 + \frac{V_0}{V_c}\right)^{-1} - 1 \approx -\frac{V_0}{V_c} \quad (\text{E.7})$$

where the final approximation is to keep terms only to first order in V_0/V_c (one can always use the full expression since this is the first appearance of this approximation).

E.3 Estimates

The compressibility can be estimated from the formula for the speed of sound in a fluid:

$$c_1 = \sqrt{\frac{B}{\rho}} = \sqrt{\frac{1}{\kappa\rho}} \quad (\text{E.8})$$

which gives

$$\kappa = \frac{1}{c_1^2\rho} \sim \frac{1}{220^2 \times 146} \sim 1.42\text{e-}7 \frac{1}{Pa} \quad (\text{E.9})$$

We recall that the characteristic volume is $V_c \equiv \frac{A^2}{\kappa k}$.

E.3.1 Old single weak link cell

- $V_0 \sim 5\text{e-}8 \text{ m}^3$
- $A = 4.95\text{e-}5 \text{ m}^2$
- $\frac{A^2}{k} = 1.76\text{e-}12 \text{ s}^2\text{m}^4/\text{kg}$ (from WL6 calibration)

This gives $V_c \sim 1.24\text{e-}5 \text{ m}^3$ and $\frac{V_0}{V_c} \sim 4\text{e-}3 \ll 1$ which doesn't lead us to expect much difference from the incompressible case. From Eq. (E.7) the fractional decrease in the displacement signal compared to the incompressible ideal is $\sim 1/1.004 - 1 = -0.00398 \approx -0.4\%$.

E.3.2 New interferometer

- $V_0 \sim 3.6\text{e-}7 \text{ m}^3$
- $A = 7.13\text{e-}5 \text{ m}^2$

- $\frac{A^2}{k} = 2.53\text{e-}12 \text{ s}^2\text{m}^4/\text{kg}$ (scaled from WL6 calibration for new area, assuming same k)

This gives $V_c \sim 1.8\text{e-}5 \text{ m}^3$ and $\frac{V_0}{V_c} \sim 2\text{e-}2$. So, the fractional deviation from Eq. (E.7) is $\sim -0.0198 \sim 2\%$.

Note that to first order in $\frac{V_0}{V_c}$, the fractional decrease in the displacement signal compared to the incompressible ideal is simply $\sim \frac{V_0}{V_c}$.

E.4 Compression with changing density

The story proceeds exactly the same way as in the previous section, except we add an additional step where the density of the fluid is now allowed to change in response to the compression of the entering fluid. We note (presciently) that the additional correction term due to this is (in a very non-intuitive way) unimportant for the typical regimes we work in. This section is therefore included only for completeness. It is worthwhile to note however that the extra correction actually *reduces* the overall signal decrease due to the finite compressibility of the fluid.

The modified density is

$$\rho \equiv \frac{M}{V} = \frac{M_0 + \delta m}{V_0 + \delta V_{obs}} = \frac{M_0}{V_0} \frac{\left(1 + \frac{\delta m}{M_0}\right)}{\left(1 + \frac{\delta V_{obs}}{V_0}\right)} = \rho_0 \frac{\left(1 + \frac{\delta m}{M_0}\right)}{\left(1 + \frac{Ax_{obs}}{V_0}\right)} \quad (\text{E.10})$$

where M_0 and V_0 are the inner cell fluid mass and volume respectively when the diaphragm is perfectly relaxed and $\delta V_{obs} = Ax_{obs}$ as usual.

The expression for the pressure change in response to the entering fluid is unchanged from Eq. (E.2):

$$\delta P = \frac{kx_{obs}}{A} \quad (\text{E.11})$$

Eq. (E.5) describing the observed volume change is modified by replacing the normal density ρ_0 with the density at the new equilibrium state -

$$\delta V_{obs} = \frac{\delta m}{\rho} - \kappa V_0 \delta P \quad (\text{E.12})$$

Note that the approximation for the differentials in Eq. (E.3) is still good since the typical volume changes due to the diaphragm motion are much smaller than the relaxed volume². Dividing the master equation (E.12) for the volume change by V_0 and using Eq. (E.11) for

²For comparison, the largest displacements are $\mathcal{O}(\text{nm})$ and the cell height H is $\mathcal{O}(0.1\text{'})$ so that the ratio $\delta V/V_0$ is $\mathcal{O}(10^{-7})$

the pressure, we obtain

$$\begin{aligned}
\delta V_{obs} &= \frac{\delta m}{\rho} - \kappa V_0 \delta P \\
\frac{Ax_{obs}}{V_0} &= \frac{\delta m}{\rho V_0} - \kappa \frac{kx_{obs}}{A} \\
\frac{Ax_{obs}}{V_0} + \kappa \frac{kx_{obs}}{A} &= \frac{\delta m}{\rho V_0} \\
\frac{Ax_{obs}}{V_0} \left(1 + V_0 \frac{\kappa k}{A^2}\right) &= \frac{\delta m}{M_0} \frac{M_0}{\rho V_0} \\
\frac{Ax_{obs}}{V_0} \left(1 + \frac{V_0}{V_c}\right) &= \frac{\delta m}{M_0} \frac{\rho_0}{\rho}
\end{aligned} \tag{E.13}$$

As before, we have defined the characteristic volume $V_c \equiv \frac{A^2}{\kappa k}$ in the last line. Using the density equation (Eq. (E.10)) allows us to simplify this further -

$$\begin{aligned}
\frac{Ax_{obs}}{V_0} \left(1 + \frac{V_0}{V_c}\right) &= \frac{\delta m}{M_0} \left(1 + \frac{Ax_{obs}}{V_0}\right) \\
\frac{Ax_{obs}}{V_0} \left(1 + \frac{V_0}{V_c}\right) &= \gamma \left(1 + \frac{Ax_{obs}}{V_0}\right)
\end{aligned}$$

where we have defined (for convenience) -

$$\frac{\delta m}{M_0} \frac{1}{\left(1 + \frac{\delta m}{M_0}\right)} \equiv \gamma \tag{E.14}$$

Simplifying further -

$$\begin{aligned}
\frac{Ax_{obs}}{V_0} \left(1 + \frac{V_0}{V_c}\right) &= \gamma + \frac{A\gamma}{V_0} x_{obs} \\
\frac{Ax_{obs}}{V_0} \left(1 + \frac{V_0}{V_c}\right) - \frac{A\gamma}{V_0} x_{obs} &= \gamma \\
\frac{Ax_{obs}}{V_0} \left(1 + \left[\frac{V_0}{V_c} - \gamma\right]\right) &= \gamma \\
x_{obs} &= \frac{\gamma V_0}{A} \left(1 + \left[\frac{V_0}{V_c} - \gamma\right]\right)^{-1}
\end{aligned} \tag{E.15}$$

No assumptions up to this point (except that any volume changes are small in comparison to the relaxed volume so that it can be approximated by the differential as before).

To get an idea of how γ enters into play, we can examine it again -

$$\gamma = \frac{\delta m}{M_0} \frac{1}{\left(1 + \frac{\delta m}{M_0}\right)} = \frac{1}{\left(\frac{M_0}{\delta m} + 1\right)} \approx \frac{1}{\frac{M_0}{\delta m}} = \frac{\delta m}{M_0} \quad (\text{E.16})$$

where the approximation is that $M_0 \gg \delta m$ (which allows us to neglect the 1 in the denominator)³.

With this approximation, Eq. (E.15) becomes

$$x_{obs} = \frac{\delta m V_0}{M_0 A} \left(1 + \left[\frac{V_0}{V_c} - \frac{\delta m}{M_0}\right]\right)^{-1} = \frac{\delta m}{\rho_0 A} \left(1 + \left[\frac{V_0}{V_c} - \frac{\delta m}{M_0}\right]\right)^{-1} \quad (\text{E.17})$$

Comparing this to our first approximation (Eq. (E.6)), since $V_0/V_c \sim \mathcal{O}(10^{-3})$ (for single weaklink cell run 6) or $\mathcal{O}(10^{-2})$ (for modular SHeQUID cell) and $\delta m/M_0 \sim \mathcal{O}(10^{-7})$ (from Section E.3), we can see that the extra term due to the density change is much smaller than the correction term. Our previous estimates are therefore quite robust.

³From footnote 2 on p. 341: $\frac{\delta V}{V_0} \sim \mathcal{O}(10^{-7})$. So, $\frac{\delta m}{M_0} = \frac{\rho \delta V}{\rho_0 V_0} \sim \mathcal{O}(10^{-7})$ since $\frac{\rho}{\rho_0} \sim \mathcal{O}(1)$

Appendix F

Inertial effects and resonant frequencies

This chapter collects several short theoretical discussions and derivations of concepts used in this dissertation (not addressed elsewhere).

F.1 Inertial effects

Here we discuss inertial effects of the mass of the diaphragm as they pertain to the signal loss due to compressibility that was discussed in Section 3.2.1 and the previous Appendix E. Reconsidering the compressibility of the helium in the inner cell κ (as defined near Eq. (E.3)):

$$\kappa \equiv \frac{-1}{V} \frac{\partial V}{\partial P} = V \frac{\partial(1/V)}{\partial P} = \frac{V}{M} \frac{\partial(M/V)}{\partial P} = \frac{1}{\rho} \frac{\partial \rho}{\partial P}$$

where ρ is the density of and P is the net pressure seen by the inner cell fluid (what we call ΔP elsewhere).

The time-rate of change of the density therefore becomes:

$$\dot{\rho} = \frac{\partial \rho}{\partial t} = \frac{\partial \rho}{\partial P} \frac{\partial P}{\partial t} = \rho \kappa \dot{P} \quad (\text{F.1})$$

Considering the archetypal cell in Fig. E.1 again, we can write the equation of motion for the flexible diaphragm (with area A , mass m_d , spring-constant k , and mass density per unit area $\sigma = m_d/A$) as:

$$\begin{aligned} F_{net} = m_d \ddot{x} &= \sigma A \ddot{x} = PA - kx \\ \Rightarrow P &= \frac{k}{A} x + \sigma \ddot{x} \end{aligned} \quad (\text{F.2})$$

The total fluid mass in the inner cell is $M = \rho V_{in} = \rho A(x + H)$, where H , as before, is the equilibrium height of the cell with a relaxed diaphragm (so that the equilibrium inner cell

volume is $V_0 = AH$). The total mass current into the inner cell is then the time-derivative of the inner cell fluid mass:

$$\begin{aligned}\dot{M} \equiv I_t &= \rho A \dot{x} + A(x + H) \dot{\rho} \\ &= \rho A \dot{x} + A(x + H) (\rho \kappa \dot{P}) \\ &= \rho A \dot{x} + \rho A \kappa H \left(1 + \frac{x}{H}\right) \left(\frac{k}{A} \dot{x} + \sigma \ddot{x}\right)\end{aligned}$$

where, in the second line, we used Eq. (F.1) and in the third line, we used Eq. (F.2). From footnote 2 on p.341, the maximum values of x/H are $\mathcal{O}(10^{-7})$, so that we can neglect it in the last line above to obtain:

$$I_t = \rho A \left(\dot{x} + \kappa H \left[\frac{k}{A} \dot{x} + \sigma \ddot{x} \right] \right) \quad (\text{F.3})$$

Now, what we measure in experiments is the integrated whistle peak in the frequency spectrum of the displacement signal $x(t)$ and use it to deduce the peak mass current amplitude. We therefore take the Fourier transform¹ of both sides above, while remembering that the Fourier transform of the n 'th time-derivative goes as $\mathcal{F}(d^n x/dt^n) = (i\omega)^n \mathcal{F}(x)$ (where $\mathcal{F}(x)$ is the Fourier transform of the displacement signal):

$$\begin{aligned}\mathcal{F}(I_t) &= \rho A \left((i\omega) \mathcal{F}(x) + \kappa H \left[\frac{k}{A} (i\omega) \mathcal{F}(x) + \sigma (i\omega)^3 \mathcal{F}(x) \right] \right) \\ &= (i\omega) \rho A \left(1 + \kappa H \left[\frac{k}{A} - \sigma \omega^2 \right] \right) \mathcal{F}(x) \\ &= (i\omega) \rho A \left(1 + \kappa \frac{kV_0}{A^2} \left[1 - \sigma \frac{A}{k} \omega^2 \right] \right) \mathcal{F}(x) \\ &\equiv (i\omega) G(\omega) \mathcal{F}(x)\end{aligned}$$

where, in the last line we have used $H = V_0/A$ and defined a new function $G(\omega)$. Rearranging, we finally obtain:

$$\boxed{\mathcal{F}(x) = \frac{\mathcal{F}(I_t)}{i\omega G(\omega)} \quad \dots \text{ where, } G(\omega) = \rho A \left(1 + \kappa \frac{kV_0}{A^2} \left[1 - \sigma \frac{A}{k} \omega^2 \right] \right)} \quad (\text{F.4})$$

F.1.1 Discussion

There is a wealth of information compressed into this one equation. The first thing to note is that this expression is the frequency spectrum that we will observe in experiments, and it

¹The general idea of looking at the Fourier transform of the equation of motion to investigate frequency dependence of transfer functions is based on examples found in an unpublished report by Emile Hoskinson.

is always attenuated by a factor of ω irrespective of the compressibility and inertial terms. This is simply the attenuation one would expect from using a displacement sensor to measure velocities (currents) - a $1/\omega$ falloff in sensitivity over frequency.

