# ADVANCES IN FETAL MAGNETOCARDIOGRAPHY USING SERF ATOMIC MAGNETOMETERS

by

Zachary J DeLand

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

Doctor of Philosophy

(Physics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2017

© Copyright by Zachary J DeLand 2017

All Rights Reserved

**DISCARD THIS PAGE**

# TABLE OF CONTENTS

Page

**LIST OF TABLES** . . . . . v

**LIST OF FIGURES** . . . . . vii

**ABSTRACT** . . . . . xix

**1 Introduction** . . . . . 1

1.1 Biomagnetism and the SQUID . . . . . 1
1.2 Biomagnetism with atomic magnetometers . . . . . 3
1.3 This work . . . . . 5

**2 SERF Magnetometer Theory and Extensions** . . . . . 7

2.1 Basic atomic magnetometry and limitations . . . . . 8
2.2 Spin-exchange collisions, spin-temperature, and the SERF regime . . . . . 10
2.3 DC-mode SERF magnetometry . . . . . 15
2.3.1 Optical pumping . . . . . 16
2.3.2 Spin-relaxation mechanisms . . . . . 19
2.3.3 Larmor precession . . . . . 23
2.3.4 Equations of motion and steady-state magnetometer response . . . . . 23
2.3.5 Dynamic field response . . . . . 29
2.3.6 Light shift . . . . . 32
2.3.7 Faraday rotation detection . . . . . 35
2.4 Diffusion mode . . . . . 38
2.4.1 Motivation . . . . . 38
2.4.2 Conceptual basis for diffusion-mode . . . . . 38
2.4.3 Optimizing cells for diffusion mode . . . . . 40
2.5 Z-mode . . . . . 43
2.6 Active feedback . . . . . 48

Page

**3 Apparatus** . . . 51

3.1 Magnetometer unit . . . 53
3.1.1 Vapor cell . . . 54
3.1.2 Heaters . . . 55
3.1.3 Shell and insulation . . . 57
3.1.4 Shell-mounted coils . . . 59
3.2 Lasers and fibers . . . 60
3.3 Magnetometer array . . . 61
3.3.1 Fiber splitters . . . 63
3.3.2 Pump optics . . . 65
3.3.3 Probe optics . . . 67
3.3.4 Detectors . . . 68
3.3.5 Feedback coil . . . 72
3.4 Magnetically shielded room . . . 72
3.4.1 MSR magnetic field coils . . . 73
3.5 Electronics . . . 76
3.5.1 Heater drivers . . . 76
3.5.2 Laser drivers and TEC controllers . . . 77
3.5.3 Coil drivers . . . 77
3.5.4 Magnetic field decoupling . . . 81
3.5.5 Photodiodes and amplifiers . . . 83
3.5.6 Feedback circuitry . . . 84
3.5.7 Data acquisition . . . 86

**4 Data Collection and Analysis** . . . 88

4.1 Initial Setup . . . 88
4.2 Alkali density measurement . . . 91
4.3 Residual field nulling . . . 92
4.4 Magnetometer calibration and characterization . . . 95
4.4.1 Chirp signal synthesis . . . 96
4.4.2 Coil driver circuit calibration . . . 98
4.4.3 Magnetometer calibration . . . 99
4.4.4 Environmental magnetic field noise measurements . . . 100
4.4.5 Additional noise measurements . . . 100
4.5 Magnetometer optimization . . . 104
4.6 Z-mode measurements . . . 107
4.6.1 Field nulling and modulation optimization . . . 109
4.6.2 Z-mode data acquisition . . . 110

Appendix

Page

4.7 Cardiographic measurements . . . 113
4.8 Noise Analysis . . . 115
4.8.1 Magnetic noise . . . 116
4.8.2 Technical noise . . . 121
4.8.3 Fundamental noise . . . 123

**5 Gradiometry** . . . 128

5.1 Introduction . . . 128
5.2 Gradiometer theory and figure of merit . . . 131
5.3 Previous Work . . . 136
5.4 Experimental methods . . . 138
5.4.1 Open-loop gradiometry . . . 138
5.4.2 Closed-Loop Gradiometry . . . 142
5.5 CMRR Stability . . . 147
5.6 Conclusion and outlook . . . 149

**6 Magnetocardiography** . . . 150

6.1 Cardiac electrophysiology . . . 150
6.1.1 Heart cycle . . . 151
6.1.2 Need for MCG . . . 153
6.2 MCG and fMCG characteristics . . . 155
6.3 MCG and fMCG measurement procedure . . . 157
6.3.1 Array setup and positioning . . . 157
6.3.2 Magnetocardiogram measurement . . . 158
6.4 MCG signal processing . . . 160
6.5 Single-channel processing . . . 162
6.6 Multi-channel processing . . . 163
6.6.1 An overview of ICA . . . 165
6.6.2 ICA Implementation . . . 167
6.6.3 Splined Independent Component Subtraction . . . 167
6.7 Averaging . . . 169
6.8 Discussion . . . 171
6.9 Comparison with SQUID . . . 173

**7 Conclusion and Future Directions** . . . 175

**APPENDICES**

Appendix

Page

Appendix A: Circuits . . . . . 187
Appendix B: Magnetometer Cell Fabrication . . . . . 198
Appendix C: Loop Shaping for Feedback . . . . . 212
Appendix D: Gradiometry Figure of Merit . . . . . 228
Appendix E: Adult MCG results . . . . . 232

**DISCARD THIS PAGE**

# LIST OF TABLES

Table Page

3.1 Design considerations for the magnetometer array. The two (sometimes competing) objectives of the apparatus are a large SNR for MCG detection while at the same time maintaining portability so the device can be easily maneuvered to an optimal position for measurement, as well as taken down and removed from the shielded chamber when necessary. . . . . 52

3.2 Shell-mounted coil geometries and calculated on-axis coil calibration at the center of the coil pair. . . . . 59

4.1 The three noise measurements commonly made to characterize the magnetometers. By making these three separate measurements, we can deduce what, if anything, is limiting the magnetometers' performances. . . . . 103

5.1 Summary of published results of atomic gradiometers. The figure of merit is calculated for $r_s = 5$ cm, and $p = 2$, typical values of $p$ and $r_s$ for fetal magnetocardiography applications **?**. The values of $\xi$ for previously-published work should be considered lower bounds, and the resulting figure of merit is calculated for an ambient environment sufficiently noisy such that $\left(\frac{\delta Bu}{\delta Bc}\right)^2 \ll \xi^{-2}$. Specifications for a commercial SQUID gradiometer in use at University of Wisconsin-Madison are included for comparison. . . . . 137

Appendix
Table

B.1 Stoichiometry for the chemical reactions in Equations (B.2a) and (B.2b) for a sample cell-making run. The $H_2O$ required indicates the *minimum* amount of water required to dissolve the given amount of rubdium chloride. Additional water was added to ensure the rubidium. Similarly, excess $Ba(N_3)_2$ solution was added to ensure that all of the rubidium reacted. . . . . 202

B.2 Issues and resulting modifications to cell-making procedure . . . . . 211

Appendix
Table Page

C.1 Laplace domain representations of the four loop-shaping building blocks. $s\,(=i\omega)$ is the independent variable related to the frequency and $\omega_0$ is the "center frequency" of the particular building block. . . . . 213

C.2 Phase complementary blocks . . . . . 215

C.3 Result of curve fitting the phase-response to the 3-pole, 1-zero model first approximated in Equation C.5. . . . . 220

C.4 Result of curve-fitting the phase-response to the 3-pole, 1-zero model first approximated in Equation C.5. The resistors and capacitors are off-the-shelf components avaliable at the department electronics shop. The resultant filter frequency using those components is shown in the last column. . . . . 223

**DISCARD THIS PAGE**

# LIST OF FIGURES

Figure Page

2.1 Diagram of the feedback-locked $M_x$ magnetometer used in [Kitching et al., 2011], reprinted with author's permission. In this setup, the alkali atoms are polarized with a circularly-polarized pump laser whose transmission is monitored by a photodiode (PD). A detected magnetic field $B_0$ will cause the atoms to precess and modulate the pump intensity on the photodiode. The atoms are driven with an oscillating field which can be locked to the signal detected by the photodiode. When the drive frequency of the coils matches the Larmor frequency $B_0$, the device becomes self-oscillating. A measurement of that Larmor frequency allows the field strength to be directly calculated. 9

2.2 A spin-exchange collision between two alkali atoms, each initially in the $|F, m_f\rangle = |2, 1\rangle$ (from [Walker and Happer, 1997]). The total spin and angular momentum projection of the atoms is conserved, but hyperfine level populations and individual magnetic sublevels are redistributed by the collision. The two hyperfine levels have opposite $g$ factors, and atoms in different sublevels will precess in opposite directions in a magnetic field $B$, leading to rapid decoherence. . . . . 10

2.3 Spin-exchange regimes, from [Happer and Tam, 1977]. In the following subfigures, $\omega_a$ corresponds to the Larmor frequency in the lower ($F = 1$, A) hyperfine state and $\omega_b$ to the frequency in the upper ($F = 2$, B) state. The plots in the right column are the frequency spectrum representations of the plots in the left column (i) In the low-SE regime, an atom will precess at its Larmor frequency until a spin exchange collision (vertical dotted line) causes it to change hyperfine levels. Note the different $g$ values of each level cause the atom to precess in the other direction after a SE collision. The atoms are allowed to precess for a time $T > 2\pi/\omega_{a,b}$ in each state. (ii) A maximally decoherent state with $\Gamma_{\text{SE}} \sim \omega_{a,b}$. (iii) In the high-SE limit ($\Gamma_{\text{SE}} \gg \omega_{a,b}$), the atom will undergo rapid SE collisions before it can precess very far. Given the statistical preference for the upper hyperfine state, over time, the atom will appear to precess with that state's phase at a slowed-down rate $\bar{\omega}$. The "jaggedness" of the precession phase smooths out as the spin exchange rate increases. . . . . 12

Figure Page

2.4 (a) The slowing down factor $q$ as a function of the atomic polarization $P$ for $^{87}$Rb ($I = 3/2$). (b) Spin-temperature distributions (from [Wyllie, 2012]) at several polarizations, with accompanying values of $q$. Note the equal population of equivalent $m_f$ substates in each hyperfine level. The $P = 50\%$ condition is optimum for SERF magnetometery (see Figure 2.9). . . . . 14

2.5 A plot of the spin-exchange relaxation rate (Equation (2.12)) at several cell temperatures and polarizations. In the SERF regime, the relaxation rate $\Gamma_{\text{SE}}$ due to spin-exchange collisions can be made arbitrarily small by appropriately "nulling" the transverse fields $B_\perp$. . . . . 15

2.6 A diagram of the magnetometer arrangement. Atoms are pumped by $\sigma^+$ light and their spins are aligned along $\hat{z}$. Any atomic precession into the $\hat{x}$ direction (caused by an applied $B_y$, for example) will cause a rotation in the polarization of a linearly polarized probe beam propagating perpendicularly to the pump. This polarization change can be detected at the far side of the atomic ensemble. . . . . 16

2.7 Optical pumping scheme for our experiment. An unpolarized ensemble of atoms initially contains roughly equal populations in the ground state ($5S_{1/2}$) magnetic sublevels. Optical pumping with resonant circularly-polarized D1 (795 nm) light drives atoms in $m_s = -1/2$ sublevel to the excited state ($5P_{1/2}$), $m_s = +1/2$ sublevel. An angular momentum selection rule prevents atoms in the $m = +1/2$ ground state level from absorbing the $\sigma^+$ photons (with $s = 1$). Rapid collisional mixing in the excited state coupled with non-radiative quenching collisions with nitrogen means that one half of the atoms pumped out of the $m_s = -1/2$ state will eventually end up in each ground sublevel. Thus, on average, one photon absorption adds 1/2 unit of angular momentum to the vapor. Spin-destruction collisions (Section 2.3.2) tend to return atoms to the $m_s = -1/2$ sublevel. The balance of the optical pumping rate $R$ and the spin-relaxation rate $\Gamma_{\text{sd}}$ determines the population in each sublevel. . . . . 18

Appendix
Figure Page

2.8 The relaxation rates $\Gamma_{sd}$ and $\Gamma_{diff}$ as functions of $N_2$ buffer gas pressure (a) and cell temperature (b). Each shows the dependence of the spin-polarization $P$, which influences $\Gamma_{diff}$ through $q$ (refer to Equation (2.32)). At low buffer gas pressures, collisions with the walls dominate the total relaxation rate. As buffer gas pressure increases, the spin-destruction collisions between the Rb and $N_2$ molecules become the dominant relaxation mechanism. Increasing temperature most strongly influences the spin-destruction collision rate, as the rate depends on Rb density and the Rb–$N_2$ spin-destruction rate, both of which grow exponentially with temperature (Equations 2.13 and 2.21). Not included in the total relaxation rate is the contribution of $\Gamma_{pr}$, which can be made arbitrarily large or small by varying probe laser parameters. The vertical dashed lines indicate the temperature and buffer gas pressure values in the other figure: the conditions at which we commonly operate in our experiment. As noted in Section 2.3.2.4, the *actual* relaxation rates in optically-pumped cells can be much higher due to the relaxation caused by the pumping rate $R$. . . . . 24

2.9 Normalized response (left axis, solid red line) and longitudinal polarization (right axis, dashed green line) as a function of the normalized pumping rate $\tilde{R} = R/\Gamma$. Note the response is maximized when $\tilde{R} = 1$ and the polarization is $P_z = 50\%$. . . . . 26

2.10 Magnetometer swept field response . . . . . 28

2.11 The frequency response of the magnetometer (Equation (2.48)) with various $\hat{z}$ field imbalances at $R = \Gamma$ ($\tilde{R} = 1$) and $\Omega_{1y} = 1$. Note the decreasing sensitivity, non-lorentzian lineshapes, and broadening as $\Omega_{0z}$ increases. In practice, the frequency response must be measured every time the magnetometer is calibrated, so the shape of these response curves provide convenient clues as to potential field imbalances. Imbalences can be corrected by re-shimming the local magnetic field until a response resembling the $\Omega_{0z} = 0$ curve is obtained. . . . . 31

2.12 The pump beam-induced light shift (Equation (2.54) in the magnetic field, normalized to a 1 cm waist and 1 mW pump power. The horizontal axis in this plot is the normalized detuning of the pump beam from the D1 resonance ($x = \Delta\nu/(\Delta\Gamma/2)$). The light shift is maximum at a detuning equal to half of the absorption linewidth ($x = 1$). The light shift can be eliminated entirely by tuning the laser exactly on resonance ($x = 0$). However if the vapor is very optically thick, tuning on resonance can make it difficult to optically pump the entire cell. . . . . 34

Appendix
Figure Page

2.13 Faraday rotation signal as a function of normalized detuning $x = \Delta\nu/(\Delta\Gamma/2)$. Without accounting for absorption (dashed line), the detuning which maximizes the rotation is $x = 1$ ($\Delta\nu = (\Delta\Gamma/2)$. With absorption (solid line), a larger detuning must be used to avoid having the probe light be completely absorbed by the vapor. The absorption curve uses an on-resonance optical depth of 100, which is common for our measurements. Probe detuning optimization scans (Section 4.5) help us choose the correct detuning to maximize signal for a particular measurement. . . . . 37

2.14 (a) The differential response (Equation (2.65)) as a function of the distance $\rho$ from the center of the pump beam at several different values of $N_2$ buffer gas pressure. The temperature of the atoms is 160°C. The atoms at a radius $\rho \leq \rho_0 = 0.05$ cm are "pinned" to the $\hat{z}$-axis. They are highly polarized ($P_z \sim 1$), but they have negligible response due to the extreme optical pumping. Only when the atoms diffuse to a radius $\rho > \rho_0$ is there any increase in sensitivity ($dP_x/d\Omega_y$). (b) The result of integrating Equation (2.65) out to $\rho = 0.4$ cm (roughly the size of the probe beam). The optimum response with these pump and probe parameters is realized with a pressure of 150-200 torr $N_2$. The response falls off faster on the low-pressure side of this maximum, so erring on the side of over-filling the cells with nitrogen is desirable compared to under-filling them. . . . . 42

2.15 The response-reducing prefactors from operating the SERF magnetometers in Z-mode. We generally try to operate at the point indicated by the black dot ($u \approx .90$). Here, both $\Omega_x$ and $\Omega_y$ sensitivity are at roughly 65% of the DC-mode level. If one desired to use the $n = 2$ harmonic to measure $\Omega_y$, the optimum location would use roughly $u = 1.5$, resulting in a similar level of $\Omega_x$ response, but the sensitivity to $\Omega_y$ would suffer greatly. . . . . 47

2.16 (a) The system in "open loop" mode. The input signal $r(s)$ is first modified by the prepared servo $K(s)$ before being passed to the system $G(s)$. (b) The system in "closed-loop" mode. The output signal $y(s)$ is fed back to the input of the system, and the servo $K(s)$ works to hold the output as close to the control signal $r(s)$ as possible, thereby minimizing the error signal $e(s)$. . . . . 50

3.1 A single magnetometer channel. The white 3-D printed shell contains the cell, heaters, and insulation. Holes in the shell allow for optical access for both pump and probe lasers. Triaxial coils (copper-colored) are wrapped directly onto the cell framework, enclosing the active region. At top left, the various electronic connections can be seen (rainbow ribbon: field cancelling and modulation currents; red/white twisted pairs: heater signals; brown/yellow twisted pair: thermistor signal. A thin layer of Mylar tape is often added to secure the shell in the magnetometer array holder. . . . . 53

Appendix
Figure Page

3.2 Heater panel drawing and photo . . . . . 56

3.3 Assembly of a magnetometer channel (clockwise from top left). First, the heater panels are arranged around the cell in a box formation and two layers of FEP heat-shrink tubing are placed around the assembly for structural support. The circular holes in the heater panels align with the optical paths the longitudinal pump and transverse probe beams take through the cell. The polyether ether ketone (PEEK) shell pieces are arranged around the perimeter of the assembly, and tape is applied to keep the pieces in place. Next, 2-3 layers of $\sim$2 mm-thick Aerogel is wrapped around the assembly. Note how the circular protrusions in the PEEK pieces keep the optical paths clear of the black Aerogel insulation. Finally, the cell and insulation are placed inside a 3-D printed shell which lends additional structural support and provides a mount for the magnetic field coils. . . . . 58

3.4 Magnetometer shell with the shell mounted $\hat{x}$, $\hat{y}$, and $\hat{z}$ coils highlighted in red, green, and blue, respectively. The slot at the bottom allows the heater and RTD leads to pass out of the shell. . . . . 60

3.5 Arrangement of the magnetometers in the array holder (orange dashed lines). Each cell (light blue) is enclosed in a four-walled PCB heater oven (dark green), surrounded by Aerogel insulation and placed inside a shell. Light is coupled into the array via fiber splitters connected to pump and probe tubes (black dashed lines). All beams are collimated and passed through high-extinction linear polarizers (brown). Pump light is circularly polarized with zero-order half waveplates (magenta). Probe light is passed through half waveplates (bright green), which provides equal division of light intensity between the two arms of polarizing beamslitting cubes. Light passing through the first beamsplitting cube is rotated an additional 90$^{o}$, so it is completely reflected by the second cube and sent through the cell. Immediately after the light is reflected off of a PBS, it is passed through a high-extinction polarizer (brown) to convert any remaining rotation noise into intensity noise. Rotation of the probe beam (primarily caused by detected magnetic fields) is detected by polarimeters after the light passes through the cell. . . . . 62

3.6 Photo of the entire apparatus connected to the wood frame, which allows for mounting on the SQUID gradiometer. RJ-45 sockets mounted on the frame provide connections for the ethernet cables (not connected) which carry all electronic signals between the apparatus and the electronics outside of the magnetically-shielded room. Fiber sockets provide interfaces between the polarizing fibers (also not connected) and the fiber splitters which break out the pump (blue fibers connecting to the apparatus) and probe (yellow fibers) laser light. . . . . 64

Appendix
Figure Page

3.7 A more detailed view of the main part of the apparatus. The four magnetometer channels (black shells) are enclosed within a structure which holds the optics: probe tubes (connected to the end of the yellow fibers), pump tubes (connected to the ends of the blue fibers), beamsplitting optics (placed between channels "A" and "E" and between channels "C" and "G") and the polarimeter detectors (wrapped with the green electrical tape, far left and far right) . . . . . 65

3.8 Fiber and ferrule holder designed by Mike Bulatowicz. Two separate clamps (green) hold the ferrule and fiber jacket, providing strain relief on the vulnerable bare fiber (lavender). A removable end cap allows for easy cleaning or polishing of the fiber ferrule should it become dirty or damaged. The end cap is secured by plastic screws (red) to the end of the main body of the holder. . . . . 66

3.9 Advanced detector . . . . . 71

3.10 Cartoon drawing of the magnetically shielded room, with the locations of the room-mounted coils indicated (red, $\hat{x}$ coils; green, $\hat{y}$ coils; blue, $\hat{z}$ coils). The approximate position of the door is indicated by the black rectangle. . . . . 74

3.11 Frequency-dependent calibration of the MSR-mounted Y coils. The red data point was obtained with a DC fluxgate, while the remaining values were determined by a differential calibration. All values agree well with the calibration given by Equation (3.6). . . . . 75

3.12 Sulai supply Noise . . . . . 80

4.1 Absorption scan . . . . . 93

4.2 Data Collection Block Diagram . . . . . 97

4.3 Chirp Calibration . . . . . 101

4.4 DC-mode noise . . . . . 105

4.5 Pump and Probe Optimization . . . . . 108

4.6 Z-mode noise . . . . . 112

4.7 heartposition . . . . . 114

Appendix
Figure Page

5.1 Time series generated using the ECGSYN program, along with their power spectral densities. To achieve a 10:1 fMCG:mMCG signal ratio, suppression of the maternal signal by at least a factor of 100 is required. . . . . 130

5.2 Configuration of source and magnetometers. The source, a fetal heart is modeled to be a current dipole $Id\vec{l}$. In the cartoon, the dipole is oriented perpendicular to the axis of the magnetometers, but it should be noted the intra-heart currents that generate the detected fields flow along all three axes. . . . . 131

5.3 $(g^p - 1) / (g^p + 1)$ as a function of the geometric parameter $g$ for various values of $p$. $p = 3$ is appropriate for a purely magnetic dipole source, while $p = 2$ is appropriate for a purely current dipole source. $g = 2$ for the case $L = r_s$. . . . . 135

5.4 A schematic of the apparatus. Two magnetometers are separated by a distance of $L$ = 4 cm. The magnetometers are operated in the Spin-Exchange Relaxation Free (SERF) regime. A circularly polarized pump beam slightly detuned from the Rb D1 line spin polarizes a rubidium vapor along the $\hat{z}$ axis. A linearly polarized probe beam detuned from the D2 line is directed perpendicularly ($\hat{x}$) to the pump beam. Magnetic fields in the $\hat{y}$ direction cause precession of the atoms's polarization, the direction of polarization of the transmitted probe light via Faraday rotation. This rotation is detected by a balanced polarimeter on the far side of the cell and recorded on a computer. In an alternative operation mode ("Z-mode"), parametric modulation is applied via the "local" $z$-coils. In this mode, the magnetometer is also sensitive to fields in the $\hat{x}$ direction. The two channels of the gradiometer (shown in dark red) are placed at the center of the magnetically shielded room ($right$). "Common" cancellation coils wrapped around the perimeter of the innermost shield are used to null the $x$ (red), $y$ (green), and $z$ (blue) fields at one of the magnetometers, as well as apply calibration and common-mode noise signals. . . . . 139

5.5 Calibrating the gradiometer. During calibration, the two channels of the were calibrated simultaneously by applying a time dependent field of the form $B(t) = (B_0 + B_1 t) \sin\left(Ate^{Ct}\right)$ (a.) to a common coil. The responses to the calibration field are recorded. Panel (b.) shows the responses of the two magnetometers in open-loop Z-mode operation. Panel (c.) shows the responses in closed-loop DC-mode operation. . . . . 141

5.6 Determination of $\xi$. White noise is added to the magnetically shielded room via the large coils mounted on the inside surface of the room itself. The individual channels operating in closed-loop DC mode are shown in the top panel, (a.) In panel (b.) the difference channel is shown multiplied by a factor of 10 for clarity. . . . . 143

Appendix
Figure Page

5.7 Power spectra of the data shown in Figure 5.6. (a.) Individiual channels. (b.) Sum and difference channels. Note the baseline noise on the difference signal of 10 fT/$\sqrt{\text{Hz}}$. 144

5.8 Comparison of the common-mode rejection ratio for various operating modes. $\xi$ was determined by taking the ratio of the sum and difference power spectra, for example the ones shown in Figure 5.7. The use of active feedback shows superior performance in suppressing common-mode noise. The large magnitude of uncorrelated 60 Hz line noise limits the performance of the gradiometer at that frequency. . . . . 146

5.9 The stability of $\xi$ over time was measured in DC-mode with and without active feedback. With active feedback ("closed-loop" operation), the device was able to maintain an acceptably high level of common-mode noise suppression over a time period comparable to a fMCG measurement. When operating without the feedback, however, the device suffers both in the absolute value of $\xi$ as well as its stability. Most notably, $\xi$ value drops below the critical value of $\xi \approx 100$ after approximately 3 minutes. This evidence strongly supports using active feedback in future work. . . . . 148

6.1 Electrophysiology of the human heart. The gold shading indicates the approximate location of the "pulse" that causes the contraction of the various chambers of the heart. These contractions pump blood between chambers as well as to the lungs and to the rest of the body. Each subfigure also contains an inset indicating part of the EKG corresponding to each stage of the cycle. The diastolic (resting) phase of the cycle is indicated at the top followed by the systolic sequence of (clockwise) *P, Q, R, S,* and *T* waves. Subfigures adapted from [NIH, 2011]. . . . . 152

6.2 Two consecutive systolic cycles with a period of diastole in between [not to scale]. The intervals most useful for cardiologists in determining arrhythmias are indicated. The heart rate (in beats per minute) is equal to $60/T_{RR}$, where $T_{RR}$ is the $RR$ interval time in seconds. The $QT$ interval is of particular interest for fMCG. . . . . 154

Appendix
Figure Page

6.3 Top: Time series of the maternal and fetal heartbeats created using the ECGSYN program [McSharry and Clifford, 2015]. The magnitudes and heart rates for each were based on commonly-acquired maternal and fetal MCG data with our apparatus positioned appropriately for a fetal measurement. Bottom: The frequency-spectrum representation of the the maternal and fetal signals from the top panel. A power spectral density of a real fMCG measurement is shown in gray. Note the total signal is dominated by the maternal signal between DC and 30 Hz, by the fetal signal over the range 30-50 Hz, and then by environmental noise (which was at a level of roughly 10 fT/$\sqrt{\text{Hz}}$ during this measurement) above that. To capture as much of the cardiographic signal as possible, environmental noise must be kept to a minimum, especially at low frequencies where most of the cardiographic information exists. . . . . 156

6.4 The array in position to make a fetal measurement. After the magnetometer array is mounted on the end of the out-of-use SQUID gradiometer (beige cylinder, top right), the mother lies down on the bed inside of the magnetically-shielded room. An ultrasound is performed to locate the fetal heart, and the array is translated and tilted so that the fetal heart lies roughly in the center of the four-unit array. Though laser light is completely contained within the apparatus, the patient wears infrared safety goggles to protect herself should unexpected stray laser light be present in the room. . . . . 159

6.5 Calibrated z-mode data collected during a fMCG measurement with no additional signal processing applied. The top four traces are the recorded $B_x$ (Channels 1, 2, 3, and 4, respectively), while the bottom four are the measurement of $B_y$. The prominent peaks occurring every $\sim$1 second are the *maternal QRS*-complexes. The fetal data is mostly obscured by noise, though some indications of it are visible on traces 5 and 6 ($B_{1y}$ and $B_{2y}$, respectively). Channels 1 and 2 were farthest from the maternal heart, and during this segment of data, likely closest to the fetal heart. The quality of this raw data shows the need for further signal processing. . . . . 161

6.6 The data from Figure 6.5 after each channel is calibrated by its own calibration function $\zeta_i(\omega)$ and two digital filters are applied: a notch filter at 60 Hz to remove electronic line noise and a 1-100 Hz bandpass filter that removes all signal outside of our band of interest. The fetal *QRS*-peaks (occurring at roughly twice the rate of the maternal peaks) are now apparent, especially on Channel $1y$, and small indications of those peaks are visible on other channels as well. Note the $y$ axes scales for each channel are different. . . . . 164

Figure Page

6.7 The eight ICA components after applying the CubICA algorithm to the filtered data in Figure 6.6. The maternal signal is primarily in component 2, while parts the *QRS* complex are also quite visible in component 3. Most of the fetal signal is confined to component 5, though evidence can also be seen in other components, including a small amount in the primary maternal component (2). The individual components were calculated using the *entire* 60 second time series of each magnetometer channel, not just the 10 second sample shown in Figures 6.5 and 6.6. . . . . 168

6.8 The eight magnetic field time series (the $\hat{x}$ and $\hat{y}$ components of the four channels of the array), after spICS reconstruction. The maternal features have been suppressed and some of the noise has been removed during the reprojection of the ICA components. The fetal *QRS* complexes are now clear on many channels, though the other structures in the cardiogram are still somewhat difficult to discern. The . . . . . 170

6.9 The averaged waveforms of the eight channels of processed fMCG data. The $y$ axis of this graph has units of femtotesla. Variations in the fetal heart rate mean that the averaging tends to "smear out" the features in the peripheral beats (such as the "QRS" complexes near .07 s and .93 s), but the central features are well-defined. The relevant cardiographic intervals are indicated by the red lines. The *QTc* intereval time is *not* a true measurement in milliseconds; rather, it is a value that combines the lengths of the *QT* and *RR* intervals to help fetal cardiologists determine if the fetus may be at risk for long-*QT* syndrome. This image was produced using the Wakai group's *Saki* software. . . . . 172

7.1 Cylindrical shield . . . . . 178

Appendix
Figure

A.1 Magnetometer swept field response . . . . . 188

A.2 PA107DP Heater . . . . . 191

A.3 Sulai coil driver . . . . . 193

A.4 Shell Coil Feedback . . . . . 196

A.5 Shell Coil Feedback . . . . . 197

B.1 Manifold . . . . . 200

Figure Page

B.2 Vacuum System . . . . . 203

B.3 Azide decomosition reaction. The temperature was slowly increased until we began to see evidence of the $N_2$ gas production (Equation (B.2a)). The many pressure spikes seen in this figure were a result of each cell (and even different parts of a single cell) being at different temperatures and thus undergoing the azide decomposition at different times. Note the difference in temperatures between Cell 1 (which as placed near the oven) and Cell 2 (which was closer to the middle). For this particular run, we used procedure 2 from Section B.2.3, turning off the heat to the oven as soon as we were confident that all of the azide had decomposed. The barium/rubidium reaction (Equation (B.2b)) took place at a later time. . . . . 205

C.1 The amplitude (left column) and phase (right column) response for the building blocks. From the top, pole, zero, inverted pole, inverted zero. Note that the frequency has been normalized to the "center frequency" $\omega_0$. For the amplitude response, the gain $p$ has been made unity. . . . . 214

C.2 The analog circuit representations for (clockwise from top left): pole, zero, inverted zero, inverted pole. For the pole and inverted zero, the center frequency is given by $\omega_0 = \frac{1}{R_2 C}$. For the zero and inverted pole, $\omega_0 = \frac{L}{R_2}$. The gain for the stage can be controlled by making $R_1$ a variable resistor, as the overall gain for the stage is $\frac{Z_{fb}}{R_1}$, where $Z_{fb}$ is the impedance in the "feedback" part of the stage (between the inverting input and output). . . . . 216

C.3 A simple two-stage circuit designed to phase-complement a system with a pole at 100 Hz and zero at 1000 Hz. The . . . . . 217

C.4 Amplitude and phase response for the DC-SERF magnetometer. . . . . 219

C.5 Successive approximations to curve-fitting the phase response data. From top: 1-pole fit; 1-pole, 1-zero fit; 2-pole, 1-zero fit; 3-pole, 1-zero fit; 3-pole, 1-zero fit with center frequencies $f_{0_i}$ optimized by a least-squares algorithm. . . . . 224

C.6 Curve fit to the amplitude response using a 3-pole, 1-zero fit with the center frequencies from the phase response fit. . . . . 225

C.7 Servo circuit created with the phase complementary pieces determined by the curve-fitting procedure. The component values can be found in Table C.4. The voltage follower between the poles and zero stages serves to impedance-isolate the stages from each other. . . . . 225

Figure Page

C.8 Phase and amplitude response of the system $G$ and the servo $K$. Note the oppposite behaviors of the phase responses of $G$ and $K$. . . . . 226

C.9 Phase and amplitude response complete system $GK$. . . . . 227

D.1 Configuration of source and magnetometers. The source, a fetal heart is modeled to be a current dipole $Id\vec{l}$. In the cartoon, the dipole is oriented perpendicular to the axis of the magnetometers, but it should be noted the intra-heart currents that generate the detected fields flow along all three axes. . . . . 228

D.2 $(g^p - 1) / (g^p + 1)$ as a function of the geometric parameter $g$ for various values of $p$. $p = 3$ is appropriate for a purely magnetic dipole source, while $p = 2$ is appropriate for a purely current dipole source. $g = 2$ for the case $L = r_s$. . . . . 231

E.1 Adult Data . . . . . 233

# ABSTRACT

The human body is a complex source of electromagnetic signals which can provide clinicians with critical information about a person's health. Well-established methods, such as measuring electrocardiograms (ECGs or EKGs) and electroencephalograms (EEGs) are often inadequate for measuring particularly small signals or those which are shielded by tissue or bone. In these cases, measurements of the associated magnetic fields from these electronic processes provide the only insight. One particular example is fetal magnetocardiography (fMCG), the measurement of the magnetic fields produced by electrical signals regulating the heartbeat of an *in utero* fetus. The measurement of these signals is critical for clinicians to diagnose heartbeat arrhythmias, which can be treated if detected early enough in gestation. Previously, the only method of detecting these small $\sim$1 pT signals was using magnetometers based on superconducting quantum interference devices (SQUIDs). While these devices demonstrated excellent ($\sim$1 fT/$\sqrt{\text{Hz}}$ ) sensitivity, they are expensive to purchase, operate, and maintain and must be operated in magnetically shielded chambers large enough to house the large cryogenic hardware required for their operation. These drawbacks make them unattainable for most hospitals and clinics, so alternate methods of measuring these small biomagnetic fields have been desired.