Turning our attention now to $G(\omega)$, in the case of a perfectly incompressible fluid ($\kappa = 0$), the inertial effects of the diaphragm mass are killed off. Conversely, for a perfectly massless diaphragm ($\sigma = 0$), the compressibility term still plays a role in attenuating the displacement signal. In fact, we recognize the massless diaphragm case $G(\omega) \approx \rho A \left(1 + \kappa \frac{kV_0}{A^2}\right) = \rho A \left(1 + \frac{V_0}{V_c}\right)$ as identical to the result (Eq. (E.6)) that we obtained in our quasi-static treatment of compressibility effects in Appendix E (where we also defined a characteristic volume $V_c \equiv \frac{A^2}{\kappa k}$).

In the real scenario where neither effect is *a priori* negligible, we note (from Section F.2) that the (fundamental) vacuum resonant frequency² of the diaphragm is $\omega_0 = \frac{y_0}{b} \sqrt{\frac{k}{8\pi\sigma}}$, where b is the diaphragm radius and $y_0 = 2.40483$ is the first zero of the Bessel function of the first kind. We can re-write this in terms of the diaphragm area as: $\frac{\sigma A}{k} \sim \frac{0.72}{\omega_0^2} \sim \frac{1}{(1.18\omega_0)^2} \sim \frac{1}{\omega_0^2}$. Eq. (F.4) then becomes:

$$G(\omega) = \rho A \left(1 + \frac{V_0}{V_c} \left[1 - \left(\frac{\omega}{\omega_0}\right)^2\right]\right) \quad (\text{F.5})$$

This finally allows us to establish three distinct regimes of interest (note that $\omega/\omega_0 = f/f_0$ and $f_0 \sim 10\text{kHz}$ for typical diaphragms that we have used in the past):

$\omega \ll \omega_0$: $G(\omega) \approx \rho A \left(1 + \frac{V_0}{V_c}\right)$. This is essentially the massless diaphragm case and only the compressibility plays a role in attenuating the displacement signal in a way that we saw in Appendix E.

$\omega \sim \omega_0$: $G(\omega) \approx \rho A$. This is the most surprising regime. Apparently, near resonance, the two effects nullify each other and we are left with a good approximation of a massless, incompressible case. In practice however, (for judicious cell designs where the compressibility effect is much reduced), this happy accident doesn't help much when compared to the always present $1/\omega$ attenuation discussed earlier.

$\omega \gg \omega_0$: This is a regime unlikely to be useful for many other reasons. In this context as well, we are left with the full expression for $G(\omega)$ and the attenuation is now enhanced by both effects together (in addition to the $1/\omega$ attenuation).

The optimal working frequency (considering only the inertial and compressibility effects described here) can be found by minimizing the denominator of Eq. (F.4) $\rightarrow \omega G(\omega)$. This

²This is only for the simple displacement sensor with a superconducting metal coating. The analysis from this point forward is not valid for the magnet-based displacement sensor.

optimal frequency is:

$$f_{opt} = f_0 \sqrt{\frac{1}{3} \left(\frac{V_c}{V_0} + 1 \right)} \quad (\text{F.6})$$

For the interferometer numbers given in Section E.3.2, $f_{opt} \sim 4f_0 \sim 40\text{kHz}$.

However, practically speaking, being close to the resonant frequency is a better option since we don't lose much sensitivity but the non-trivial transfer function in Eq. (F.4) becomes the simplest possible: $\mathcal{F}(x) = \frac{\mathcal{F}(I_t)}{i\omega A\rho}$. Calibrations and theoretical treatments can therefore be carried out in extremely simple terms by merely assuming incompressibility and ignoring the inertial effects of the diaphragm mass. All things considered, working at the diaphragm resonance would seem like the most optimal solution all round, except that in the work described in this dissertation, we reach a limit (at around 4 kHz) where the signal to noise (due to the $\sim 1/\omega$ attenuation) drops to $\sim 1:1$ and we cease to be able to measure the whistle.

In practical terms, the conclusion of this section is clear - stay as close as possible to the diaphragm vacuum resonance within the constraints of the signal to noise. In practice, we have found that the results of this section (particularly the inertial effects of the diaphragm) impose weaker constraints than those set by the simple $1/\omega$ sensitivity drop that kills the whistle (given our ambient noise). However, any changes in cell design should be checked against these constraints as we are within an order of magnitude of when we would expect them to become important.

F.2 Diaphragm resonant frequency

A circular membrane that is loaded with a mass at its center behaves quite differently from an unloaded one. For a membrane with mass per unit area σ and radius b under tension T that is loaded by a disk (magnet in this case) of mass m and radius a (where obviously $a < b$), the modified eigenfrequencies are given by

$$f_n = \frac{x_n}{2\pi b} \sqrt{\frac{T}{\sigma}} = \frac{x_n}{2\pi b} \sqrt{\frac{k}{8\pi\sigma}} \quad (\text{F.7})$$

where x_n is the n 'th eigenvalue of the equation of a freely (i.e. unforced) vibrating membrane and T is the surface tension. The tension is related to the effective spring constant of the membrane by [26, pp. 256-7]:

$$k = 8\pi T \quad (\text{F.8})$$

These eigenvalues x_n are the solutions to the equation [93]:

$$\frac{\alpha x J_0(\alpha x) - \gamma \alpha^2 J_1(\alpha x)}{\alpha x N_0(\alpha x) - \gamma \alpha^2 N_1(\alpha x)} = \frac{J_0(x)}{N_0(x)} \quad (\text{F.9})$$

where $J_m(z)$ and $N_m(z)$ are the Bessel functions of the first and second kind respectively, $\alpha = a/b$ and $\gamma = 2M/m$ (M is the membrane mass $\sigma\pi b^2$). In the limiting case where the

magnet disappears, Eq. F.9 reduces to the usual equation for a simple vibrating membrane found in acoustics books (such as [94, p. 189]) -

$$J_0(x) = 0 \quad (\text{F.10})$$

Note that Eq. F.9 actually depends on the various parameters of the membrane and magnet while Eq. F.10 does not. This means that for mass-loaded membranes, the eigenvalues must be recalculated each time a parameter is changed.

We note in passing that the effective spring constant from Ref. [26, pp. 256-7] should be recalculated for the magnet-loaded diaphragm. We do this using Poisson's equation for the shape of the diaphragm by a procedure similar to that in the quoted reference (with the boundary conditions that the membrane is clamped at the boundary and its slope is smoothly 0 at the edge of the glued magnet (no sharp edges)) and find that for the disk-loaded diaphragm,

$$k = \frac{8\pi T}{1 - \alpha^4 + 4\alpha^4 (\ln \alpha)} \quad (\text{F.11})$$

This expression reduces to Eq. (F.8) in the limit that the magnet disappears ($\alpha \rightarrow 0$). For the magnet and diaphragm dimensions used here ($a \sim 1/32''$ and $b \sim 3/16''$), the error in using Eq. (F.8) is around 0.6%. Since we are already at the point where larger magnets become unwieldy, it is doubtful that this will be important but we provide this information here for completeness.

F.3 Helmholtz resonance

The Helmholtz mode in superfluid cells has been discussed thoroughly in previous work so we merely quote the theoretical results here. Essentially, the cell dynamics equations in Appendix G can be manipulated [46, p. 75] to yield a "small-angle" pendulum-like equation in the chemical potential difference between the inner and outer cells. The resonant frequency of this linearized model can be simply obtained by inspection as:

$$f_h^2 = \frac{\rho_s \beta_s}{4\pi^2} \left(\frac{1}{\rho^2} \frac{k}{A^2} + \frac{s^2 T}{\rho c_p V_{cell}/M_4} \right) \quad (\text{F.12})$$

The energy decay time constant τ_H and quality factor Q can be similarly obtained -

$$\tau_H = \frac{(\rho c_p V_{cell}/M_4)^2 R}{\rho_s \beta_s T s^2} \omega_H^2 \quad (\text{F.13})$$

$$Q = \omega_H \tau_H \quad (\text{F.14})$$

F.4 Cavity acoustic resonances

The calculation of modal frequencies for acoustic resonances in a cylindrical (or box) cavity is a very old, solved problem and can be found in acoustics textbooks such as Morse [94, p. 398]. These modes are indexed by one axial (nz) and two tangential mode indices (m , n), whose precise physical meanings are discussed in the referenced text (essentially having to do with longitudinal, azimuthal and radial modes of fluid vibration). The cylindrical cavity radius is a and height is L and c is the speed of first sound in the fluid medium. The frequency is given by:

$$f(nz, m, n, a, L, c) = \frac{c}{2} \sqrt{\left[\frac{nz}{L}\right]^2 + \left[\frac{\alpha(m, n)}{a}\right]^2} \quad (\text{F.15})$$

F.4.1 Calculated modes for SHeQUID 3

Using the scripts provided in the next section, we calculate all possible standing wave modes in the cell (defining various, roughly cylindrical cavities defined by cell parts) and present the results in Table F.1 for modes nominally below 6 kHz (which is the maximum frequency in the experimental sweeps shown in Figs. 2.1 and 2.2). The first sound speed is set to $c = 218.0$ m/s for $T = 2.17$ K. The cell being modeled is the same cell (SHeQUID 3) used for the experimental frequency sweeps. The cavity descriptions (which can be easily matched to the various cell components described elsewhere in this dissertation) are also provided. Note that $d = 2a$ is the cavity diameter.

Table F.1: Cavity resonant modes: SHeQUID 3

cavity	f (hz)	nz	m	n
Outer cell (from bottom to ICP) [d = 1.850000 "] [L = 2.000000 "]	2145.669	1	0	0
	2718.932	0	1	0
	3463.595	1	1	0
	4291.339	2	0	0
	4510.284	0	2	0
Outer cell (From bottom to D-ring) [d = 1.850000 "] [L = 1.910000 "]	2246.774	1	0	0
	2718.932	0	1	0
	3527.121	1	1	0
	4493.548	2	0	0
	4510.284	0	2	0
Outer cell (From bottom to E-ring) [d = 1.850000 "] [L = 1.460000 "]	2718.932	0	1	0
	2939.273	1	0	0
	4003.988	1	1	0
	4510.284	0	2	0
	5383.492	1	2	0

... Continued on next page

Table F.1 – continued from previous page

cavity	f (hz)	nz	m	n
Inner cell [d = 0.375000 "] [L = 0.128200 "]	13413.397	0	1	0
	22250.735	0	2	0
	27914.754	0	0	1
	30606.513	0	3	0
	33473.780	1	0	0
Sense arm: V tunnel [d = 0.062500 "] [L = 0.846000 "]	5072.504	1	0	0
	10145.008	2	0	0
	15217.513	3	0	0
	80480.382	0	1	0
	80640.078	1	1	0
Sense arm: side arm [d = 0.086000 "] [L = 1.130000 "]	3797.645	1	0	0
	7595.290	2	0	0
	11392.934	3	0	0
	58488.650	0	1	0
	58611.810	1	1	0
Sense arm: heat tube [d = 0.086000 "] [L = 1.489000 "]	2882.027	1	0	0
	5764.055	2	0	0
	8646.082	3	0	0
	58488.650	0	1	0
	58559.612	1	1	0
Sense arm: V tunnel + side arm (use avg. radius and combined length) [d = 0.074000 "] [L = 1.980000 "]	2167.343	1	0	0
	4334.685	2	0	0
	6502.028	3	0	0
	67973.295	0	1	0
	68007.840	1	1	0
Between E and D [d = 1.000000 "] [L = 0.002992 "]	5030.024	0	1	0
	8344.026	0	2	0
	10468.033	0	0	1
	11477.442	0	3	0
	14565.240	0	1	1
Sense arm: return path 1 [d = 0.062500 "] [L = 0.688000 "]	6237.411	1	0	0
	12474.821	2	0	0
	18712.232	3	0	0
	80480.382	0	1	0
	80721.727	1	1	0
Sense arm: return path 2 [d = 0.081000 "] [L = 0.135000 "]	31787.693	1	0	0
	62099.060	0	1	0
	63575.386	2	0	0
	69762.101	1	1	0

... Continued on next page

Table F.1 – continued from previous page

cavity	f (hz)	nz	m	n
	88871.385	2	1	0
	1430.446	1	0	0
Sense arm: return path1+2((length-weighted avg. d)) [d = 0.065000 "] [L = 3.000000 "]	2860.892	2	0	0
	4291.339	3	0	0
	77384.982	0	1	0
	77398.202	1	1	0
	14304.462	1	0	0
Sense arm: tube flange well [d = 0.062300 "] [L = 0.300000 "]	28608.924	2	0	0
	42913.386	3	0	0
	80738.746	0	1	0
	81996.114	1	1	0

F.4.2 Matlab code

We present below a complete system for generating sorted lists of mode frequencies for a set of cylindrical cavities. The front-end script is modified by the user and executed. It uses function file “freq.m” for calculating the frequencies, which uses “alpha.m”, which finally uses “zerobess.m”, which is a Bessel zero calculator (see note in the “alpha.m” code-listing where it is called). The alpha script calculates the tangential and azimuthal mode contributions in the above expression.

Listing F.1: front_end.m

```
% 8/8/12 Aditya Joshi
% frontend (version 4) -SHEQUID 4 run1(new dimensions)
% Script to build table of standing wave modes in ascending order
% Function zerobess.m is used to calculate zeros of the derivate of the
% BesselJ function. Function freq.m is used by this front-end to call the
% zerobess function with the appropriate params.

%% Init
clear all; close all; clc;

%% Constants
% din = 1.85; %cylinder diameter in inches
% Lin = 2; %cylinder height in inches
% partlabel = 'outer cell (from bottom to ICP)'; %For saving in file
% filename = 'outer_cell.txt';
% a = (din*0.0254)/2; %cylinder radius in m
% L = Lin*0.0254; %cylinder height in m

c = 218; %Sound speed m/s (Donnelly value near Tlambda)

% Add more params in the format:
% params{n} = {'File header label' cylinder_ID(inches) cylinder_height(inches) 'Save_filename'}
sfol = '3_shequid4_run1'; %subfolder to save files
% If save folder doesn't exist, create it
if (exist(sfol, 'dir') ~= 7)
    mkdir(sfol);
end;
```

```

general_header = 'Shequid_4_run_1_[Cell_can_1.1379]_deep!'; %Appears as block at start of all files

params{1}={'outer_cell_(from_bottom_to_ICP)' 1.85 1.1379 'outer_cell_full.txt'};
params{2}={'Outer_cell_(From_bottom_to_D-ring)' 1.85 1.0297 'outer_cell_D.txt'};
params{3}={'Outer_cell_(From_bottom_to_E-ring)' 1.85 0.5817 'outer_cell_E.txt'};
params{4}={'Inner_cell' 0.375 0.0802 'Inner_cell.txt'};
params{5}={'Sense_arm:_V_tunnel' 0.0625 0.846 'Sense_arm_V_tunnel.txt'};
params{6}={'Sense_arm:_side_arm' 0.086 1.13 'Sense_arm_side_arm.txt'};
params{7}={'Sense_arm:_heat_tube' 0.086 1.489 'Sense_arm_heat_tube.txt'};
params{8}={'Sense_arm:_V_tunnel+_side_arm_(gentle_slope_may_make_it_one_single_tube)_(use_avg_radius_and_
    combined_length)' 0.074 1.98 'Sense_arm_V_tunnel_plus_side_arm.txt'};
params{9}={'Between_E_and_D' 1 0.0029921 'E_D_spacing.txt'};
params{10}={'Sense_arm:_return_path1' 0.0625 0.688 'Sense_arm_return_path1.txt'};
params{11}={'Sense_arm:_return_path2' 0.081 0.135 'Sense_arm_return_path2.txt'};
params{12}={'Sense_arm:_return_path1+2((length-weighted_avg_d))' 0.065 3.3 'Sense_arm_return_path_1_and_2.txt'};
    % 0.081 0.135 and 0.065 (length-weighted avg. d) 0.829
%Note that there is a small gap between the two sections of the return path (the gap defined by the indium seal)
params{13}={'Sense_arm:_tube_flange_well' 0.0623 0.3 'Sense_arm_tubeFL_well.txt'}; %Smaller d because tube flange lip
    takes up some space when assembled

%% max indices (all combinations will be looked at) to be generated
nzmax = 3;
mmax = 3;
nmax = 3;

%% smallest N frequencies to be put in consolidated file for each part
nFmax = 5;

%% Main seq
consol_data = cell(length(params),1);

for i = 1:length(params)
    sprintf('This_is_iteration_%d_of_%d!', i, length(params));

    %Get params
    temp_par = params{i};

    din = temp_par{2}; %cylinder diameter in inches
    Lin = temp_par{3}; %cylinder height in inches
    partlabel = sprintf('%s_[d=_%f_]_[H=_%f_]', temp_par{1}, din, Lin); %Header for saving in file
    filename = temp_par{4};

    a = (din*0.0254)/2; %cylinder radius in m
    L = Lin*0.0254; %cylinder height in m

    % Initialize arrays
    ntot = (nzmax+1)*(mmax+1)*(nmax+1);
    nzA = zeros(ntot,1);
    mA = nzA;
    nA = nzA;
    fA = nzA;

    ctr = 0; %vector index (data is flattened into a single dimension vector)

    for nz = 0:nzmax
        for m = 0:mmax
            for n = 0:nmax
                ctr = ctr+1;
                nzA(ctr) = nz;
                mA(ctr) = m;
                nA(ctr) = n;
                fA(ctr) = freq(nz,m,n,a,L,c);
            end
        end
    end
end

```



```
f = (c/2)*sqrt((nz/L)^2 +(alpha/a)^2);
```

```
end
```

```
%Modified 1/10/12 to be accurate (n was offset by 1)
```

Listing F.3: alpha.m

```
function x = alpha(m,n)
```

```
% alpha(m,n) are the roots of the equation d/dx(J_m[pi*x])=0 for the Bessel
```

```
% Functions J_m (n is the root index).
```

```
% m and n index the tangential mode indices (0,1,2...)
```

```
x = zerobess('DJ', m,n+1)/pi;
```

```
%the function zerobess('DJ', a, nu) should output the first a positive zeros of the derivative of the Bessel function of the first kind J_nu(x)
```

```
%One can use the pre-fab function file (zerobess.m) for this purpose, written by Jonas Lundgren (2010) and found at http://www.mathworks.com/matlabcentral/fileexchange/26639-zerobess
```

```
x = x(n+1);
```

```
end
```

```
%Modified 1/10/12 to be accurate (n was offset by 1)
```

Appendix G

Cell dynamics derivations

In this chapter, we sketch a derivation of the dynamical equations for a typical “dummy cell” shown in Fig. G.1(a) (based on work by Hoskinson [39]). Note that this is a very generic cell and the number of weak-links shown do not matter. What matters in this chapter is the following set of assumptions:

1. There is a small inner cell and a much larger outer cell separated by some flow impedance.
2. The outer cell is thermally sunk to the large outer helium bath, which is temperature regulated so that the outer cell temperature T is under our control and doesn’t change much for small temperature changes in the inner cell.
3. Finally, there is a flexible diaphragm with area A and an effective spring constant k that caps the inner cell.

Since these very general conditions apply to all our cells (single or multiple weak-links), the equations derived here will be of universal use throughout this dissertation.

G.1 Normal and total currents

G.1.1 Normal current

Here, we just quote a relatively straightforward result from [39, p. 16] for the normal current through the aperture array(s):

$$I_n = \rho_n \frac{\beta_n}{\eta} \left(\frac{\rho_n}{\rho} \Delta P + s \rho_s \Delta T \right) \quad (\text{G.1})$$

This flow form, derived from the Navier-Stokes equations, is Poiseuille-like with an empirically determined geometric factor β_n (the normal flow conductance). This factor is included

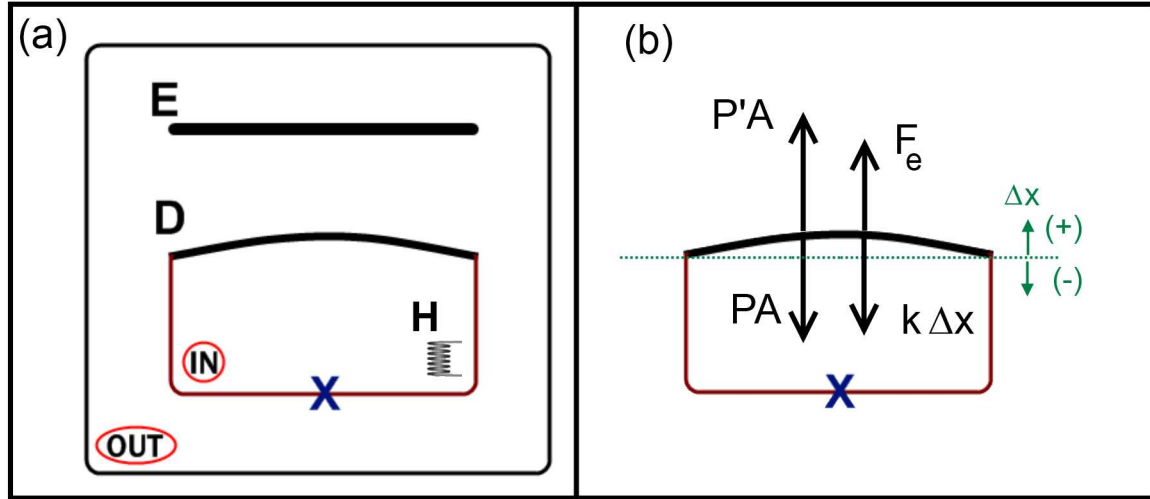


Figure G.1: (a) Simple cell schematic for cell dynamics. See introductory note at the beginning of this chapter for applicability of this dummy cell to the real cells used in experiments. D is a flexible, metallized diaphragm and E is a fixed (metal) electrode used to exert electrostatic force on the diaphragm. H is a resistive heater used to inject heater power \dot{Q} into the inner cell. (b) Force diagram for diaphragm - see Section G.1.2.

to automatically take care of flow field effects from multiple apertures with variable aspect ratios and is obtained during Fountain calibration (Section 10.8.3). A model for the conductance β_n for an array with N apertures of diameter d and channel length l (from [39, p. 78] and [95])

$$\beta_n = \frac{\pi d^4 N}{8(3\pi d + 16l)} \quad (\text{G.2})$$

can be used to deduce the aperture diameter from the experimentally obtained value of β_n (during the normal flow calibration procedure described in Section 10.8.2).

G.1.2 Total current

Referring to Fig. G.1(b), the mass in the inner cell is $M = \rho V' = \rho A(x_0 + \Delta x)$, where x_0 is the flat position of the diaphragm and Δx denotes diaphragm displacements. Differentiating this equation, we see that the total current can be described by

$$\dot{M} \equiv I_t = \rho A \dot{\Delta x} + V' \dot{\rho} \approx \rho A \dot{\Delta x} \quad (\text{G.3})$$

where we have neglected any density changes and approximated the fluid as incompressible (see Appendix E for details about the limits of validity of this approximation).

It is instructive at this point to carefully think about quasi-static processes and the relation between the pressure difference between the inner and outer cells and the diaphragm

displacement Δx . Consider Fig. G.1(b), where we have shown a force diagram for the diaphragm, assuming a point mass and a simple, effective spring constant k . F_e is any electrostatic force that may be applied to the diaphragm by putting a voltage across the capacitor defined by the diaphragm (D) and electrode (E). The net force (using the sign convention shown in the figure) can be written as:

$$F_{net} = P'A - PA + F_e - k\Delta x = A\Delta P + F_e - k\Delta x$$

If we assume a quasi-static process so that the diaphragm is always in (approximate) equilibrium, the net force is approximately zero at each instant and we can write:

$$\Delta P = \frac{F_e}{A} - \frac{k}{A}\Delta x \quad (\text{G.4})$$

Assuming an electrostatic force constant in time (which is usually the case for electrostatically induced pressure steps), we can differentiate this equation to get:

$$\dot{\Delta P} = -\frac{k}{A}\dot{\Delta x} \quad (\text{G.5})$$

Using this in Eq. (G.3), we finally obtain:

$$I_t = \rho \frac{A^2}{k} \dot{\Delta P} \quad (\text{G.6})$$

We note again that:

$$I_t = I_s + I_n \quad (\text{G.7})$$

SQUID signal and current

The total current I_t (flowing into or out of the inner cell), inflates or deflates the flexible diaphragm and changes its average position Δx away from equilibrium. The SQUID-based displacement sensor¹ puts out a voltage signal V_{SQ} in response to diaphragm displacements. For the analysis presented here, it is unimportant which design is actually used. The only assumption we need here is that the change in SQUID output voltage ΔV_{SQ} from the sensor is linearly proportional to the diaphragm displacement:

$$\Delta x \equiv \Delta V_{SQ}/\alpha \quad (\text{G.8})$$

where α (in V/m) is called the displacement sensitivity (or ‘‘SQUID sensitivity’’ in some places in program code, notes, etc.)