More recently, atomic magnetometers operating in the spin-exchnage relaxation-free (SERF) regime have shown promise in supplementing or perhaps replacing SQUIDs for biomagnetic applications. SERF magnetometers have demonstrated an even higher sensitivity (¡ 1 fT/$\sqrt{\text{Hz}}$ ) than SQUID magnetometers and do so while rectifying many of the shortcomings (notable cost and size) of the SQUID systems. SERF magnetometers have already been used to measure many

biomagnetic signals, including most recently, fMCG. In this thesis, we present a compact, four-channel array of vector SERF magnetometers, each of which can measure two components of the local magnetic field with a sensitivity of 5-10 fT/$\sqrt{\text{Hz}}$ . Though the magnetometers are fixed, the entire array is portable and can be set up and taken down in minutes and provides a flexible platform to measure biomagnetic signals. We employ several extensions to "traditional" SERF magnetometry, including longitudinal field modulation, diffusive-mode operation, and active field-stabilization feedback to enhance our device's ability to reliably measure a fetal magnetocardiogram. We present a thorough investigation of using a pair of adjacent magnetometers in order to create a hardware gradiometer and suppress unwanted common-mode magnetic noise which interferes with our ability to measure the fetal heart's magnetic field. Finally, we present an eight-channel acquisition of a 28 week gestational age fetus and implement sophisticated signal processing techniques to isolate the magnetic signal of interest.

# Chapter 1

# Introduction

The study of magnetic fields began in ancient China. People discovered that allowing a hanging piece of lodestone (a naturally-magnetized ore) to freely rotate always resulted in the lodestone oriented in the same direction: along a north-south axis in alignment with the field produced by the earth. Later, the use of iron bars in compasses allowed explorers to cross continents and oceans without the aid of landmarks to guide them. Despite their usefulness for navigation, compasses could only measure the *direction* of the local magnetic field. The absolute *magnitude* of the magnetic field was a mystery.

In C.F. Gauss' 1832 paper "The Intensity of the Earth's Magnetic Force Reduced to Absolute Measurement", he described an apparatus not unlike the hanging piece of lodestone. In his experiment, a bar magnet suspended horizontally from a gold fiber was purposely misaligned from the earth's magnetic field. When released, the torque exerted by the earth's field caused the bar to oscillate around the axis of the field. By comparing the oscillation frequencies of a magnetized and demagnetized bar, Gauss was able to deduce for the first time the magnitude of the earth's magnetic field. His experiment pioneered the field of *magnetometry*, the precision measurement of magnetic fields.

## 1.1 Biomagnetism and the SQUID

Human biomagnetism (the magnetic fields produced by the human body) has become one important area in magnetometry research. Biomagnetic fields are tiny compared to the earth's field; the largest such field is produced by the electrical currents which control the beating of

the human heart. With a magnitude of 50-100 pT at the surface of the skin, the human heart's magnetic field is a million times smaller than the Earth's field. These cardiographic magnetic fields could not be accurately measured until the advent of the superconducting quantum interference device (SQUID) magnetometer in the 1970s [Cohen et al., 1970]. Even with this development, the measurement of the adult magnetocardiogram (MCG) was practically nothing more than a complementary procedure; nearly all of the cardiographic data obtained in an MCG could already much more easily be obtained with an *electro*cardiogram (ECG or EKG), a well-established and commonly-used medical procedure.

On the other hand, much of the information regarding processes in the human body can *only* be accessed through the resulting magnetic fields. Tissue, bone, and other internal structures tend to distort or suppress the transmission of electrical signals from source to skin, but these same materials do little to attenuate the simultaneously-produced magnetic signals. One example in which magnetometry provides the only insight into the internal processes of the body is the fetal magnetocardiogram (fMCG): the measurement of the magnetic fields produced by the beating heart of an *in utero* fetus. During gestation, the fetus is encased in the *vernix caseosa*, a material which strongly attenuates electrical signals, but has a lesser effect on magnetic ones [Wakai et al., 2000]. Thus, while a pair of ECG leads would have little hope of detecting the cardiographic electric signal, a sufficiently sensitive magnetometer can detect the fetal magnetocardiographic signals at the surface of the mother's abdomen.

fMCG can be used to diagnose a host of fetal cardiac arrhythmias (abnormalities in the heartbeat), one of which is long-$QT$ syndrome (LQTS) [Cuneo et al., 2003]. It has been postulated that this arrhythmia a mechanism of sudden cardiac death for affected fetuses. Work published in Cuneo's paper suggests that the condition is treatable and reversible *if the condition can be detected at a suitably early gestational age*. The diagnosis of the arrhythmia is in turn conditional on the ability to measure the fMCG. This task is considerably more difficult than adult MCG. The small size of the fetal heart results in a generated magnetic field that is 10-100 times smaller than an adult heart at a comparable distance.

SQUIDs have proved useful in the diagnosis of LQTS, including progress made at the University of Wisconsin-Madison [Cuneo et al., 2013]. The UW apparatus consisted of 20-40 SQUID gradiometers (devices sensitive magnetic field gradients). Uniform background environmental noise was largely suppressed as the gradiometers were only sensitive to fields that varied in magnitude or direction over the gradiometer baseline (the distance between the two "arms" of the gradiometer); 5 cm for the device used at UW-Madison). Until ten years ago, these SQUID apparatus were the only devices that had proved capable of measuring an fMCG.

Despite their excellent sensitivity and background-noise suppression abilities, SQUIDs are not without their drawbacks. The devices themselves are very costly, on the order of $10,000/channel in some cases. Because of the small size of the magnetic signal being detected compared to the size of the "noise" magnetic fields generated by electronics, large metal objects, and the earth itself, fMCG measurements must be performed in magnetically-shielded rooms. The shielding is prohibitively expensive, so small magnetic field-sensing devices would be ideal as the size of the shielded chamber could be made correspondingly small. Unfortunately, SQUIDs require large dewars of liquid helium to keep the coils superconducting. Not only is the liquid helium itself costly, but the size of the dewar also sets a lower limit on the size of the shielded room. The cost of purchasing, installing, and maintaining these large shielded rooms means that fMCGs are generally only measured by SQUIDs in specialized research facilities like the one at UW-Madison. Smaller hospitals simply cannot afford the expense of a SQUID fMCG device, despite the lifesaving measurements it can provide.

## 1.2 Biomagnetism with atomic magnetometers

More recently atomic magnetometers have shown the potential for replacing costlier SQUID sensors in biomagnetism experiements. Spin-Exchange Relaxation Free (SERF) magnetometers, invented by the Romalis Group at Princeton in the early 2000s [Allred et al., 2002], have shown an even higher sensitivity than SQUIDs: as low as .160 fT/$\sqrt{\text{Hz}}$ at 40 Hz [Dang et al., 2010]. The fundamental limit to the sensitivity of these devices is another two and a half times lower. Furthermore, SERF magnetometers can offer this performance at a fraction of the cost of expensive

SQUID systems. Far from relying on cryogenics to keep its sensors at superconducting temperatures, SERF magnetometers generally operate in the 140-200°C temperature range, which can be achieved with gentle resistive heating and a few layers of thermal insulation. Without the need to accommodate a large dewar of liquid helium, SERF magnetometers can operate in much smaller magnetically-shielded environments at a drastically reduced cost.

Several biomagnetism experiments using atomic magnetometers have been previously performed. Bison et al. [2009] built an experiment in which an adult MCG could be measured using a tri-level array of 25 cesium atomic magnetometers. In this experiment, (performed in an eddy-current shield) the magnetometers were arranged with three above and three below a central plane of nineteen sensors. This configuration acted as second-order gradiometers, as feedback from the satellite units was used to stabilize the central units. With this strategy, the experimenters were able to suppress ambient noise by a factor of 1000 and successfully acquire an adult MCG. However, the sensitivity of this device was limited to 300 fT/$\sqrt{\text{Hz}}$ , insufficient for a fMCG measurement.

The key features of all of these experiments were unified in the work of Bob Wyllie at UW-Madison [Wyllie, 2012]. Wyllie constructed an array of four independent SERF magnetometers, arranged in a square with a side length of $\sim$7 cm. Working in the same magnetically-shielded chamber as Shah and Wakai, Wyllie demonstrated a sensitivity of $\sim$10 fT/$\sqrt{\text{Hz}}$ at frequencies above 10 Hz. This four-channel device was first used to measure the MCG of 13 adult subjects [Wyllie et al., 2012a]. This work hinted at the possibility of measurement on fetal subjects, and soon after, the device yielded the first measurement of a fetal MCG using an atomic magnetometer [Wyllie et al., 2012b]. In this second paper, the fMCG of a 31 weeks gestational age fetus was successfully measured, and an analysis of the data (similar to the ones that would be used in a clinical setting to diagnose an arrhythmia) was performed. SQUID measurements of the same fetus, taken shortly after the SERF measurement, showed comparable quality and the interval timings between the two fMCGs matched to within 10%.

Even more recently, Shah and Wakai [2013] designed a compact single-channel atomic magnetometer with a sensitivity of $<$10 fT/$\sqrt{\text{Hz}}$ . They acquired sequential adult MCG measurements with the atomic magnetometer and a commercial SQUID magnetic sensing device in a $\mu$-metal

shielded chamber. The data obtained from the atomic magnetometer was nearly indistinguishable from the SQUID-collected data.

## 1.3 This work

In this work, we will build on the success of Wyllie's experiment and address several areas of improvement. We maintain four individual magnetometer channels, but have decreased the channel spacing (and the overall size of the apparatus) for improved fMCG signal detection and background noise suppression. Each magnetometer channel of the array exhibits a noise floor sensitivity of 5-10 fT/$\sqrt{\text{Hz}}$ in the 10-100 Hz band–the frequency range that contains much of the critical data in an fMCG. We also extend the functionality of these units by operating in a parametrically-modulated mode [Li et al., 2006], which allows us to detect a second directional component of the magnetic field at the location of each sensor. Operating in this mode also suppresses the effect of "technical noise"–noise generated by equipment and detection–in that second direction. This noise is not magnetic in nature but appears as magnetic noise in measurements. For only a modest decrease in the magnetometers' responses to magnetic fields, we are able to effectively double our channel count from four to eight.

Concern about elevated environmental noise levels led us to experiment with using two channels of our device as a gradiometer, and we measured its ability to subtract away noise common to both channels. For this work, we developed an active feedback scheme to help stabilize the magnetometers. Operating in this mode helps with gradiometry, as the cell-to-cell differences which often limit the ability to effectively calibrate the units against each other are mitigated. Furthermore, it suppresses technical noise that can arise if the environmentals fields drift with time.

Ultimately, we will present the fetal magnetocardiogram measurements acquired from a fetus at under 30 weeks gestational age. Because these measurements were captured in our parametrically-modulated mode, we were able to implement more sophisticated signal processing techniques when analyzing the fMCG–techniques that were previously only used to analyze SQUID-generated fMCG data.

The thesis is presented as follows: Chapter 2 provides the theory behind SERF magnetometers and demonstrate the important features that set them apart from traditional magnetometers that are limited in their sensitivity by spin-exchange collisions between the alkali atoms. The theory behind three extensions of basic SERF magnetometry (the aforementioned parametrically-modulated mode, active feedback operation, and a configuration which uses diffusion of optically-pumped atoms to reduce laser-generated magnetic fields) are also discussed. Chapter 3 describes the apparatus, highlighting the upgrades made over the past several years, notably an improvement to the magnetic field nulling system. By orthogonalizing the twelve magnetic field coils used to null the four channels, the device can now be set up to measure fields in a fraction of the time that it once took. Chapter 4 contains the details of the data-taking process, including the procedures used to zero the local magnetic field at each sensor and to calibrate the sensors. A discussion of several of the prominent magnetic field noise sources and their effect on measurements is included at the end of this chapter. In Chapter 5, we discuss our work using two nearby magnetometer channels as a magnetic gradiometer and demonstrate our ability to suppress common-mode magnetic field noise by a factor of over 400. Chapter 6 begins with a discussion of the signal of interest for this work, the magnetocardiogram, the cardiac processes that generate it, and its properties. The chapter concludes with our 8-channel measurement of a fetal magnetocardiogram, and proceeds step-by-step through the advanced signal processing techniques we use to transform our raw data into something that could be clinically-viable. Finally, we conclude the thesis with our plans for the project going forward, using what we have learned to steer the project into the future.

# Chapter 2

# SERF Magnetometer Theory and Extensions

In this chapter, we present a short introduction to atomic magnetometry. We focus on how the fundamental limitations of atomic magnetometers motivated the earliest work with Spin-Exchange Relaxation Free (SERF) magnetometers. SERF magnetometers are essential in our efforts to capture adult and fetal magnetocardiograms (MCG and fMCG, respectively). We next provide an overview of the theory behind SERF magnetometers, to complement other literature [Allred et al., 2002, Budker and Romalis, 2007, Kominis et al., 2003, Kornack, 2005, Ledbetter et al., 2008, Seltzer, 2008, Wyllie, 2012]. We then discuss the theory behind three extensions to standard "DC-mode" SERF operation:

1. *diff-SERF*: A variation of a SERF magnetometer that takes advantage of diffusion to reduce the effects of a spatially-varying AC Stark shift field, thereby increasing the magnetometer's response and reducing noise. [Sulai et al., 2013]

2. *Z-mode*: A modified SERF magnetometer which employs a large (on the scale of normal SERF operation), high-frequency oscillating field along the direction of the pump laser which allows two-component detection of the magnetic field (versus a single direction for normal SERF magnetometers). [Li et al., 2006]

3. *Active Feedback*: Employing feedback to stabilize the sensors against fluctuations in the magnetic field. This mode can theoretically be used in conjunction with other any of the other magnetometer modes.

## 2.1 Basic atomic magnetometry and limitations

Magnetometers can be constructed out of any object with a non-zero magnetic moment $\mu$. An applied magnetic field $B$ will cause a torque $\tau$ on the the object, where $\tau = \mu \times B$. The field thus causes the object to precess in the field at its Larmor frequency $\Omega = \gamma B$, where $\gamma$ is the gyromagnetic ratio of that object. For a free electron, the gyromagnetic ratio is

$$\gamma_e = g_e \frac{\mu_B}{\hbar} \tag{2.1}$$

where $g_e$ is the electron's g factor and $\mu_B$ is the Bohr magneton.

For an electron bound to an atom, the atom's nuclear spin must also be taken into account and the precession frequency in a field is

$$\begin{aligned} \Omega &= \gamma B/q \\ &= \frac{g}{q}\frac{\mu_B}{\hbar}B. \end{aligned} \tag{2.2}$$

$g$ is the Landé g-factor for the particular atomic level being measured. $q$ is the *slowing-down factor*: this factor accounts for the hyperfine interaction between the electron and the nucleus and the fact that the nuclear spin increases the rotational inertia of the atom and *slows down* the atomic precession in a given magnetic field. $q$ is dependent on both the nuclear spin of the atom and, for a system of atoms, m-level distribution of the electrons in the system. Because $\gamma$ is assembled from fundamental constants and known atomic quantum numbers, by measuring $\Omega$, one can deduce the magnitude of $B$.

For a single atom, the sensitivity of such a measurement is given by $\delta B \simeq 1/\gamma\sqrt{T_2 t}$, where $T_2$ is the coherence time (the time during which the atom's phase is well-defined) for the precession and $t$ is the measurement time. The sensitivity can be increased by observing the coherent precession of $N$ atoms. If we write $N = nV$ where $n$ is the atomic density and $V$ is the measurement volume, then

$$\delta B \simeq \frac{1}{\gamma\sqrt{nVT_2 t}}, \tag{2.3}$$

which is the fundamental atomic shot-noise-limited sensitivity for any atomic magnetometer [Kominis et al., 2003]. To observe this coherent precession, it i first necessary to align the spins of the

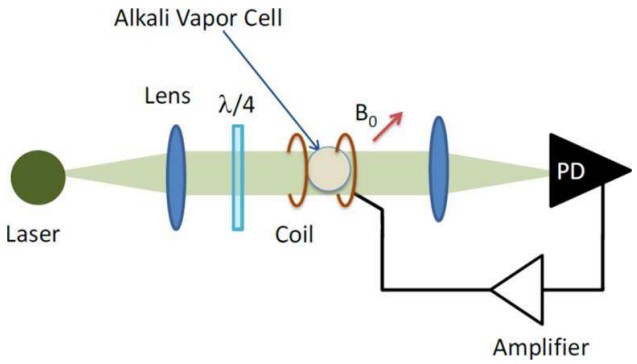



**Figure 2.1:** Diagram of the feedback-locked $M_x$ magnetometer used in [Kitching et al., 2011], reprinted with author's permission. In this setup, the alkali atoms are polarized with a circularly-polarized pump laser whose transmission is monitored by a photodiode (PD). A detected magnetic field $B_0$ will cause the atoms to precess and modulate the pump intensity on the photodiode. The atoms are driven with an oscillating field which can be locked to the signal detected by the photodiode. When the drive frequency of the coils matches the Larmor frequency $B_0$, the device becomes self-oscillating. A measurement of that Larmor frequency allows the field strength to be directly calculated.

$nV$ atoms. This can be accomplished by optical pumping the ensemble with resonant laser light, and the precession frequency of the atoms can be detected by monitoring the amount of absorbed pump light. This type of magnetometer, referred to as an $M_x$ magnetometer, was first developed by Bell and Bloom [Bell and Bloom, 1957], and it is still a popular choice for high-sensitivity magnetometery [Aleksandrov et al., 1995, Schwindt et al., 2007], including biomagnetic detection [Groeger et al., 2006]. A simple diagram of the scheme is shown in Figure 2.1.

For a fixed volume, it would appear that to increase the sensitivity of the device, one would want to increase the atomic density $n$. This is only true, however, if this does not also proportionally decrease the coherence time $T_2$. In fact, the coherence time of an ensemble of precessing atoms will often be limited by spin-exchange collisions between the atoms, such that

$$T_2^{-1} \sim T_{SE}^{-1} = \Gamma_{SE} = n\bar{v}\sigma_{SE} \tag{2.4}$$

where $\bar{v}$ is the thermal velocity of the atoms and $\sigma_{SE}$ is the spin-exchange collision cross section between the optically-pumped atoms ($1.9 \times 10^{-14}$ cm$^{-2}$ for $^{87}$Rb-$^{87}$Rb atoms [Gibbs and Hull, 1967]). Comparison of Equation (2.3) and Equation (2.4) shows that simply increasing $n$ will not increase the magnetometer's sensitivity if spin-exchange collisions dominate the decoherence. The desire to circumvent the traditional limits placed on sensitivity by this relationship was the driving

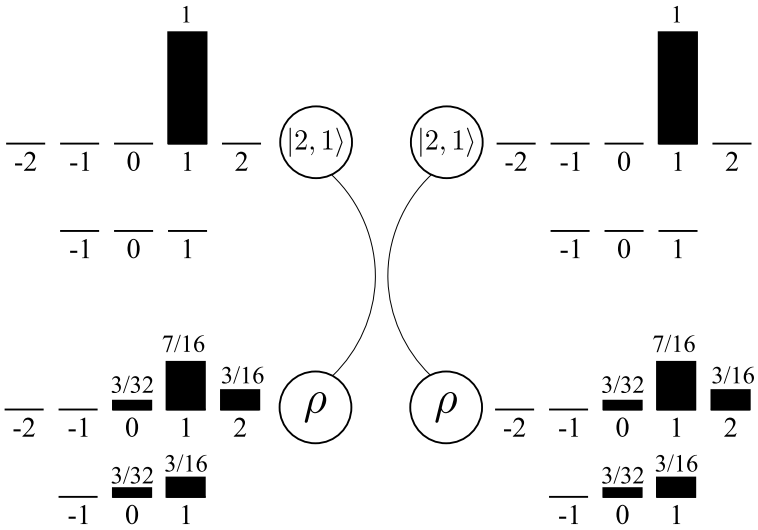



**Figure 2.2:** A spin-exchange collision between two alkali atoms, each initially in the $|F, m_f\rangle = |2, 1\rangle$ (from [Walker and Happer, 1997]). The total spin and angular momentum projection of the atoms is conserved, but hyperfine level populations and individual magnetic sublevels are redistributed by the collision. The two hyperfine levels have opposite $g$ factors, and atoms in different sublevels will precess in opposite directions in a magnetic field $B$, leading to rapid decoherence.

force for the development of the Spin-Exchange Relaxation Free (SERF) magnetometer [Allred et al., 2002], and the basis for the magnetometers used in the work described in this thesis.

## 2.2 Spin-exchange collisions, spin-temperature, and the SERF regime

An alkali-alkali spin exchange collision is shown in Figure 2.2. During a spin exchange collision, the electron spins of the atoms are coupled together, though the nuclear spin and total spin of the atom are conserved. This results the atoms leaving the collision with their total angular momentum projection conserved (*i.e.*, $m_{f1} + m_{f2} = m'_{f1} + m'_{f2}$), but in a statistical distribution of $|F, m_f\rangle$ states. For a $^{87}$Rb atom in its ground state $5^2S_{1/2}$, $I = 3/2$ and $J = S = 1/2$, so the Landé $g$-factors for each hyperfine level can be calculated as

$$\begin{aligned} g_F &= g_s \frac{F(F+1) - I(I+1) + J(J+1)}{2F(F+1)} \\ &= \begin{cases} -1/2 & (F = 1) \\ +1/2 & (F = 2) \end{cases} \end{aligned} \tag{2.5}$$

Recall that $\Omega \propto g$, so this difference in sign means that the redistribution to the different $F$ levels will cause the colliding atoms to precess at their Larmor frequency in opposite directions after the collision. This leads to a rapid decoherence of the atomic spins *if* the spin-exchange rate is sufficiently slow such that the atoms are allowed to precess a significant amount through a Larmor rotation between spin-exchange collisions.

This picture changes considerably, however, if the spin-exchange collisions occur at a rate much higher than the Larmor precession frequency, (that is, $\Omega T_{SE} \ll 2\pi$). In other words, the electron does not precess very far before it undergoes another spin-exchange collision. In this regime, it is not appropriate to think of each atom precessing in an individual hyperfine state. Rather, the atoms precess in a coherent manner in the direction that they would in the upper hyperfine state ($F = 2$) (given the higher statistical weight of that state and higher electron-spin polarization), but at a slowed rate compared to a free atom in that $F = 2$ state. The low-SE, rapid-SE, and a maximally decoherent "medium"-SE regimes are shown in Figure 2.3.

When the spin-exchange collision rate between alkali-metal atoms is rapid compared to all other rates in the alkali-metal spin evolution (most importantly, the pumping rate), the population distribution of the atoms into their magnetic sublevels satisfies the spin-temperature distribution [Anderson et al., 1960, Lancor, 2011, Savukov and Romalis, 2005], such that

$$n_{m_f} \propto \exp[-\beta m_f]. \tag{2.6}$$

Note that in this distribution, the corresponding $m_f$ sublevel populations in the two hyperfine levels are equal. Additionally, the population ratio of neighboring sublevels with angular momentum projections $m$ and $m-1$ is given by

$$\frac{n_m}{n_{m-1}} = \exp[-\beta]. \tag{2.7}$$

The spin temperature parameter $\beta$ is defined as ([Savukov and Romalis, 2005])

$$\beta = \ln\left(\frac{1+P}{1-P}\right), \tag{2.8}$$

or conversely,

$$P = \tanh(\beta/2), \tag{2.9}$$

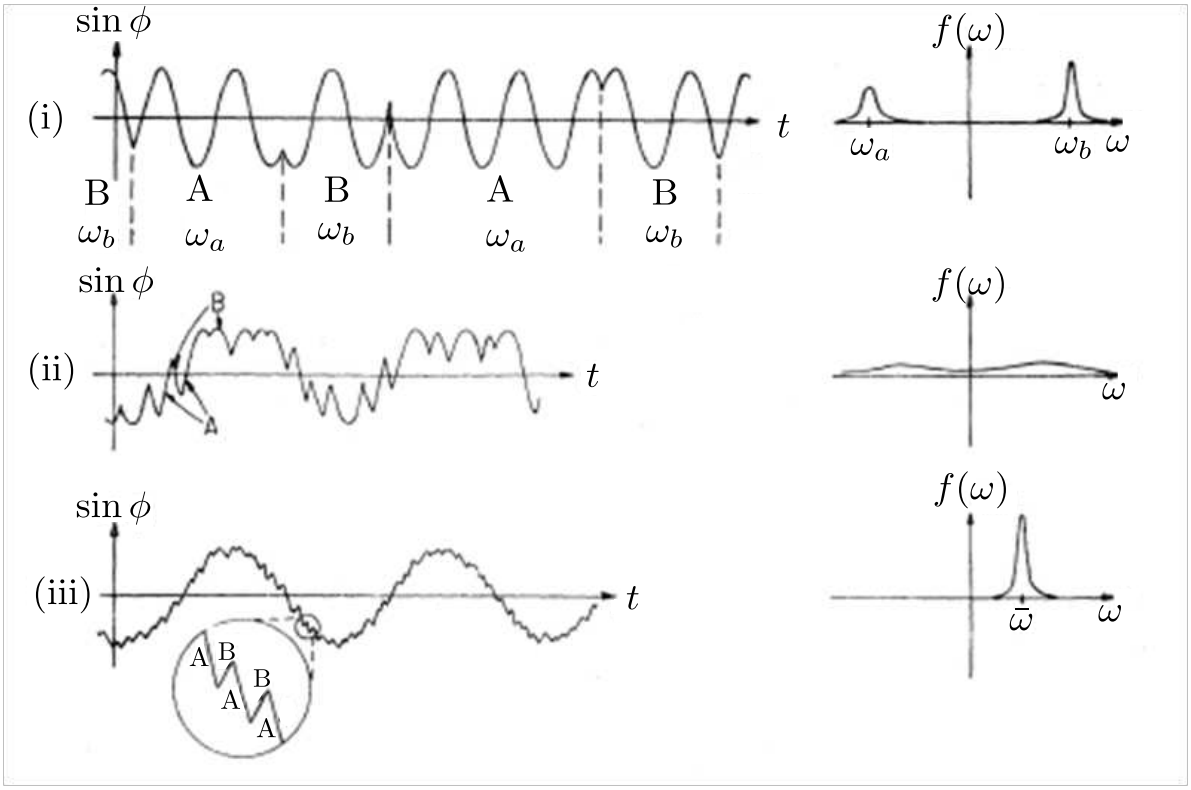



**Figure 2.3:** Spin-exchange regimes, from [Happer and Tam, 1977]. In the following subfigures, $\omega_a$ corresponds to the Larmor frequency in the lower ($F = 1$, A) hyperfine state and $\omega_b$ to the frequency in the upper ($F = 2$, B) state. The plots in the right column are the frequency spectrum representations of the plots in the left column (i) In the low-SE regime, an atom will precess at its Larmor frequency until a spin exchange collision (vertical dotted line) causes it to change hyperfine levels. Note the different $g$ values of each level cause the atom to precess in the other direction after a SE collision. The atoms are allowed to precess for a time $T > 2\pi/\omega_{a,b}$ in each state. (ii) A maximally decoherent state with $\Gamma_{\text{SE}} \sim \omega_{a,b}$. (iii) In the high-SE limit ($\Gamma_{\text{SE}} \gg \omega_{a,b}$), the atom will undergo rapid SE collisions before it can precess very far. Given the statistical preference for the upper hyperfine state, over time, the atom will appear to precess with that state's phase at a slowed-down rate $\bar{\omega}$. The "jaggedness" of the precession phase smooths out as the spin exchange rate increases.

where $\mathbf{P} = 2\langle \mathbf{S} \rangle$ is the alkali electron spin-polarization produced by the optical pumping (Section 2.3.1). In the spin-temperature distribution, the total atomic angular momentum is related to the electron angular momentum by the slowing-down factor $q(I, P)$, where

$$\begin{aligned} q(I,P) &= \frac{\langle F_z \rangle}{\langle S_z \rangle} \\ &= 1 - \coth(\beta/2) + (2I+1)\left(\coth\left[\beta/2\right]\right)\left(\coth\left[\beta(I+1/2)\right]\right) \end{aligned} \tag{2.10}$$

where $I$ is the atom's nuclear spin. Substituting Equation (2.8) into Equation (2.10) and using $I = 3/2$ (for $^{87}$Rb), we obtain

$$q(P) = \frac{6 + 2P^2}{1 + P^2}, \tag{2.11}$$

a function that decreases smoothly from 6 to 4 as the electron spin polarization increases from 0 to 1 (Figure 2.4a). Several spin-temperature distributions at various alkali polarizations are shown in Figure 2.4b. Happer and Tam [1977] and Allred et al. [2002] showed that in the rapid spin-exchange regime, the spin-exchange relaxation rate decreases quadratically with decreasing transverse magnetic field and vanishes in zero field:

$$\Gamma_{\text{SE}} = \Omega_{\perp}^2 T_{\text{SE}} \left( \frac{1}{2} - \frac{(2I+1)^2}{2q^2} \right) \tag{2.12}$$

Here, $T_{\text{SE}} = (n_{\text{Rb}} \bar{v} \sigma)^{-1}$ represents the time between rubidium-rubidium spin-exchange collisions. Using the standard method for calculating thermal velocity $\left(\bar{v} = \sqrt{3k_B T/m}\right)$, the cross section $\sigma$ for $^{87}$Rb-$^{87}$Rb collisions ($\approx 1.9 \times 10^{-14}$ cm$^2$ [Gibbs and Hull, 1967]) and the rubidium vapor pressure model from Alcock et al. [1984]

$$n_{\text{Rb}}(T) = \left( \frac{7.25 \times 10^{16}}{T} \right) 10^{9.318 - 4040/T}, \tag{2.13}$$

we can evaluate the spin exchange rate $T_{SE}^{-1}$ and then use that to calculate the corresponding relaxation rate $\Gamma_{\text{SE}}$. Figure 2.5 shows this relaxation rate at several different temperatures and alkali polarizations as a function of the transverse field. As long as the magnetic fields are kept sufficiently low (*i.e.* $T_{SE}^{-1} \gg 2\pi\gamma B$), the spin-exchange relaxation rate will not be the primary source of decoherence. The parameter space for which this condition is met is known as the Spin-Exchange Relaxation Free, or SERF regime.

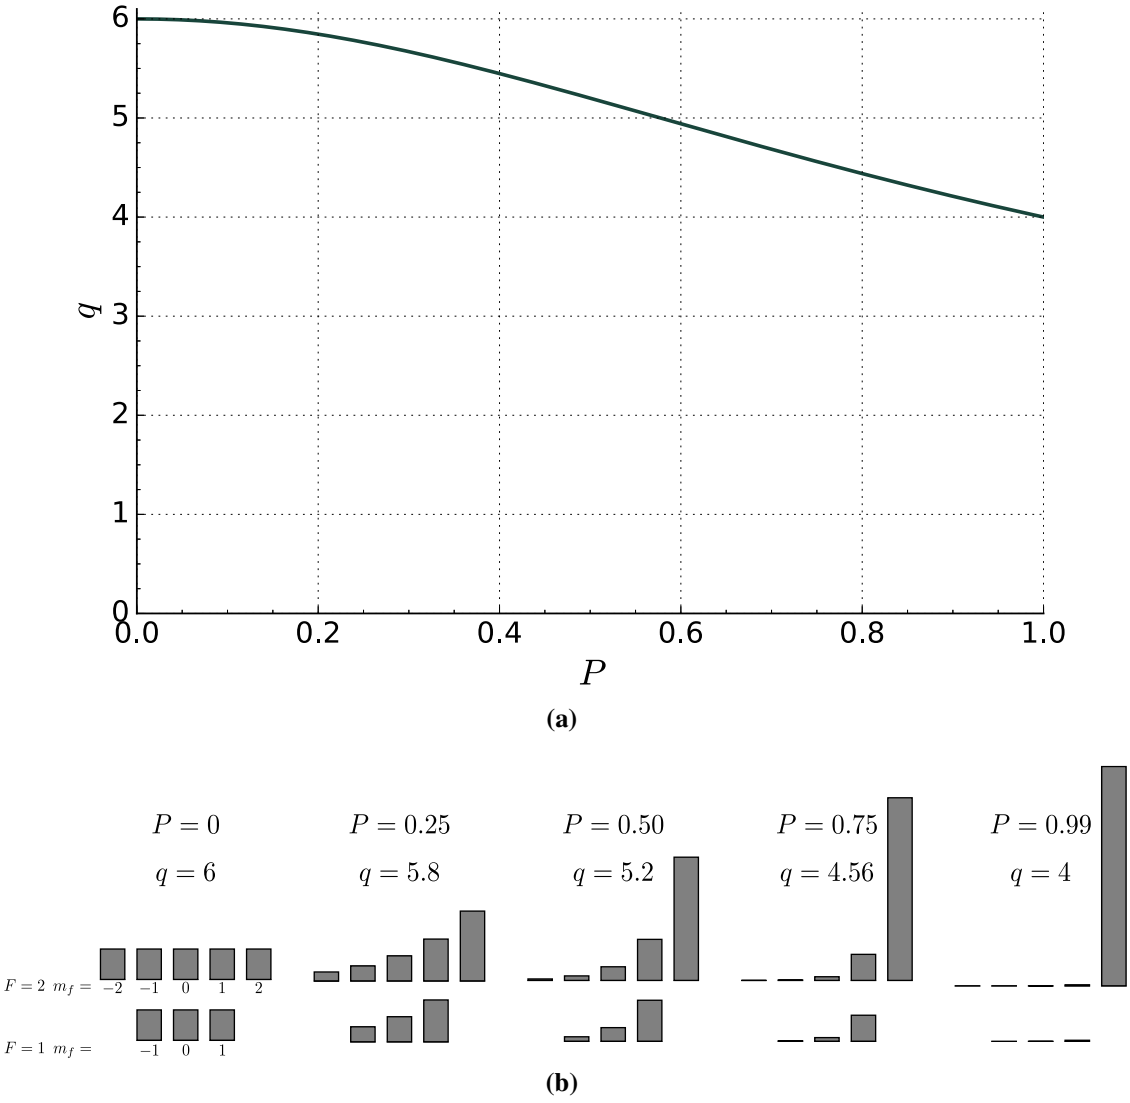



**Figure 2.4:** (a) The slowing down factor $q$ as a function of the atomic polarization $P$ for $^{87}$Rb ($I = 3/2$). (b) Spin-temperature distributions (from [Wyllie, 2012]) at several polarizations, with accompanying values of $q$. Note the equal population of equivalent $m_f$ substates in each hyperfine level. The $P = 50\%$ condition is optimum for SERF magnetometery (see Figure 2.9).