Taking the derivative of this equation over time and using it with Eq. (G.5), we see that

$$\dot{\Delta P} = \frac{k}{A\alpha} \dot{\Delta V}_{SQ} \equiv \gamma_1 \dot{\Delta V}_{SQ} \quad (\text{G.9})$$

¹The displacement sensor is summarized in Section 3.1.2 and described in detail in Chapter 5.

where γ_1 (in Pa/V) is known as the pressure calibration factor that translates between SQUID voltage and pressure difference between the inner and outer cells. In the case where the externally applied electrostatic force F_e is constant during the diaphragm displacement, the above equation is also valid for ΔP and ΔV_{SQ} themselves (and not just their time derivatives). We will use this fact for some of the calibration derivations. What we typically calibrate is γ_1 . The result above ($\frac{k}{A\alpha} = \gamma_1$) is used to obtain (with some uncertainty in k and especially A) the SQUID sensitivity, which is needed for establishing displacement noise figures for the displacement sensor.

Of course, what is important to us (nearly always), is the sensitivity to *current* and the current noise (not displacement). Using the current from Eq. (G.6) with Eq. (G.9), we obtain the final calibration of the total mass current against SQUID voltage:

$$I_t = \rho \frac{A^2}{k} \gamma_1 \Delta \dot{V}_{SQ} \quad (\text{G.10})$$

G.2 Supercurrent and inductance

The supercurrent behaves differently in the various regimes described in Section 1.2.2. We can identify two extreme regimes where the supercurrent can be written out in a straightforward way (see Section 1.2.2 for descriptions of physics in these regimes).

G.2.1 Weakly coupled (Josephson) regime

In the weakly coupled (ideal Josephson) regime, we can write the Schrödinger equations² for the left and right volumes in Fig. 1.3:

$$\begin{aligned} i\hbar \frac{\partial \psi_L}{\partial t} &= \mu_L \psi_L + K \psi_R \\ i\hbar \frac{\partial \psi_R}{\partial t} &= \mu_R \psi_R + K \psi_L \end{aligned}$$

where μ is the chemical potential, ϕ is the order parameter phase and K is a coupling energy that describes the amount of order parameter “leakage” (the subscripts L and R refer to the left and right volumes in the figure).

Substituting Eq. (1.1) for the order parameters in the above equations and simplifying (with the notation $\dot{z} \equiv \partial z / \partial t$), we get³:

$$\begin{aligned} i\hbar \dot{\rho}_L &= 2\hbar \rho_L \dot{\phi}_L + 2\mu_L \rho_L + 2K \sqrt{\rho_L \rho_R} e^{i(\phi_R - \phi_L)} \\ i\hbar \dot{\rho}_R &= 2\hbar \rho_R \dot{\phi}_R + 2\mu_R \rho_R + 2K \sqrt{\rho_R \rho_L} e^{i(\phi_L - \phi_R)} \end{aligned}$$

²After Feynman [96, Ch. 21].

³Note that all densities and currents here refer to the superfluid - we are omitting the subscript ‘s’ to avoid notational chaos.

These 2 complex equations yield 4 real equations (where we define $\Delta\phi \equiv \phi_R - \phi_L$):

$$\begin{aligned}\hbar \dot{\rho}_L &= 2K\sqrt{\rho_L\rho_R}\sin\Delta\phi \\ \hbar \dot{\rho}_R &= -2K\sqrt{\rho_L\rho_R}\sin\Delta\phi \\ -\hbar \dot{\phi}_L &= \mu_L + K\sqrt{\rho_R/\rho_L}\cos\Delta\phi \\ -\hbar \dot{\phi}_R &= \mu_R + K\sqrt{\rho_L/\rho_R}\cos\Delta\phi\end{aligned}$$

We can interpret the time derivative of the (superfluid) density as the (superfluid) mass current flowing through the aperture: $\dot{\rho}_L = -\dot{\rho}_R \equiv I$. Then the first two equations above become:

$$\boxed{I = I_c \sin \Delta\phi} \quad (\text{G.11})$$

where we have defined a critical current $I_c \equiv 2K\sqrt{\rho_L\rho_R}/\hbar$. This equation is called the Josephson equation (after Brian Josephson, who first predicted this effect for superconducting weak-links [23]).

With the assumption of incompressibility⁴, we can assume that the ratio $\rho_L/\rho_R \approx 1$ and the last two equations above become:

$$\boxed{\frac{\partial\Delta\phi}{\partial t} = -\frac{\Delta\mu}{\hbar}} \quad (\text{G.12})$$

This is known as the Josephson-Anderson phase evolution equation and it can be shown to be valid even in the strongly coupled regime [24].

The two boxed equations describe the evolution of the supercurrent in a parametric way given a known chemical potential difference $\Delta\mu$.

G.2.2 Aperture array inductance

Here, we consider the idea of hydrodynamic inductance in an aperture array. In electronic systems, we define the inductance L using $LdI/dt \equiv -V$, where we can interpret the inductance as the system's intrinsic inertia (against changes in current). Analogously, we can define the superfluid hydrodynamic inductance L as:

$$L\frac{dI_s}{dt} \equiv -\frac{\Delta\mu}{m_4} \quad (\text{G.13})$$

⁴See Fig. 1.20: the actual single weak-link cell is composed of an inner cell placed inside a larger, outer cell. One of the inner cell walls has the aperture (in practice, an array of apertures) that weakly connects the inner and outer cells. Because there is a flexible diaphragm capping one of the faces of the inner cell, mass flow through the aperture causes the diaphragm to bulge in a way that maintains the inner and outer cell densities approximately constant during such flows, within limits discussed in Section 3.2.1. Fig. 1.21 shows a double weak-link cell, where similar considerations apply.

This can be re-written in terms of derivatives with respect to the phase-difference across the aperture array $\Delta\phi$:

$$-\frac{\Delta\mu}{m_4} = L \frac{dI_s}{dt} = L \frac{dI_s}{d\Delta\phi} \frac{d\Delta\phi}{dt} = L \frac{dI_s}{d\Delta\phi} \left(-\frac{\Delta\mu}{\hbar} \right)$$

where we have used the phase-evolution equation Eq. (G.12) in the last step. We finally obtain for the inductance of the aperture array:

$$L = \frac{\hbar}{m_4} \left(\frac{dI_s}{d\Delta\phi} \right)^{-1} \quad (\text{G.14})$$

Using this expression with the current phase-phase relation in the Josephson regime (Eq. (G.11)), we can calculate the (phase-dependent) aperture array inductance in the Josephson regime (L_J) as:

$$L_J(\Delta\phi) = \frac{\hbar}{m_4 I_c \cos \Delta\phi} \quad (\text{G.15})$$

G.2.3 Strongly coupled (phase-slip) regime

In the strongly coupled regime, we have bulk superfluid flow with a velocity described by Eq. (1.2), whose time variation can be written as:

$$\mathbf{v}_s = \frac{\hbar}{m_4} \nabla\phi \Rightarrow \dot{v}_s = \frac{\hbar}{m_4} \frac{\partial \nabla\phi}{\partial t} = -\frac{\hbar}{m_4} \frac{\nabla\mu}{\hbar} \quad (\text{G.16})$$

where a term proportional to the chemical potential gradient has been substituted for the time-derivative of the phase gradient across the aperture from the Josephson-Anderson phase evolution equation (Eq. (G.12), which, as previously stated, is valid in both the weak and strong coupling regimes). We finally obtain the time rate of change of the superfluid velocity, which can be thought of as Newton's 2nd law of motion for a ⁴He atom:

$$m_4 \dot{v}_s = -\nabla\mu \quad (\text{G.17})$$

where the right hand side implies that chemical potential gradients are forces acting on superfluids. For flow through an aperture of area a and effective length l_e , we can use the above equation and $I_s = \rho_s v_s a$ to obtain an equation for the superfluid current:

$$\begin{aligned} m_4 \dot{v}_s &\approx -\frac{\Delta\mu}{l_e} \\ m_4 \frac{\dot{I}_s}{\rho_s a} &= -\frac{\Delta\mu}{l_e} \\ \frac{l_e}{\rho_s a} \dot{I}_s &= -\frac{\Delta\mu}{m_4} \end{aligned}$$

Comparing the last line to Eq. (G.13), we identify the hydrodynamic inductance in the phase-slip regime as⁵:

$$L_{ps} = \frac{l_e}{\rho_s a} \quad (\text{G.18})$$

and write the equation for the supercurrent in this regime as:

$$\boxed{\dot{I}_s = -\frac{1}{L_{ps}} \frac{\Delta\mu}{m_4}} = \frac{\hbar}{m_4} \frac{\dot{\Delta\phi}}{L_{ps}} \quad (\text{G.19})$$

where the last re-statement comes from substituting for $\Delta\mu$ from the phase-evolution equation Eq. (G.12). We can trivially integrate this equation (with the self-evident condition that $I_s \propto v_s \propto \Delta\phi = 0$ for $\Delta\phi = 0$) to obtain a linear current-phase relationship for the strongly coupled regime:

$$\boxed{I_s = \frac{\hbar}{m_4} \frac{\Delta\phi}{L_{ps}}} \quad (\text{G.20})$$

Either of the two boxed equations above can be used to represent the evolution of the current in this regime.

Array inductance (strong coupling)

Note that this argument can be carried through for an array of N apertures in a similar way, with the only difference being that we now have N inductances, each equal to $\frac{l_e}{\rho_s a}$ in parallel (where the inductance formula is valid for long channels such that $l \gg \sqrt{a}$). Naively, this should give an effective inductance N times smaller: $\frac{l_e}{N\rho_s a}$, but flow distortions due to proximity of the apertures require corrections in this formula. As stated in Eq. (1.12) (and the text following it), this flow ambiguity is folded into an empirical, geometric factor $\beta_s \sim Na/l_e$ (but with corrections), such that:

$$L_{ps} = \frac{1}{\rho_s \beta_s} \quad (\text{G.21})$$

Here, β_s has dimensions of length and is described in [39, p. 79] as the superfluid conductance. This quantity can be found experimentally by measuring the temperature dependence of the so-called Helmholtz frequency (see Section 10.8.7). A flow model⁶ for a square array of N apertures with diameter d and inter-aperture spacing D , gives a predicted value for β_s :

$$\beta_s = N \left(\frac{l + 8d/(3\pi)}{\pi d^2/4} + 2\sqrt{N}/D \right)^{-1} \quad (\text{G.22})$$

This can be used with the experimentally obtained value (from the Helmholtz frequency) to get an independent measurement of the aperture diameter (with N and D known).

⁵The same result for the inductance can be obtained by using instead Eq. (G.14) with $I_s = \rho_s v_s a$.

⁶See [39, p. 81]. There is an error in Eq. (C.11) of this reference, corrected in the equation given here (with a missing N inserted).

G.3 Chemical potential difference

We simply quote here the expression for the chemical potential difference from Eq. (1.5) for completeness:

$$\Delta\mu = m_4 \left(\frac{\Delta P}{\rho} - s\Delta T \right) \quad (\text{G.23})$$

G.4 The temperature equation

Throughout this section⁷, we use a naming convention where primed quantities refer to the inner cell and unprimed quantities refer to the outer cell. All differences are defined as ($\Delta \equiv$ inner cell value - outer cell value). The thermodynamic identities and definitions used here may be found in any elementary thermodynamics textbook, such as Ref. [97, pp. 43-45].

The inner and outer cell volumes are V' and V respectively. The pressure, temperature and number of atoms in the inner cell are P' , T' and N' respectively. The corresponding quantities for the outer cell are P, T and N , which are related to the corresponding inner cell quantities via the respective differences as: $P' = P + \Delta P$, $T' = T + \Delta T$ and $N' = N + \Delta N$.

G.4.1 Change in inner cell entropy

We can write a thermodynamic description of the entropy change $\delta S'$ of the inner cell fluid in terms of the inner cell quantities P', T' and N' :

$$\delta S' = \left(\frac{\partial S'}{\partial T'} \right)_{P', N'} \delta T' + \left(\frac{\partial S'}{\partial P'} \right)_{T', N'} \delta P' + \left(\frac{\partial S'}{\partial N'} \right)_{P', T'} \delta N'$$

Dividing by δt and in the limit that $\delta t \rightarrow 0$, we obtain the rate of change of inner cell entropy:

$$\dot{S}' = \left(\frac{\partial S'}{\partial T'} \right)_{P', N'} \dot{T}' + \left(\frac{\partial S'}{\partial P'} \right)_{T', N'} \dot{P}' + \left(\frac{\partial S'}{\partial N'} \right)_{P', T'} \dot{N}' \quad (\text{G.24})$$

The outer cell pressure (P) and temperature (T) are essentially constant in time as a consequence of our initial assumptions: the pressure, because the outer cell is significantly larger than the inner cell and the temperature, because the outer cell is thermally sunk to a large, temperature-regulated bath⁸.

⁷This derivation is a detailed version based on a briefly sketched out map found in Ref. [38]. Since the chemical potential battery method can easily require much larger inner cell powers than Hoskinson, *et al.* ever dealt with in that reference (or anyone else that we know of), it is necessary to understand more fully the assumptions that go into deriving the temperature equation and re-check their limits of validity.

⁸These assumptions will need to be re-assessed if there is a significant change in the cell design (for instance, using thermally insulating materials like plastics instead of metal for the cell can or changing the cell volumes to a point where they start becoming comparable).

Therefore, we can write the time-derivatives above as $\dot{T}' \approx \dot{\Delta T}$ and $\dot{P}' \approx \dot{\Delta P}$. The quantity \dot{N}' is the rate of change of the particle number in the inner cell. If this is positive, we can interpret it a scenario with a net flow of particles *into* the inner cell. Since each particle is a ${}^4\text{He}$ atom, the total mass current can be written as $I_t = m_4 \dot{N}'$ (where a positive I_t automatically implies a net flow *into the inner cell*).

Further, we have the following elementary thermodynamic relations:

$$T' \left(\frac{\partial S'}{\partial T'} \right)_{P', N'} = C_p$$

where C_p is the total isobaric heat capacity in J/K of the inner cell fluid (and $C_p = c_p V' \rho$, where ρ is the total mass density of He-II and c_p is its specific heat capacity per unit mass - this relation will only be used in the actual scripts used to perform calculations. For now, we will continue to use the total heat capacity C_p).

The quantity $\left(\frac{\partial S'}{\partial N'} \right)_{P', T'}$ is simply the entropy density per particle, which can be written in terms of the entropy density per unit mass (s) as:

$$\left(\frac{\partial S'}{\partial N'} \right)_{P', T'} = m_4 s$$

The quantity $\left(\frac{\partial S'}{\partial P'} \right)_{T', N'}$ is not readily recognizable, but a Maxwell relation turns it into

$$\left(\frac{\partial S'}{\partial P'} \right)_{T', N'} = - \left(\frac{\partial V'}{\partial T'} \right)_{P', N'} = -V' \alpha_p$$

where α_p is the (isobaric) coefficient of linear expansion (in K^{-1}).

Putting all this together with Eq. (G.24), we finally obtain the rate at which the inner cell entropy is changing in time:

$$\dot{S}' = \frac{C_p}{T'} \dot{\Delta T} - V' \alpha_p \dot{\Delta P} + s I_t \quad (\text{G.25})$$

G.4.2 Entropy flows

We can list the possible reasons why the inner cell entropy can change:

Superflow

According to the two-fluid model (see Section 1.2.1), the superfluid does not carry any entropy.

Normal flow

There's a subtle point to be noted here. While the normal fluid transports any entropy in the system, the entropy must be thought of as “belonging” to the entire volume of He-II. The entropy density per unit volume is therefore $s_v = s\rho$ (where it is the total density ρ and not just the normal fluid density ρ_n that is used).

The change in inner cell entropy $\delta S'$ due to entropy transport by normal flow is then $\delta S' = s_v \delta V'_n$, where $\delta V'_n$ is the increase in normal fluid in the inner cell. The entropy rate of change is then: $\dot{S}' = s_v \dot{V}'_n = s\rho \dot{V}'_n$, where \dot{V}'_n must now be interpreted as the volume flow rate (m^3/s) of normal fluid *into* the inner cell (since an outflow would carry entropy away from the inner cell, thereby *reducing* its entropy). The normal current is $I_n = \rho_n \dot{V}'_n$. Finally, the inner cell entropy rate of change due to normal flow is:

$$\dot{S}'_{norm} = \frac{\rho}{\rho_n} s I_n \quad (\text{G.26})$$

Inner cell heater power

If a heater power \dot{Q}_{in} is injected into the inner cell through the resistive heater H (see Fig. G.1)(a), it raises the inner cell entropy at a rate given by:

$$\dot{S}'_H = \frac{\dot{Q}_{in}}{T'} \quad (\text{G.27})$$

Boundary conduction

Heat is lost from the inner cell through the inner cell walls and the diaphragm (D). We assume an effective boundary resistance R_K in an Ohm's law analogue of Fourier's heat conduction equation, so that the heat loss becomes: $\dot{Q}_{in} = -\Delta T/R_K$. Since $\Delta T \equiv T' - T$, the additional minus sign enforces a heat *loss* from the inner cell if $\Delta T > 0$ (i.e. if the inner cell is hotter). The resulting entropy change in the inner cell is:

$$\dot{S}'_{cond} = \frac{-\dot{Q}_{in}}{T'} = \frac{-\Delta T}{R_K T'} \quad (\text{G.28})$$

Putting all three contributions (Eqs. (G.26), (G.27) and (G.28)) together, we obtain the total inner cell entropy rate of change due to the various entropy flows:

$$\dot{S}' = \frac{\rho}{\rho_n} s I_n + \frac{\dot{Q}_{in}}{T'} - \frac{\Delta T}{R_K T'} \quad (\text{G.29})$$

G.4.3 Putting it all together with approximations

The sum of all entropy flows into the inner cell (Eq. (G.29)) must account for the thermodynamically predicted changes in inner cell entropy (Eq. (G.25)). Equating these two expressions yields:

$$\begin{aligned} \frac{C_p}{T'} \dot{\Delta T} - V' \alpha_p \dot{\Delta P} + s I_t &= \frac{\rho}{\rho_n} s I_n + \frac{\dot{Q}_{in}}{T'} - \frac{\Delta T}{R_K T'} \\ C_p \dot{\Delta T} - T' V' \alpha_p \dot{\Delta P} + s T' I_t &= \frac{\rho}{\rho_n} s T' I_n + \dot{Q}_{in} - \frac{\Delta T}{R_K} \\ C_p \dot{\Delta T} - T' V' \alpha_p \dot{\Delta P} &= s T' \left(\frac{\rho}{\rho_n} I_n - I_t \right) + \dot{Q}_{in} - \frac{\Delta T}{R_K} \end{aligned}$$

We recall that the inner cell temperature is $T' = T + \Delta T$. Since the temperature differences are typically $\sim nK$ (and not more than $\sim \mu K$), while the working temperature (outer cell) is $T \sim 2K$, T' can be immediately replaced with T , with an error of no more than $\mathcal{O}(10^{-6})$ for the largest temperature differences used. This gives us:

$$C_p \dot{\Delta T} - T V' \alpha_p \dot{\Delta P} = s T \left(\frac{\rho}{\rho_n} I_n - I_t \right) + \dot{Q}_{in} - \frac{\Delta T}{R_K}$$

The α_p term is roughly (at least) two orders of magnitude smaller than the terms on the right hand side, so it can be neglected. Finally, we obtain the temperature equation for the cell:

$$\boxed{C_p \dot{\Delta T} = s T \left(\frac{\rho}{\rho_n} I_n - I_t \right) + \dot{Q}_{in} - \frac{\Delta T}{R_K}} \quad (\text{G.30})$$

Re-writing the total current as $I_t = I_n + I_s$, where I_s is the superfluid current, and re-arranging a bit gives us a more suggestive form:

$$C_p \dot{\Delta T} = s T \frac{\rho_s}{\rho_n} I_n + \dot{Q}_{in} - s T I_s - \frac{\Delta T}{R_K} \quad (\text{G.31})$$

This is essentially a statement of energy conservation. The inner cell gains energy via heater power and normal flow inward and loses it via superflow into the cell (the latter of which has the net effect of ‘‘cooling’’ the inner cell fluid) and wall conduction. The net energy left goes into raising the temperature of the inner cell (the left-hand side of the equation). What we neglected a while back is the heat that goes into thermal expansion of the fluid.

G.5 Summary

The boxed equations (G.30), (G.23), (G.1), (G.6) and (G.7) (for all coupling regimes) and the supercurrent equation for the respective regime constitute a set of coupled differential

equations whose solutions in certain scenarios are used in applications like the fountain calibration, aperture size measurements using normal flow and superflow and in modeling the chemical potential battery.

Appendix H

Flow simulations for a superfluid grating

H.1 Introduction

A superfluid placed in a box that resides in a non-inertial (specifically, rotating) reference frame is analyzed and its behavior is simulated at increasing levels of complexity of rotation. This work is partly based on previous work by K. Schwab [60].

The goal is to simulate the response of a SHeQUIG (Superfluid Helium 4 Quantum Interference Grating) to rotations and optimize that response to enable its use as a highly sensitive gyroscope. A common first test of such a gyroscope is to detect and measure the rotation of the Earth in the reference frame of the fixed stars. In a relativistic universe, this frame (defined by the “fixed” stars) is the closest thing to an “absolute” inertial reference frame that we can imagine. We will call this frame S .

H.2 Single-axis rotation

H.2.1 Reference frames and coordinate systems

- The main assumption in this work is that all flow fields set up in the box add linearly so that rotations of the box around different axes can be analyzed separately. Accordingly, we define a reference frame R attached to the rotating box with a coordinate system aligned along the sides of the box. The linear dimensions of the box are L_x and L_y (respectively aligned along the X and Y axes of the reference frame attached to the box).
- A rotation around the Z -axis should produce flow that would look the same in an XY plane at any Z -coordinate, if the box were infinitely long along Z . If we look at XY planes close to the Z -center of the box, this would still be a reasonable starting point.
- We can now begin to solve for the flow field. We imagine that the box has been rotating at a constant angular speed Ω around one of the axes of the box (as observed in the




Figure H.1: A snapshot of the system at equilibrium. Two reference frames have two distinct coordinate systems that coincide perfectly at *this* moment of time. All coordinates can therefore be written in either primed or unprimed coordinates as long as time remains fixed.

inertial frame S). Fig. H.1 shows the box face normal to this rotation axis at a point of time when the flow has reached steady state.

- There are two distinct reference frames for the box - the inertial (S) and the non-inertial (R). We now freeze the picture at the moment of time shown in Fig. H.1 and define a coordinate system in each frame (XYZ and $X'Y'Z'$) such that their origins and axes coincide. Clearly, these coordinate systems will coincide *only at this instant*. However, as long as we do not take any *time derivatives*, the two coordinate systems (but *not* the frames themselves) are identical and we may use either one to describe the fluid. We will use the unprimed coordinates henceforth.

H.2.2 Flow equations

- A fluid can be well-approximated as incompressible if its flow speed is much less than the speed of sound within it [98]¹. Since this is generally true for superfluids in the regimes we work in², we can start by assuming incompressibility, so that $\frac{\partial \rho_s}{\partial t} = 0$ (where ρ_s is the superfluid density).
- From the equation of continuity, $\frac{\partial \rho_s}{\partial t} = \nabla \cdot \mathbf{v}_s$, where v_s is the flow velocity in the inertial frame S . We therefore have:

$$\nabla \cdot \mathbf{v}_s = 0 \quad (\text{H.1})$$

¹This reference will be used for all elementary fluid dynamics unless otherwise specified

²1st sound speed is $\sim 200\text{m/s}$ and flow velocities stemming from rotation are of the order of microns/sec.

- In the inertial frame, the superfluid is described by an order parameter ($\psi = \psi_0 e^{i\phi}$), with a quantum phase ϕ [16]. Acting the probability current operator on it and identifying the result with the superfluid mass current gives us the quantum mechanical relation:

$$\mathbf{v}_s = \frac{\hbar}{m_4} \nabla \phi \quad (\text{H.2})$$

where \hbar is Planck's constant and m_4 is the mass of a helium-4 atom.

- The curl of this flow field is:

$$\nabla \times \mathbf{v}_s = \frac{\hbar}{m_4} \nabla \times \nabla \phi = 0 \quad (\text{H.3})$$

Superflow is therefore irrotational (as seen in the inertial frame).