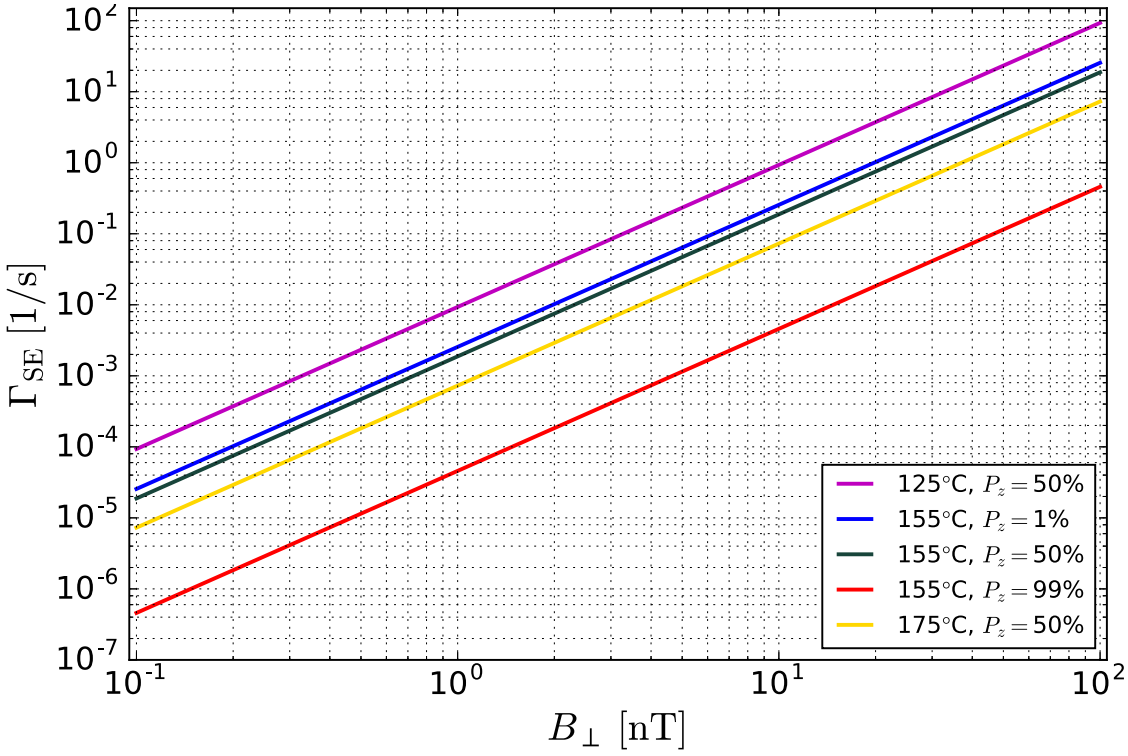



**Figure 2.5:** A plot of the spin-exchange relaxation rate (Equation (2.12)) at several cell temperatures and polarizations. In the SERF regime, the relaxation rate $\Gamma_{\mathrm{SE}}$ due to spin-exchange collisions can be made arbitrarily small by appropriately "nulling" the transverse fields $B_\perp$.

## 2.3 DC-mode SERF magnetometry

The first SERF magnetometers were described by Allred et al. [2002], and the simplest version of our experimental design is based on their work. We use a resonant, circularly-polarized pump laser to optically pump the Rb valence electrons to a state of nonzero spin polarization $P$ along the $\hat{z}$ axis in the absence of external magnetic fields. An external magnetic field along $\hat{y}$ causes the atoms to Larmor precess in the $\hat{x} - \hat{z}$ plane. The resulting spin-polarization in the $\hat{x}$ direction is detected by the Faraday rotation of a linearly polarized probe beam along that axis (perpendicular to the pump beam). A diagram of this arrangement is shown in Figure 2.6. This operating mode, in which we are mostly sensitive to a single component of the magnetic field ($B_y$, is referred to as "DC-mode" in this work. Other measurement schemes (differing from Allred) that allow for the detection of other magnetic field components will be discussed in later sections. In general, the

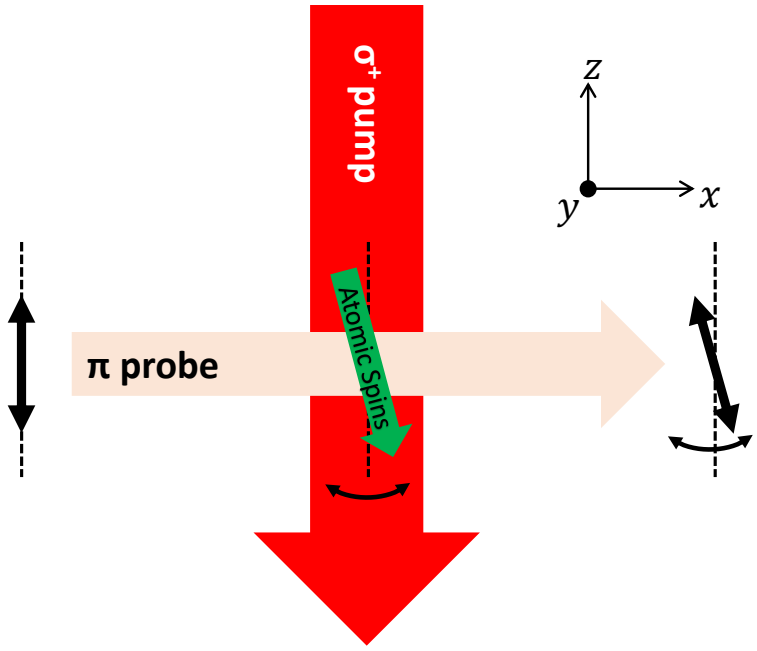



**Figure 2.6:** A diagram of the magnetometer arrangement. Atoms are pumped by $\sigma^+$ light and their spins are aligned along $\hat{z}$. Any atomic precession into the $\hat{x}$ direction (caused by an applied $B_y$, for example) will cause a rotation in the polarization of a linearly polarized probe beam propagating perpendicularly to the pump. This polarization change can be detected at the far side of the atomic ensemble.

time-evolution equations governing the dynamics of a magnetometer can be quite complicated, but under the SERF regime, the description is simplified greatly. When the atoms assume the spin-temperature distribution required for SERF operation, only three effects contribute to the evolution of the atoms [Ledbetter et al., 2008]: optical pumping (from the pump photons, Section 2.3.1), atomic spin relaxation (which occurs via inter-atomic collisions, absorption of probe laser light, and collisions between atoms and the cell walls, Section 2.3.2), and Larmor precession (which occurs at the slowed-down rate characteristic of the spin-temperature regime, Section 2.3.3). These three effects will be covered in depth in the following sections. The dynamics and performance of the magnetometer are determined by interplay of these terms.

### 2.3.1 Optical pumping

A full treatment of the optical pumping process, including the hyperfine levels of the ground and excited states, is complicated [Happer and Van Wijngaarden, 1987]. However, the picture is

simplified considerably if those hyperfine levels are pumped equally by the light. In high-pressure cells with buffer gas pressures of $\sim$.5 amg or more, the pressure-broadening of the resonance lines (17.8 GHz/amg for Rb atoms in $N_2$ [Romalis et al., 1997]) is sufficient to make the two hyperfine levels (separated by 6.835 GHz [Steck, 2001]) unresolved. In our cells, with room-temperature nitrogen pressures of only 175 torr (=.21 amg), the pressure-broadened widths are 3.8 GHz, which is roughly half of the hyperfine splitting. However, rapid spin-exchange collisions bring the system to the spin-temperature regime. This ensures that if optical pumping depletes either of the hyperfine levels, its population will be quickly restored. Thus, we can treat the situation as though the two levels are unresolved to the pump photons.

Our optical pumping scheme is shown in Figure 2.7. Circularly polarized ($\sigma^+$) light near the rubidium D1 ($5^2S_{1/2}$ to $5^2P_{1/2}$, $\lambda \approx 795$ nm) resonance line excites atoms in one of the ground state magnetic sublevels ($m_s = -1/2$). Angular momentum selection rules forbid atoms in the other sublevel ($m_s = +1/2$) from absorbing light of this polarization. Rapid inter-atomic collisions between atoms in the excited state and buffer gas molecules equalize the excited state sublevel populations before the atoms are quenched though non-radiative collisions with $N_2$. This quenching process, while preserving the angular momentum of the excited-state atom, returns the atom to the corresponding ground state. Thus, every pumping cycle, atoms are pumped out of the $m_s = -1/2$ sublevel and roughly half of them are deposited into the $m_s = +1/2$. In other words, each absorbed photon adds, on average, 1/2 unit of its angular momentum to the atoms per cycle. The nuclear spin, which is not randomized by the excited state collisions, tends to become polarized as well via the hyperifne interaction, and the ensemble evolves into a spin-temperature distribution by Equation (2.14):

$$\left.\frac{d\langle \mathbf{F}\rangle}{dt}\right|_{\text{OP}} = \frac{R}{2}(1 - \mathbf{s}\cdot\mathbf{P}). \tag{2.14}$$

Here, $\mathbf{P}$ is the electron spin-polarization along the pumping direction (=$2\langle \mathbf{S}\rangle$), $\mathbf{s}$ is the angular momentum carried by the pump photon ($|\mathbf{s}| = 1$ for perfectly circularly polarized light), and $R$ is the optical pumping rate, the rate at which nominally *un*polarized atoms absorb the circularly-polarized pump light. $R$ is a convolution of the photon flux $\Phi_0(\nu)$ and the unpolarized absorption

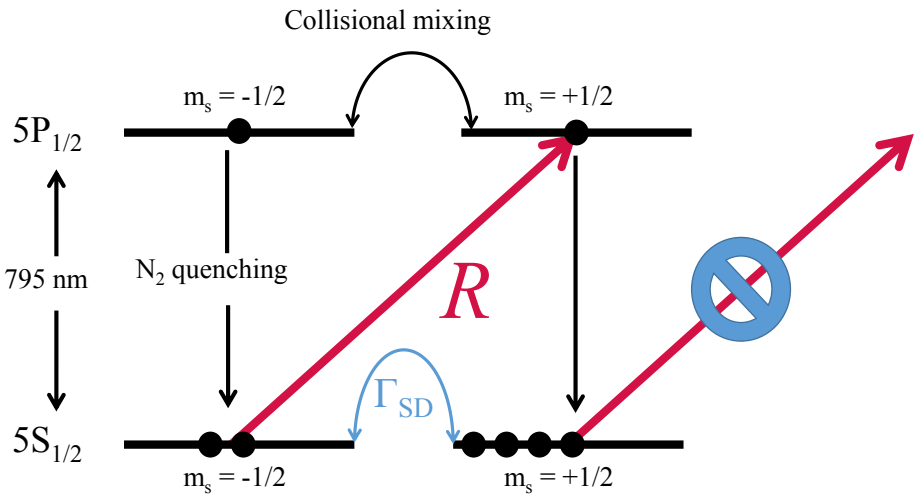



**Figure 2.7:** Optical pumping scheme for our experiment. An unpolarized ensemble of atoms initially contains roughly equal populations in the ground state ($5S_{1/2}$) magnetic sublevels. Optical pumping with resonant circularly-polarized D1 (795 nm) light drives atoms in $m_s = -1/2$ sublevel to the excited state ($5P_{1/2}$), $m_s = +1/2$ sublevel. An angular momentum selection rule prevents atoms in the $m = +1/2$ ground state level from absorbing the $\sigma^+$ photons (with $s = 1$). Rapid collisional mixing in the excited state coupled with non-radiative quenching collisions with nitrogen means that one half of the atoms pumped out of the $m_s = -1/2$ state will eventually end up in each ground sublevel. Thus, on average, one photon absorption adds 1/2 unit of angular momentum to the vapor. Spin-destruction collisions (Section 2.3.2) tend to return atoms to the $m_s = -1/2$ sublevel. The balance of the optical pumping rate $R$ and the spin-relaxation rate $\Gamma_{sd}$ determines the population in each sublevel.

cross section $\sigma_0(\nu)$:

$$R = \int_{-\infty}^{\infty} \Phi(\nu' - \nu)\sigma_0(\nu' - \nu_0)d\nu' \tag{2.15}$$

$$R \approx \Phi_0\sigma_0(\Delta\nu) \tag{2.16}$$

$$\Phi_0 = \frac{\mathcal{I}}{h\nu} \tag{2.17}$$

$$\sigma_0(\Delta\nu) = r_e f_{\mathrm{D1}} c\mathcal{L}(\Delta\nu) \tag{2.18}$$

$$\mathcal{L}(\Delta\nu) = \frac{\Delta\Gamma/2}{(\Delta\nu)^2 + (\Delta\Gamma/2)^2}. \tag{2.19}$$

Here, $\nu$ is the center frequency of the laser and $\nu_0$ is the center frequency of the D1 absorption [Steck, 2001]. $r_e$ is the classical electron radius, $f_{\mathrm{D1}}$ is the oscillator strength of the D1 transition, $\mathcal{I}$ is the pump laser intensity, and $\Delta\nu = \nu - \nu_0$ is the pump laser detuning from the D1 line. Equation (2.16) is true as long as the laser linewidth is much smaller than the atomic absorption linewidth, so $\Phi(\nu' - \nu) \approx \Phi_0\delta(\nu' - \nu)$, with $\Phi_0$ defined in Equation (2.17). In this experiment, the pump laser is a single-mode distributed feedback diode laser with a typical FWHM spectral width of 2 MHz and the pressure-broadened absorption widths are on the order of a few GHz.

### 2.3.2 Spin-relaxation mechanisms

There are three primary spin-relaxation mechanisms of a SERF magnetometer: spin-destruction collisions that occur when an alkali atom collides with another alkali atom or nitrogen molecule, relaxation due to the absorption of probe photons, and spin destruction collisions that occur at the cell walls. The first two mechanisms are quite straightforward (a spin-destruction collision or photon absorption causes an immediate loss of spin coherence), while the third one involves the diffusion of atoms through the buffer gas in the cell. Since collisions with the cell walls immediately depolarize the atoms, it would seem that using high buffer gas pressures would be advantageous, and in fact this strategy has been employed with the several-amagat pressure cells used in [Allred et al., 2002]. An extension of SERF magnetometry which takes advantage of this diffusion proved to be a vital part of this experiment and will be discussed in Section 2.4. Thus, our magnetometer cells operate with more modest buffer gas pressures.

#### 2.3.2.1 Atom-atom spin-destruction collisions

The spin destruction rates between the atoms present in our magnetometers are (from [Chen et al., 2007])

$$\Gamma_{\text{Rb-Rb}} = \left(4.2 \times 10^{-13}\,\frac{\text{cm}^3}{\text{s}}\right) n_{\text{Rb}} \tag{2.20}$$

$$\Gamma_{\text{Rb-N}_2} = \left(1.3 \times 10^{-25}\,\frac{\text{cm}^3}{\text{s K}^3}\right) n_{\text{N}_2}\, T^3 \tag{2.21}$$

where $n_{\text{Rb}}$ and $n_{\text{N}_2}$ are the number densities of $^{87}$Rb and $N_2$, respectively. Though only explicit in Equation (2.21), the Rb-Rb rate is also temperature-dependent through the rubidium vapor pressure (Equation (2.13)). The contribution of these collisions to the atoms' total spin can be written as

$$\begin{aligned}\left.\frac{d\langle\mathbf{F}\rangle}{dt}\right|_{\text{sd}} &= -\left(\Gamma_{\text{Rb-Rb}} + \Gamma_{\text{Rb-N}_2}\right)\mathbf{S} \\ &\equiv -\Gamma_{\text{sd}}\mathbf{S}.\end{aligned} \tag{2.22}$$

#### 2.3.2.2 Probe absorption

Absorption of probe photons (incident perpendicular to the optical pumping along the $\hat{x}$ axis) on the Rb D2 transition ($\lambda \approx 780$ nm) acts as a second simple relaxation mechanism. The relaxation rate $\Gamma_{\text{pr}}$ is simply the absorption rate $R_{\text{pr}}$ of the probe photons, similar to the rate defined in Equation (2.15). We can make the same narrow-linewidth laser approximation as we did in Equation (2.16) and write

$$\Gamma_{\text{pr}} = R_{\text{pr}} \approx \Phi_{\text{pr}}\sigma_{\text{pr}}(\Delta\nu), \tag{2.23}$$

where the probe photon flux $\Phi_{\text{pr}}$, the cross section $\sigma_{\text{pr}}$ and detuning $\Delta\nu$ are all defined as they were in Equations (2.15) through (2.19), using the appropriate values for the D2 transition line. Much like Equation (2.22), we write

$$\left.\frac{d\langle\mathbf{F}\rangle}{dt}\right|_{\text{pr}} = -\Gamma_{\text{pr}}\mathbf{S}. \tag{2.24}$$

Unlike the pumping rate, which in general must be quite high to create the electron spin-polarization necessary for magnetometry, $\Gamma_{\text{pr}}$ can be made arbitrarily small by detuning the probe laser far away from the transition or reducing its intensity. This, however, will cause a reduction in magnetometer

signal size and an increase in the photon shot noise. These effects will be explained in more detail in Sections 2.3.7 and 4.8.3.2.

#### 2.3.2.3 Diffusion

Diffusion plays a critical role in the time evolution of the magnetometer, as the movement of rubidium atoms in the cell carries angular momentum obtained through optical pumping or spin-exchange collisions to other parts of the cell. The time-evolution of the angular momentum under diffusion is [Ghosh, 2009]

$$\left.\frac{d\left\langle \mathbf{F}\right\rangle }{dt}\right|_{\text{diff}}=D\nabla^{2}\left\langle \mathbf{F}\right\rangle , \tag{2.25}$$

where the diffusion constant $D$ is defined as

$$D=\frac{1}{3}\lambda\bar{v}, \tag{2.26}$$

and $\lambda$ is the mean free path and $\bar{v}$ is the thermal velocity of the alkali atoms. These two variables are dependent on the specifics of the cell (alkali atom, buffer gas composition, pressure, and temperature). We use a measured $\text{Rb}-\text{N}_2$ diffusion constant $D_0$ from [Ishikawa and Yabuzaki, 2000], and modify it appropriately for our cell conditions:

$$D(P,T)=D_{0}\left(\frac{p_{0}}{p}\right)\left(\frac{T}{T_{0}}\right)^{3/2} \tag{2.27}$$

$$D_{0}(\text{Rb-N}_{2})=.159\text{ cm}^{2}\text{ s}^{-1}\qquad\qquad T_{0}=60^{\circ}\text{C},\ p_{0}=760\text{ torr} \tag{2.28}$$

The used hypothesis is that collisions between a polarized alkali atom and the cell wall will completely destroy the angular momentum carried by that atom (both electronic and nuclear) [Wagshul and Chupp, 1994, Walker and Happer, 1997]. Thus we enforce the Dirichlet boundary condition at the walls ($\left\langle \mathbf{F}\right\rangle |_{\text{walls}}=0$) when solving Equation (2.25).

A thorough treatment of diffusion is presented by Happer et al. [2010], but in this work we will focus on two primary effects:

1. Diffusion increases the relaxation rate in the cell by allowing the atoms to depolarize at the cell walls. In turn, this requires a higher pumping rate to maintain a steady-state atomic polarization.

2. Diffusion alters the spatial dependence of cell's polarization. This effect will be critical when we examine an extension to SERF magnetometery ("diff-SERF" operation) in Section 2.4.

For now, let us examine the first effect, the extra relaxation caused by atoms diffusing to the walls. Following Wagshul and Chupp [1994] and Ghosh [2009], we write the additional relaxation caused by diffusion to the cell walls as

$$\left.\frac{d\langle\mathbf{F}\rangle}{dt}\right|_{\text{diff}} = -\Gamma_{\text{diff}}\langle\mathbf{F}\rangle \tag{2.29}$$

$$= -q\Gamma_{\text{diff}}\langle\mathbf{S}\rangle\,, \tag{2.30}$$

where $q$ is the slowing-down factor and

$$\Gamma_{\text{diff}} \approx D(P,T)\times 3\left(\frac{\pi}{l}\right)^2. \tag{2.31}$$

Equation (2.31) uses the normal mode expansion to the diffusion equation in a cube of side length $l$ to calculate the final term. The inclusion of the slowing down factor $q$ in Equation (2.30) emphasizes an important point: unlike the spin relaxation due to probe photon absorption and atom-atom spin-destruction collisions, the collision of an atom with the cell wall will destroy *both* electronic and nuclear spins.

#### 2.3.2.4 Spin-relaxation summary

Combining the three terms from the preceding sections, we can write the total spin-destruction rate, and its contribution to the evolution of $\langle\mathbf{F}\rangle$ as

$$\Gamma_{\text{total}} \equiv \Gamma = \Gamma_{\text{sd}} + \Gamma_{\text{pr}} + q\Gamma_{\text{diff}} \tag{2.32}$$

$$\frac{d\langle\mathbf{F}\rangle}{dt} = -\Gamma\langle\mathbf{S}\rangle \tag{2.33}$$

Figures 2.8a and 2.8b show how these relaxation rates depend on nitrogen pressure and temperature, respectively. To increase optical pumping efficiency and magnetometer performance, we try to work at an optimum point which minimizes the total spin-relaxation rate. The cells used in our experiment actually have somewhat lower nitrogen pressure than the optimal pressure that Figure

2.8a would suggest; however this lower pressure allows us to work in diff-SERF mode (Section 2.4). Diffusion is impeded and diff-SERF performance is degraded if the buffer gas pressure is too high.

It's important to note that the optical pumping rate $R$ also contributes significantly to the alkali relaxation rate. We will fold its contribution into an effective relaxation rate $\Gamma'$ in Section 2.3.4.

### 2.3.3 Larmor precession

In between spin relaxations and photon absorptions, the atoms will undergo Larmor precession in the presence of a magnetic field $\mathbf{\Omega} = (\gamma/q)\mathbf{B}$,

$$\left.\frac{d\left\langle \mathbf{F}\right\rangle }{dt}\right|_{\text{Larmor}} = \mathbf{\Omega} \times \left\langle \mathbf{S}\right\rangle . \tag{2.34}$$

### 2.3.4 Equations of motion and steady-state magnetometer response

Combining the terms of the previous sections and assuming that the spin-exchange rate is sufficiently high to create a spin-temperature distribution of the atoms, we can write the Bloch equation for the time evolution of the angular momentum as in [Ledbetter et al., 2008, Savukov and Seltzer, 2013]

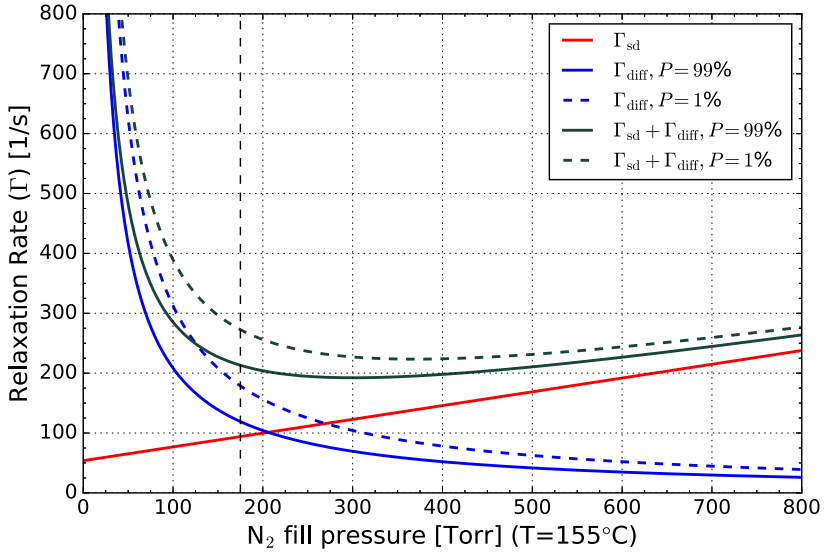
$$q(P)\frac{d\mathbf{P}}{dt} = q(P)D\nabla^2\mathbf{P} + \mathbf{\Omega} \times \mathbf{P} + R\mathbf{s} - R\mathbf{P} - (\Gamma_{\text{sd}} + \Gamma_{\text{pr}})\mathbf{P}, \tag{2.35}$$

where $P = |\mathbf{P}|$ and $\mathbf{s}$ is the angular momentum of the pump photons. Additionally, we have used the relationships $\left\langle \mathbf{F}\right\rangle = q(P)\left\langle \mathbf{S}\right\rangle$ and $\mathbf{P} = 2\left\langle \mathbf{S}\right\rangle$. If we assume that the diffusion term contributes to the relaxation as in Equation (2.29) and assume the pump light is perfectly circularly polarized along the $\hat{z}$ direction (*i.e.*, $R\mathbf{s} = R\hat{z}$), then we can write the steady state solution of Equation (2.35) as

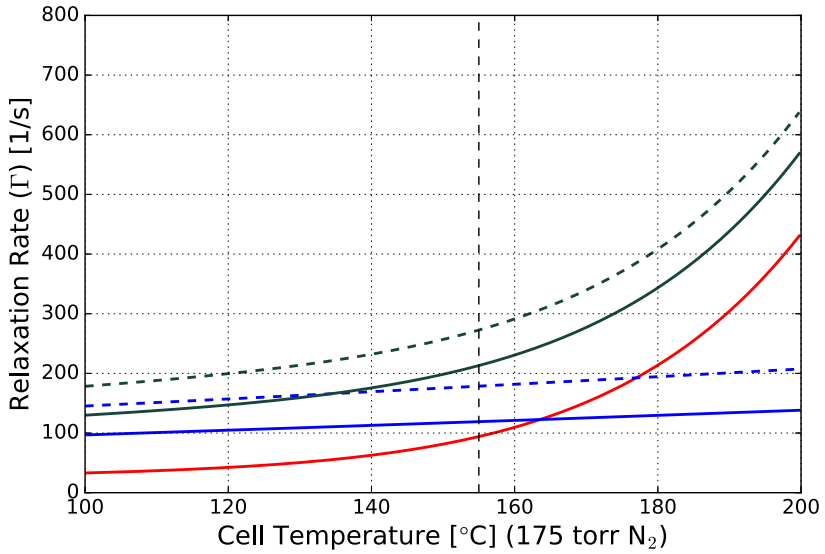
$$\mathbf{P} = \frac{R}{\Gamma'\left(\Gamma'^2 + \Omega^2\right)} \begin{pmatrix} \Gamma'\Omega_y + \Omega_x\Omega_z \\ -\Gamma'\Omega_x + \Omega_y\Omega_z \\ \Gamma'^2 + \Omega_z^2 \end{pmatrix}, \tag{2.36}$$

where $\Gamma'$ is the *total effective relaxation rate*

$$\Gamma' = R + \Gamma = R + \Gamma_{\text{sd}} + \Gamma_{\text{pr}} + \Gamma_{\text{diff}}, \tag{2.37}$$




**(a)**




**(b)**

**Figure 2.8:** The relaxation rates $\Gamma_{sd}$ and $\Gamma_{diff}$ as functions of $N_2$ buffer gas pressure (a) and cell temperature (b). Each shows the dependence of the spin-polarization $P$, which influences $\Gamma_{diff}$ through $q$ (refer to Equation (2.32)). At low buffer gas pressures, collisions with the walls dominate the total relaxation rate. As buffer gas pressure increases, the spin-destruction collisions between the Rb and $N_2$ molecules become the dominant relaxation mechanism. Increasing temperature most strongly influences the spin-destruction collision rate, as the rate depends on Rb density and the Rb$-N_2$ spin-destruction rate, both of which grow exponentially with temperature (Equations 2.13 and 2.21). Not included in the total relaxation rate is the contribution of $\Gamma_{pr}$, which can be made arbitrarily large or small by varying probe laser parameters. The vertical dashed lines indicate the temperature and buffer gas pressure values in the other figure: the conditions at which we commonly operate in our experiment. As noted in Section 2.3.2.4, the *actual* relaxation rates in optically-pumped cells can be much higher due to the relaxation caused by the pumping rate $R$.

and $\Omega = |\mathbf{\Omega}| = \sqrt{\Omega_x^2 + \Omega_y^2 + \Omega_z^2}$. For our experimental arrangement (Figure 2.6), the $\hat{x}$ propagation direction of the probe beam means we are sensitive to $P_x$ (top line of Equation (2.36)). We can use the expression for $P_z$ (the bottom line of Equation (2.36)) to rewrite $P_x$ as

$$P_x = P_z \left( \frac{\Gamma' \Omega_y + \Omega_x \Omega_z}{\Gamma'^2 + \Omega_z^2} \right) \tag{2.38}$$

In the limit where $\Omega \ll \Gamma'$

$$P_z \to \frac{R}{\Gamma'} \tag{2.39}$$

$$P_x \to P_z \frac{\Omega_y}{\Gamma'} \tag{2.40}$$

The magnetic field regime where this is true depends on the pumping and relaxation rates of the magnetometer. If we assume a typical $\Gamma = 500$ s$^{-1}$ and let $R \approx \Gamma$ (which we can experimentally accomplish by varying pump laser power and/or detuning), then this is true for magnetic fields $B$ on the order of 1 nT or smaller.

Using Equation (2.36), we see that if we make $\Omega_x$ and $\Omega_z$ small, then the magnetometer is primarily sensitive to fields in the $\hat{y}$ direction. A sweep of the field $\Omega_y$ gives a dispersive lineshape

$$P_x = \frac{R\Omega_y}{\Gamma'^2 + \Omega_y^2}. \tag{2.41}$$

Transverse fields $\Omega_x$ and $\Omega_z$ shift the zero-crossing of this dispersion and reduce the slope $dP_x/d\Omega_y$ of the dispersive curve, as seen in Figure 2.10. The shape of the dispersion gives some insight into the sizes of the bias fields present in the magnetometer. In theory, if one wished to null the transverse magnetic fields to maximize the slope of the dispersion, he or she could continuously sweep $\Omega_y$ while adjusting $\Omega_x$ and $\Omega_z$ until a dispersive curve matching Equation (2.41) was found. This is time-consuming in practice, though, and an alternate method is employed to null out those fields (Section 4.3).

It is often useful to rewrite the steady-state equations 2.36 to simplify notation and to understand the way in which varying controllable parameters (notably, $R$ and $\Omega_i$) change the behavior of the magnetometer. One normalization is provided by Wyllie [2012], in which the controllable parameters are normalized against the total spin-relaxation rate $\Gamma$. Once $\Gamma_{\text{pr}}$ is made to be small

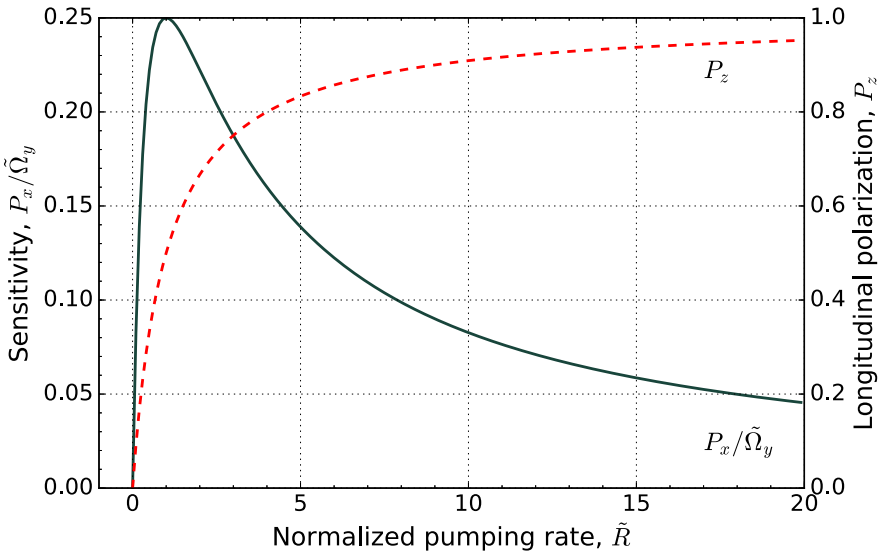



**Figure 2.9:** Normalized response (left axis, solid red line) and longitudinal polarization (right axis, dashed green line) as a function of the normalized pumping rate $\tilde{R} = R/\Gamma$. Note the response is maximized when $\tilde{R} = 1$ and the polarization is $P_z = 50\%$.

(either by reducing probe power or increasing probe detuning), $\Gamma$ is primarily dependent on the cell properties (cell geometry, cell temperature, buffer gas pressure) with only a weak dependence on the cell polarization through $q(P)$. If we introduce the variables $\tilde{R} \equiv R/\Gamma$ and $\tilde{\Omega} \equiv \Omega/\Gamma$, then we can rewrite the $P_x$ and $P_z$ components of Equation (2.36)

$$P_z = \frac{\tilde{R}}{\tilde{R}+1}\left(\frac{\left(\tilde{R}+1\right)^2 + \tilde{\Omega}_z^2}{\left(\tilde{R}+1\right)^2 + \tilde{\Omega}^2}\right) \tag{2.42}$$

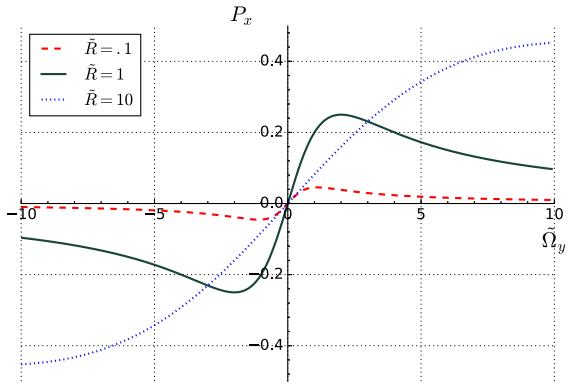
$$P_x = \frac{\tilde{R}}{\tilde{R}+1}\left(\frac{\left(\tilde{R}+1\right)\tilde{\Omega}_y + \tilde{\Omega}_x\tilde{\Omega}_z}{\left(\tilde{R}+1\right)^2 + \tilde{\Omega}^2}\right), \tag{2.43}$$

and the relationship

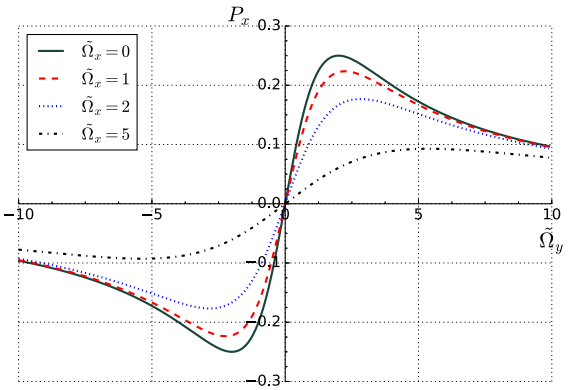
$$P_x = P_z\left(\frac{\left(\tilde{R}+1\right)\tilde{\Omega}_y + \tilde{\Omega}_x\tilde{\Omega}_z}{\left(\tilde{R}+1\right)^2 + \tilde{\Omega}_z^2}\right) \tag{2.44}$$

is similar to that derived in Equation (2.38). Plotting Equation (2.43) normalized to the detected magnetic field $\tilde{\Omega}_y$ at small $\tilde{\Omega}_x$ and $\tilde{\Omega}_z$ (Figure 2.9), indicates that maximum sensitivity is realized when $\tilde{R} = 1$ (*i.e.*, $R = \Gamma'$), corresponding to a longitudinal low-field polarization of $P_z = 50\%$. A naïve guess might have suggested that a polarization closer to 100% (like what is used for spin-exchange optical pumping [Walker and Happer, 1997]) would be ideal, as that distribution contains more atoms coherently precessing in the same direction. However, "overpumping" the atoms with $R \gg \Gamma$ tends to simply pin the atoms' spins along $\hat{z}$ and reduces the sensitivity. It should be stressed, though, that observing the behavior of $P_x/\tilde{\Omega}_y$ near $\tilde{R} = 1$ shows the function falls off faster for $\tilde{R} < 1$ than $\tilde{R} > 1$. In other words, pumping too hard is normally preferable to not pumping hard enough. In practice, the proper pumping rate is determined by tuning laser parameters to maximize the response $P_x/\tilde{\Omega}_y$ to a slowly varying applied $\hat{y}$ field.

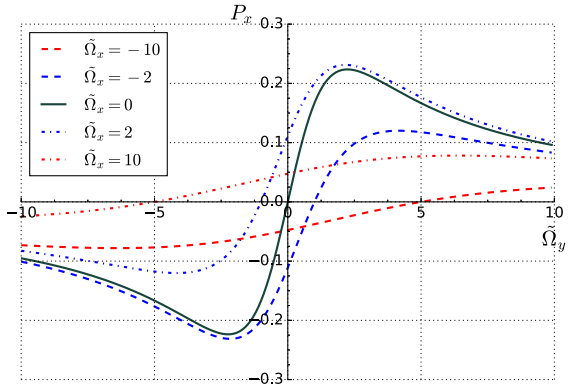
It is important to note here that because of the spatial dependence of the pump beam (longitudinally due to the high optical thickness of the vapor and transversely due to the beam's Gaussian profiile), it is impossible to have a uniform pumping rate $\tilde{R}$ at all points in the cell. Therefore, it is similarly impossible to maintain the ideal $P_z = 1/2$ condition across the entire cell volume. Rather, what we hope to achieve is an *averaged* value of $\tilde{R} \approx 1$ (and $P_z = 1/2$) across the cell,



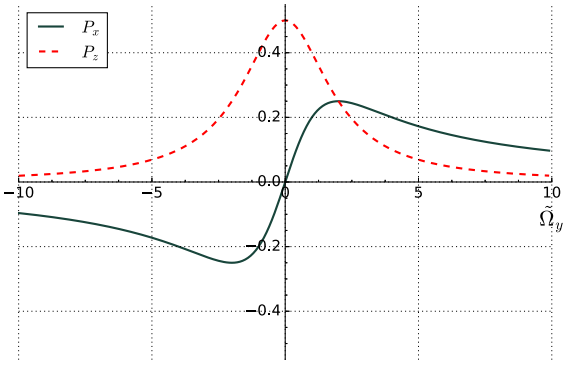

**(a)** Swept $\tilde{\Omega}_y$ at various $\tilde{R}$ with $\tilde{\Omega}_x = \tilde{\Omega}_y = 0$




**(b)** Swept $\tilde{\Omega}_y$, $\tilde{R} = 1$, at various $\tilde{\Omega}_x$ with $\tilde{\Omega}_z = 0$




**(c)** Swept $\tilde{\Omega}_y$, $\tilde{R} = 1$, at various $\tilde{\Omega}_x$ with $\tilde{\Omega}_z = 1$




**(d)** Swept $\tilde{\Omega}_y$ with $\tilde{R} = 1$ and $\tilde{\Omega}_x = \tilde{\Omega}_y = 0$

**Figure 2.10:** Adiabatically swept fields $\tilde{\Omega}_y$ with various pumping rates and transverse bias fields. The SERF magnetometer is most sensitive (has the largest value $dP_x/d\tilde{\Omega}_y$) at $\tilde{R} = 1$ with the transverse fields nulled (a). Introducing one transverse field will reduce the sensitivity of the device (b), while having two transverse fields will not only reduce sensitivity, but also induce asymmetries and alter the zero-crossings of the dispersion (c). These shapes can thus help one correctly zero out the fields on the magnetometer to obtain the curves shown in (a). In (d), the longitudinal spin-polarization $P_z$ is shown to be maximum (and optimal) at $\tilde{\Omega}_y = 0$ and $\tilde{R} = 1$.

or at least in the volume being interrogated by the transverse probe beam. In Section 2.4 we will discuss the strides we have made to reduce the longitudinal dependence of the pump beam.

### 2.3.5 Dynamic field response

The biomagnetic signals we are attempting to detect with our SERF magnetometers are not confined to the slow field sweeps that form the curves in Figure 2.10. It is equally important to understand how the magnetometer responds to fields at all frequencies, which will become critical during the magnetometer calibration (Section 4.4.3). Consider a field $\mathbf{\Omega}' = \mathbf{\Omega}_0 + \mathbf{\Omega}_1 e^{i\omega t}$, in which the DC component of this field is much larger than the AC component (*i.e.* $|\mathbf{\Omega}_0| \gg |\mathbf{\Omega}_1|$). We can similarly express the spin-polarization $\mathbf{P}'(t) = \mathbf{P}_0 + \mathbf{P}_1(t)$. Plugging these expressions into Equation (2.35) and carefully grouping terms yields

$$q(P)\frac{d\mathbf{P}_1(t)}{dt} = \mathbf{\Omega}_0 \times \mathbf{P}_1(t) + e^{i\omega t}\mathbf{\Omega}_1 \times \mathbf{P}_0 + e^{i\omega t}\mathbf{\Omega}_1 \times \mathbf{P}_1 - \Gamma'\mathbf{P}_1(t) \\ + \mathbf{\Omega}_0 \times \mathbf{P}_0 + R\mathbf{s} - \Gamma'\mathbf{P}_0 - q(P)\frac{d\mathbf{P}_0}{dt}. \quad (2.45)$$

Note that the second line of (Equation (2.45)) is the steady-state equation given in Equation (2.35), which is unaffected by the AC magnetic field as long as $|\mathbf{\Omega}_0| \gg |\mathbf{\Omega}_1|$. Thus we will focus only on the time-varying part of the solution, the first line of Equation (2.45). If we let $\mathbf{P}_1(t) = \mathbf{P}_1 e^{i\omega t}$, we can write the linear response to signal at frequency $\omega$ as

$$-iq(P)\omega\mathbf{P}_1 = \mathbf{\Omega}_1 \times \mathbf{P}_0 + \mathbf{\Omega}_0 \times \mathbf{P}_1 - \Gamma' P_1. \quad (2.46)$$

Note that we have discarded the third term on the right-hand side of Equation (2.45), which the product of two small quantities $\mathbf{\Omega}_1$ and $\mathbf{P}_1$. It is possible to solve Equation (2.46) exactly by transforming into a rotating basis of $P_\pm = P_x \pm iP_y$. We will employ such a strategy when we discuss actively modulating the atomic polarization with a driving field (Section 2.5). For now, however, Wyllie [2012] notes that the solutions are simplified greatly when the DC part of the magnetic field only exists along $\hat{z}$ (*i.e.* $\mathbf{\Omega} = \Omega_{0z}\hat{z}$) and the AC part only contains terms along $\hat{x}$

and $\hat{y}$ ($\mathbf{\Omega}_1 = \Omega_{1x}\hat{x} + \Omega_{1y}\hat{y}$ ). If we let $\omega' = q(P)\omega$, then

$$P_{1x} = P_{0z}\frac{(\Gamma' - i\omega')\Omega_{1y} + \Omega_{0z}\Omega_{1x}}{(\Gamma' - i\omega')^2 + \Omega_{0z}^2} \tag{2.47}$$

$$|P_{1x}| = P_{0z}\sqrt{\frac{(\Gamma'^2 + \omega'^2)\Omega_{1y}^2 + \Omega_{0z}^2\Omega_{1x}^2}{(\Gamma'^2 + \Omega_{0z}^2)^2 + \omega'^2(\omega'^2 - (\Omega_{0z}^2 - \Gamma'^2))}}. \tag{2.48}$$

This response decreases at low frequency with increasing $\Omega_0 z$, as shown in Figure 2.11. When the fields are well-nulled ($\mathbf{\Omega}_0 = 0$, the regime which maximizes the value of $dP_x/d\tilde{\Omega}_y$ as seen in the previous section), Equation (2.48) becomes

$$|P_{1x}| = P_{0z}\sqrt{\frac{\Omega_{1y}^2}{\Gamma'^2 + \omega'^2}}, \tag{2.49}$$

which gives the response a characteristic square-root of a Lorentzian lineshape centered at DC. It is also only sensitive to fields in the $\hat{y}$ direction. The bandwidth of the magnetometer is set by $\Gamma'/q(P)$, though there are some subtleties to that claim that Wyllie expounds on in his thesis [2012].

It is also worth noting the contribution to $q(P)$ in the bandwidth is one of the pertinent reasons for choosing isotopically-pure $^{87}$Rb alkali gas over a natural-abundance mixture of isotopes (72% $^{85}$Rb and 28% $^{87}$Rb) for most of our magnetometers. The $^{85}$Rb nuclear spin of $I = 5/2$ means the slowing-factor for natural-abundance rubidium varies from 10.8 to 5.44 as $P$ increases from 0 to 1 (compared to the $6 \rightarrow 4$ range for $^{87}$Rb, Figure 2.4a). The increased bandwidth for $^{87}$Rb aids detection of the magnetocardiographic signals at higher frequencies. A discussion of the frequency content of these signals appears in Chapter 6

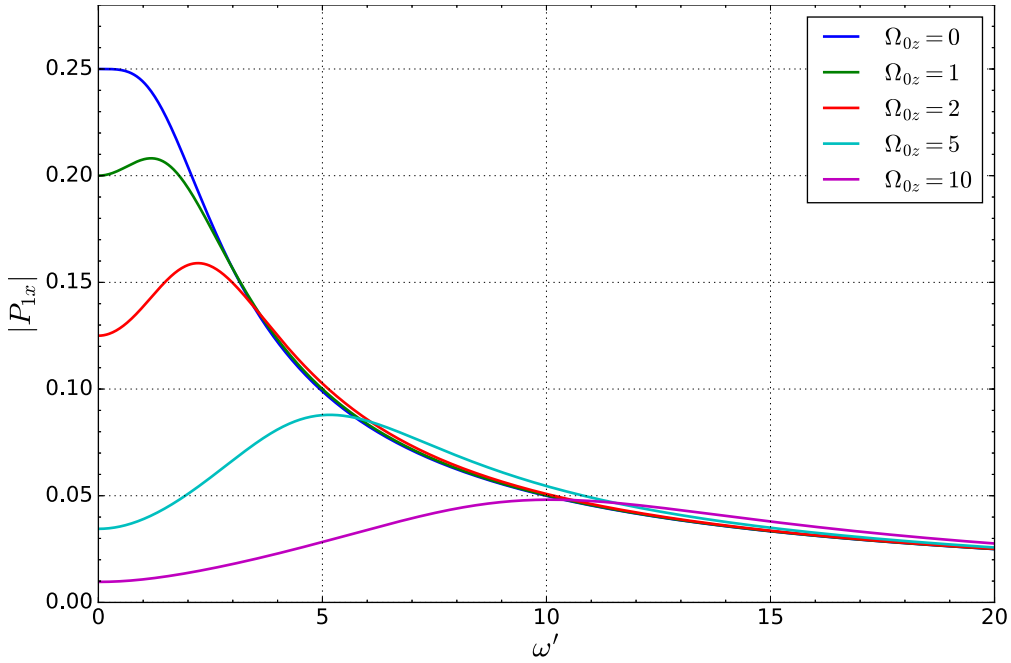



**Figure 2.11:** The frequency response of the magnetometer (Equation (2.48)) with various $\hat{z}$ field imbalances at $R = \Gamma$ ($\tilde{R} = 1$) and $\Omega_{1y} = 1$. Note the decreasing sensitivity, non-lorentzian lineshapes, and broadening as $\Omega_{0z}$ increases. In practice, the frequency response must be measured every time the magnetometer is calibrated, so the shape of these response curves provide convenient clues as to potential field imbalances. Imbalences can be corrected by re-shimming the local magnetic field until a response resembling the $\Omega_{0z} = 0$ curve is obtained.

### 2.3.6 Light shift

From Figures 2.10 and 2.11 (and more generally, the requirements of the SERF regime), it is apparent that the elimination of transverse magnetic fields $\Omega_x$ and $\Omega_z$ is particularly important for the optimization of a SERF magnetometer. Most of the work of eliminating the earth-scale ($\sim 50\,\mu$T) magnetic fields is accomplished by working in a magnetically-shielded room, which reduces $\mathbf{B}_{\text{earth}}$ to the $\sim 100$ nT level. The remainder of this field can be removed through the use of magnetic-field canceling coils which are wrapped around the perimeter of the interior of the room. The use of magnetic materials in construction of the magneteomters must be kept to an absolute minimum to avoid generating additional fields.

However, due to the atomic polarizability of the alkali atoms [Bonin and Kresin, 1997], an effective magnetic field is created when these atoms are optically pumped. This effect (known as the "light shift" or "AC Stark shift") is well-understood [Appelt et al., 1999, Bonin and Kresin, 1997, Shah and Romalis, 2009, Wyllie, 2012], but is an obstacle that must be accounted for and corrected for in SERF magnetometry.

The effective field created by a circularly-polarized, $\hat{z}$-propagating pump beam near the D1 resonance line is [Shah and Romalis, 2009, Wyllie, 2012]

$$\mathbf{\Omega}_{\text{LS}} = r_e f_{\text{D1}} c \Phi_0 \mathcal{D}(\Delta\nu)\hat{z} \tag{2.50}$$

$$= R\frac{\Delta\nu}{\Delta\Gamma/2}\hat{z}, \tag{2.51}$$

where

$$\mathcal{D}(\Delta\nu) = \frac{\Delta\nu}{(\Delta\nu)^2 + (\Delta\Gamma/2)^2} \tag{2.52}$$

$$= \frac{\Delta\nu}{\Delta\Gamma/2}\mathcal{L}(\Delta\nu) \tag{2.53}$$

is the dispersive lineshape. $\mathcal{L}(\Delta\nu)$ was defined in Equation (2.19). Also recall from Section 2.3.1 that $r_e$ is the classical electron radius, $f_{D1}$ is the frequency of the Rb D1 transition on which the atoms are optically pumped, $\Phi_0$ is the incident photon flux at the laser's center frequency. The narrowband laser approximation allows for the magnitude of the light shift to be related directly to

the pumping rate $R$ (Equation (2.51)). To get an idea of the actual size of the light shift fields, it is useful to plot Equation (2.50) in real magnetic units, but normalize it to a pump power of 1 mW and a pump beam waist of 1 cm (recall $\Phi_0 = \mathcal{I}/h\nu = \mathcal{P}_{\text{pu}}/w_0^2 h\nu$). While these are not normally the conditions we use in our experiment, it does allows us to quickly estimate the value of the light shift for our particular conditions by multiplying or dividing by the appropriate factor. Thus,

$$\frac{B_{\text{LS}}}{\mathcal{P}/w_0^2} = \frac{4\gamma r_e f c}{\pi h\nu} \frac{1}{\Delta\Gamma/2} \frac{x}{1+x^2}, \tag{2.54}$$

where $x = \Delta\nu/(\Delta\Gamma/2)$ is a normalized detuning from resonance. A plot of this function is shown in Figure 2.12. Using reasonable estimates of pump parameters for our experimental conditions of $\mathcal{P}_{\text{pu}} = 5$ mW and $w_0 = .2$ cm, light shifts of $B_{\text{LS}} \geq 500$ nT are easily attainable. Since these light shifts are just DC fields which vary depending on pump parameters, they can theoretically be canceled by simply using a set of magnetic field coils to shim away the induced light shift field. However, in practice, there are several issues to consider:

1. Because the pump light is absorbed as it passes through the cell, the steady-state intensity (and hence, the pumping rate $R$ and light shift $\mathbf{\Omega}_{\text{LS}}$) will vary as a function of position along the $\hat{z}$-axis of the cell, as well as in the transverse directions, as the beam has a gaussian intensity profile. Thus, it is impossible to cancel the light shift at all locations in the cell; the best that is normally possible is to cancel the *average* light shift across the cell.

2. The problem is magnified with multiple magnetometer channels. By careful arrangement of the cells and adjustment of each pump beam's circularity, it is possible to make all of the light shifts in the same *direction*. However, because power and spatial profile variations exist between each cell's pump laser beam, equalizing the *magnitude* of the light shift across all cells is extraordinarily difficult. For this reason, the use of a single set of "common" magnetic field coils is usually not desirable, as light shift-enhanced field gradients will make it difficult to bring all cells into an optimum SERF magnetometry condition ($|\mathbf{\Omega}| = 0$) simultaneously. We have countered this issue by using a set of tri-axial field-canceling coils for each channel in addition to the large set of coils wrapped around the perimeter of the magnetically shielded room. These coils will be discussed in more detail in Chapter 3.

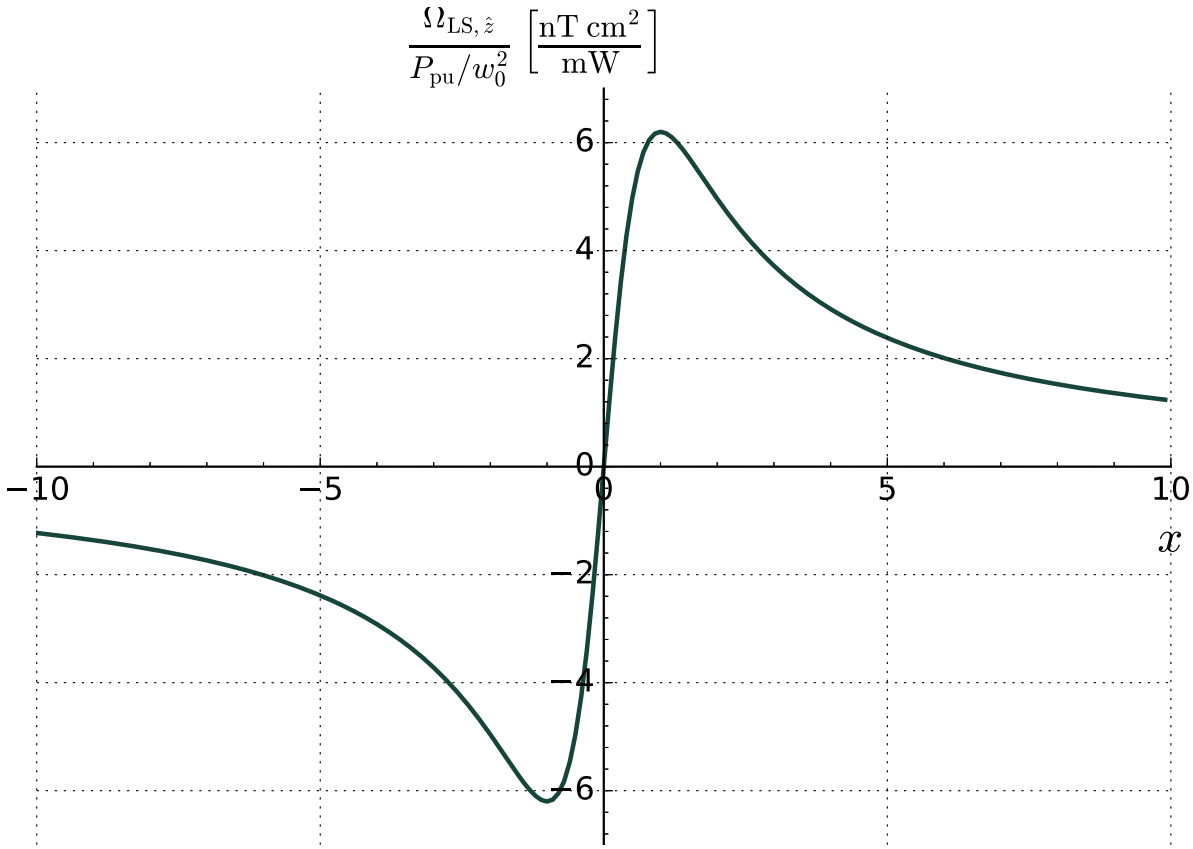



**Figure 2.12:** The pump beam-induced light shift (Equation (2.54) in the magnetic field, normalized to a 1 cm waist and 1 mW pump power. The horizontal axis in this plot is the normalized detuning of the pump beam from the D1 resonance ($x = \Delta\nu/(\Delta\Gamma/2)$). The light shift is maximum at a detuning equal to half of the absorption linewidth ($x = 1$). The light shift can be eliminated entirely by tuning the laser exactly on resonance ($x = 0$). However if the vapor is very optically thick, tuning on resonance can make it difficult to optically pump the entire cell.

The non-uniformity of the light shift in the cells led to the development of a new type of SERF magnetometer that mitigates the effect of the light shifts. This method will be discussed in Section 2.4.

### 2.3.7 Faraday rotation detection

We have shown how in our SERF magnetometer, circularly polarized light creates an atomic polarization along the $\hat{z}$ axis of the cell ($P_z$), and magnetic fields with a component in the $\hat{y}$ direction cause those polarized atoms to precess in the $xz$ plane, producing $P_x$. The only remaining question is how to detect that $P_x$. Again, we utilize the atomic polarizability of the atoms in the presence of laser light, except this time we will take advantage of it! To detect $P_x$, we send a linearly polarized probe beam detuned from the D2 transition down the $\hat{x}$ axis of the cell. Any change in $P_x$ (say, caused by an applied magnetic field $\Omega_y$) alters the polarizability $\overleftrightarrow{\boldsymbol{\alpha}}$ of the alkali vapor. This causes a slight difference in the refraction index for $\sigma^+$ and $\sigma^-$ light, a linear combination of which describe the linearly polarized light of the probe. As the probe propagates through the medium, the difference in refracive indicies causes the polarization of the beam to rotate by an angle $\phi$, given by

$$\phi = -\frac{1}{4} n_{\mathrm{Rb}} l r_e f_{\mathrm{D2}} c \mathcal{D}(\Delta\nu) P_x \tag{2.55}$$

where $l$ is the path length of the light through the vapor. All other variables are as previously defined, although values corresponding to the rubidium D2 transition should be used for $f_{\mathrm{D2}}$ and $\Delta\nu$. (A complete derivation of $\phi$ follows the method of Bonin and Kresin [1997] and is available in Appendix A of Wyllie [2012]). Given Equation (2.55), one would be tempted to tune the probe laser to a detuning of $\Delta\nu = \pm\Delta\Gamma/2$, as these frequencies are the local maxima of $|\mathcal{D}(\Delta\nu)|$. However, this neglects to take the absorption of the probe photons into account. In Section 2.3.2.2, we discussed how absorption of probe photons inccreases the spin-relaxation rate $\Gamma$, but now we also take into account that the rotation signal $\mathcal{S}$ depends on the probe photon flux incident on the photodiodes of the rotation detector. Absorption of the probe decreases the number of probe photons counted by the photodiodes and hence *increases* the effective photon shot noise of the measurement. Taking the probe absorption ($\propto \exp[-n_{\mathrm{Rb}} l \sigma_{\mathrm{D2}}]$) into account, we can write the

rotation signal as the product of the probe power $\mathcal{P}_{\text{pr}}$ and the rotation $\phi$, or

$$\mathcal{S} = \mathcal{P}_{\text{pr}}\phi \tag{2.56}$$

$$\mathcal{S} = \frac{\mathcal{P}_{\text{pr}}}{2} n_{\text{Rb}} l \sigma_{\text{D2}} \frac{\Delta\nu}{\Delta\Gamma/2} \exp\left[-n_{\text{Rb}} l \sigma_{\text{pr}}\right] P_x \tag{2.57}$$

$$= \frac{\mathcal{P}_{\text{pr}}}{2} x\tau(x) \exp\left(-\tau(x)\right) P_x \tag{2.58}$$

where $x$ is again the normalized detuning $\Delta\nu/(\Delta\Gamma/2)$, $\sigma_{\text{D2}}$ is the absorption cross section for probe photons by the vapor, and and $\tau(x) = n_{\text{Rb}} l \sigma_{\text{D2}}(x)$ is the optical depth of the medium at the detuning $x$. Figure 2.13 shows the signal size $\mathcal{S}$ both with and without the effect of probe absorption. Detuning much farther away from resonance than Equation (2.55) would suggest is necessary to optimize the Faraday rotation signal when absorption is taken into account. If we express $P_x$ in the low-field limit $\omega \ll \Gamma'$ then, using Equations (2.39) and (2.40), we can write

$$P_x = \frac{R}{\Gamma'^2}\Omega_y = \frac{R\gamma}{\Gamma'^2} B_y \tag{2.59}$$

and thus,

$$\mathcal{S} = \frac{\mathcal{P}_{\text{pr}}}{2} x\tau(x) \exp\left[-\tau(x)\right] \frac{B_y}{B_0} \tag{2.60}$$

where $B_0 = (R\gamma/\Gamma'^2)^{-1}$. Equation (2.60) gives us a good "target" to aim for when optimizing the magnetometer, and plugging in common experimental values gives a value of $\phi/B_y = 1 \times 10^{-7}$rad/fT. Generally, if we measure a level within a factor of two of this value (in response to an adiabatic field sweep), we say that our laser detunings have been well-optimized.

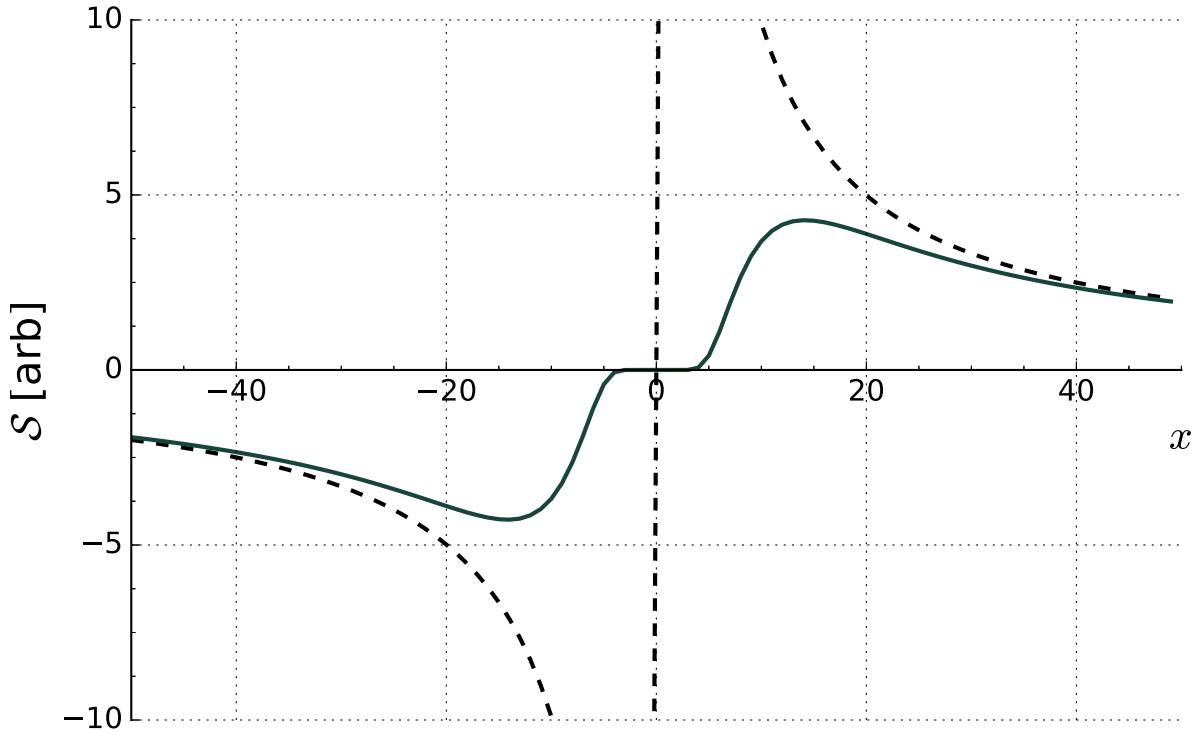



**Figure 2.13:** Faraday rotation signal as a function of normalized detuning $x = \Delta\nu/(\Delta\Gamma/2)$. Without accounting for absorption (dashed line), the detuning which maximizes the rotation is $x = 1$ ($\Delta\nu = (\Delta\Gamma/2)$. With absorption (solid line), a larger detuning must be used to avoid having the probe light be completely absorbed by the vapor. The absorption curve uses an on-resonance optical depth of 100, which is common for our measurements. Probe detuning optimization scans (Section 4.5) help us choose the correct detuning to maximize signal for a particular measurement.

## 2.4 Diffusion mode

### 2.4.1 Motivation

The light shift (Section 2.3.6) presents an interesting problem in high-sensitivity magnetometry measurements. Recall from Equation (2.51) that the magnitude of the light shift is directly proportional to the pumping rate $R$. In theory, we would set $R = \Gamma$ everywhere in the cell (to maintain $P_z = 1/2$) and would therefore also have a uniform light shift that could be canceled through the application of a uniform $\hat{z}$ bias field. In practice, though, this is impossible due to the high optical depth (50-200 on resonance, depending on density) of the vapor, which results in significant absorption of the pump as it propagates across the cell. This in turn leads to a longitudinally-varying pumping rate (in addition to the transversely-varying rate due to the gaussian intensity profile of the pump), which leads to a spatially-varying light shift $\mathbf{\Omega}_{\mathrm{LS}}$.

There are several avenues to attempt to circumvent the effect of the light shift. As shown in Figure 2.12, the light shift goes to zero when the pump laser is tuned exactly to resonance, regardless of the laser intensity. In this situation, however, the light is quickly absorbed at the front of the cell, and the average cell polarization is low and the sensitivity suffers. Simply increasing the pumping rate to a large value (say, $R = 1000\Gamma$) to reach atoms in the rear of the cell is not ideal either, because in this case, many of the atoms simply become "pinned" to the $\hat{z}$ axis and the sensitivity to $\Omega_y$ is low. Tuning the pump away from resonance allows one to reach the back of the cell with a relatively low pumping rate, but this will result in increased light shifts (especially if $\Delta \approx \Gamma/2$), according to Equation (2.51). In short, it is difficult to maintain the three optimum conditions ($R = \Gamma$, $\mathbf{\Omega}_{\mathrm{LS}} = 0$, and pumping the entire cell) simultaneously with a "standard" SERF magnetometer. However, a proposal, first presented in Wyllie [2012] and first published by Sulai et al. [2013], utilizes the diffusion of the polarized atoms as a tool to do just that.

### 2.4.2 Conceptual basis for diffusion-mode

In diffusion mode pumping, the pump beam is shrunk down to a tight waist $w_{\mathrm{pu}} \sim 0.3$ mm, a factor of ten or more smaller than non-diffusion-mode experiments. We maintain (or even increase)

the pump power in the narrow waist to increase intensity. At a radial distance of $\rho$ from the center of the beam, the pumping rate obeys the Gaussian profile of the pump intensity,

$$R(\rho) = R_0 \exp\left[-2\frac{\rho^2}{w_{\text{pu}}^2}\right], \tag{2.61}$$

with

$$R_0 = \frac{2\Phi_{\text{pu}}}{\pi w_{\text{pu}}^2}\sigma_0(\Delta\nu), \tag{2.62}$$

where $\Phi_{\text{pu}}$ is the pump photon flux and $\sigma_0(\Delta\nu)$ is the D1 absorption cross section for detuning $\Delta\nu$. We can define a critical radius $\rho_0$ such that inside this critical radius, the atoms will be pumped to a high polarization ($P_z \approx 1$). We define this critical radius by the location where the pumping rate approximately equals the atom-atom spin-destruction rate $\Gamma_{\text{sd}}$, or

$$\rho_0 = w_{\text{pu}}\sqrt{\frac{1}{2}\ln\left(\frac{R_0}{\Gamma_{\text{sd}}}\right)}. \tag{2.63}$$

When the atoms are highly polarized, they primarily occupy the $m = +1/2$ "dark state" of Figure 2.7. In this state, they do not absorb the $\sigma^+$ light, and this region of the cell can be pumped to a state where the pump light is linearly (rather than exponentially) attenuated along the longitudinal $\hat{z}$ axis of the cell [Lancor, 2011]. In this way, we ensure a more uniform pumping rate longitudinally through the cell, a goal mentioned in the introduction to this section.