- We can therefore write the flow velocity as the gradient of a velocity potential u_s (which, from Eq. H.2, is clearly proportional to the phase ϕ):

$$v_s = \nabla u_s \quad (\text{H.4})$$

We have defined this velocity potential in the inertial frame S (this distinction will be important later).

- From Eqs. H.2 and H.4, the phase is

$$\phi = \frac{m_4}{\hbar} u_s \quad (\text{H.5})$$

- From Eqs. H.1 and H.4, the velocity potential therefore satisfies Laplace's equation:

$$\nabla^2 u_s = 0 \quad (\text{H.6})$$

with the boundary condition that the flow cannot penetrate the walls. This means that the normal component v_{sn} of the flow velocity at each point \mathbf{r}_w of the wall must match the normal component of the wall velocity \mathbf{v}_w as seen in the frame S .

$$v_{sn} \equiv \mathbf{v}_s \cdot \hat{n} = \mathbf{v}_w \cdot \hat{n} = (\boldsymbol{\Omega} \times \mathbf{r}_w) \cdot \hat{n} \quad (\text{H.7})$$

- $\boldsymbol{\Omega} = (0, 0, \Omega)$ and $\mathbf{r}_w = (x_w, y_w, z_w)$ gives:

$$v_{sn} = \Omega(x_w n_y - y_w n_x) \quad (\text{H.8})$$

where $\hat{n} \equiv (n_x, n_y, n_z)$ is the surface normal at the boundary point \mathbf{r}_w .

H.2.3 Solutions

Inertial frame S

The problem is now essentially 2-dimensional and Laplace's equation (Eq. H.6) can be readily solved with the boundary condition in Eq. H.8 using a partial differential equation solver. We use Matlab's PDE toolkit [99] to do this calculation and obtain the velocity potential u_s in the inertial frame. The velocity field follows simply from Eq. H.4 as $\mathbf{v}_s = \nabla u_s$

Rotating frame R

The flow field (\mathbf{v}_R) at all points \mathbf{r} inside the box as seen by an observer in the rotating frame R (fixed to the box) can be found by boosting from S to R :

$$\mathbf{v}_R = \mathbf{v}_s - \boldsymbol{\Omega} \times \mathbf{r} = \mathbf{v}_s - (-\Omega y, \Omega x, 0) \quad (\text{H.9})$$

H.2.4 Phase ambiguity

In the rotating frame, the flow field is no longer irrotational, as may be easily seen by taking the curl of Eq. H.9 and using Eq. H.3:

$$\nabla \times \mathbf{v}_R = \nabla \times \mathbf{v}_s - \nabla \times (\boldsymbol{\Omega} \times \mathbf{r}) = (0, 0, -2\Omega) \quad (\text{H.10})$$

Therefore, there can be no scalar function u_R whose gradient ∇u_R gives the flow field \mathbf{v}_R . Does this mean then that the superfluid is no longer described by a simple order parameter with a quantum phase ϕ ? Or does it merely mean that the probability current has a different form in a non-inertial frame so that the classical fluid dynamics and the quantum descriptions stay consistent?

The answer is simply that the intrinsic (or kinetic) velocity referenced to some inertial frame is associated with phase textures of the superfluid. Coordinate transformations cannot change the intrinsic phase texture (they are just gauge transformations). See, for instance, Ref. [30].

The idea is that information about the angular velocity is encoded in the inertial frame velocity since it is the walls that push on the fluid to make it move. If, instead of a rectangular box, we had a circular cylinder, the fluid velocity in the inertial frame would NOT contain information about the angular velocity. In essence, the phase textures would be insensitive to any rotation and so would the grating. In a non-circular geometry, the walls DO push on the fluid to an extent that depends on the wall rotation rate. The phase textures will therefore change in response.

The phase texture change can be found as we have described above. The velocity in the inertial frame is $\mathbf{v}_s = \nabla u_s$ and u_s is found by solving Laplace's equation in the inertial frame with boundary conditions that encode the wall motion and therefore the rotation rate. u_s is therefore the actual phase texture associated with this motion. The further coordinate boost to the rotating frame ($\mathbf{v}_R = \mathbf{v}_s - \boldsymbol{\Omega} \times \mathbf{r}$) does not contribute to the phase.

This is valid because experimentally (e.g. [100]), the circulation in the rotating frame is no longer quantized [30] (since it IS quantized in the inertial frame and we are adding to it an unquantized term $\oint_{L1} \boldsymbol{\Omega} \times \mathbf{r} \cdot d\mathbf{l} = \int_{S1} \boldsymbol{\Omega} \cdot d\mathbf{S}$ corresponding to the boost, where S1 is the surface enclosed by the loop L1).

In conclusion, u_s will be the phase texture *as seen by the grating*.

H.2.5 Results

Simulation results, including phase textures and interference curves are available in Section 13.4.4 of Chapter 13.

H.3 Matlab code

H.3.1 Generating the phase data for each value of rotation rate

Main front-end script

This is the main script file that should be edited with user parameters and run from Matlab. It requires two function files to run: “rampgen.m” and “box6.m”, provided in the code listings that follow this one.

Listing H.1: front_end_w_6.m

```
%Aditya Joshi 10/22/11
%v6 -frontend for flow PDE calc
%Uses box6.m for pde calcs (same as box5) only front end (this file) is
%changed. Refinemesh changed from 4 to 8.

clear all;
close all;
clc;
set(0,'DefaultFigureWindowStyle','docked'); %Docks all figures in one container
%This script will probe the dependence of flow velocities on the angular
%velocity W and plot a modulation curve for a given box vs. W

%% Input params
inch = 0.0254;
Lx=4*inch;
Ly = 4*inch;

%W=2*pi/(3600*24); %Max value of earth's angular speed (W = 2pi/T)
%latberk = Berkeley's latitude in radians
latberk = 37.8717*pi/180;
WEarth = 2*pi/(3600*24);
%Wp = Earth's rotation parallel to earth's surface (that can create large
%flows parallel to the apertures.
Wp = WEarth*cos(latberk);

WI = (0.001)*Wp; %0.1% of W_Earth
WF = 10*Wp; %W_Earth parallel to surface (max)
WN = 1000; %steps

rmcount=5; %Refine mesh iterations (3 is good)
```

```

savpathstr = 'C:\SIMS\box\interf_W_6\'; %All plots will be saved here
cd(fullfile(savpathstr, '6\'));

%% Preliminary calcs
W = rampgen(WI,WF,WN); %Ly grid line

%% Initialize data collectors for main loop

%Init cell array. Each cell is a vector of data from box6.m
xh = cell(WN,1);
uh = xh;

%% Main loop
for j = 1:WN
    sprintf('Started_iteration_%d_of_%d', j, WN)
    sprintf('Lx=_%0.4f_in,_Ly=_%0.4f_in,_W=_%0.4e_rad/s', Lx/inch, Ly/inch, W(j))

    [xhtemp uhtemp] = box6(Lx, Ly, W(j), Ly/2, rmcount);

    xh{j} = xhtemp;
    uh{j} = uhtemp;
    clear xhtemp uhtemp;

    sprintf('Finished_iteration_%d_of_%d', j, WN)
end;

sprintf('Last_iteration_completed_%d_of_%d', j, WN)
sprintf('Finished_main_loop_Saving_data.')

%% Save all workspace vars up to this point
%(doesn't include interference curves)
%Does not contain ANY information about aperture array
%This is the data for a flow field setup in a rotating box. Only quantum
%thing is the rescaling by the box6 function of the velocity potential to
%the phase using hbar and m4.
cd(savpathstr);
save('data.mat');
sprintf('Saved_all_data')
sprintf('SCRIPT_FINISHED')

%Do not need to redo simulation after this point. The simulated flow data has been saved.
%Think of it as actual "DATA" pertaining to superflow generated in the box.

```

Ramp generator (helper script)

Listing H.2: rampgen.m

```

function y = rampgen(yi, yf, n)
%RAMPGEN generate ramp (row vector) from yi to yf for n equi-spaced samples
% yi = initial value, yf = final value, n = num samples in ramp vector
%  $y(k) = yi + (k-1)*(yf-yi)/(n-1)$ 
y = zeros(n,1);
dy = (yf-yi)/(n-1);
for k = 1:n
    y(k) = (k-1)*dy;
end;
y = y+yi;
end

```


Main PDE solver function (customized)

Listing H.3: box6.m

```

function [xh uh] = box6(Lx, Ly, W, yp0, rmcount)
%Lx and Ly are x and y dimensions of box, W is angular speed, yp0 is
%the y-coord of the horizontal line for which to obtain the phases
%that are output in the matrix ap.
%A box with sides Lx and Ly are created centred at 0.
%rmcount is the number of times to refine the mesh (3 is
%good enough most times). savpathstr is a string for the path to save
%the solution images (e.g. 'C:\temp\pdeplots')

%xh si a vector of x-coords along the chosen horizontal line (y=yp0)
%uh = real phase (in the inertial frame) = 2*pi/kap0*velocity potential

%Aditya Joshi 6/20/11
%v6

%% Step 1: geometry description (see help on the decsg function in step 2)
gd = [3; 4; -Lx/2; Lx/2; Lx/2; -Lx/2; Ly/2; Ly/2; -Ly/2; -Ly/2];
sf = 'R1';
ns = [82;49];

%% Step 2: Get decomposed geometry matrix
dl = decsg(gd,sf,ns);

%% Step 3: Create mesh and refine it based on input 'rmcount'
[p,e,t]=initmesh(dl);

%Refine mesh 'rmcount' number of times
for rmc = 1:rmcount
    [p,e,t]=refinemesh(dl,p,e,t);
end;

%% Step 4: Define PDE coefficients
c = 1;
a = 0;
f = 0;

%% Step 5: Define boundary conditions
%There are 10 boundary segments in all. 4 for each curved section and 1
%each for the two septum walls.
%Construct the b.c. strings for the two different kinds of segments
%There are the 2 ends of the heatpipe in one qualitative set of boundaries
%and then the remaining segments in another set.
%The first set will have v.n = W*(x*ny -y*nx) (normal flow velocity
%forced to be the wall velocity).
%The second set will have this PLUS a term for the heat current induced
%flow: vH*nx. This is not done yet so ignore this one.

bstrR=[" num2str(W,'%11f') '(x.*ny_-y.*nx)' "];

%Construct the columns of the boundary condition matrix -b
%See Help entry of "assemb" function for details on the meaning of each
%row. There must be one column for each boundary segment. The order of the
%segments must be found by running this sort of problem at least once in
%the pdetool gui. Alternately, look at the decomposed geometry matrix and
%using its format, try to figure out from scratch which boundaries are
%being referred to (x and y coords are included in there).

%The two types of boundary conditions constructed.
%(These rows will become columns in the end (when b is transposed)).

```

```

bR = [1 0 1 length(bstrR) '0' bstrR];

%Compile all the segments' boundary conditions together into a matrix b
%There are 4 segments.
b = char(bR,bR,bR,bR);

%Note the transpose for the matrix b
b = b';

clear bR bH bstrR bstrH;

%% Step 6: Assemble the PDE problem and solve it
u=asempde(b,p,e,t,c,a,f);

x = p(1,:)' ; %first row of p converted to column vector
y = p(2,:)' ; %second row of p converted to column vector

%% After getting all velocity fields, rescale the velocity potential
kap0 = 9.969287862274535e-8; %Quantum of circulation = h/m4
u=(2*pi/kap0).*u; %Rescale velocity potential to get the quantum phase in real units (phi = u*2pi/kappa0)
%From here on, u is actually the phase phi.

%% Find phase and velocity values at horizontal slice (y = yp0)
yp0 = y(dsearchn(y,yp0)); %find nearest value to yp0, save THAT as new yp0
%Use THAT new yp0 for boolean search

%Find indices in the main table corres. to y=yp0
inds1=(y==yp0); %Bool array with 1 if equal and 0 if not. Can be used to extract the "1" elements from other matched
arrays (like x and u)
uh = u(inds1); %Phases
xh = x(inds1); %x-coord col vector
[xh, sorter1] = sort(xh); %for properly connected graph
uh = uh(sorter1); %Re-align to keep the coordinate points matched

end

```

H.3.2 Using the saved phase data to obtain interference curves and maximum sensitivities

Front-end script (set grating parameters)

Listing H.4: reanalyze_series_frontend_v4.m

```

%reanalyze_series_frontend_v4.m
%10/22/11
%Uses v4 of the reanalyze function

%% Init sys
clear all; close all; clc;

%% Common params (not iterated)
inch = 0.0254;
apextent = 2; %aperture array grating total extent in inches (fixed)

dataroot = 'C:\SIMS\box\interf_W_6\'; %directory of "data" saved by front end and box function
datfile = fullfile(dataroot, 'data.mat');
savpathstr3 = fullfile(dataroot,num2str(apextent, 'Lap=%din'),'results\'); %To save stuff in

```

```

apextent = apextent*inch; %meters
apcols = 1;
lonehole = 60/5000;
%current in ng/s for one hole (reference -max amplitude seen in SHeQUID3
%run1)
Irow = lonehole*apcols;

N = 200; %Max num rows to test (starting with 2, not 1)
%Note that for 3in, 22500 rows gives ~ 3.39 micron spacing between rows
%This is the # holes in 150x150 array. So, about 3 times this number is the
%largest practical N for a 3in max grating length.

%% Do series analysis
aprows = 2:N; aprows = aprows';
MSE = zeros(N-1,1);
MSN = MSE;
MSW = MSE;
MSE2 = MSE;
MSN2 = MSE;
MSW2 = MSE;
MSI = MSE;
MSI2 = MSE;

%See function file: "reanalyze_v4_func.m" for definitions of these max sensitivities
for j = 1:(N-1)
    tempdat = reanalyze_v4_func(apextent, aprows(j), apcols, lonehole, datfile);
    MSI(j) = tempdat(1);
    MSE(j) = tempdat(2);
    MSN(j) = tempdat(3);
    MSW(j) = tempdat(4);
    MSI2(j) = tempdat(5);
    MSE2(j) = tempdat(6);
    MSN2(j) = tempdat(7);
    MSW2(j) = tempdat(8);
end;

%% plot
labelstring = sprintf('array_extent=_%0.1f_in,_num_cols=_%d', apextent/inch, apcols);

close all;

figure(1) ;clf;
plot(aprows, abs(MSE), 'r' , aprows, abs(MSE2), 'xb');
ylabel('|Max_Sensitivity|_((ng/s)/rad)');
xlabel('Phase_between_grating_ends_(rad)');
title({'Spacing_comparison_[Grating_ends]', labelstring});
legend('equispaced','optimal', 'Location', 'Best');
saveas(1,fullfile(savpathstr3,'spacing_comparison_E.fig'));
saveas(1,fullfile(savpathstr3,'spacing_comparison_E.png'));

figure(2) ;clf;
plot(aprows, abs(MSN), 'r' , aprows, abs(MSN2), 'xb');
ylabel('|Max_Sensitivity|_((ng/s)/rad)');
xlabel('Phase_between_adjacent_slits_(rad)');
title({'Spacing_comparison_[nearest_neighbor_slits]', labelstring});
legend('equispaced','optimal', 'Location', 'Best');
saveas(2,fullfile(savpathstr3,'spacing_comparison_N.fig'));
saveas(2,fullfile(savpathstr3,'spacing_comparison_N.png'));

figure(3) ;clf;
plot(aprows, abs(MSW), 'r' , aprows, abs(MSW2), 'xb');

```

```

ylabel('Max_Sensitivity|_((ng/s)/rad)');
xlabel('Angular_velocity_(rad/s)');
title({'Spacing_comparison_[angular_velocity]', labelstring});
legend('equispaced','optimal', 'Location', 'Best');
saveas(3,fullfile(savpathstr3,'spacing_comparison_W.fig'));
saveas(3,fullfile(savpathstr3,'spacing_comparison_W.png'));

%% save
mkdir(savpathstr3); %results
cd(savpathstr3);
save('results1.mat');
sprintf('Saved_all_data')
sprintf('SCRIPT_FINISHED')

```

Function file to create interference curves

Listing H.5: reanalyze_v4_func

```

function [maxsensdata] = reanalyze_v4_func(apextent, aprows, apcols, Ionehole, datfile)
%apextent is the total grating length in meters
%aprows and apcols are rows and cols (rows are along the grating length)
%Ionehole is the current amplitude for one hole (from an experiment)
%datfile is a mat file containing the data saved by "front_end_W_6.m"

%Output:
%maxsensdata = [max_sens_ind, MaxSensE, MaxSensN, MaxSensW, max_sens_ind2,
%MaxSensE2, MaxSensN2, MaxSensW2];
%E: vs. phase diff across grating ends
%N: vs. phase diff between nearest neighbors (average or exact[2])
%W: vs. angular velocity
%suffix 2 is for optimized array and no suffix is equispaced array

% debug values
% clear all;
% close all;
% clc;
% inch = 0.0254;
% apextent = 2*inch; %Length of the aperture array (the grating length)
% % NOTE that this is not the window length (window should be made slightly bigger)
% apcols = 1; %num of columns (holes along short dimension) -just to raise signal -these don't interfere since they are
% spaced 1 micron apart
% aprows = 10; %typically 60-100micron spacing
% Ionehole = 60/5000; %current in ng/s for one hole (reference -max
% % amplitude seen in SHeQUID3 run1)
% datfile = 'C:\SIMS\box\interf_W_6\data.mat';

set(0,'DefaultFigureWindowState','docked'); %Docks all figures in one container

load(datfile); %Load saved data file
%MUST have been saved by "front_end_W_6.m"

%Save current dir and restore it when function is done
tempcd = cd;

%All interference plots and data will be saved here
savpathstr2 = [savpathstr,num2str([apextent/inch, aprows], 'Lap=%din\\Nrows=%08d\\')];
cd(savpathstr);
mkdir(savpathstr2); %Create new directory (nothing happens if it exists)

% apextent = apextent*inch;

```

```

labelstring = sprintf('array_extent=_%0.1f_in,_num_rows=_%d', apextent/inch, aprows);

% It0 = Ionehole*apcols*aprows; %Expected max current (to normalize interference pattern with)
Irow = Ionehole*apcols; %Amplitude for one row. (each row interferes with every other row)
% looks like this: [ | | ] for aprows=4. apextent is from [ to ]
% [ ] are the 1st and last rows. The window should be slightly bigger than
% this of course.
% So, apextent is divided into aprows-1 equal divisions
% This division size is:
dz = apextent/(aprows-1);

% The x-coord of each aperture row is then:
z = zeros(aprows,1);
for k=1:aprows
    z(k) = (k-1)*dz - apextent/2;
end;
clear k;

%% Create mod curve for equi-spaced aperture rows

% since u returned by the box function is already scaled to phi
% kap0 = 9.969287862274535e-8; %Quantum of circulation = h/m4

% Find indices of x-positions for equi-spaced apertures (these indices will
% match the phase-indices).
% Note that the mesh is the same but to allow for future jiggling and other
% param sweeps, we will find indices for each W separately

inds1 = cell(WN,1); %Each cell is a vector of indices [s] such that xh(s) is the location of the s'th aperture row
ItotS = zeros(WN,1); %Total current amplitude for each W
% ItotR = ItotS;
Pe = ItotS; %phase diff between grating ends (one for each W)
Pn = ItotS; %average (or exact) nearest neighbor phase diff (one for each W)

clear j;

% Find ItotS for each W: ItotS = Irow*SQRT(Sum_m(Sum_k( cos(uk-um))))
% Note: initially ItotS is used to sum all the cos(uk -um) terms. Later it
% is scaled and square rooted, etc.
for j=1:WN
    % Find index-locations of each aperture row from the vector xh (matched
    % to uh)

    % Note that z is an array generated up near the top of this script
    % z is the array of x-coords of the equi-spaced holes.
    tempinds = dsearchn(xh{j}, z); % #ok < *USENS > % Find the aperure row location's index in xh{j} (the horizontal slice)
    % This is a unique match (as opposed to searching through phases)
    % z is also in ascending order (from left to right of the aperture
    % array) so that tempinds is also in that order.

    clear k;
    % Now, tempinds is a vector containing indices for each aperture row location
    % in xh{j} for all j

    tempuhS = uh{j}; %phase vector for slice (for j'th iteration in W)
    % tempuhR = uh{j} + (2*pi/kap0).*W(j).*(Ly/2).*xh{j}; %phase vector for slice for j'th iteration

    % Find ItotS for each W: ItotS = Irow*SQRT(Sum_m(Sum_k( cos(uk-um))))
    % phase at each aperture row is obtained using the index of that
    % aperture row's x-position that we found earlier.
    for k=1:aprows
        for m=1:aprows

```

```

        ItotS(j) = ItotS(j) + cos(tempuhS(tempinds(k)) -tempuhS(tempinds(m)));
        % ItotR(j) = ItotR(j) + cos(tempuhR(tempinds(k)) -tempuhR(tempinds(m)));
    end;
end;
clear k m;
inds1{j} = tempinds; %Save for later use

%Calculate phase diff (ends or nearest neighbor) corresponding to W
Pe(j) = tempuhS(tempinds(aprows)) -tempuhS(tempinds(1));
Pn(j) = Pe(j)/(aprows -1);
%NOTE that this Du is equal to what I call theta is the writeup (the
%nearest neighbor phase diff. for a perfectly linear grating like we are
%building here or the average nearest neighbor phase diff. for a non-linear
%grating).
%Optimized grating == perfectly linear in phase (not sapcing).
%unoptimized == perfectly linear in spacing but not phase.
%NOTE: the phase I'm talking about here is the nearest neighbor Dphi -call
%this Pn (OR the Dphi between the ends of the grating -call this Pe) -
%this will be one value for each W.

clear tempinds;
end;
clear j k;
ItotS = Irow*sqrt(ItotS); %sqrt and scale by real current size
%ItotR = Irow*sqrt(ItotR); %sqrt and scale by real current size
%Now, ItotS(j) is the current for each W(j) (should be an interference
%pattern)

% Find sensitivity against 3 different x-vars (equispaced)

% phase diff across ends
dPe=Pe(2:WN)-Pe(1:WN-1); %Offsets array by 1 and subtracts
dI=ItotS(2:WN)-ItotS(1:WN-1); %same
SensE = dI./dPe;

% nearest neighbor phase diff (average since it is different for equispaced
% grating)
dPn=Pn(2:WN)-Pn(1:WN-1); %Offsets array by 1 and subtracts
SensN = dI./dPn;

% angular speed W
dW = W(2:WN) -W(1:WN-1); %Offsets array by 1 and subtracts
SensW = dI./dW;

% Max sensitivity (same index for all 3) (equispaced)
%since all 3 x-vars are proportional to each other
[MaxSensE max_sens_ind] = max(abs(SensE));
MaxSensN = SensN(max_sens_ind);
MaxSensW = SensW(max_sens_ind);

%% Create mod curve for optimized spacing; with inertial phase u

clear j;

%Pick an W >0 (say Wmax)
JM = dsearchn(W,max(W)); %the index of Wmax
%inds1{JM} is a set of indices for equi-spaced x positions of the aprows
%Take the first and last of these and fix them in place.
%Du is then the diff between the phase at these positions divided by the
%number of divisions needed (aprows -1). Du is therefore the phase step.
uI = uh{JM}(inds1{JM}(1));
uF = uh{JM}(inds1{JM}(aprows));
Du = (uF -uI)/(aprows-1);

```

```

%Generate a sequence of equi-spaced phases between these two endpoints
%including the endpoints (this is the phase value for each ap-row when it's
%placed at an "optimal position"). Note that these optimal positions are
%limited by the e-beam capability in terms of tolerance.

%Limit the search to the index range for the aperture row positions (xh).
%Phase can repeat so we must proceed carefully.

uE = zeros(aprows,1);
for k=1:aprows
    uE(k) = uI + (k-1)*Du;
end;
%These are the phases ONLY FOR THIS W. Use these to find positions for the
%aperture rows such that the phase drop between adjacent rows is the same
%for all rows. NOTE that this is *guaranteed* by construction ONLY ONLY ONLY
%for this chosen W and no others. It may or may not remain optimal when you
%change W and the phases readjust -finding what happens is the point of
%this simulation).

%Now, find the positions xE to satisfy the above goal.
%To be explicit, xE are the positions where u = uE as defined above
%k = dsearchn(X,XI) returns the indices k of the closest points in
%X for each point in XI

ut1 = uh{JM}; %temp var, will be cleared in the end
ut1 = ut1(inds1{JM}(1):inds1{JM}(aprows)); %ONLY the subset within the grating

xEinds = dsearchn(ut1, uE); %ONLY for the chosen W (this is indices in the above subset)
xEinds = xEinds + inds1{JM}(1) - 1; %index referred to the original uh,xh arrays
xE = xh{JM}(xEinds);
%Now, we fix the positions of the aperture rows at these positions (xE).

%A lot of the following is copied from the equi-spaced apertures code.
%Find indices of these x-positions for each j cell (each iteration of W)
%Note that the mesh is the same but to allow for future jiggling and other
%param sweeps, we will find indices for each W separately

%Initialize arrays
inds2 = cell(WN,1); %Each cell is a vector of indices [s] such that xh(s) is the location of the s'th aperture row
ItotS2 = zeros(WN,1); %Total current amplitude for each W
% ItotR2 = ItotS2;

Pe2 = ItotS2; %average (or exact) phase diff between grating ends (one for each W)
Pn2 = ItotS2; %average (or exact) nearest neighbor phase diff (one for each W)

clear j;

%Find Itot for each W: Itot = Irow*SQRT(Sum_m(Sum_k( cos(uk-um))))
%Note: initially Itot is used to sum all the cos(uk -um) terms. Later it
%is scaled and square rooted, etc.
for j=1:WN
    tempinds = dsearchn(xh{j}, xE); %Find the aperure row location indices in xh{j} (the horizontal slice)
    %Now, tempinds is a vector containing indices for each aperture row location
    %in xh{j} for all j

    tempuhS = uh{j}; %phase vector for slice for j'th iteration
    % tempuhR = uh{j} + (2*pi/kap0).*W(j).*(Ly/2).*xh{j}; %phase vector for slice for j'th iteration
    %the second one is the pseudo-phase in the rotational frame

    %Find Itot for each W: Itot = Irow*SQRT(Sum_m(Sum_k( cos(uk-um))))
    for k=1:aprows