The problem is that these highly pumped atoms are not very useful for magnetometry. Overlapping the $\hat{x}$ probe beam precisely with this volume of highly pumped atoms would result in a poor magnetometer whose constituent atoms would be pinned along $\hat{z}$ and not able to precess into the $\hat{x}$ direction to be detected by the probe. Instead, we expand the probe beam until its waist matches (or slightly exceeds) the diffusion length $z_D$ of the cell where

$$z_D = \sqrt{\frac{qD(P,T)}{\Gamma_{\text{sd}}}}. \tag{2.64}$$

With this large waist, the probe still samples the low-response "pinned" atoms inside the pump beam, but also samples a much larger volume of surrounding atoms. These atoms were polarized in the high-pumping rate region of the vapor and have since diffused outside of it. Because the pump

light intensity is small outside of the narrow region in the center, the atoms that have diffused experience a pumping rate of nearly zero and are thus not subject to the strong light shift fields. Atoms inside the pump beam will see still see these enormous light shift fields, but they do not contribute much to the magnetometer signal because their response to applied $\Omega_y$ is so small with the large pumping rate (Figure 2.9).

Operating in diffusion mode not only decreases the size of the average light shift field [Sulai et al., 2013], but also decreases the light shift *gradients* down the longitudinal axis of the cell due to the increased uniformity of the pumping rate on that axis. Consequently, the mean value of $\mathbf{\Omega}_{\mathrm{LS},z}$ is much more representative of the entire cell, and performance is improved by an increased ability to null that field by applying a uniform $B_z$ shim field to the atoms. With enough pump light, it is even possible to operate the pump beam right on resonance and eliminate the light shift completely. Any excess pump light present at the back of the cell can be used to pump another cell if the units are arranged appropriately (see Figure 3.5).

### 2.4.3 Optimizing cells for diffusion mode

The first experiments with diffusion mode operation were performed with cells that had originally been filled with approximately 760 Torr $^4$He background gas to limit diffusion to the walls [Wyllie, 2012]. Over time, the $^4$He diffused through the cell walls, leaving only a small amount of the gas plus the 50 Torr $N_2$ that was always present. The low total pressure in the cells allowed for a sufficiently long $z_D$ that diffusion mode operation could be attempted. It was noted in these experiments, though, that the increased wall relaxation probably limited the performance of the magnetometers.

For this work, we purchased new cells and even experimented with fabricating our own (Appendix B). In both cases, we had the ability to choose a $N_2$ buffer gas pressure, and we optimize that pressure for diffusion mode. To do so, we assume a critical pump radius of $\rho_0 = .05$ cm, and posit that the atoms inside this radius are spin-polarized to unity, but completely insensitive to $\hat{y}$ fields. We then assume that once the atoms are outside this critical radius, they spatially diffuse

uniformly, with each particle moving at a constant velocity, carrying the spin-polarization it obtained in the highly-pumped regime with it. All this time, the atoms undergo spin-relaxation due to the contributions mentioned in Section 2.3.2 ($\Gamma = \Gamma_{\text{sd}} + \Gamma_{\text{pr}} + \Gamma_{\text{diff}}$). Note that in this volume (outside of the pump beam), we assume $R = 0$ so $\Gamma' = \Gamma$.

With these conditions in mind, we can write the differential response to a $\hat{y}$ field as

$$d\left(\frac{dP_x}{d\Omega_y}\right) \propto \begin{cases} 0 & \rho \leq \rho_0 \\ t_D \exp\left[-\Gamma t_D\right] dV. & \rho > \rho_0 \end{cases} \tag{2.65}$$

Here, $dV = 2\pi\, \rho d\rho\, l$ is the cylindrical differential volume element and $t_D$ is a characteristic diffusion time given by

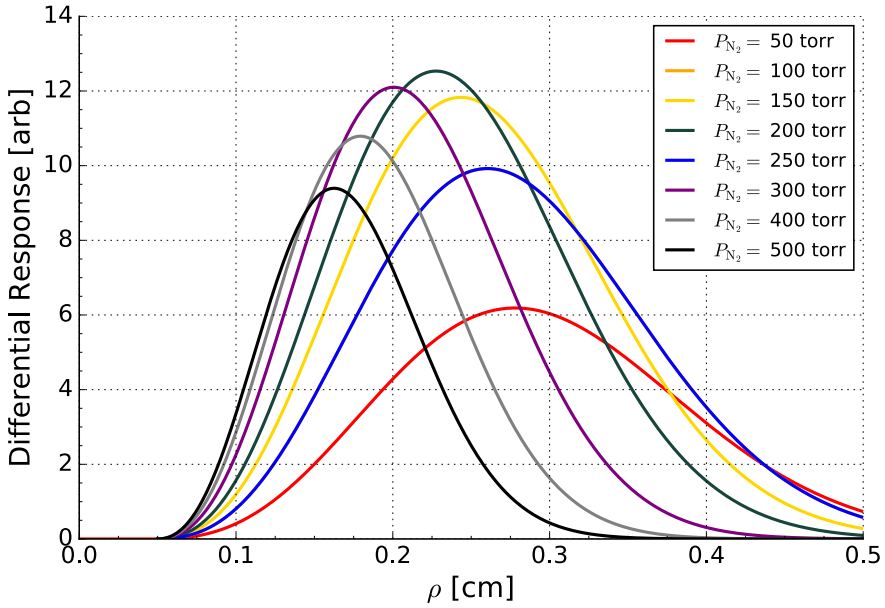
$$t_D = \frac{(\rho - \rho_0)^2}{D(P_{\text{N}_2}, T)}, \tag{2.66}$$

with $D(P_{\text{N}_2}, T)$ defined in Equation (2.27). Though not explicitly stated in Equation (2.65), the relaxation rate $\Gamma$ is also a function of buffer gas pressure and cell temperature (Figure 2.8). Figure 2.14a shows a plot of Equation (2.65) at several different values of $P_{\text{N}_2}$. To get an idea of the signal we obtain from this polarization distribution, we integrate over the volume that the probe beam samples. The calculation to do this with a perpendicular probe beam (as in our setup) is quite complicated. To simplify matters somewhat, we assume that instead the probe laser is collinear with the pump laser, and we integrate outwards from the center of the pump beam. We believe this to be an acceptable approximation as

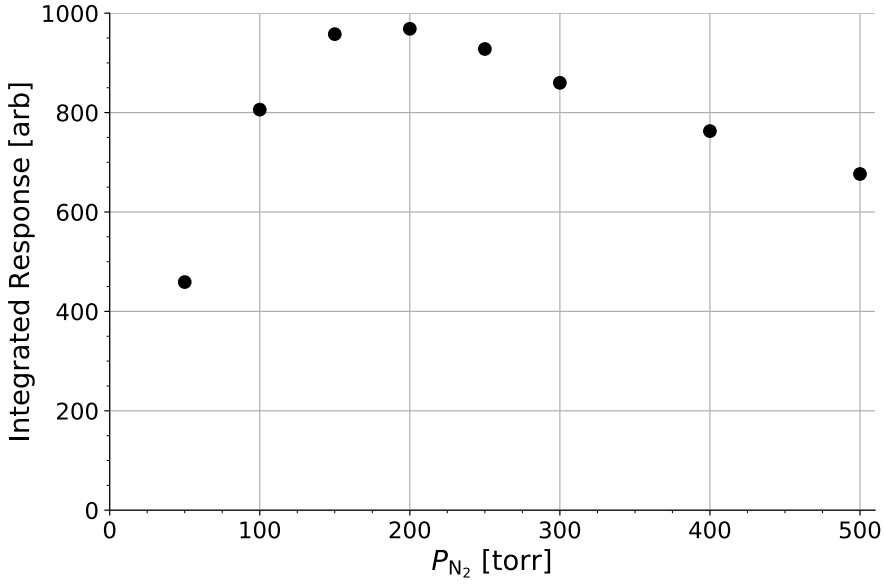
1. the cell has the same length dimension in both the $\hat{x}$ (probe) and $\hat{z}$ (pump) directions, so the *number* of atoms sampled will be identical; and

2. although the region of "pinned" atoms will have a different appearance depending on the orientation of the probe beam (a longitudinal "column" for a collinear beam or a horizontal "stripe" for a transverse beam), the actual number of probed atoms inside this "pinned zone" will be similar as long as the probe beam is expanded to a sufficient size.

Figure Figure 2.14b shows the results of integrating the curves of Figure 2.14a from $\rho = 0$ to $\rho = .3$ cm ($w_{\text{pr}} \approx .6$ cm). Using these data points, we estimated that the ideal buffer gas pressure was




**(a)**




**(b)**

**Figure 2.14:** (a) The differential response (Equation (2.65)) as a function of the distance $\rho$ from the center of the pump beam at several different values of $N_2$ buffer gas pressure. The temperature of the atoms is 160°C. The atoms at a radius $\rho \leq \rho_0 = 0.05$ cm are "pinned" to the $\hat{z}$-axis. They are highly polarized ($P_z \sim 1$), but they have negligible response due to the extreme optical pumping. Only when the atoms diffuse to a radius $\rho > \rho_0$ is there any increase in sensitivity ($dP_x/d\Omega_y$). (b) The result of integrating Equation (2.65) out to $\rho = 0.4$ cm (roughly the size of the probe beam). The optimum response with these pump and probe parameters is realized with a pressure of 150-200 torr $N_2$. The response falls off faster on the low-pressure side of this maximum, so erring on the side of over-filling the cells with nitrogen is desirable compared to under-filling them.

approximately 175 torr $N_2$, which is what we specified when we ordered new cells from Precision Glassblowing.

## 2.5 Z-mode

One of the limiting factors in the detection of magnetic fields with a SERF magnetometer is non-magnetic "technical" noise, originating from the laser diodes, optics, data acquisition devices, and other electronics. The optical fibers used in our experiment in particular can be quite a nuisance because their noise contribution is often worst at low frequency ($< 10$ Hz), a band that is critical to accurately measuring our cardiographic signals of interest. A similar issue was present in the early SERF magnetometers of Li [2006]. In that case, the low frequency noise was due to polarization fluctuations in the probe beam caused by air currents in the forced-air oven which heated the magnetometer cell.

Despite our best efforts to limit these technical noise sources, including switching to polarizing (rather than polarization-maintaining) fibers, our low-frequency detection is often limited by this non-magnetic noise. Therefore, a scheme, first presented by Li et al. [2006], was implemented that uses high-frequency parametric modulation of the longitudinal ($\hat{z}$) magnetic field and lock-in detection of the resultant magnetometer signal to suppress these technical noise sources. The scheme is built on the principle that the only signals near the modulation frequency (which includes the modulated magnetic fields we wish to detect) will be present in the measurement, and the technical noise at this frequency is of negligible size.

The second advantage of using a parametrically-modulated scheme is that the lock-in detected signal is proportional to a second component of the magnetic field: in this case, the magnetic field along the probe direction, $\Omega_x$. The ability to detect two components of the magnetic field ($\Omega_y$ and $\Omega_x$) is useful in biomagnetism, as it allows us to effectively double our channel count and utilize more sophisticated signal processing techniques to extract the MCG or (especially) fMCG signal of interest.

In our scheme, which we refer to as "Z-mode", a large modulation field is applied along the pump ($\hat{z}$) direction and the resultant signal is proportional to both $\Omega_x$ (by demodulation at the modulation frequency $\omega_z$) and $\Omega_y$ (either by lowpass filtering the modulated signal or by demodulating the output signal at the second harmonic $2\omega_z$). We present a derivation here that mostly follows the derivation presented in Wyllie [2012], which itself follows the original derivation in Li [2006].

Let us begin with the SERF magnetometer explained in Section 2.3. Again, we ignore the hyperfine structure of the ground state and in addition, and ignore the effect of the spin polarization $P$ on the slowing-down factor $q$ (*i.e.* $q(P) \to q$). We apply a large modulating field along the $\hat{z}$ direction such that the total field can be expressed as

$$\mathbf{\Omega} = \Omega_x \hat{x} + \Omega_y \hat{y} + (\Omega_1 + \Omega_0 \cos(\omega_z t))\hat{z}, \tag{2.67}$$

with $\Omega_0 \gg \Omega_1$. We then rewrite the steady-state solution (Equation (2.36)) in terms of rotating bases of polarization ($P_\pm = P_x \pm iP_y$) and magnetic field ($\Omega_\pm = \Omega_x \pm i\Omega_y$):

$$q\frac{dP_+}{dt} = -\Gamma' P_+ - i\Omega_+ P_z + i\Omega_z(t) P_+ \tag{2.68}$$

$$q\frac{dP_z}{dt} = \frac{1}{2i}\left(P_+\Omega_- - P_-\Omega_+\right) - \Gamma' P_z + R, \tag{2.69}$$

where $\Omega_z(t) = \Omega_z + \Omega_0 \cos(2\omega_z t)$. As long as the transverse fields $\Omega_x$ and $\Omega_y$ are small compared to $\Omega_z$, then the steady-state longitudinal spin-polarization is still $P_z = R/\Gamma'$. For the polarization along $\hat{x}$ and $\hat{y}$, we try the rotating solution $P_+ = A_+ e^{-i\phi(t)}$. Inserting this into Equation (2.68) and simplifying yields

$$\frac{dA_+}{dt} = -\left(\frac{\Gamma'}{q} + i\left(\frac{d\phi}{dt} - \frac{\Omega_z(t)}{q}\right)\right) A_+ - i\frac{\Omega_+ P_z}{q} e^{-i\phi(t)}. \tag{2.70}$$

To make the first term equal to zero, we can choose $\phi$ such that

$$\frac{d\phi}{dt} = \frac{1}{q}\left(i\Gamma' + \Omega_z(t)\right), \tag{2.71}$$

which we can integrate such to find $\phi$:

$$\phi = \frac{1}{q}\left((i\Gamma' + \Omega_z)t + \frac{\Omega_0}{\omega_z}\sin\omega_z t\right). \tag{2.72}$$

Assuming $\Omega_+ \ll d\phi/dt$, $A_+$ is approximately constant, and we can integrate Equation (2.70) to find an expression for $A_+$:

$$A_+ = -i\frac{\Omega_+ P_z}{q}\int dt' e^{-i\phi(t')} \tag{2.73}$$

$$= -i\frac{\Omega_+ P_z}{q}\int dt' e^{(\Gamma' - i\Omega_z)t'/q - iu\sin\omega_z t'} \tag{2.74}$$

$$= i\Omega_+ P_z \sum_{m=-\infty}^{\infty} \frac{\Gamma' + i(\Omega_z + mq\omega_z)}{\Gamma'^2 + mq\omega_z} J_m(u) e^{(\Gamma' - i\Omega_z)t/q - im\omega_z t}. \tag{2.75}$$

Here, $u = \Omega_0/q\omega_z$ is a dimensionless parameter that is proportional to the magnitude of the applied longitudinal field and inversely proportional to its frequency. In the last line, we use the Jacobi-Anger expansion to express the sine term in the exponential as a sum of Bessel functions $J_m$:

$$e^{iu\sin(\omega_z t)} = \sum_{m=-\infty}^{\infty} J_m(u) e^{im\omega_z t}. \tag{2.76}$$

Recalling that $P_+ = A_+ e^{i\phi(t)}$, and applying the Jacobi-Anger expression in a similar way as we did in Equation (2.75), we write

$$P_+ = i\Omega_+ \frac{P_z}{\Gamma'} \sum_{m=-\infty}^{\infty}\sum_{m'=-\infty}^{\infty} F_m J_m(u) J_{m'}(u) e^{-i(m-m')\omega_z t}, \tag{2.77}$$

where the coefficients $F_m$ are given by

$$F_m = \frac{1 + i(\Omega_z + mq\omega_z)/\Gamma'}{1 + (\Omega_z + mq\omega_z)^2/\Gamma'^2}. \tag{2.78}$$

Next, we expand the Fourier components of Equation (2.77) to see the response $P_+$ to DC fields, as well as fields at each harmonic $n\omega_z$. For DC, we simply multiply the expression by the Kronecker $\delta_{m,m'}$. We employ a similar trick to find the harmonic response, though instead we multiply Equation (2.77) by $\delta_{m,n\pm m'} e^{i\pm n\omega_z t}$. After using Euler's formula to expand the complex exponentials ($e^{i\omega_z t} = \cos(\omega_z t) + i\sin(\omega_z t)$), we can write (using $J_i(u) \to J_i$ for compactness)

$$P_+(\mathrm{DC}) = i\Omega_+ \frac{P_z}{\Gamma'} \sum_m F_m J_m^2 \tag{2.79}$$

$$P_+(n\omega_z) = \Omega_+ \frac{P_z}{\Gamma'} \sum_m F_m J_m((J_{m-n} - J_{m+n})\sin\omega_z t + i(J_{m+n} + J_{m-n})\cos\omega_z t) \tag{2.80}$$

As we are still using the same pump-probe alignment scheme as in Figure 2.6, we are most interested in $P_x = \Re\{P_+\}$, as that is the signal our Faraday rotation detection scheme interrogates (Section 2.3.7). Replacing the coefficients $F_m$ and evaluating the real parts of Equations 2.79 and 2.80 gives

$$P_x(\mathrm{DC}) = -\frac{P_z}{\Gamma'}\Omega_y \sum_m \frac{\Gamma'^2}{\Gamma'^2 + (\Omega_z + mq\omega_z)^2} J_m^2 \tag{2.81}$$

$$\begin{aligned} P_x(n\omega_z) = &\sum_m \frac{P_z}{\Gamma'^2 + (\Omega_z + mq\omega_z)^2} J_m \times \\ &[\Omega_x \left(\Gamma'(J_{m-n} - J_{m+n}) \sin(n\omega_z t) - (\Omega_z + mq\omega_z)(J_{m+n} + J_{m-n}) \cos(n\omega_z t)\right) + \\ &\Omega_y \left(-(\Omega_z + mq\omega_z)(J_{m-n} - J_{m+n}) \sin(n\omega_z t) + \Gamma'(J_{m+n} + J_{m-n}) \cos(n\omega_z t)\right)] . \end{aligned} \tag{2.82}$$

Now, if we work in a regime where the DC bias fields are small and the modulation frequency is large ($\Omega_z \ll \Gamma' \ll \omega_z$) [Li et al., 2006], then only the $m = 0$ term is important, as the $F_m$ terms fall off rapidly for $m \geq 1$. Finally, we make use of the Bessel function relationship $J_{-m} = (-1)^m J_m$ to find

$$P_x(\mathrm{DC}) \approx -\frac{P_z}{\Gamma'}\Omega_y J_0(u)^2 \tag{2.83}$$

$$P_x(n\omega_z) \approx -\frac{P_z}{\Gamma'} 2J_0(u)J_n(u) \begin{cases} -\Omega_x \sin(n\omega_z t) & n \text{ odd} \\ \Omega_y \cos(n\omega_z t) & n \text{ even} \end{cases} \tag{2.84}$$

The DC response to an applied $\hat{y}$ field is the same as the DC-mode magnetometer, except modified by an additional factor $J_0^2(u)$. The response at higher harmonics alternates between $\hat{x}$ and $\hat{y}$ sensitivity, modified by the Bessel product $2J_0(u)J_n(u)$. In practice, this means we have a magnetometer sensitive to two components of the field simultaneously, with $\Omega_x$ sensitivity at the first harmonic of the modulation (using a lock-in amplifier, for instance) and $\Omega_y$ sensitivity *either* by simply low-pass filtering the magnetometer signal or by detecting at the second harmonic.

A justification for the former method is provided by Figure 2.15. In an effort to maximize sensitivity in both components simultaneously, we generally work with a value of $u$ $(= \Omega_0/q\omega_z) \approx .90$, which is the intersection of the $n = 1$ $(J_0^2(u))$ and $n = 2$ $(2J_0(u)J_1(u))$ curves, near the

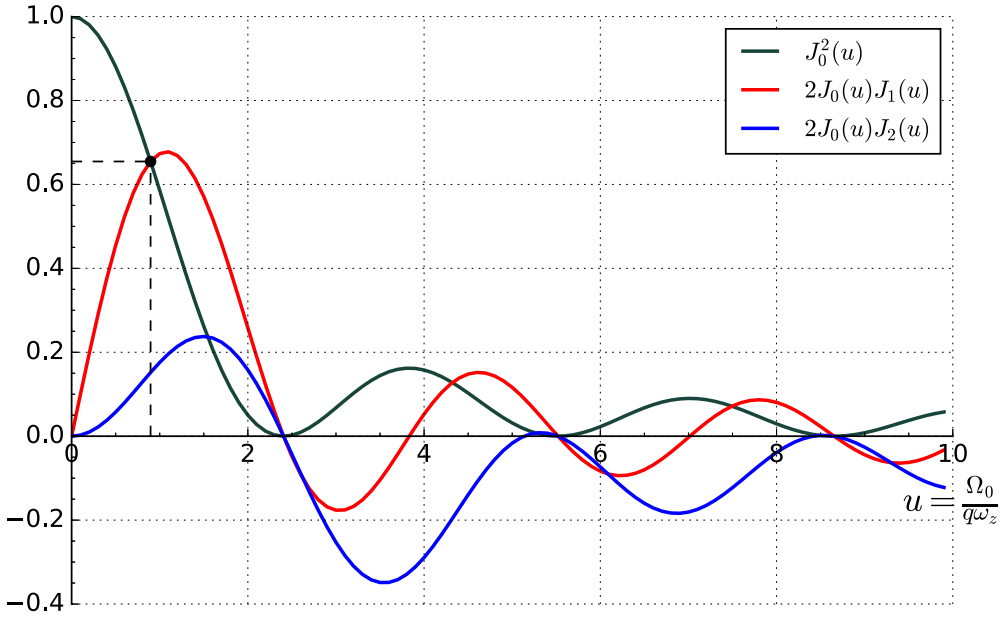



**Figure 2.15:** The response-reducing prefactors from operating the SERF magnetometers in Z-mode. We generally try to operate at the point indicated by the black dot ($u \approx .90$). Here, both $\Omega_x$ and $\Omega_y$ sensitivity are at roughly 65% of the DC-mode level. If one desired to use the $n = 2$ harmonic to measure $\Omega_y$, the optimum location would use roughly $u = 1.5$, resulting in a similar level of $\Omega_x$ response, but the sensitivity to $\Omega_y$ would suffer greatly.

local maximum of the $\hat{x}$ sensitivity. At this value of $u$, the response to $\Omega_y$ is 65% of the DC-mode sensitivity, an acceptable decrease given the extra component of the magnetic field ($\Omega_x$) which can be measured (also at 65% of the DC-mode $\Omega_y$ sensitivity). Alternatively, the second-order harmonic could be used to detect $\Omega_y$ [Li et al., 2006]. This is an attractive strategy as both components of magnetic field would use lock-in detection to suppress low-frequency noise. However, the common optimum for the $n = 1$ and $n = 2$ Bessel functions results in reduced sensitivities to both $\Omega_x$ and, to a more severe degree, $\Omega_y$, and thus is not often implemented.

To satisfy $\Omega_0 = .90q\omega_z$, we ostensibly have an unlimited set of $(\Omega_0, \omega_z)$ pairs to choose from. However, we must keep in mind the relationship $\Omega_z \ll \Gamma' \ll \omega_z$ when selecting a pair. Large $\omega_z$ would seem to be beneficial for moving away from the low-frequency noise which can trouble a DC-mode measurement. At the same time, this requires a larger magnitude of the modulation field ($\Omega_0$), which may become large enough to take the magnetometer out of the SERF regime.

In general we operate at a frequency of $\Omega_z = 2\pi \times 1$ kHz to strike a balance between these two extremes.

A complete description of operating the magnetometers in Z-mode, as well as some performance data, is given in Section 4.6. MCG data acquired using Z-mode are discussed in Chapter 6 and Appendix E.

## 2.6 Active feedback

Active feedback has been used since the earliest days of atomic magnetometry. The $M_x$ magnetometer shown in Figure 2.1 uses positive feedback to make a self-oscillating device when the drive frequency matches the Larmor frequency of the precessing atoms. Negative feedback can be used to stabilize a magnetometer against unwanted variances in the environmental magnetic field, such as those produced by nearby moving magnetic objects (such as elevators, vehicles, construction equipment, and in the case of Wyllie [2012] a ∼6 Hz HVAC fan that could not always be disabled during measurements). An exceptional example of the use of active negative feedback in magnetometry was Bison et al. [2009] who employed 25 magnetometers (19 field sensors plus 6 dedicated solely to active feedback) to make an adult MCG measurement in an *unshielded* enviroement. Seltzer and Romalis [2004] presented another unshielded magnetometer with three axis feedback with ∼1 pT/$\sqrt{\text{Hz}}$ sensitivity.

In its most basic sense, the goal of active feedback is to have a system output follow a control signal as closely as possible [Bechhoefer, 2005]. This is accomplished by measuring the output of the system (call it $y(t)$) and comparing it to the control signal, $r(t)$. The result of this comparison is known as the error signal, $e(t)$, where

$$e(t) = r(t) - y(t). \tag{2.85}$$

A correction signal is generated based on the error signal, and this corrective term is fed back into the feedback servo, which is connected to the system input. The output of the servo is added to the system. The dynamics of the servo, $K$, are adjusted to minimize the error signal, thus holding the output $y(t)$ close to the desired signal $r(t)$.

It is often easier to describe these dynamics in the frequency domain. Consider a system whose transfer function is given by $G(s)$, where $s = i\omega$. $i$ is the imaginary unit and $\omega$ is the frequency of interest. The system takes an input signal $\tilde{u}(s)$ and outputs a signal $y(s)$ such that

$$\tilde{y}(s) = G(s)\tilde{u}(s). \tag{2.86}$$

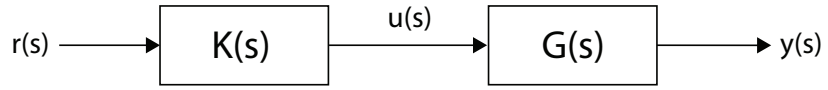
Often $G$ is a function over which we have little to no control, for example, the atomic magnetometer response, or the transfer function for a hard-wired electronic system. Now consider adding a second transfer function $K(s)$ that we *do* control, inserted into the system before the "uncontrollable" function $G(s)$, such that the output of $K$ is the input to $G$. Now, $\tilde{u}(s) = K(s)\tilde{r}(s)$ as shown in Figure 2.16a. The output of the new system is

$$\begin{aligned} \tilde{y}(s) &= G(s)\tilde{u}(s) \\ &= G(s)K(s)\tilde{r}(s). \end{aligned} \tag{2.87}$$

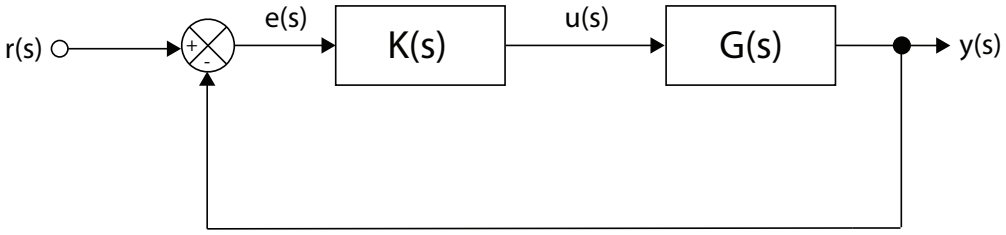
We define the "open-loop" gain as the product of the two transfer functions $G(s)$ and $K(s)$. We can introduce feedback into the system by "closing the loop", as shown in Figure 2.16b. In this scheme, we compare the output signal $\tilde{y}(s)$ to the desired control signal $\tilde{r}(s)$, the result of this comparison being the error signal $\tilde{e}(s)$. It is this error signal that is fed into the servo $K(s)$. Thus,

$$\begin{aligned} \tilde{y}(s) &= G(s)K(s)\tilde{e}(s) \\ &= G(s)K(s)\left(\tilde{r}(s) - \tilde{y}(s)\right) \\ &= \frac{G(s)K(s)}{1 + G(s)K(s)}\tilde{r}(s) \\ &\equiv T(s)\tilde{r}(s). \end{aligned} \tag{2.88}$$

The hope is that the dynamics of $T(s)$ are superior to those of $G(s)$. In other words, $T(s)$ forces the output of the combined system to more closely follow the control signal $\tilde{r}(s)$ than does$G(s)$. This is accomplished through careful design of $K(s)$. A discussion of our specific implementation of active feedback is found in Section 3.5.6, and a detailed explanation of how we tailored our servo $K(s)$ to fit the needs of our experiment is discussed in Appendix C.




**(a)** Open loop




**(b)** Closed loop

**Figure 2.16:** (a) The system in "open loop" mode. The input signal $r(s)$ is first modified by the prepared servo $K(s)$ before being passed to the system $G(s)$. (b) The system in "closed-loop" mode. The output signal $y(s)$ is fed back to the input of the system, and the servo $K(s)$ works to hold the output as close to the control signal $r(s)$ as possible, thereby minimizing the error signal $e(s)$.

# Chapter 3

# Apparatus

The goal of this project is the measurement of adult and fetal magnetocardiograms with an array of SERF magnetometers. A large signal-to-noise ratio is necessary to discriminate the small cardiographic signals from the rest of the environment; thus all measurements are taken inside a magnetically shielded chamber. Measuring inside this chamber dramatically increases our SNR, but logistical concerns from the shared use of the chamber (with another biomagnetism experiment) require modifications to the experimental design. These additional design considerations, which would not be necessary with a permanent setup and may seem burdensome at first, are in the end good motivation for constructing a device that could one day be deployed in a healthcare environment with limited magnetic-shielding resources. Table 3.1 provides an overview of the design requirements and considerations that were taken into account when constructing our device.

A complete MCG measurement system was constructed by Wyllie [2012]. While the fundamental design of the apparatus described in this chapter comes from Wyllie, nearly every component setup has been redesigned and improved. A main goal of this redesign was to increase reliability, compactness, and user-friendliness. This chapter will discuss each component of our magnetometer array apparatus, highlighting the improvements made from the first generation device.

| Req. type | Requirement | Design solution |
|---|---|---|
| SNR | Reduce magnetic bias fields and environmental noise in measurement bandwidth of interest | Operate apparatus in magnetically-shielded room or chamber |
| | Minimize distance from magnetometer to source (patient) | Use quality thermal insulation to reduce required insulation thickness |
| | Minimize metal near cell | Use non-magnetic ceramics and plastics for structural parts |
| | Optimize for uniform background noise removal | Use multiple magnetometer channels, set channel spacing roughly equal to source depth for optimal Garachtchenko et al. [1999] gradiometer performance |
| Practical | Array must be portable | Modular design:<br>• Pump and probe single lasers fiber coupled, fiber split to interface with individual channels<br>• Electric signals carried in bulk by Cat5 cables<br>• Individual tri-axial coils for each magnetometer for local field adjustment<br>• Resistive heating<br>Optics and electronics fit on a movable electronics rack, can be set up and packed up in $<$15 minutes |
| | Array must be easily repositioned to minimize distance from detectors to cardiac magnetic field source | Easily attached to commercial SQUID with overhead gantry |
| | Compact size to fit inside of smaller shielded chamber | Channels mounted in fixed positions with shared probe optics |

**Table 3.1:** Design considerations for the magnetometer array. The two (sometimes competing) objectives of the apparatus are a large SNR for MCG detection while at the same time maintaining portability so the device can be easily maneuvered to an optimal position for measurement, as well as taken down and removed from the shielded chamber when necessary.

## 3.1 Magnetometer unit

A single magnetometer channel is shown in Figure 3.1. The components of the magnetometer are the alkali-metal vapor cell, heaters to warm the cell to a temperature suitable for SERF magnetometery, thermal insulation to limit heat loss to the environment (and simultaneously protect the test subject from the significantly elevated temperature of the glass vapor cell), and a 3-D printed "shell," which serves the dual purpose of maintaining the structure of the magnetometer channel and providing a structure for the triaxial magnetic field coils that are used to zero the field inside of the vapor cell and apply modulation fields when appropriate.

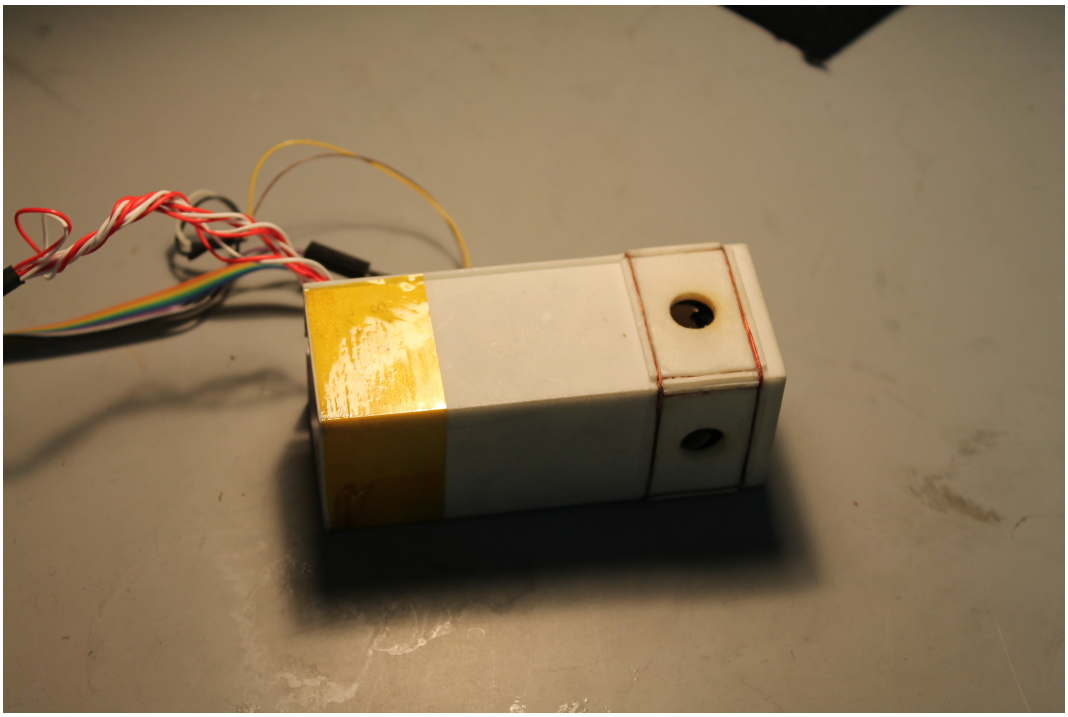


**Figure 3.1:** A single magnetometer channel. The white 3-D printed shell contains the cell, heaters, and insulation. Holes in the shell allow for optical access for both pump and probe lasers. Triaxial coils (copper-colored) are wrapped directly onto the cell framework, enclosing the active region. At top left, the various electronic connections can be seen (rainbow ribbon: field cancelling and modulation currents; red/white twisted pairs: heater signals; brown/yellow twisted pair: thermistor signal. A thin layer of Mylar tape is often added to secure the shell in the magnetometer array holder.

### 3.1.1 Vapor cell

The vapor cells used in this thesis were 1 cm x 1 cm x 5 cm rectangular fluorometer cells made of Pyrex glass obtained from Starna Cells (part number 3-PX-10). Five of the six walls of the cell were float-surfaced, whereas the sixth was formed into a tip-off during the filling process. Due to the glassblowing process, the final length of the cells varied between 3.5 and 7 cm. The cells contained enough isotopically-pure rubidium-87 (generally 3-10 mg) to obtain a saturated alkali vapor up to temperatures of at least 180°C.

In addition to the rubidium, each cell contained $N_2$ buffer gas. As explained in Section 2.4.3, to facilitate diffusion-SERF mode, this nitrogen buffer gas pressure should not be made too large, else it will hamper the redistribution of angular momentum in the cell and reduce the device's performance. Too small a buffer gas density, however, and radiation trapping Happer [1972] will make the vapor cell unusable as a magnetometer. (We observed with the latter case when we took delivery of cells that contained little to no $N_2$. Re-radiated photons from optically-pumped atoms were so pervasive that a purplish glow was visible with the naked eye.) Cells with pressures from 170 to 500 torr of $N_2$ pressure were used in experiments; after significant amounts of modeling and experimentation, it was determined that pressures at the low end of that spectrum, $\sim$175 torr, gave the best diff-SERF magnetometer performance.

An effort, described in its entirety in Appendix B, was made to produce these vapor cells in-house. Slow turn-around times and high costs motivated a shift away from cells produced by third party sources such as Opthos and Triad Technologies. While some excellent cells were created by our group, overall, the work was plagued by an abnormally high number of cells which either performed well for only a short time or never worked at all. A decision was made to return to professionally made cells, and the most recently-obtained cells were four units obtained from Precision Glassblowing of Colorado. These cells were isotopically-pure $^{87}$Rb with 175 torr of $N_2$ buffer gas. Having all four channels of magnetometer array with nominally identical cells has simplified our optimization procedure considerably, as each cell can be elevated to the same temperature and exhibit similar performance to its neighbors.

### 3.1.2 Heaters

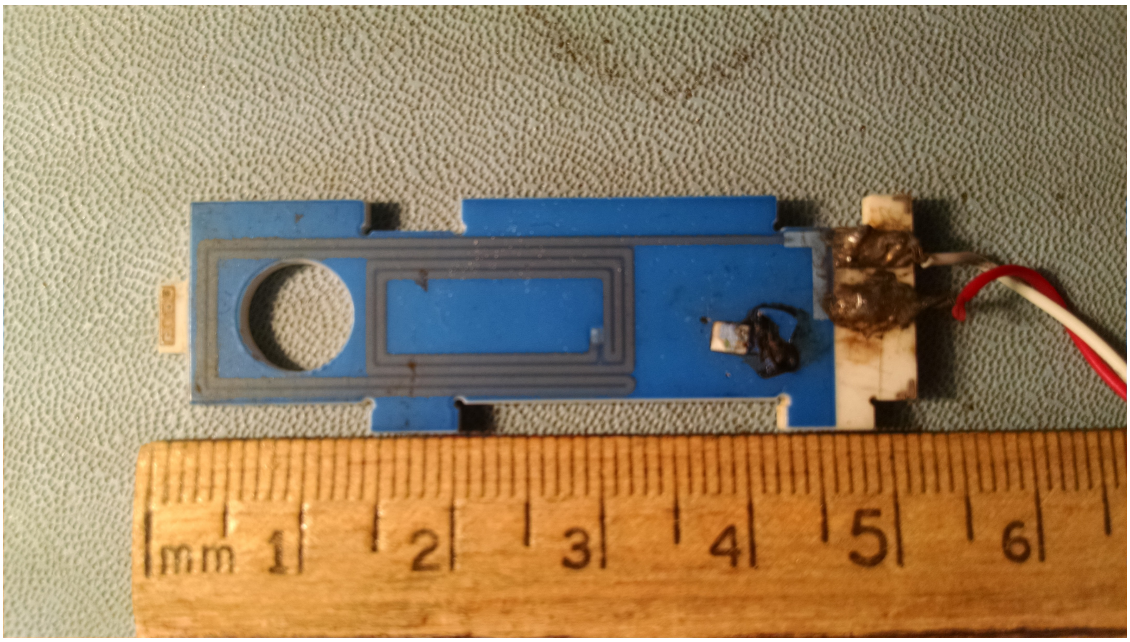
To obtain a sufficiently high rubidium density for SERF magnetometry, the cells must be heated to temperatures between 140-180°C. The previous iteration of the apparatus accomplished this with the cell mounted inside a thermally-conductive Shapal-M heatsink with resistive heaters attached to the opposite end of the sink [Wyllie, 2012]. In an effort to more evenly distribute the heat across the cell walls, we designed resistive heaters printed directly on a thermally-conductive substrate. The heater is shown in Figure 3.2. The 2.1 mm-thick ceramic is a 96% alumina composite material which is laser cut into a jigsaw shape, which allows for four interconnected pieces to form a box oven around the cell. An 8 mm hole is cut at one end of the panel for optical access.

Because of the close proximity of the heaters to the cells, great care must be taken to limit the magnetic fields produced by the resistive traces in the heaters. To accomplish this, the resistive traces on each heater panel were printed counterwound, one on top of each other, separated by two layers of a thin dielectric (ESL 4917) and connected by a via through the dielectric. A layer of dielectric is placed over the top resistive layer to protect the trace from the environment The result was the current in every part of the trace was effectively canceled that from a neighboring trace, which almost completely negates any magnetic field contribution from these traces. A small magnetic field arises from the non-aligned solder pads at the bottom of the panel where wire leads are attached, but this location is designed to be as far as possible from the measurement region of the vapor cell. Using platinum ($\rho = 10.6 \times 10^{-8}\,\Omega$·m) traces of width 1.2 mm, the panels have a resistance of $\sim$200 $\Omega$ each; four such panels connected in parallel form a 50 $\Omega$ oven. The current through the oven is driven by a op-amp circuit described briefly in Section 3.5.1 and in greater detail in Appendix A. Two layers of fluorinated ethylene propylene (FEP) heat-shrink tubing placed around the exterior of the oven give the cell and oven assembly some structural rigidity, while also allowing for quick disassembly if a cell or heater panel should fail.

A platinum resistive temperature detector (RTD, Heraeus L220 model) is affixed to one of the four heater panels, giving us a real-time temperature readout of the oven assembly. Because the



**(a)** Heater panel drawing




**(b)** Heater panel photo

**Figure 3.2:** CAD drawing and photo of the heater panel. Current enters from the pad in the top right corner (connected to the white wire in the photo) and flows through the platinum printed trace around the perimeter of the panel, working its way to the center. A via (square to the right of the vertical centerline) provides access through a thin dielectric at which point the current traces back the same path–this time flowing in the opposite direction–until it reaches other pad (which the red wire has been soldered to). "Counterwinding" the traces in this manner reduces magnetic fields produced when current flows through them. Visible in the photo near the 4 cm mark of the ruler is one of the platinum RTD temperature sensors that failed and is no longer in use.

sensor is resistive, however, it is necessary to disconnect it from the ohm-meter while magnetometer measurements are underway, or else small but noticeable magnetic fields can be produced by the currents running through the RTD. Appendix A contains more details on the RTD.

### 3.1.3 Shell and insulation

Thermally insulating the cell provides two benefits.

1. Reducing the power (and thus, the current) required to maintain the cell at SERF operation temperature.

2. Protecting the patient's skin from the hot cell during a measurement.

The insulation of choice is Aerogel, a silica-based fabric with an ultra-low thermal conductivity of 0.021 W/ m·K. The cell is enclosed on all sides by 4-8 mm of Aerogel sheets, with holes cut in the fabric aligning with the holes in the heater panels for pump and probe beam access. The chalky, flexible composition of the Aerogel makes the windows particularly vulnerable to obstructing Aerogel fabric and dust which can degrade measurements. An additional polyether ether ketone (PEEK) structure was 3D printed which keeps the optical path clear. The structure can be seen in detail in Figure 3.3.

The final piece of the magnetometer channel assembly is a 3D printed shell that holds the assembled magnetometer and insulation securely in place. In addition, it serves as a mount for the magnetic field coils which are used to null the fields and apply modulation fields. These coils will be discussed in-depth in the next section.

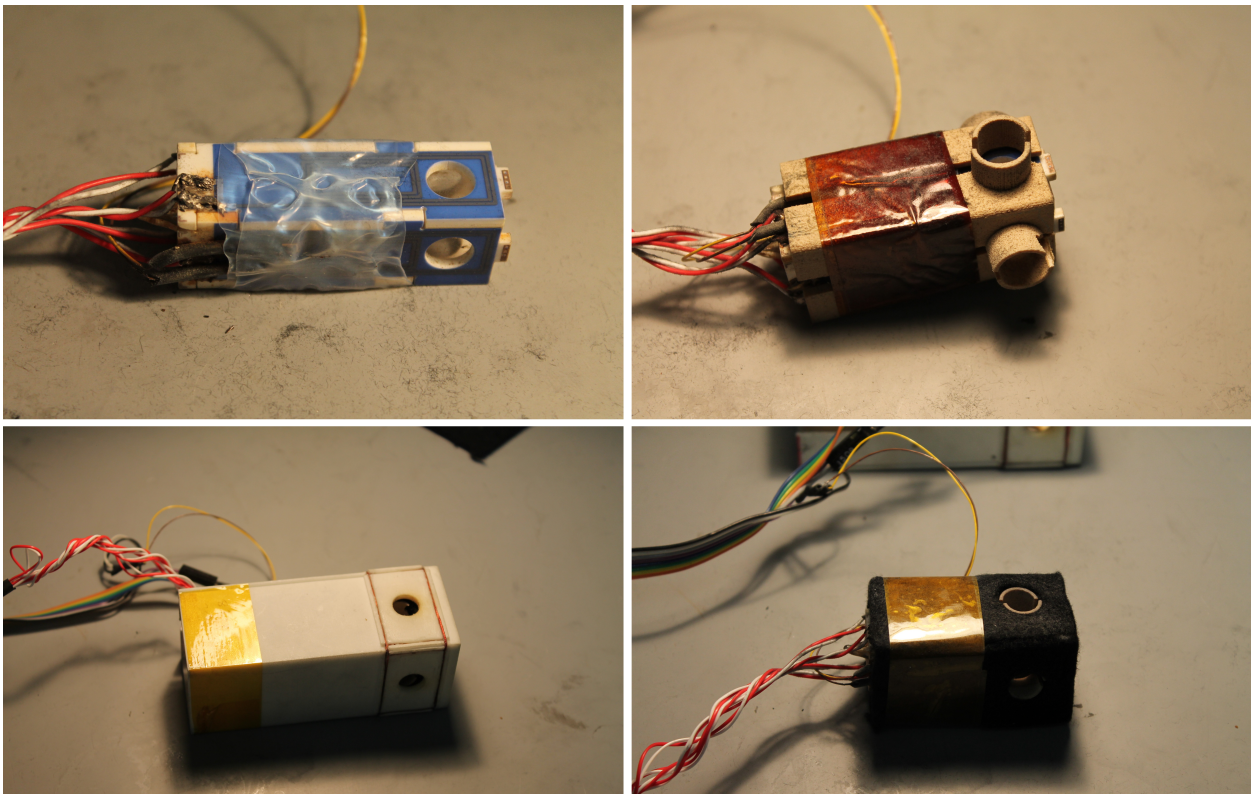


**Figure 3.3:** Assembly of a magnetometer channel (clockwise from top left). First, the heater panels are arranged around the cell in a box formation and two layers of FEP heat-shrink tubing are placed around the assembly for structural support. The circular holes in the heater panels align with the optical paths the longitudinal pump and transverse probe beams take through the cell. The polyether ether ketone (PEEK) shell pieces are arranged around the perimeter of the assembly, and tape is applied to keep the pieces in place. Next, 2-3 layers of $\sim$2 mm-thick Aerogel is wrapped around the assembly. Note how the circular protrusions in the PEEK pieces keep the optical paths clear of the black Aerogel insulation. Finally, the cell and insulation are placed inside a 3-D printed shell which lends additional structural support and provides a mount for the magnetic field coils.

### 3.1.4 Shell-mounted coils

The magnetic field coils mounted on the shell consist of three single-turn pairs of rectangular coils, as shown in Figure 3.4. Previous iterations had used three sets of square coils mounted symmetrically around the sensitive region of the magnetometer. In an effort to move this sensitive region even closer to the magnetic field source while still keeping it centered on the axis of the coils, the horizontal $\hat{x}$- and $\hat{z}$-coils were reduced in their y-dimension. This brought the $\hat{y}$ coils into a pseudo-Helmholtz configuration with the coil spacing equal to roughly half the coil "diameter" (the length of one of edge of the square coil). This provides superb field uniformity in the $\hat{y}$ direction. The $\hat{x}$ and $\hat{z}$ fields produced by this tri-axial set of coils suffer somewhat in their uniformity, though. They are still constant to better than 10% transversely across the entire vapor in a plane halfway between the coil pair.

| Coil set | Side 1 Length [cm] | Side 2 Length [cm] | Separation [cm] | Coil Calibration (*calculated*) [fT/nA] |
|---|---|---|---|---|
| $\hat{y}$ | 3.8 | 3.8 | 2.1 | 42.4 |
| $\hat{x}, \hat{z}$ | 3.6 | 2.1 | 3.6 | 21.1 |

**Table 3.2:** Shell-mounted coil geometries and calculated on-axis coil calibration at the center of the coil pair.

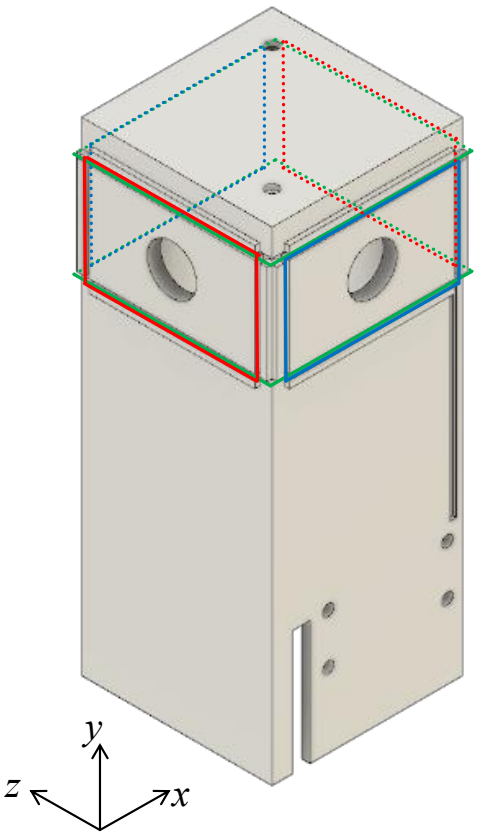



**Figure 3.4:** Magnetometer shell with the shell mounted $\hat{x}$, $\hat{y}$, and $\hat{z}$ coils highlighted in red, green, and blue, respectively. The slot at the bottom allows the heater and RTD leads to pass out of the shell.

## 3.2 Lasers and fibers

To minimize electronic interference from near the magnetometers, the lasers, optics, and accompanying electronics to generate the pump and probe light are placed outside of the magnetically shielded room on an electronics rack. The sources for the probe and pump light are two Eagleyard distributed feedback (DFB) lasers, with the pump centered at 795 nm and the probe at 780 nm. The DFB lasers were chosen for their large temperature coefficient of laser frequency (28.5 GHz/°C) and high output power (100 mW at peak current). The laser is safely tunable over a temperature range of 0-55°C, which is several nanometers of detuning bandwidth. While we don't require nearly this range of detunings for SERF magnetometry, it is a nice convenience.

Despite the high output powers of the DFB lasers, power losses from optics–in particular Faraday isolators (which prevent back-reflections from the fiber from destabilizing the laser diode) and linear polarizers–and fiber coupling left us with insufficient light at the magnetometer array. We therefore decided to increase the laser power by installing two tapered amplifiers (TAs): A Newport 7600 fiber-in-fiber-out amplifier for the pump and a ThorLabs TPA780P20 fiber-in amplifier for the probe. Polarization maintaining fibers couple the light from the DFB laser diodes into the TAs. Free-space output light from the probe tapered amplifier is re-coupled into a fiber launch, while a PM fiber-fiber connector links the pump TA to the light entering the shielded room.

To transmit the light over the $\sim$5 m distance from the electronics rack to the shielded room, we use polarizing fibers purchased from IXfiber. We had long had issues with standard polarization-maintaining fibers giving us poorly polarized or noisy-polarization light at the output, despite our best efforts to couple the light down the fibers' slow axes. The polarizing fibers, on the other hand, strongly attenuate light not coupled along the fiber's slow axis, resulting in better than 40 dB extinction ratio at the output. We verified this ratio a polarization analyzer placed at the fiber output. The fibers connect directly to two- and four-way fiber splitters mounted near the magnetometers. These splitters allow us to minimize the number of fibers running to the magnetically shielded room whiile still allowing us to use multiple magnetometer channels simultaneously. Details of the array of magnetometers will be presented in the next section.

## 3.3 Magnetometer array

The previous iteration of the magnetometer array was comprised of four structurally independent magnetometers, each with its own set of pump and probe lasers, as well as a polarimeter to detect each probe's Faraday rotation. Each channel could be positioned independently of the other three, which gave flexibility to arrange the devices in a variety of spatial patterns. The price paid for the added flexibility was a larger device, which is incompatible with the smaller-sized magnetic shields that we hope to use for patient measurements in the future. We have transitioned to a more compact, permanent setup that creates the framework for the addition of more magnetometers.

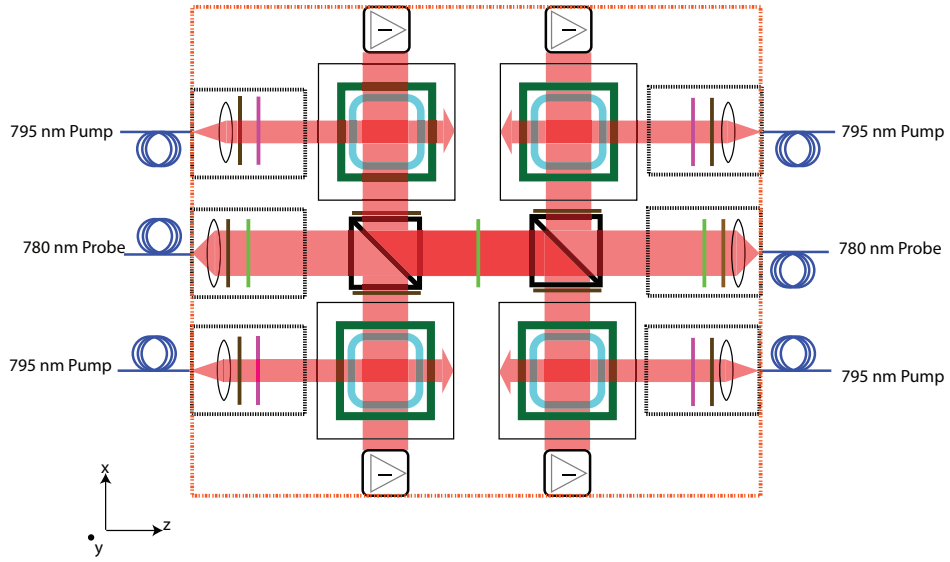



**Figure 3.5:** Arrangement of the magnetometers in the array holder (orange dashed lines). Each cell (light blue) is enclosed in a four-walled PCB heater oven (dark green), surrounded by Aerogel insulation and placed inside a shell. Light is coupled into the array via fiber splitters connected to pump and probe tubes (black dashed lines). All beams are collimated and passed through high-extinction linear polarizers (brown). Pump light is circularly polarized with zero-order half waveplates (magenta). Probe light is passed through half waveplates (bright green), which provides equal division of light intensity between the two arms of polarizing beamslitting cubes. Light passing through the first beamsplitting cube is rotated an additional $90^{o}$, so it is completely reflected by the second cube and sent through the cell. Immediately after the light is reflected off of a PBS, it is passed through a high-extinction polarizer (brown) to convert any remaining rotation noise into intensity noise. Rotation of the probe beam (primarily caused by detected magnetic fields) is detected by polarimeters after the light passes through the cell.

This design allows for magnetometers that are not too far from the heart source that we wish to detect.

With those points in mind, we constructed a new structure to house four magnetometer channels in fixed positions. The four channels are arrayed in a rectangular pattern, as shown in Figure 3.5, with the sensitive regions of the channels separated by 5.0 cm on the short side of the rectangle and 6.8 cm on the long side. The slight radial asymmetry of the arrangement arises from the need for probe laser optics between the two pairs of magnetometers. In addition to the higher degree of compactness offered by this setup, there are two notable advantages:

1. The four-way fiber splitter which was previously required for the probe laser has been replaced by a two-way fiber splitter, which the further splitting of the beam handled by off-the-shelf optical elements. This minimizes our reliance on the splitters, which are notoriously lossy. The pump laser, which still requires a four-way splitter in this setup, experiences a roughly 50% insertion loss despite careful alignment and polishing of the fiber connectors.

2. The pump lasers are arranged in such a way that the excess light from one magnetometer can optically pump its neighbor. This is especially important when the cells are at particularly high temperature and the optical thickness is high. Using diffusion-SERF mode mitigates this issue, as the cells are relatively uniformly polarized. However, empirically we have found that allowing the remaining light from one channel to enter the next improves magnetometer response, suggesting that perhaps the rubidium atoms in the cells are not as well polarized as the theory would suggest.

The structure is connected to a wooden platform, which can then be attached to one of the SQUID devices in our shielded room (Section 3.4), as shown in Figures 3.6 and 3.7. The wooden platform also provides a place for the fiber sockets and RJ-45 ports that interface with the fibers and Cat6A cables that carry light and electronic signals in and out of the room. The wooden structure is somewhat large (a diameter of  75 cm and a height of 30 cm), though its large size is necessary only to match the SQUID that our device is mounted on. The wooden structure can be moved freely after disconnecting electronic and fiber connections - the rest of this section will deal with the components held by the platform.

### 3.3.1 Fiber splitters

To limit the number of optical fibers between the interior and exterior of the magnetically shielded room, we split the incoming light using two polarization-maintaining OZ Optics fiber splitters. For the pump, the same 5-output (1 port is a 5% monitor port while the remaining light is divided among the other four) from the previous iteration of the experiment [Wyllie, 2012]. A

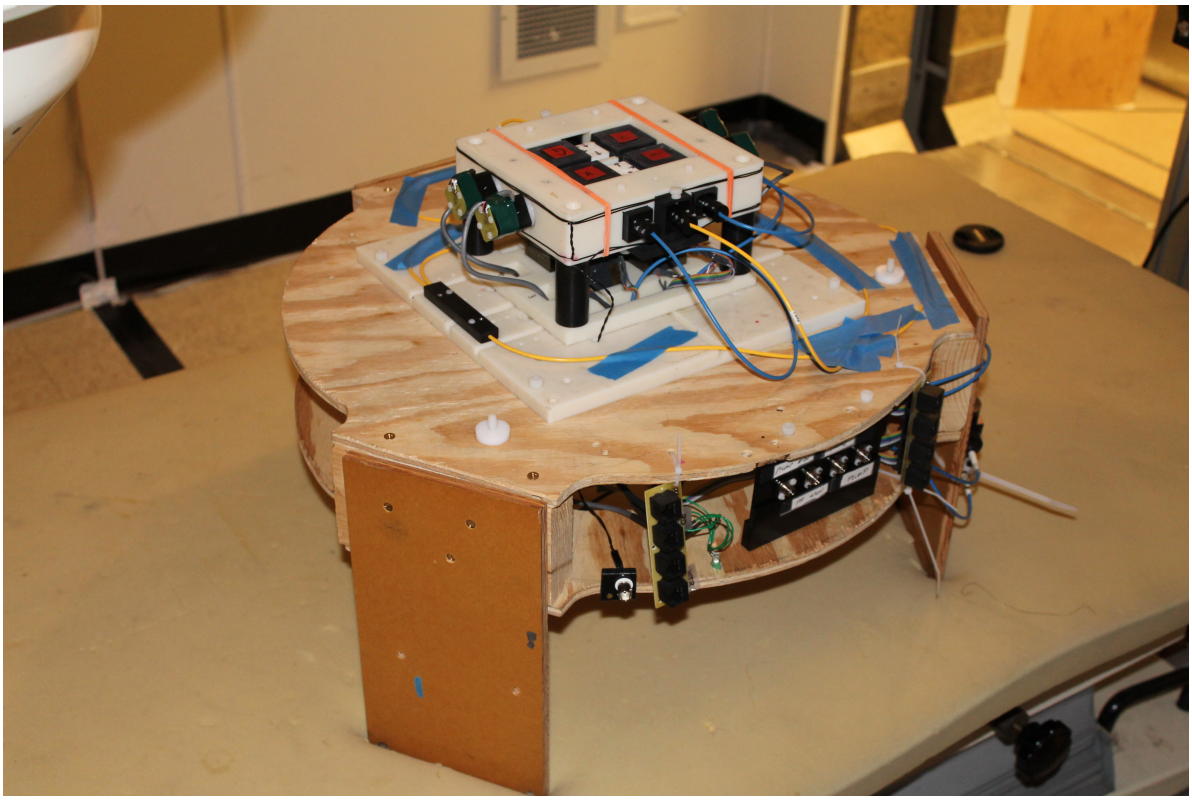


**Figure 3.6:** Photo of the entire apparatus connected to the wood frame, which allows for mounting on the SQUID gradiometer. RJ-45 sockets mounted on the frame provide connections for the ethernet cables (not connected) which carry all electronic signals between the apparatus and the electronics outside of the magnetically-shielded room. Fiber sockets provide interfaces between the polarizing fibers (also not connected) and the fiber splitters which break out the pump (blue fibers connecting to the apparatus) and probe (yellow fibers) laser light.

2-output fiber splitter was employed for the probe beam, with each of those beams further divided by beamsplitters in order to individually probe the four magnetometer channels.

The outputs of the fiber splitters were standard FC/APC connectors that had been known to become magnetized over time. Their close proximity to the magnetometers' detection regions made the sensors vulnerable to stray fields and field gradients produced by these connectors. Thus, custom plastic connectors were designed to hold the fiber ferrule and jacket in place and screw into the optics tubes discussed in the next section. A diagram of the connector highlighting the important parts is shown in Figure 3.8.

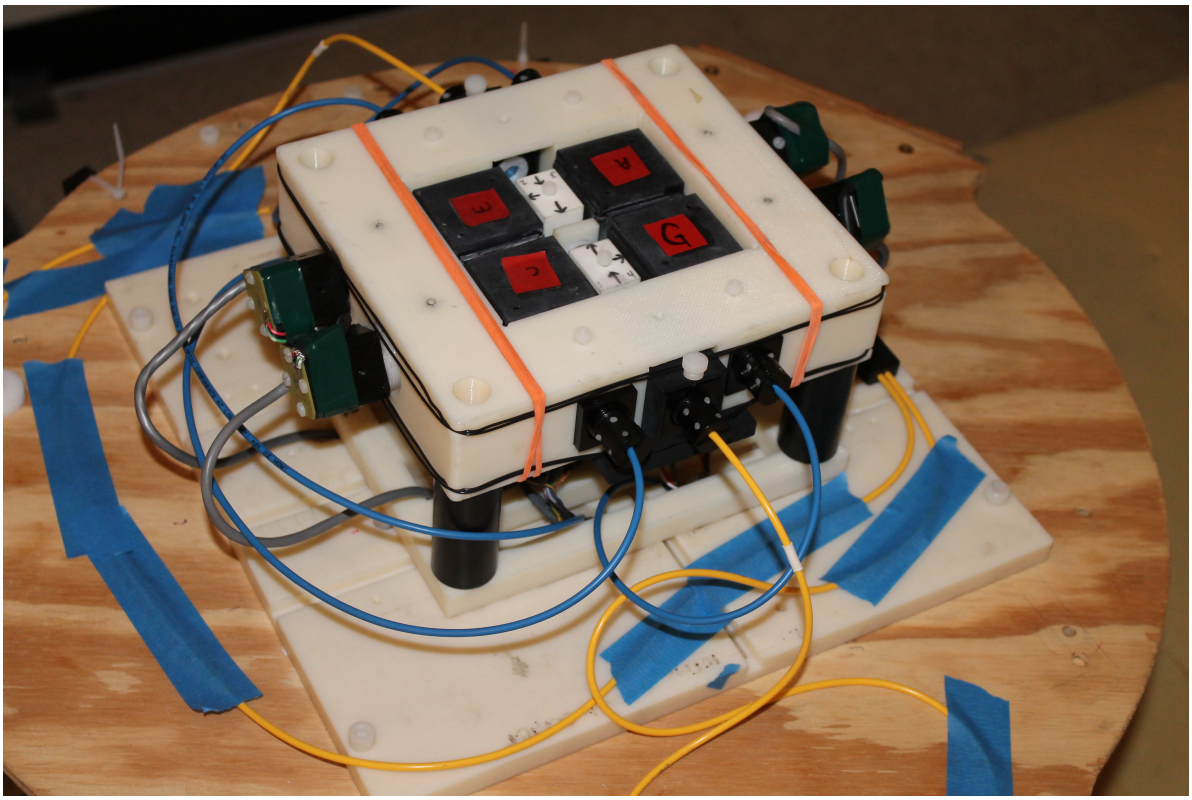



**Figure 3.7:** A more detailed view of the main part of the apparatus. The four magnetometer channels (black shells) are enclosed within a structure which holds the optics: probe tubes (connected to the end of the yellow fibers), pump tubes (connected to the ends of the blue fibers), beamsplitting optics (placed between channels "A" and "E" and between channels "C" and "G") and the polarimeter detectors (wrapped with the green electrical tape, far left and far right)

### 3.3.2 Pump optics

Placing all of the pump optics in removable tubes allows us to assemble several tubes (of differing pump waists) and swap them if new cells with different characteristics are used, or to maximize performance of an individual channel. The optics are mounted in square Delrin "pump tubes" with a ⅞" side length to match the size of the holes. A ⅝"-24 tapped hole is drilled along the long axis of the tube and all optics are placed in custom plastic holders with ⅝"-24 external threads. While using commercially-available metal components would have been easier, the use of these materials near magnetometers can be detrimental to performance, especially if the pieces become accidentally magnetized. A decision was made early on to abandon the standard ThorLabs

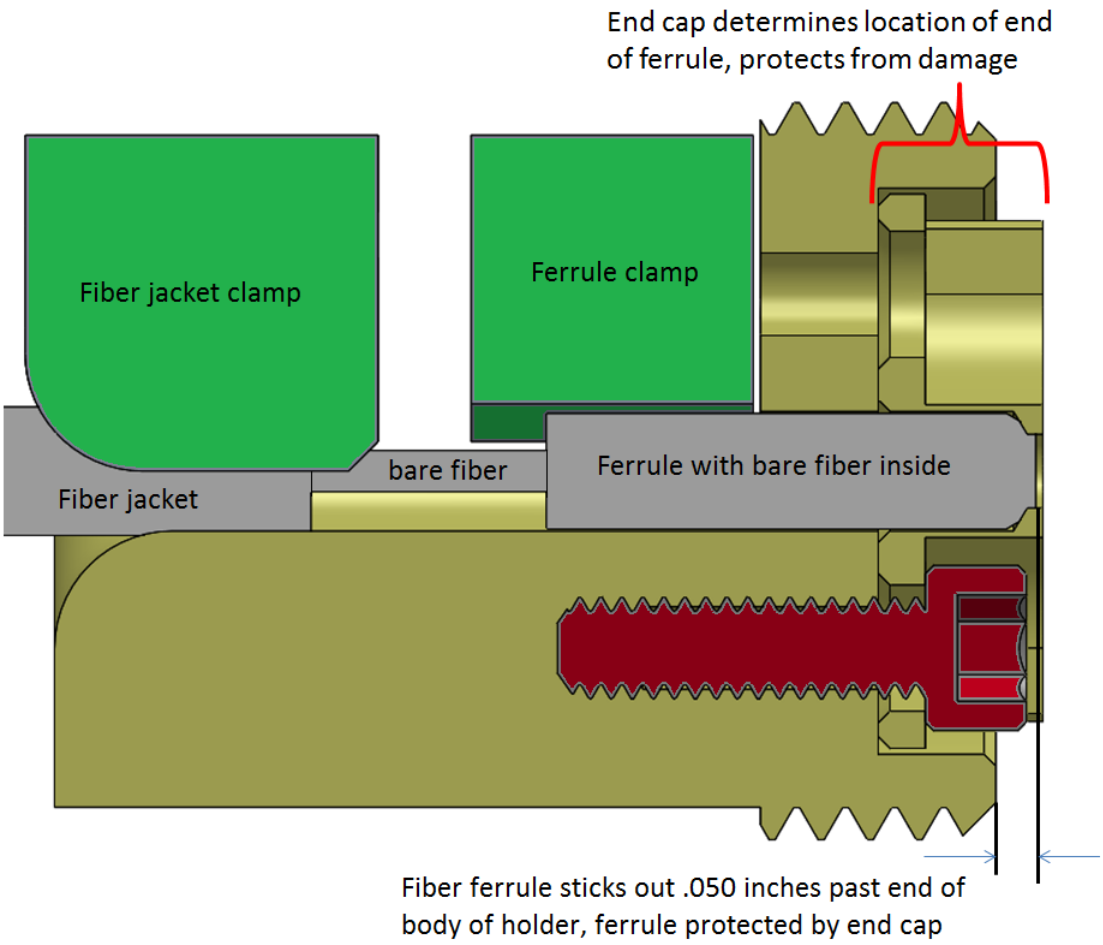



**Figure 3.8:** Fiber and ferrule holder designed by Mike Bulatowicz. Two separate clamps (green) hold the ferrule and fiber jacket, providing strain relief on the vulnerable bare fiber (lavender). A removable end cap allows for easy cleaning or polishing of the fiber ferrule should it become dirty or damaged. The end cap is secured by plastic screws (red) to the end of the main body of the holder.