```

```

    for m=1:aprows
        ItotS2(j) = ItotS2(j) + cos(tempuhS(tempinds(k)) -tempuhS(tempinds(m)));
% ItotR2(j) = ItotR2(j) + cos(tempuhR(tempinds(k)) -tempuhR(tempinds(m)));
    end;
end;
clear k m;

%Calculate phase diff (ends or nearest neighbor) corresponding to W
Pe2(j) = tempuhS(tempinds(aprows)) -tempuhS(tempinds(1));
Pn2(j) = Pe2(j)/(aprows -1);
%NOTE that this Du is equal to what I call theta is the writeup (the
%nearest neighbor phase diff. for a perfectly linear grating like we are
%building here or the average nearest neighbor phase diff. for a non-linear
%grating).
%Optimized grating == perfectly linear in phase (not spacing).
%unoptimized == perfectly linear in spacing but not phase.
%NOTE: the phase I'm talking about here is the nearest neighbor Dphi -call
%this Pn (OR the Dphi between the ends of the grating -call this Pe) -
%this will be one value for each W.

inds2{j} = tempinds; %Save for later use
clear tempinds; %clear tempinds;

end;
clear j k;
ItotS2 = Irow*sqrt(ItotS2); %sqrt and scale by real current size
% ItotR2 = Irow*sqrt(ItotR2); %sqrt and scale by real current size
%Now, Itot(j) is the current for each W(j) (should be an interference
%pattern)

% Find sensitivity against 3 different x-vars -optimized spacing

% phase diff across ends
dPe2=Pe2(2:WN)-Pe2(1:WN-1); %Offsets array by 1 and subtracts
dI2=ItotS2(2:WN)-ItotS2(1:WN-1); %same
SensE2 = dI2./dPe2;

% nearest neighbor phase diff (average since it is different for equispaced
% grating)
dPn2=Pn2(2:WN)-Pn2(1:WN-1); %Offsets array by 1 and subtracts
SensN2 = dI2./dPn2;

% angular speed W
dW2 = W(2:WN) -W(1:WN-1); %dW is already defined but redo it in case
%this is moved somewhere
SensW2 = dI2./dW2;

% Max sensitivity (same index for all 3) -optimspaced
%since all 3 x-vars are proportional to each other
[MaxSensE2 max_sens_ind2] = max(abs(SensE2));
MaxSensN2 = SensN2(max_sens_ind2);
MaxSensW2 = SensW2(max_sens_ind2);

% Collect all max sens data into an array
maxsensdata = [max_sens_ind, MaxSensE, MaxSensN, MaxSensW, max_sens_ind2, MaxSensE2, MaxSensN2, MaxSensW2
];

% Plot x-positions of equispaced and optimized slits to compare

postemp = zeros(length(z),1); %bar height

figure(3);clf;

```



```

hold on;
st1 = stem(z/inch,postemp+1);
st2 = stem(xE/inch,postemp-1);
axis([-Lx/(2*inch) Lx/(2*inch) -3 3]);
set(st1,'MarkerFaceColor','blue')
set(st2,'MarkerFaceColor','red','Marker','square');
set(get(st1,'BaseLine'),'LineStyle','-','LineWidth', 5);
xlabel('Position_(inches)');
title({'Slit_positions:_equispaced_vs._optimized', labelstring});
legend('Equispaced','Optimized','Location','NorthEast');
rectangle('Position', [-apextent/(2*inch),-0.15,apextent/inch,0.3], 'FaceColor', 'Black');
hold off;
saveas(3,fullfile(savpathstr2,'slit_positions.fig'));
saveas(3,fullfile(savpathstr2,'slit_positions.png'));

% Plot equispaced grating mod curves

%phase ends
figure(4); clf; hold on;
[AX,H1,H2] = plotyy(Pe(1:WN-1), ItotS(1:WN-1), Pe(1:WN-1),abs(SensE), 'plot');
set(get(AX(1),'YLabel'),'String','Amp_(ng/s)') ;
set(get(AX(2),'YLabel'),'String','Sens_((ng/s)/rad)') ;
set(AX(1),'YColor','r') ;
set(AX(2),'YColor','b') ;
set(H1,'LineStyle','-','Color','r');
set(H2,'LineStyle','.','Color','b');
xlabel('Phase_diff._between_grating_ends_(rad)');
title({'I_vs._Pe_[Equi-spaced_grating]', labelstring});
legend('Amplitude','Sensitivity', 'Location', 'Best'); hold off;
saveas(4,fullfile(savpathstr2,'equispaced_E.fig'));
saveas(4,fullfile(savpathstr2,'equispaced_E.png'));

% phase nearest neighbor (average)
figure(5); clf; hold on;
[AX,H1,H2] = plotyy(Pn(1:WN-1), ItotS(1:WN-1), Pn(1:WN-1),abs(SensN));
set(get(AX(1),'YLabel'),'String','Amp_(ng/s)') ;
set(get(AX(2),'YLabel'),'String','Sens_((ng/s)/rad)') ;
set(AX(1),'YColor','r') ;
set(AX(2),'YColor','b') ;
set(H1,'LineStyle','-','Color','r');
set(H2,'LineStyle','.','Color','b');
xlabel('Phase_diff._between_nearest_neighbors_-_avg_(rad)');
title({'I_vs._Pn_[Equi-spaced_grating]', labelstring});
legend('Amplitude','Sensitivity', 'Location', 'Best'); hold off;
saveas(5,fullfile(savpathstr2,'equispaced_N.fig'));
saveas(5,fullfile(savpathstr2,'equispaced_N.png'));

%angular velocity
figure(6); clf; hold on;
[AX,H1,H2] = plotyy(W(1:WN-1), ItotS(1:WN-1), W(1:WN-1),abs(SensW));
set(get(AX(1),'YLabel'),'String','Amp_(ng/s)') ;
set(get(AX(2),'YLabel'),'String','Sens_((ng/s)/(rad/s))') ;
set(AX(1),'YColor','r') ;
set(AX(2),'YColor','b') ;
set(H1,'LineStyle','-','Color','r');
set(H2,'LineStyle','.','Color','b');
xlabel('W_(rad/s)');
title({'I_vs._W_[Equi-spaced_grating]', labelstring});
legend('Amplitude','Sensitivity', 'Location', 'Best'); hold off;
saveas(6,fullfile(savpathstr2,'equispaced_W.fig'));
saveas(6,fullfile(savpathstr2,'equispaced_W.png'));

```

```

% Plot optimized (spacing) grating mod curves

%phase ends
figure(7); clf; hold on;
[AX,H1,H2] = plotyy(Pe2(1:WN-1), ItotS2(1:WN-1), Pe2(1:WN-1),abs(SensE2), 'plot');
set(get(AX(1),'YLabel'),'String','Amp_(ng/s)');
set(get(AX(2),'YLabel'),'String','Sens_((ng/s)/rad)');
set(AX(1),'YColor','r');
set(AX(2),'YColor','b');
set(H1,'LineStyle','-','Color','r');
set(H2,'LineStyle','.', 'Color','b');
xlabel('Phase_diff._between_grating_ends_(rad)');
title({'I_vs._Pe_[Optimized_grating]', labelstring});
legend('Amplitude','Sensitivity', 'Location', 'Best'); hold off;
saveas(7,fullfile(savpathstr2,'optimspaced_E.fig'));
saveas(7,fullfile(savpathstr2,'optimspaced_E.png'));

% phase nearest neighbor (average)
figure(8); clf; hold on;
[AX,H1,H2] = plotyy(Pn2(1:WN-1), ItotS2(1:WN-1), Pn2(1:WN-1),abs(SensN2));
set(get(AX(1),'YLabel'),'String','Amp_(ng/s)');
set(get(AX(2),'YLabel'),'String','Sens_((ng/s)/rad)');
set(AX(1),'YColor','r');
set(AX(2),'YColor','b');
set(H1,'LineStyle','-','Color','r');
set(H2,'LineStyle','.', 'Color','b');
xlabel('Phase_diff._between_nearest_neighbors_-_avg_(rad)');
title({'I_vs._Pn_[Optimized_grating]', labelstring});
legend('Amplitude','Sensitivity', 'Location', 'Best'); hold off;
saveas(8,fullfile(savpathstr2,'optimspaced_N.fig'));
saveas(8,fullfile(savpathstr2,'optimspaced_N.png'));

%angular velocity
figure(9); clf;hold on;
[AX,H1,H2] = plotyy(W(1:WN-1), ItotS2(1:WN-1), W(1:WN-1),abs(SensW2));
set(get(AX(1),'YLabel'),'String','Amp_(ng/s)');
set(get(AX(2),'YLabel'),'String','Sens_((ng/s)/(rad/s))');
set(AX(1),'YColor','r');
set(AX(2),'YColor','b');
set(H1,'LineStyle','-','Color','r');
set(H2,'LineStyle','.', 'Color','b');
xlabel('W_(rad/s)');
title({'I_vs._W_[Optimized_grating]', labelstring});
legend('Amplitude','Sensitivity', 'Location', 'Best'); hold off;
saveas(9,fullfile(savpathstr2,'optimspaced_W.fig'));
saveas(9,fullfile(savpathstr2,'optimspaced_W.png'));

% Final comparison: optimal vs. non-optimal

%phase ends
figure(10); clf;
plot(Pe(1:WN-1), abs(SensE),'.r', Pe2(1:WN-1),abs(SensE2), 'xb');
xlabel('Phase_diff._between_grating_ends_(rad)');
ylabel('Sensitivity_((ng/s)/rad)');
title({'Spacing_comparison_[phase:_grating_ends]', labelstring});
legend('equispaced','optimal', 'Location', 'Best');
saveas(10,fullfile(savpathstr2,'spacing_comparison_E.fig'));
saveas(10,fullfile(savpathstr2,'spacing_comparison_E.png'));

%phase nearest neighbor (average)
figure(11); clf;

```

```

plot(Pn(1:WN-1), abs(SensN),'.r', Pn2(1:WN-1),abs(SensN2), 'xb');
xlabel('Phase_diff_between_nearest_neighbors_-_avg_(rad)');
ylabel('Sensitivity_((ng/s)/rad)');
title({'Spacing_comparison_[phase:_grating_nearest_neighbor]', labelstring});
legend('equispaced','optimal', 'Location', 'Best');
saveas(11,fullfile(savpathstr2,'spacing_comparison_N.fig'));
saveas(11,fullfile(savpathstr2,'spacing_comparison_N.png'));

%angular velocity
figure(12); clf;
plot(W(1:WN-1), abs(SensW),'.r', W(1:WN-1),abs(SensW2), 'xb');
ylabel('Sensitivity_((ng/s)/(rad/s))');
xlabel('W_(rad/s)');
title({'Spacing_comparison_[angular_velocity]', labelstring});
legend('equispaced','optimal', 'Location', 'Best');
saveas(12,fullfile(savpathstr2,'spacing_comparison_W.fig'));
saveas(12,fullfile(savpathstr2,'spacing_comparison_W.png'));

% Save all workspace vars up to this point
%(DOES include interference curves)
cd(savpathstr2);
% clear ut1;
% save('int_data.mat');
% cd([savpathstr, '5\']); %return to script folder
cd(tempcd); %Return to original "current dir" saved at begining of function

end

```