SM1 and SM05 threads as the high density of threads in these standards prove difficult to use with plastic pieces.

The aforementioned custom fiber connector is inserted into one end of the pump tube and secured by friction against a retaining ring screwed into the tube. To collimate the beam emerging from the fiber, a short focal length lens is gradually threaded down the opposite end of the tube while the laser light is observed on a screen some distance away to collimate the beam. Various lenses were used to achieve beam waists between 0.5-3.0 mm - the larger waists for standard SERF

operation and the narrowest ones for diff-SERF magnetometry. A high-extinction ratio (100,000:1) linear polarizer (ThorLabs LPVIS050) is aligned with the slow axis of the fiber splitter output to maximize transmission and transform any fiber polarization fluctuations into intensity fluctuations. Finally, a zero-order quarter wave plate is inserted and rotated to maximize the circularity of the pump light. The assembled pump tube is then slid into the array holder and secured with a plastic screw.

When the magnetometer is in operation, pump light transmitted through one of the magnetometers (channel 1, *e.g.*) is aligned precisely to also pass through a neighboring cell (channel 2, in this instance). Because of this, it is important that the circularly-polarized light that pumps these magnetometers have the same "handedness", else the excess pump light from the first will actually *depump* the next cell. Determining the handedness of circularly polarized light is somewhat difficult with simple optics and prisms, so the diagnosis is usually done empirically: An index card is placed between the cells to block the light and we evaluate the performance of one magnetometer. The index card is then removed and the performance is measured again. If the magnetometer performs significantly worse with the index card removed, the two pump tubes likely have different handedness. In this case, one tube's quarter waveplate is rotated 90°, giving each tube the same light polarization.

### 3.3.3 Probe optics

Identical square, ⅝"-24 tapped tubes are employed to house the probe optics as well. Again, the fiber connector is screwed into one end of the tube, and a collimation lens (this one with a longer focal length to give the probe beam a 6-8 mm collimated waist) is screwed into the opposite end. Another LPVIS050 polarizer is inserted into the tube and aligned to maximize probe transmission. Lastly a *half* waveplate is inserted into the tube, which allows the for the rotation of the linear polarization of the output light.

As Figure 3.5 illustrates, the two arms of the probe splitter each pump two cells, and the light is divided using two polarizing beamsplitting cubes placed directly between magnetometers 1 and 4 and magnetometers 2 and 3. Using the upper arm as an example, the half waveplate is adjusted

such that the light is incident on the first beamsplitter at a 45$^{o}$ angle. This splits the light into its $s$- (vertical) and $p$- (horizontal) components of roughly equal intensity. The $s$-light is reflected to probe channel 2, while the $p$-light is transmitted. Another half waveplate is placed after the cube to rotate the transmitted beam by 90 degrees, transforming it to pure $s$-light. Thus, upon striking the backside of the second beamsplitting cube, the light is totally reflected and sent to probe channel 1. Because of the relatively low extinction ratio of the reflected beams (20:1 - 100:1 for ThorLabs PBS cubes), additional LPVIS050 polarizers are placed just after the cubes to clean up the linear polarization. After passing through the cells, the light is incident on polarimeters that detect the light's Faraday rotation. Channels 3 and 4 use the same set of optics, just swapping the order in which the light hits each PBS cube.

Because of the longer distance that the probe beams must cover (17 cm from tube output to detector for the longer beam path), instead of being fixed into place like the pump tubes, the probe tubes are placed in an adjustable mount which allows for pitch and yaw adjustments of the beam to ensure that each beam passes through the cell windows and strikes its detector appropriately. The beamsplitting cubes can be adjusted as well. However, each movement (of either the probe tube or PBS cube) changes the alignment of two beams, so an iterative process is usually required to align all of the beams properly.

### 3.3.4 Detectors

To detect the rotation of the probe beam, we employed two different types of polarimeters in rotatable mounts that screw into the array holder. All polarimeters are based on a polarization beamsplitting element which divides up the light according to the relationship

$$I_1 = I_0 \cos^2(\theta) \quad I_2 = I_0 \sin^2(\theta), \tag{3.1}$$

where $I_1$ and $I_2$ are the light intensities output from each arm of the element when light of intensity $I_0$ is incident with a polarization angle $\theta$ relative to the primary axis. Measuring the intensity on one arm would be sufficient to calculate a rotation of the light polarization (assuming a constant

$I_0$), however the sensitivity can be improved by measuring the difference

$$\begin{aligned} I_1 - I_2 &= I_0 \cos^2(\theta) - I_0 \sin^2(\theta) \\ &= I_0 \cos(2\theta). \end{aligned} \tag{3.2}$$

This function is steepest (and hence, the polarimeter will have maximum sensitivity) at $\theta = 45^{\text{o}}$. If we let $\theta = 45^{\text{o}} + \alpha$, Equation (3.2) can be solved for $\alpha$ in terms of the light intensities $I_1$ and $I_2$:

$$\begin{aligned} \frac{I_1 - I_2}{I_0} &= \cos(2\theta) \\ &= \cos(2(45^{\text{o}} + \alpha)) \\ &= \cos(90^{\text{o}} + 2\alpha) \\ &= \cos(90^{\text{o}}) \cos(2\alpha) - \sin(90^{\text{o}}) \sin(2\alpha) \\ &= -\sin(2\alpha) \end{aligned} \tag{3.3}$$

For small values of $\alpha$, $\sin(2\alpha) \approx 2\alpha$, so solving for $\alpha$ yields

$$\alpha \approx \frac{I_1 - I_2}{2(I_1 + I_2)}, \tag{3.4}$$

noting that $I_0 = I_1 + I_2$. These light intensities can be measured using photodiodes, with the photocurrent difference directly proportional to the rotation.

Rotation detector performance depends on the polarization-selective element and the photodiodes which measure the light intensity. Polarimeters based on 10 mm PBS cubes with large-area photodiodes attached directly to the PBS holder provided a cheap, simple solution. However, we found that these PBS cubes suffered from etaloning effects and were extremely sensitive to the angle of the front face of the cube with respect to the incoming beam. Furthermore, the photodiodes (ThorLabs FDS1010) suffered from slow detection time (which translates to poor phase response) due to their large detecting area. Later, we switched to an advanced detector (Figure 3.9), which was comprised of a 20°-separation Wollaston prism followed by a condensing lens. The lens focused the two beams from the Wollaston to two matched photodiodes (Section 3.5.5) attached to a printed circuit board, connected to the back of the detector housing. Rather than

immediately subtracting the photocurrents, the currents were individually carried from the MSR to a box where subtraction or addition of the signals could take place. (The ability to add the photodiode signals was needed specifically for absorption or other transmission measurements.) Initial concerns of increased noise from wiring both photodiode signals out of the MSR (rather than immediately subtracting them) were allayed by repeated measurements that showed no degradation in performance.

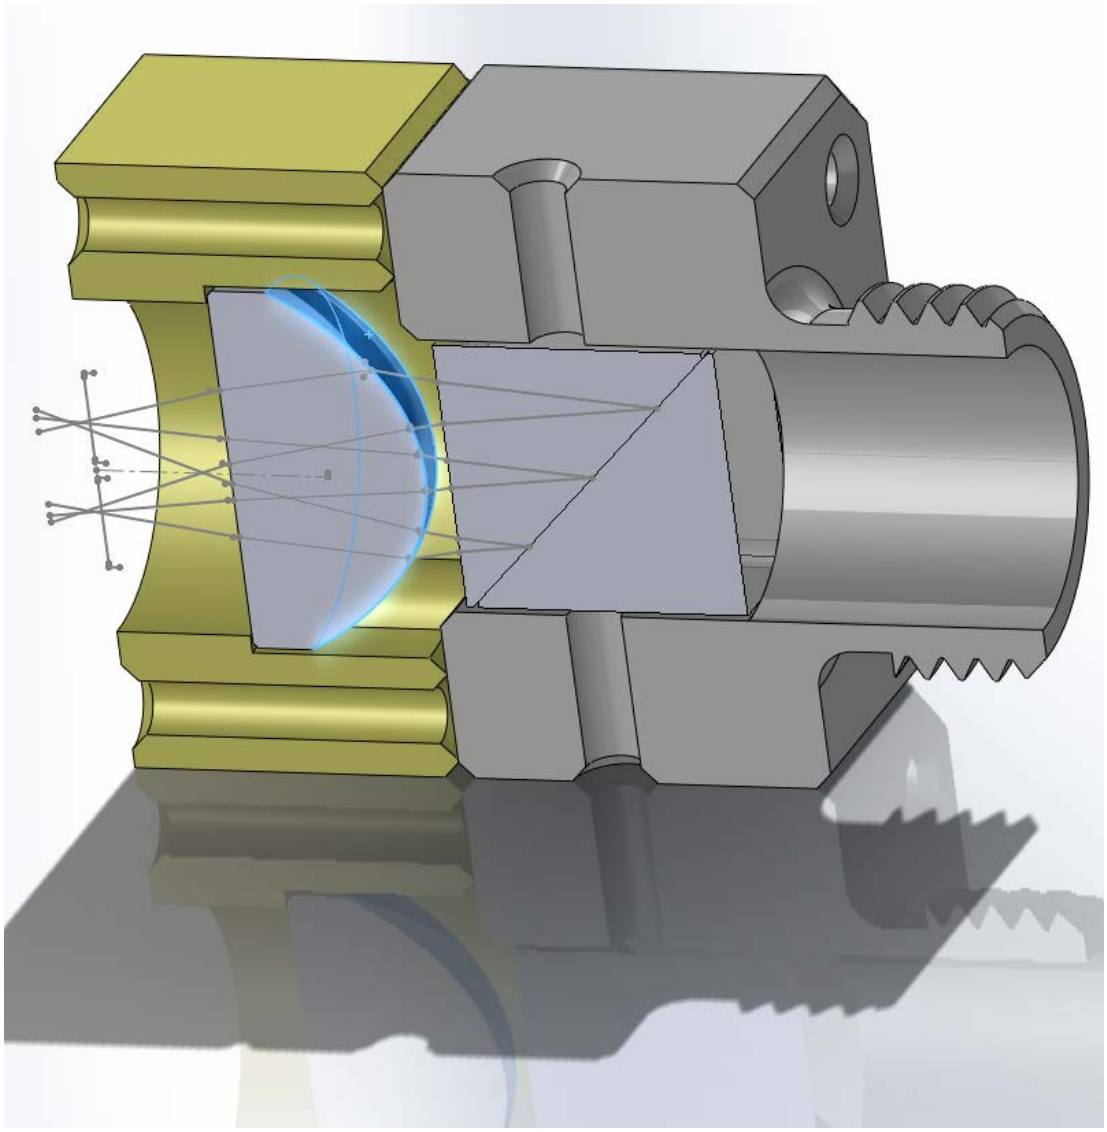


**Figure 3.9:** The advanced detector designed by Mike Bulatowicz. The linearly-polarized light enters through the right side and the $s$ and $p$ components are separated at a 20$^{o}$ angle by the Wollaston prism. The light is refocused by a condenser lens to a pair of photodiodes placed at the location of the focused light, indicated by the points where the rays cross. Not shown are the set screws which hold the optics in place. The detector screws directly into the tapped hole in the array structure (Figure **??**). Drawing courtesy of MB.

### 3.3.5 Feedback coil

In addition to the room and shell-mounted magnetic field coils, we have an additional set of coils wrapped around the array structure. Unlike the other coils in the experiment, these only generate fields in the $\hat{y}$ direction, which is the sensitive direction for DC-mode. These coils are used for common-coil feedback, as explained in Section 3.5.6, and are necessary when the array must be tilted for an MCG measurement and the $(\hat{x}, \hat{y}, \hat{z})$ axes of the large room coils no longer align with the $(\hat{x}', \hat{y}', \hat{z}')$ axes of the array. Because these coils are not used for calibration, their precise current-to-field coefficient has not been measured.

Similar "common" coils for $\hat{x}$ feedback would have been a boon for Z-mode operation, however their structure would have extended in the $\hat{y}$ direction and limit how close we could place our sensor to the adult or fetal heart of interest.

## 3.4 Magnetically shielded room

The nature of the SERF magnetometers demands zero- or near-zero-field operation; thus it is not possible to operate these magnetometers in an environment dominated by the earth's $\sim 40\,\mu$T field. Large magnetic fields can be generated by coils to counter the earth's field, but the simplest method for removing the earth-scale fields near the magnetometers is working in a magnetically shielded environment. The shielding has the added benefit of removing much of the environmental noise that is a detriment to MCG measurements.

The magnetically shielded room (MSR) in which all measurements for this thesis were made is located at the Wisconsin Institutes for Medical Research building on the UW campus. The room is constructed from a dual layer of $\mu$-metal alloy and has an interior dimension of approximately 3 m x 3 m x 3 m. Access to the room is provided by a hydraulically-operated door that is closed during a measurement. Two small (one 20 cm x 20 cm, the other 50 cm x 30 cm) portholes in the room provide access pathways for electrical signals and light for our experiment, as well as for the two SQUID gradiometers located inside of the MSR. One of the SQUIDs is still in use by

a biomagnetism research group; the other is currently inactive, but its ceiling-mounted gantry is used to position our apparatus during the measurement of an adult or fetal subject.

### 3.4.1 MSR magnetic field coils

To eliminate residual DC fields inside the MSR (which are generally on the 10-200 nT level after the door to the room is closed), magnetic field-canceling coils are wrapped around the edges of the interior of the room. For the two horizontal axes of the room (referred to as $\hat{x}$ and $\hat{z}$ to match the probe and pump directions of the magnetometer units), a ten-turn coil was wrapped along the midpoint of the room. The presence of the hydraulic door made an origin-centered N-turn coil impossible in the $\hat{y}$ direction, so a pair of ten-turn coils arere wrapped at the top and bottom of the MSR interior, as shown in Figure 3.10

The coils were calibrated using a three-axis DC fluxgate magnetometer placed at the center of the room and applying known currents to each set of coils. The DC axial calibration values are

$$\begin{aligned} \frac{dB_x}{dI_x} &= 5.16\,\text{fT/nA}, \\ \frac{dB_y}{dI_y} &= 8.15\,\text{fT/nA}, \\ \text{and} \quad \frac{dB_z}{dI_z} &= 5.22\,\text{fT/nA}, \end{aligned} \tag{3.5}$$

which agree well with values predicted by a Biot-Savart prediction of the coils' field strengths.

Because these coils sit so close to the large $\mu$-metal sides of the room, the applying magnetic fields with these coils induces eddy currents in the shields. These eddy currents obey Faraday's Law and in turn generate magnetic fields that reduce the strength of the original field from the coils. Because of the inductive nature of these fields, the true coil calibrations decrease with increasing frequency, as both the induced current and resulting eddy fields increase in amplitude.

The true magnetic field calibration was deduced using a differential calibration between magnetic fields applied to small ($\sim$3 cm), shell-mounted coils of an atomic magnetometer unit (which were assumed to induce zero eddy currents in the MSR walls) and fields applied to the same magnetometer using the large room-mounted coils. Only the vertical $\hat{y}$-coils were calibrated in this manner, though we would expect to see similar effects in the other two directions. The resulting

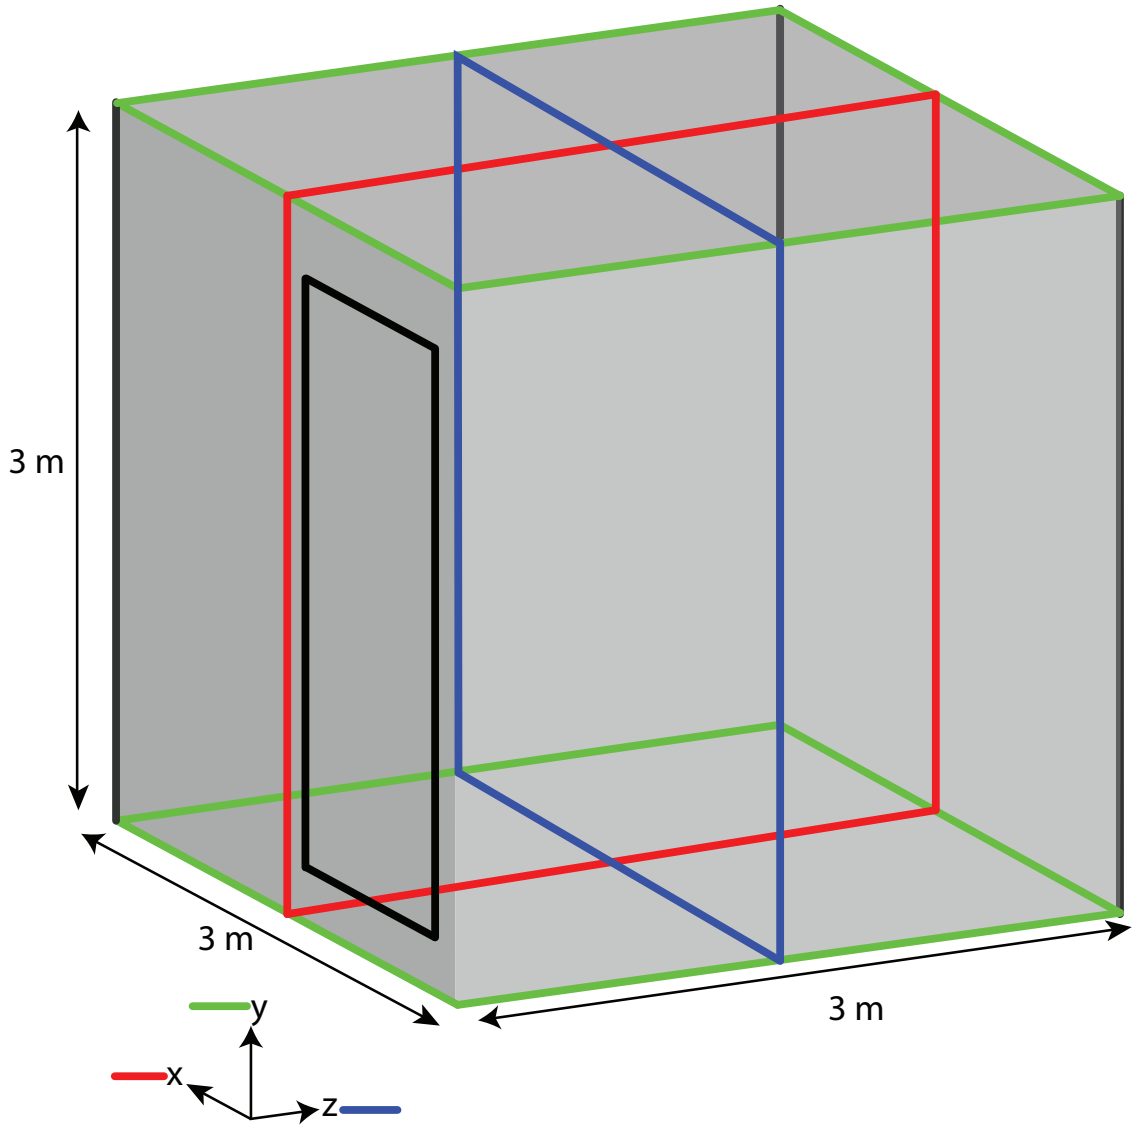



**Figure 3.10:** Cartoon drawing of the magnetically shielded room, with the locations of the room-mounted coils indicated (red, $\hat{x}$ coils; green, $\hat{y}$ coils; blue, $\hat{z}$ coils). The approximate position of the door is indicated by the black rectangle.

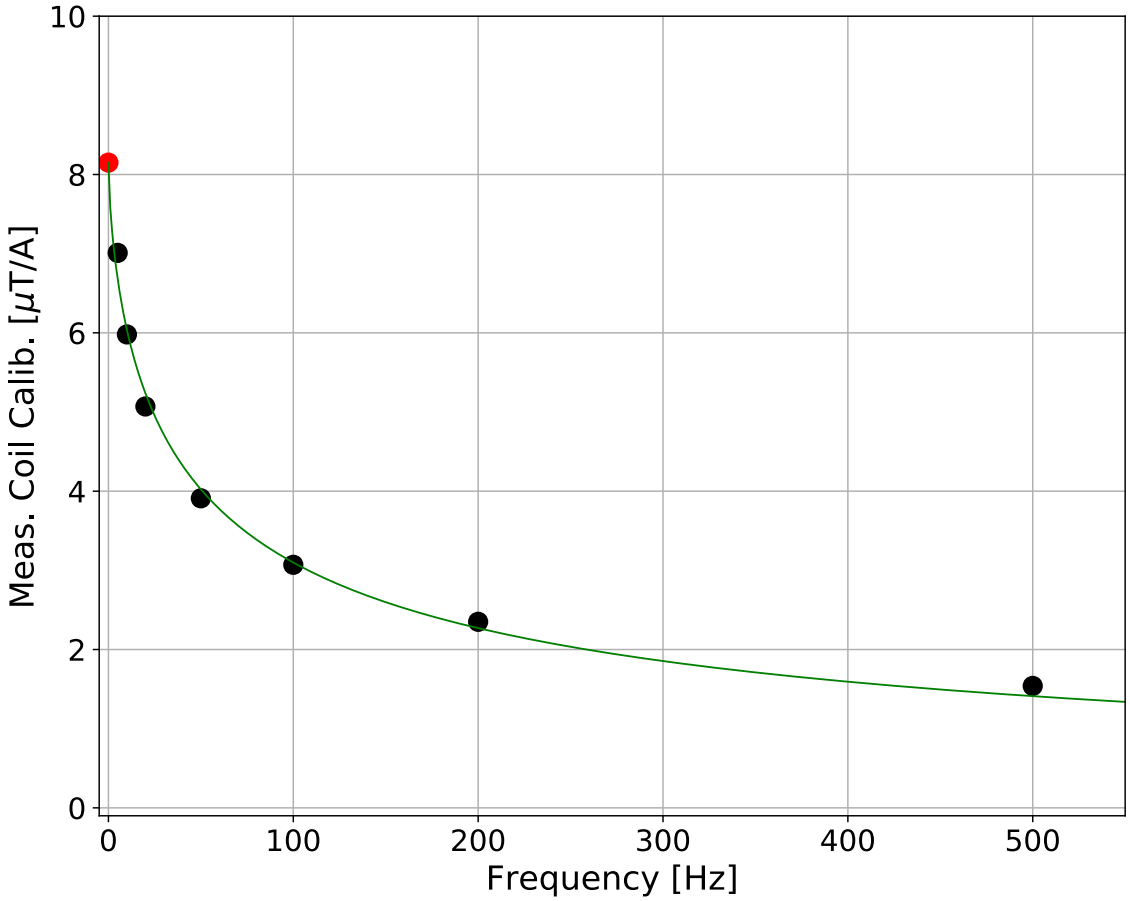



**Figure 3.11:** Frequency-dependent calibration of the MSR-mounted Y coils. The red data point was obtained with a DC fluxgate, while the remaining values were determined by a differential calibration. All values agree well with the calibration given by Equation (3.6).

calibration, shown in Figure 3.11, was curve-fit using a power function given by

$$\frac{dB_y}{dI_y}(f) = \frac{8.15\,\text{fT/nA}}{1 + (.0752\,\text{s})f^{.668}}. \tag{3.6}$$

As our magnetometer calibration is normally done with the shell-mounted coils, this calibration is not particularly critical. However, during the gradiometery studies described in Chapter 5, calibration fields were simultaneously applied to both gradiometery channels using these large coils and an exact calibration was paramount.

## 3.5 Electronics

The electronics required to run a single atomic magnetometer are six-fold: cell heaters to warm the cells, laser current drivers and thermoelectric cooler (TEC) controllers, magnetic field coils and coil drivers to apply nulling and calibration fields, photodiodes and current amplifier circuitry to measure the probe Faraday rotation, analog feedback circuits that to counteract external field drifts in the shielded room, and an analog to digital converter (ADC) system to acquire the data. Good analog electronics practice would suggest putting these devices as close to the device that the electronic component is sensing or driving as possible. However, the low-magnetic field, small magnetic gradient, and low noise requirements for SERF operation - as well as the requirement that the entire device be portable - makes this type of setup impossible. Instead, we mount all of the electronics on a rack placed outside of the room and send and receive all signals from the magnetometer array along shielded twisted pair Cat6A ethernet cables whose shields are grounded to the rack outside of the room. The shielded cables are needed to reduce electronic pickup, especially from the large amplitude signals that drive the heaters.

### 3.5.1 Heater drivers

Each of the ovens described in Section 3.1.2 have a resistance of approximately 50 $\Omega$. To maintain the 145-160°C temperatures required for SERF operation, 5-8 W of heater power is required, or 15-20 VDC. Despite the efforts to limit heater-produced magnetic fields in the heater panel design, DC heating is inadvisable due to the potential low-frequency magnetic noise production in close proximity to the measurement region. Instead, an AC voltage is applied to the heaters with a frequency much larger than the atomic Larmor precession that would be generated by the heater, *i.e.*, $\omega_H \gg \Omega_H$. To satisfy this requirement, we drive the heaters sinusoidally at a frequency of $\omega_H = 2\pi(150\,\text{kHz})$ or more. Driving the heaters at 15-20 $V_{rms}$ is non-trivial, as many conventional op-amps lack either the gain-bandwidth product or the current output levels ($\sim$500 mA) to meet our demands. We therefore designed an circuit around the Apex Electronics PA107, a high-current, high-voltage op-amp with a slew rate of 3000 V/$\mu$s. The circuit is described in detail in Appendix

A. Unfortunately, these circuits lack easy compatibility with the software feedback loop used in the past iteration of the device. Thus, we use a simple "turn-on-and-wait" strategy, and the heat-up time from room to operation temperature is generally on the order of 45 minutes. A new feedback program should be implemented in the future to shorten this warm-up time.

### 3.5.2 Laser drivers and TEC controllers

The DFB laser diodes are driven and temperature-controlled with Stanford Research Systems LDC501 controllers. The keypads on these devices allowed us to set precise drive currents and diode temperatures. This precision was helpful for making repeatable settings for the DFB diodes whose output characteristics are very sensitive to these two parameters. These SRS controllers are also programmable, which allow us to create computer programs to sweep the laser detunings and powers back and forth, expediting the time it takes to find the optimal settings for magnetometry. Pump and probe laser optimization is discussed in more detail in Chapter 4.

### 3.5.3 Coil drivers

The magnetic field coils in this experiment are connected to coil drivers that allow us to zero the magnetic fields at the locations of the magnetometers and to apply modulation fields for nulling, Z-mode, and calibration purposes. The coil drivers in this work represent two generations of devices, each named after its illustrious designer: An older circuit colloquially referred to as the "Wyllie supply" (designed by Bob Wyllie and detailed in Appendix E of his thesis [2012]), and a circuit designed by Ibrahim Sulai (known as the "Sulai supply"), a schematic of which is located in Appendix A. The main function of each of these circuits is to sum a DC input voltage with an AC modulation voltage and apply this voltage to a stable current supply with selectable output resistors which can be changed to suit the current demands for the fields being applied.

There are two principal differences between the circuit designs. The motivation to transitioning to the Sulai supplies is described below:

1. To reduce the footprint of the coil drivers, three individual driver circuits (connected to the $(x, y, z)$ magnetic field coils) were placed on one board. To accomplish this, the Sulai supplies do not incorporate the DC and AC attenuation features as well as their high-current components found on the Wyllie supplies. With a Wyllie supply, applying a 10 V signal across the 50 Ω output resistors requires a component able to source 200 mA, a requirement satisfied by the use of BUF634 buffer chips. Without these buffer chips, the new supplies are limited by the 30 mA current limit of the TL074 op amps in the circuit. To account for this change, the 50 Ω output resistors from the Wyllie supplies were replaced by 200 Ω resistors to better match the current characteristics of the new boards. The 200, 500, 1,000, and 5,000 Ω resistor selections were sufficient to apply all types of fields required for magnetometry experiments.

2. The source of applied DC voltage on the Wyllie supplies is two LM399 voltage references and a potentiometer which create a tunable voltage across a 14 V (-7 V to +7 V) range. The Sulai supplies rely on an external DC voltage input, which is then added to the AC modulation input. The source of DC voltage is now a Maxim MAX5735 16-bit Digital to Analog Converter with a -5 V to +5 V range. The DAC is controlled by an Arduino Mega microcontroller and accompanying software. The motivations for this change will be discussed fully in the next section.

Limiting the current noise output by the coil drivers is of utmost importance. Current noise produced in the coil driver circuits will become magnetic field noise when the current noise is driven through the magnetic field-nulling coils. For the shell-mounted $\hat{y}$-coils, the calibration is ∼40 fT/nA. This means the output current noise must be limited to 25 pA/$\sqrt{\text{Hz}}$ to attain a 1 fT/$\sqrt{\text{Hz}}$ magnetic noise noise level. This proved difficult, even with carefully selected low-noise components. Studies of the current noise on the new current supplies resulted in equivalent magnetic field noise levels in the 1-10 fT/$\sqrt{\text{Hz}}$ range near DC as long as the output resistors were at the 1 kΩ or 5 kΩ level, as shown in Figure 3.12. Noise increased linearly with decreasing output

resistance, leading us to believe we are limited by voltage noise. The exact source of that voltage noise, however, is elusive.

The large output resistors required to keep field noise small makes it impossible to null fields larger than $\sim$100 nT. Fortunately, the magnetically shielded room provides this type of low-field environment, and an additional set of large coils wrapped around the perimeter of the room (connected to a Wyllie supply) eliminate much of the remaining Earth field from the room (see Section 3.4.1 for more information on these coils). Thus the shell coils and coil drivers are only responsible for eliminating remaining field gradients and light-shift fields, normally a maximum of 15 nT.

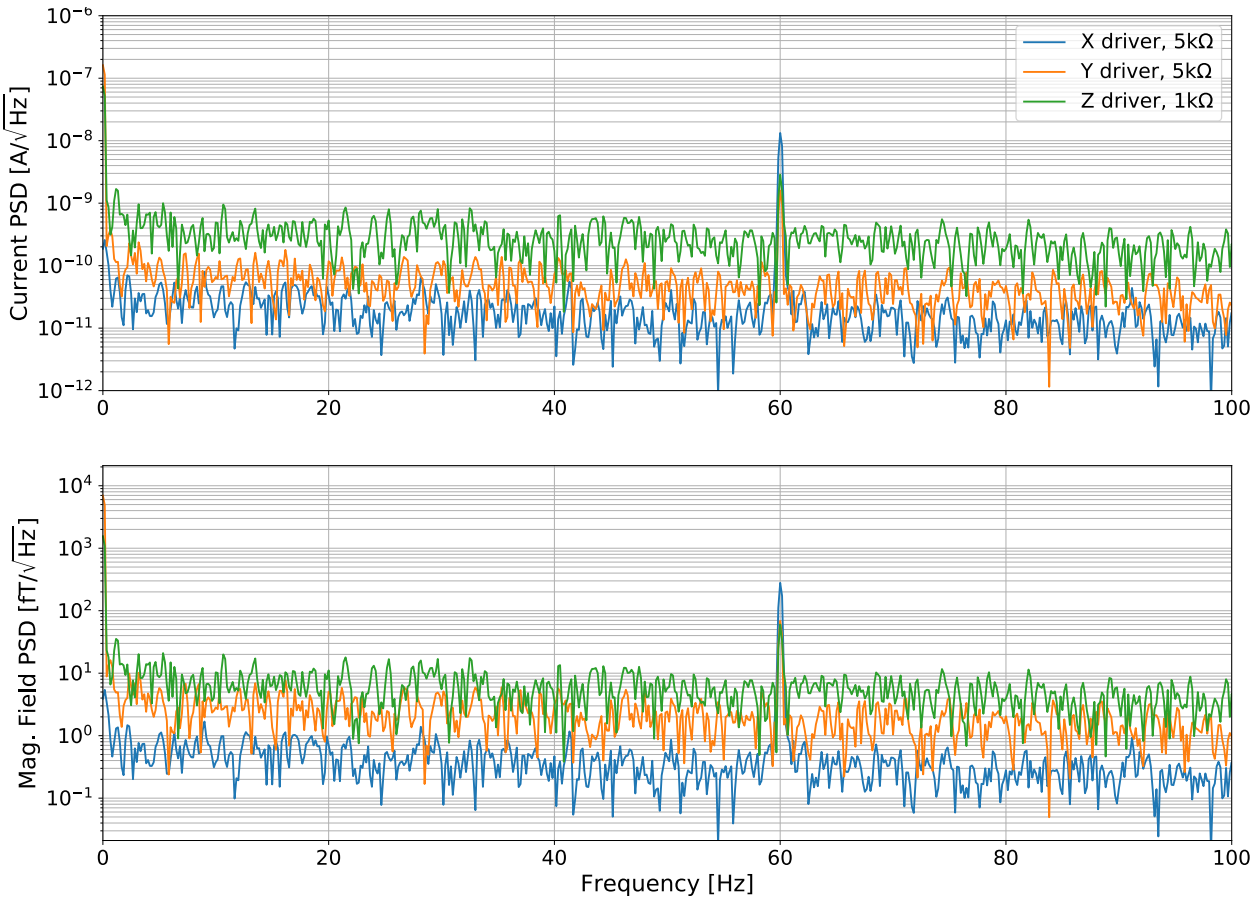



**Figure 3.12:** Power spectral densities (PSD) or the noise generated by the coil drivers. The upper plot contains the current noise output by the Channel 1 shell coil driver circuits with typical output resistors. (Note the larger output resistor for the $\hat{z}$ coils necessary for application of the large z-mode fields.) The bottom plot scales each current noise by the appropriate shell coil calibration (Table 3.2) to obtain equivalent magnetic field noise. The generated noise is $<$10 fT/$\sqrt{\text{Hz}}$ at all frequencies for the $\hat{x}$ and $\hat{y}$ coils, which represent the two directions in which we are able to detect magnetic fields. The elevated noise from the $\hat{z}$ coils would only be problematic if the magnetic field nulling is poor and the magnetometer becomes sensitive to $B_z$.

### 3.5.4 Magnetic field decoupling

#### 3.5.4.1 Motivation

By decreasing the footprint of the array and making the magnetometers closer together, we inevitably increase the crosstalk between the units. This is not particularly noticeable from a noise standpoint, as the magnetic noise generated by the magnetic field coils in each sensor is $<$10 fT/$\sqrt{\text{Hz}}$ , and the distance between the units means the noise generated at a given magnetometer by its neighbor is $<$1 fT/$\sqrt{\text{Hz}}$ . However, the application of local DC magnetic fields to null the field at one magnetometer also alters the local field of a nearby magnetometer. If that magnetometer had previously been nulled, it must be "re-nulled" before taking a measurement...which slightly alters the DC field at the other magnetometer, and so forth. This iterative process becomes even more time-consuming when four channels must be simultaneously nulled with this procedure. For this array to be a truly feasible fMCG measurement device, spending 10-15 minutes nulling magnetic fields while the patient waits is completely unacceptable.

We have devised a way to counteract this problem. Using straightforward Biot-Savart Law calculations, we calculate the magnetic field coupling terms created by each of the three sets of shell coils wrapped around a magnetometer channel as a function of position:

$$\frac{\partial B_i}{\partial I_j} = \frac{\mu_0}{4\pi}\int_C \frac{d\ell \times \mathbf{r}'}{|\mathbf{r}'|^3} \quad i,j = \{x,y,z\}, \tag{3.7}$$

where $B_j$ is the magnetic field created by a current $I$ running through the $j$ set of magnetic field coils, $d\ell$ is the differential current segment, $\mathbf{r}'$ is the displacement vector from that segment of current to the point in space where the coupling term is evaluated, and the integral runs through the entire coil pair $C$. We then evaluate each one of the $\frac{\partial B_i}{\partial I_j}$ terms at four points in space: At the center of the the magnetometer generating the fields (equivalent to the calculated coil calibration terms in Table 3.2), and at the locations of the three other magnetometers in the array. We then repeat this process for the other three magnetometers, though because of the symmetric nature of our array, the magnitude of terms are identical in the other magnetometers. With four channels, four positions, and nine $\frac{\partial B_i}{\partial I_j}$ terms per channel, this gives a total of 144 coupling terms. We arrange

these terms into a matrix $\mathbf{M}$, such that

$$\mathbf{M} = \begin{bmatrix} \frac{\partial B_{1x}}{\partial I_{1x}} & \frac{\partial B_{1x}}{\partial I_{1y}} & \frac{\partial B_{1x}}{\partial I_{1z}} & \frac{\partial B_{1x}}{\partial I_{2x}} & \cdots & \frac{\partial B_{1x}}{\partial I_{4z}} \\ \frac{\partial B_{1y}}{\partial I_{1x}} & \frac{\partial B_{1y}}{\partial I_{1y}} & \frac{\partial B_{1y}}{\partial I_{1z}} & \frac{\partial B_{1y}}{\partial I_{2x}} & \cdots & \frac{\partial B_{1y}}{\partial I_{4z}} \\ \frac{\partial B_{1z}}{\partial I_{1x}} & \frac{\partial B_{1z}}{\partial I_{1y}} & \frac{\partial B_{1z}}{\partial I_{1z}} & \frac{\partial B_{1z}}{\partial I_{2x}} & \cdots & \frac{\partial B_{1z}}{\partial I_{4z}} \\ \frac{\partial B_{2x}}{\partial I_{1x}} & \frac{\partial B_{2x}}{\partial I_{1y}} & \frac{\partial B_{2x}}{\partial I_{1z}} & \frac{\partial B_{2x}}{\partial I_{2x}} & \cdots & \frac{\partial B_{2x}}{\partial I_{4z}} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial B_{4z}}{\partial I_{1x}} & \frac{\partial B_{4z}}{\partial I_{1y}} & \frac{\partial B_{4z}}{\partial I_{1z}} & \frac{\partial B_{4z}}{\partial I_{2x}} & \cdots & \frac{\partial B_{4x}}{\partial I_{4z}} \end{bmatrix} \tag{3.8}$$

Then we can write

$$\mathbf{B} = \mathbf{M}\,\mathbf{I} \tag{3.9}$$

where $\mathbf{B} = [B_{1x}\ \ B_{1y}\ \ B_{1z}\ \ B_{2x}\ \ \ldots\ \ B_{4z}]^T$ are the magnetic fields produced in the three directions at all four sensors by applying currents $\mathbf{I} = [I_{1x}\ \ I_{1y}\ \ I_{1z}\ \ I_{2x}\ \ \ldots\ \ I_{4z}]^T$ to each magnetometer's set of three coils. In this way we can see how applying a current to, for example, the $\hat{z}$-coils of magnetometer 2 ($I_{2z}$) changes the $\hat{x}$ field of magnetometer 4 ($B_{4x}$). The fields are coupled together in a way such that changing one current may change many fields, and this is what we observe.

Alternatively, we can invert the matrix $\mathbf{M}$ so that each term now takes the form $\frac{\partial I_i}{\partial B_j}$ and write

$$\mathbf{I} = \mathbf{M}^{-1}\,\mathbf{B} \tag{3.10}$$

In this equation, the *independent* variables are the magnetic fields of each magnetometer and the *dependent* variables are the currents that create these fields. This allows us to *independently set the fields at each magnetometer*. In other words, setting a component of the magnetic field at some location to a particular value is now accomplished by adjusting the currents in several magnetic field coils. This is significant because it allows us to set one field in a way that *does not affect the fields at any other part of the array*.

### 3.5.4.2 Implementation

To accomplish this in practice, we wrote a program using Arduino software that realizes the theory above. The Arduino microcontroller takes input signals from 3 rotary encoders (representing $B_x,\ B_y,\ B_z$) and a 4-position switch which selects the magnetometer channel being adjusted.

With each click of the encoder, the program outputs an array of twelve 16-bit numbers which correspond to the twelve currents necessary to produce the desired fields. These 16-bit numbers are sent to the MAX 5235 DAC which scales them to a -5 V to +5 V range. These voltages are in turn sent to the coil drivers and output to the shell-mounted coils on the magnetometer array. Because the Arduino is controlling a voltage rather than a current, adjustments must be made if the output resistors are not all the same. The code makes allowances for this, though, and varies the coupling coefficients appropriately if different output resistors are used. An LCD screen connected to the the Arduino displays an array of twelve 16-bit numbers which correspond to the currents being applied to the twelve shell coils (four channels, three axes per channel). During the nulling process (Section 4.3), the magnetometers are nulled sequentially, and the nulling of any individual magnetometer does not disturb the field at the locations of its neighbors.

The implementation of this decoupling software has made for tremendous gains in the practicality of the device. Tests done by the author and his colleagues have shown a factor of 5 or more decrease in the time spent nulling the four channels of the array (from 12-15 minutes down to 3 minutes or less). This shortened time is critical when a patient is actually being measured and is necessary if this device is to ever see widespread clinical use.

### 3.5.5 Photodiodes and amplifiers

The photodiodes used for Faraday rotation detection are a matched pair provided by First Sensor. The diodes are mounted on a printed circuit board and attached to the end of the detector described in Section 3.3.4. Each of the four traces (anode and cathode for each diode) are wired to a Cat6A cable and carried out of the MSR. This is a departure from the previous iteration of the photodiode detection which subtracted the two photocurrents and sent only the difference signal to the detection electronics. The new arrangement allowed us to quickly switch between measuring the photocurrent sum (for probe laser coupling and probe absorption/transmission) and photocurrent difference (for normal magnetometry).

Transmitting μA signals over the $\sim$10 m distance from the array to our detection electronics does not follow the best electronics practices [Hobbs, 2000], which would suggest amplifying and

recording signals as spatially close to the source as possible. However, the limitations mentioned in Table 3.1 (specifically, to keep noisy electronics out of the magnetically shielded room and to keep the array portable) do not allow such a setup. We have found that using shielded-twisted-pair Cat6A ethernet cables provide a low-noise solution to the dilemma. The shielding on all of these cables is non-magnetic copper foil, and that shield is grounded outside of the MSR to provide additional noise attenuation. The twisted pairs themselves have varying turn ratios to limit cross-talk between channels. A single 8-conductor cable is sufficient for carrying two detectors worth of photocurrents, so two cables are required for all of the magnetometer detector signals. Identical Cat6A cables are used to carry heater voltages (one cable for all four channels) and local nulling coil currents plus RTD readings (one cable per magnetometer channel).

To convert the photocurrents to voltages we use Stanford Research Systems SR570 Low-Noise Current Preamplifiers. The preamps include amplification stages ranging from a few pA/V all the way up to 1 mA/V, and have built-in 1- and 2-pole lowpass filters, which we can set to avoid aliasing during data collection. For magnetometery measurements, typical amplifications are in the 500 nA/V to 2 μA/V range, while sum measurements use much lower amplification (100-500 μA/V) as to not saturate the amplifiers. Anti-aliasing lowpass filters are set to the smallest possible frequency that will not attenuate signals in our band of interest. For DC-mode measurements, a 300 Hz 2-pole filter is applied, while in Z-mode (which requires an un-attenuated 1 kHz modulation signal be passed), a 10 kHz 2-pole filter is implemented. The application of these filters has a noticeable effect on the phase of our measurements, an effect which is critical when operating the magnetometers in feedback mode.

### 3.5.6 Feedback circuitry

Active feedback, a theoretical basis for which is provided in Section 2.6, is a critical part of the reliable operation of our magnetometer array. The benefits of operating the magnetometers in a feedback mode is especially evident when operating the magnetometers in gradiometric operation, which will be discussed in Chapter 5. Over the course of this work, we have applied active feedback

in two different ways: using the signal of a single channel to feedback to a large common coil wrapped around the entire array, and using each channel's signal to feed back to itself.

The common-coil feedback was used mostly to eliminate the large, slow, approximately uniform (on the length-scale of our magnetometer array) field-strength drifts that result from far-away ($<$5 m) sources such as elevators traveling up and down and MRI machines in other labs turning on and off. Magnetic fields created by these sources generally have frequency components $<$1 Hz, meaning that they don't overlap with the 1-80 Hz bandwidth of interest for magnetocardiography. However, the Faraday rotation induced by these fields can cause currents in our detectors which saturate the SR570 preamps, which is catastrophic during a measurement. Thus it's important to counteract these fields with feedback. To control the common coil, we designed a simple integrator circuit with an $RC$ value of $2\pi$(0.5 Hz). The output of one of the magnetometers was amplified, sent through the integrator, and sent back to the coils with the correct polarity to create negative feedback. Because the bandwidth of this feedback does not overlap with the band of the cardiographic signal, we still record the raw signals from the magnetometer, rather than "correction" signal which is fed back to th common coil.

The approach changes somewhat when we implement individual channel feedback. In this mode, we use a separate Proportional+Integral feedback circuit (diagrammed in Appendix A) to control each channel. The bandwidth of the feedback for this loop can be made arbitrarily large: in this work, the bandwidth was made to be 500 Hz to far exceed the 1-80 Hz bandwidth of interest. The application of this broadband feedback effectively makes all of the magnetometer channels self-zeroing. Instead of recording the magnetometer signal which, outside of noise, is identically zero in time, we record an amplified copy of the feedback signal, which contains the magnetic field information we are interested in.

A more sophisticated technique for feedback design is based on so-called "loop shaping" [Doyle et al., 1990]. In loop shaping, the feedback circuit is designed by first measuring the amplitude and phase response of the system without feedback applied (so-called "open-loop" operation), and then using analog components to build a feedback loop that matches the open-loop

response of the system. This process is detailed in Appendix C. Loop-shaped feedback circuitry, or perhaps a digital realization of it, will be implemented in the next generation of the device.

### 3.5.7 Data acquisition

The final part of the apparatus is the data recording accomplished by analog to digital conversion. Because the signals we are interested in recording are composed of low-frequency components, high speed (MHz or more) acquisition equipment is unnecessary. On the other hand, the recorded signals, especially the critical components of the fetal heart signal, can be very small, so a high-precision ADC is required to digitize the voltages. To accomplish this, we settled on a National Instruments PXI-7853R FPGA card mounted in a National Instruments PXI chassis. The FPGA includes eight analog inputs, each with a 16-bit ADC for digitization.

As Wyllie notes in his thesis [2012], it is critical to have an up-to-the-minute measure of the magnetometer's performance which is based on the unit's measured noise floor. To characterize how well the magnetometers are functioning, a calibration procedure is required to convert the voltages recorded by the ADC to equivalent magnetic field units. Uncalibrated voltage measurements give little insight into magnetometer performance and leave one open to the wrath of his or her advisor. The calibration is dependent on both easily accountable parameters (such as the settings on a current-to-voltage converter) and also on the magnetometer response (Equation (2.49) which is a convolution of a multitude of parameters including pump and probe laser waists, intensities, and detunings, atomic number density, and local environmental fields. Changing any one of these parameters changes the voltage-to-field calibration and thus the ability to rapidly re-calibrate the magnetometer channels is required.

The FPGA provides us with an effective monitor of magnetometer performance. National Instruments LabVIEW software code originally developed by former ThadLab postdoc Matt Kauer and furthered by other researchers (the author included) uses the analog outputs of the FPGA to generate calibration signals and record both coil driver and atomic response to these signals. Using this calibration, further measurements of either ambient noise or cardiographic fields are immediately transformed into magnetic field units, and recalibration takes only seconds. The rapid

recalibration allows us to vary any of the experimental parameters and quickly determine whether or not the magnetometers' SNR has improved. The FPGA also generates the modulation signals required for Z-mode (described in Section 2.5) and handles the demodulation of these signals. A complete description of the calibration and data collection procedure will be given in the next chapter of this thesis.

# Chapter 4

# Data Collection and Analysis

This chapter describes the procedure for setting up, characterizing, and operating the magnetometer. While many of the intricate steps will be omitted here, this chapter will hopefully be useful as a manual for using the magnetometer array. The first several sections will assume that the magnetometers are being set up for DC-mode operation, and Section 4.6 will detail the steps required to set up the device to work in Z-mode.

First we will explain the basic initial setup of the magnetometer, nulling the DC fields of each magnetometer, and optimizing the pump and probe laser detunings to maximize each channel's response to magnetic fields. Once each channel has been optimized, they are calibrated against a known magnetic field and their amplitude and phase responses are recorded. All further magnetic field measurements (whether it is ambient background noise of the MSR or an adult or fetal cardiographic signal) use this calibration to reconstruct the magnetic field signal. Finally, we will present a brief theoretical analysis of some of the sources of noise that limit our measurement sensitivity.

## 4.1 Initial Setup

Assuming the magnetometer array has been placed inside the MSR, the electronics rack is positioned nearby outside the room, and the two have been appropriately connected together with electronics cables and fibers, the array can be brought into magnetic field-sensing operation as soon as the cells are sufficiently warm to produce a saturated rubidium vapor. The first step is to ensure the pump and probe lasers are aligned appropriately.

The construction of the array holder (Figure **??**) does not allow for much adjustment of the pump beam alignment - the pump tube is simply slid into the array structure and secured with a thumbscrew. For pumping a single cell, the close proximity (<5 cm) of the cell to the pump tube makes aligning the pump beam roughly through the circular optical cutouts in the heater panels (Figure 3.2) relatively trivial. However, if one desires to operate in a quasi-dual-sided pumping configuration, where excess pump light transmitted from one cell is used to pump the opposite side of a neighboring cell, more careful alignment is required. However, we have noted that as the pump beam does not always exit the tube precisely along the axis (though the deviation from this is $< 5^{o}$), it is possible to change the propagation direction of the beam slightly by rotating the square tube $90^{o}$ in its holder. When operating in this configuration, we can verify that the beams are counter propagating by checking for their overlap at the outputs of opposing pump tubes.

More degrees of freedom exist for alignment of the probe laser, as each probe tube sits in a "launcher," which can be adjusted in both the horizontal and vertical directions. Additionally, the polarizing beam splitters that direct the pump light to the magnetometer channels can be rotated to adjust the horizontal position of the probe beam. Because of the shared optics, some degree of flexibility is lost, as each adjustment actually changes the probe beam alignment in *two* of the magnetometer channels. Several iterations of adjustments to the launchers and PBS holders is usually required to converge to a point where each beam passes squarely through its magnetometer cell and is incident on the center of that channel's detector. Initial alignment is done by eye, and is verified by measuring the total light intensity by summing the currents from the detectors.

Once we are satisfied with the alignment of the probe beam, we proceed to balance each of the polarimeters (Section 3.3.4). To accomplish this, the pump beam is blocked so that the atomic polarizations $P_z = P_x \equiv 0$, and any rotation of the probe beam is solely due to cell birefringence, which will be present whether or not the device is operating as a magnetometer. The detector is then connected so that the difference signal is output to the current pre-amps, the pre-amps' gains are increased to a high sensitivity (.5-2 μA/V), and the detector itself is rotated until the output difference is zero. A difference signal of zero corresponds to equal light intensities on each of the detector's photodiodes, which is a consequence of the laser light being incident on the Wollaston

prism at an angle of 45° with the Wollaston's slow axis. Careful balancing of the detectors is important for two reasons:

1. Recalling Equation (3.2), the difference signal $S$ is

$$S \propto I_1 - I_2 \propto \cos(2\theta). \tag{4.1}$$

The detector will be most sensitive when the slope of $S$ is maximal, which occurs when

$$\begin{aligned} 0 &= \frac{dS}{d\theta} \\ 0 &= 2\sin(2\theta) \\ &\rightarrow 2\theta = 90^{\circ} \rightarrow \theta = 45^{\circ}. \end{aligned} \tag{4.2}$$

2. The equivalent probe noise (discussed in detail in Section 4.4.5) is directly proportional to the imbalance of the detector. Careful balancing is required to make sure this "technical noise" does not limit the device's ability to measure the the magnetic field.

Under normal measurement conditions (a total probe power of 200 μW and a preamp setting of 2 μA/V), the sensitivity of the detector is in a ∼3.5 V per degree of imbalance. A imbalance of just 2° is sufficient to saturate the detector with these settings. Without micrometers, maintaining a balance to within 2° is often difficult. We have found that coating the external thread of the detector with vacuum grease before screwing it into the array structure allows for more precise adjustment.

Even after balancing the detectors, a large amount of low-frequency noise is often visible when looking at the probe signal on the oscilloscope. The cause of this noise is often poor coupling of the amplified probe light into the polarizing fiber, which is easily rectified by a careful re-coupling of the light. The MSR also appears to have vibrational resonances near 14 and 16 Hz, which can be coupled into the measurement either due to vibrations in the fiber or vibrations in the detectors themselves. Finally, probe pumping (also discussed in Section 4.4.5) can create a small amount of $P_x$, rendering the magnetometer sensitive to magnetic field noise even when the pump is blocked. This effect is amplified when the probe detuning from resonance is small, which is often the optimum point for magnetometry with low atomic density.

## 4.2 Alkali density measurement

After alignment is complete, the alkali metal density is measured with a transmission measurement. Strictly speaking, this is not a required step for magnetometry, but is often done as a "sanity check" to verify our cells are in good working condition. Though we typically work in the 145-160°C range for SERF magnetometry, we have occasionally pushed the temperature higher to experiment with working in a regime where alkali density is high ($\geq 2 \times 10^{14}$ cm$^{-3}$). At extreme temperatures, we have noticed a degradation of cell performance which, after performing an alkali density measurement, we attributed to a loss of rubidium due to its reacting with the Pyrex cell walls. This is a well-known phenomenon, and fortunately reaction rates at normal SERF temperatures are small, though the rate approximately doubles with each 10°C increase in temperature. Performing semi-regular density measurements at identical cell wall temperatures ensures that the magnetometer cells are not losing their precious rubidium.

The density measurement is performed by measuring the transmission of the probe beam as its frequency is swept across the D2 resonance line of rubidium. The intensity $\mathcal{I}$ of transmitted laser light through an atomic vapor is given by Beer's Law:

$$\mathcal{I}(\nu) \propto \exp[-n\sigma(\nu)l] \tag{4.3}$$

where $n$ is the atomic density, $\sigma(\nu)$ is the absorption cross-section, and $l$ is the optical path length of the beam through the atomic vapor. The cross section $\sigma(\nu)$ is

$$\sigma(\nu) = \sigma_0(\nu) \sum_i r_i \mathcal{L}_i(\nu) \tag{4.4}$$

$$\begin{aligned} \sigma_0(\nu) &= \frac{3\lambda^2}{2\pi} \frac{\gamma_{\text{nat}}}{\gamma_{\text{pres}}} f_{osc} \\ &= \frac{3c^2}{2\pi\nu^2} \frac{\gamma_{\text{nat}}}{\gamma_{\text{pres}}} f_{osc} \end{aligned} \tag{4.5}$$

$$\mathcal{L}_i(\nu) = \frac{1}{1 + 4\left(\frac{\nu - \nu_i}{\gamma_{\text{pres}}}\right)^2} \tag{4.6}$$

$$\nu_i = \nu_0 - \nu_{\text{hfs},i}. \tag{4.7}$$

Here, $\gamma_{\text{nat}}$ is the natural linewidth of the D2 absorption and $\gamma_{\text{pres}}$ is the pressure-broadened absorption width, which is calculated using the cell's buffer gas pressure and the appropriate coefficient from Romalis et al. [1997]. The values for the oscillator strength $f_{osc}$, D2 transition frequency $\nu_0$ and ground-state hyperfine splittings $\nu_{\text{hfs},i}$ for the $F = 1$ and $F = 2$ levels of the $5S_{1/2}$ state are taken from [Steck, 2001]. The $r_i$ values are 5/8 for the $F = 2$ hyperfine state and 3/8 for the $F = 1$ hyperfine state.

To measure the density, each magnetometer's detector is connected in its sum configuration to measure the total laser intensity incident on the pair of photodiodes. A LabVIEW program written by the author (Laser Scan v1.1.vi) sweeps the probe DFB laser diode's frequency across a 300 GHz range roughly equally on each side of the D2 resonance. The transmission of the probe laser through each cell is recorded in 2 GHz steps. Typical results are seen in Figure 4.1. The transmission data are fit to the function in Equation (4.4), and a value for the alkali density $n$ is extracted. This value of $n$ is compared to density values taken at the same temperature on earlier days to ensure the cell is not consuming its rubidium. Once the transmission measurement is complete, the detectors are returned to difference mode where they will remain for the rest of the data collection procedure.

## 4.3 Residual field nulling

A SERF magnetometer's steady-state response is characterized by Equation (2.36), which is reprinted below:

$$P_x = P_z \left( \frac{\Gamma' \Omega_y + \Omega_x \Omega_z}{\Gamma'^2 + \Omega_z^2} \right) \tag{4.8}$$

$$\approx P_z \frac{\Omega_y}{\Gamma'} \quad \text{for} \quad \Omega_x, \Omega_z \ll \Gamma' \tag{4.9}$$

The important thing to note here is that the magnetometer will respond to magnetic fields in the $\hat{x}$ direction only if the fields in the $\hat{z}$ direction are non-zero, and vice versa. Thus, we follow the method also described in Li et al. [2006], Seltzer and Romalis [2004], and Wyllie [2012] to zero out all the three fields $B_x$, $B_y$, and $B_z$:

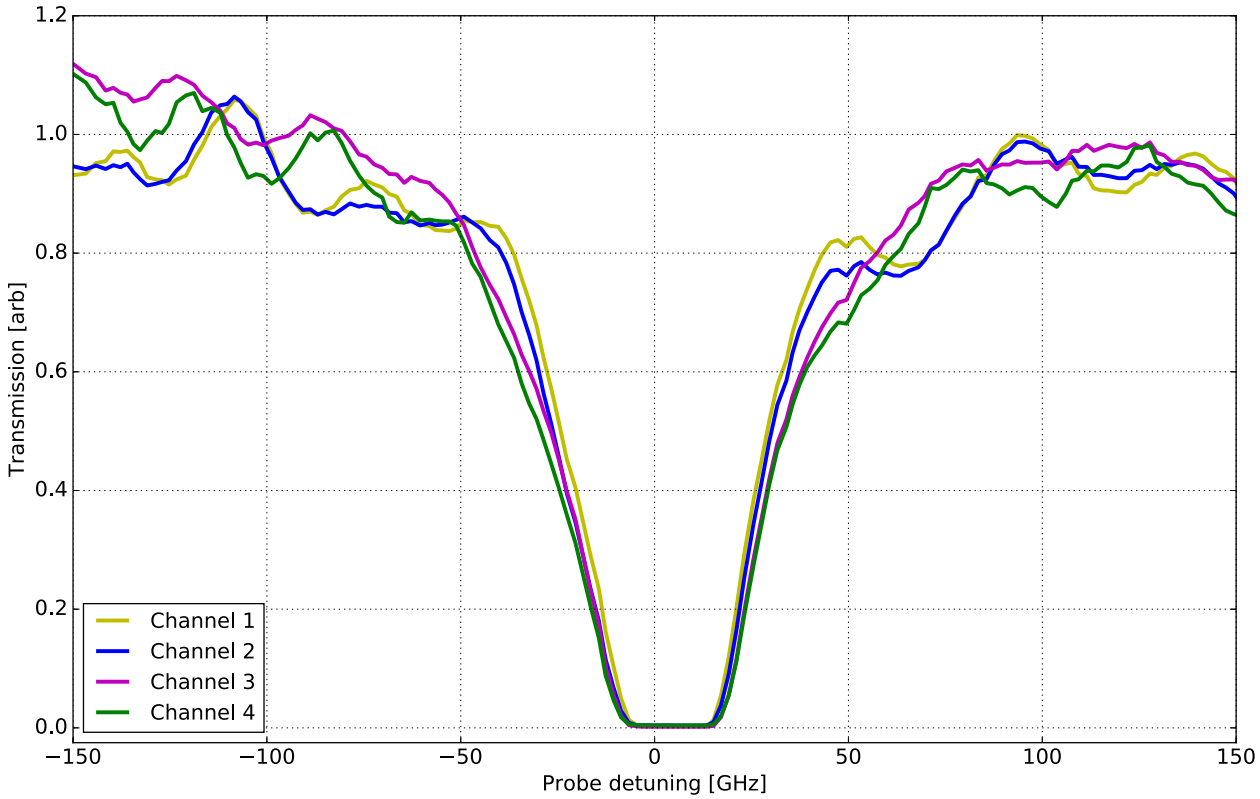



**Figure 4.1:** Results of scanning the probe beam over the D2 transition. The transmission has been normalized for each channel: normally the probe intensities can vary by $\sim$50% from channel to channel. The fluctuations in the probe intensity in the wings is a result of a frequency-dependent splitting ratio in the probe fiber splitter. Note that channels 1 and 2 (which share one "arm" of the splitter) show the same intensity behavior, as well as channels 3 and 4. Those two sets of channels are also anti-correlated with each other. To extract a density $n$, the transmission data are fit to Equation (4.3). When including the oscillating wings, the fit usually passes through the "middle" of the oscillation. We often run multiple fits, including and excluding the wings to verify those oscillations are not skewing our calculation of $n$.

1. Sweep $B_y$ to observe the entire dispersive shape, and then return to the center of curve ($B_y = 0$), where the magnetometer is most sensitive to $\hat{y}$ fields. If the transverse $\hat{x}$ and $\hat{z}$ fields are very far from their null points, it is possible a dispersive shape will not be attainable. In this case, set $B_y$ to a point where the response is the steepest with respect to $B_y$.

2. Apply a field $B_0 \sin(\omega t)\hat{x}$, with $B_0 \approx 50$ pT and $\omega \approx 2\pi\times$ 20 Hz. Adjust $B_z$ to minimize the magnetometer's response to this applied field. This will likely change the position of the magnetometer signal within the dispersion curve, so adjust $B_y$ to keep the magnetometer at its most sensitive.

3. Apply $B_0 \sin(\omega t)\hat{z}$ and adjust $B_x$ until the response is nulled. Again, adjust $B_y$ to keep the magnetometer signal at its point of maximum sensitivity. After one iteration of $B_z$ and $B_x$ adjustment, the dispersive shape of the magnetometer response should be apparent: $B_y$ should be adjusted to place the signal at the center of this pattern.

4. Iterate steps 2 and 3 until responses to both $B_x$ and $B_z$ are minimized, while constantly adjusting $B_y$ to maintain high sensitivity.

Each channel of the magnetometer must be nulled in this fashion. To expedite the process of nulling all four channels, the array is usually placed close to the center of the MSR (or, in the case of preparation for an adult or fetal cardiogram measurement, close to the location the sensor will be placed) with its $x$, $y$, and $z$ axes aligned with the corresponding axes of the room. The magnetometer signal from one of the magnetometers (usually Channel 1) is observed while the modulation signals described above are applied using either the room- or shell-mounted coils. The room coils are adjusted until the nulled condition of step 4 is met. In addition to completely nulling the fields of Channel 1, this brings the residual fields at the other magnetometers somewhat close to zero, so only small additional fields are required to null those channels. The magnetic field decoupling scheme described in Section 3.5.4 allows for the nulling of each channel to occur independently.

An important feature of this process is a result of the AC Stark shift described in Section 2.3.6. While a uniform AC Stark shift can be completely eliminated by the application of a DC field,

this is rarely the case in practice. The shift is proportional to pump intensity, and an unavoidable pump intensity gradient exists along the $\hat{z}$ axis of the cell due to pump absorption. Therefore, the best we can hope to accomplish is to cancel out the *average* AC Stark shift with the application a DC $B_z$ field. Thus, the magnetometers will normally exhibit an undesirable (at least in DC-mode) non-zero response to $\hat{x}$ fields. This effect can be mitigated by tuning the pump laser to resonance, where the pump intensity gradient is smallest, though this often sacrifices some of the desired response to $\hat{y}$ fields. Optimizing these parameters will be covered in Section 4.5.

Once all four channels of the magnetometer are individually nulled, small adjustments to cancel out uniform field drifts can be made using the room coils. However, if the array has been tilted or rotated, which is common during a patient measurement, either the array-mounted feedback or and the shell-mounted coils must be used. External field drifts of magnitude $\geq 2$ nT can cause the magnetometers to extend of the linear regime of their dispersion curves, reducing their sensitivity to the magnetic fields of interest. Active feedback can help stabilize the array against some of these drifts, though at the time of this thesis, we had only successfully implemented feedback in the most sensitive ($\hat{y}$) direction. Future versions of the experiment will expand feedback to all three axes to hopefully completely eliminate these issues.

## 4.4 Magnetometer calibration and characterization

Once the magnetometers have been brought to their operating point, their performance can be evaluated. Much of this procedure has been automated using an FPGA-based LabVIEW program originally developed by post-doc Matt Kauer and further refined and expanded by successive group members (Magnetometer 16.3.vi). As discussed in Chapter 3, the FPGA contains eight digital-to-analog converted output channels for synthesis of modulation signals and eight 16-bit analog-to-digital converted input channels for data acquisition. The data are acquired by the input channels at 500 kHz and immediately downsampled and low-pass filtered to 20 kHz by the FPGA. Further downsampling and filtering is available to match the demands of the particular type of data being acquired. The averaging done during the downsampling process reduces high frequency noise that may be aliased down to our frequency band otherwise.

The basic process is as follows: a synthesized calibration voltage (detailed in the next section) is applied to the coil driver's input while monitoring the voltage output over the coil driver's output resistors to obtain a measure of the applied calibration signal. Next, the signal is reapplied to the coil driver, but this time the output signal from the magnetometer is measured. The ratio of these measurements provides a calibration (in volts per applied field) that will be used to calibrate further magnetometer measurements. A block diagram of the measurement procedure is given in Figure 4.2.

### 4.4.1 Chirp signal synthesis

Because of the non-infinite bandwidth of the circuit components comprising the coil driver and the magnetometer itself (Equation (2.49)), it is not sufficient to measure the response of either by applying a single frequency signal. One solution is to sequentially apply several cycles of single-frequency signals to the device and measure the amplitude and phase response at each frequency. Fitting the data points with a modified version of Equation (2.49) provides a function against which future data can be calibrated.

A somewhat more sophisticated approach involves using a frequency chirp, a calibration signal whose frequency is swept continuously in time and allows for a more rapid evaluation of the system's frequency response. This approach was first implemented by Wyllie and the end of his tenure, and has been used exclusively since. The chirp function used for our calibration is

$$V(t) = A_0(t) \sin\left(\omega(t)t\right) \tag{4.10}$$

with

$$\omega(t) = \omega_0 \exp(Bt). \tag{4.11}$$

In this way, $\omega_0$ sets the initial frequency of the sweep and the sweep velocity is B, where

$$B = \frac{1}{T} \ln\left(\frac{\omega_{\text{max}}}{\omega_0}\right), \tag{4.12}$$

where $T$ is the time-length of the applied chirp and $\omega_{\text{max}}$ is the maximum desired calibration frequency. Typical chirp parameters are $\omega_0 = 2\pi \times 0.2$ Hz, $\omega_{\text{max}} = 2\pi \times 750$ Hz, and $T = 1$ s. The

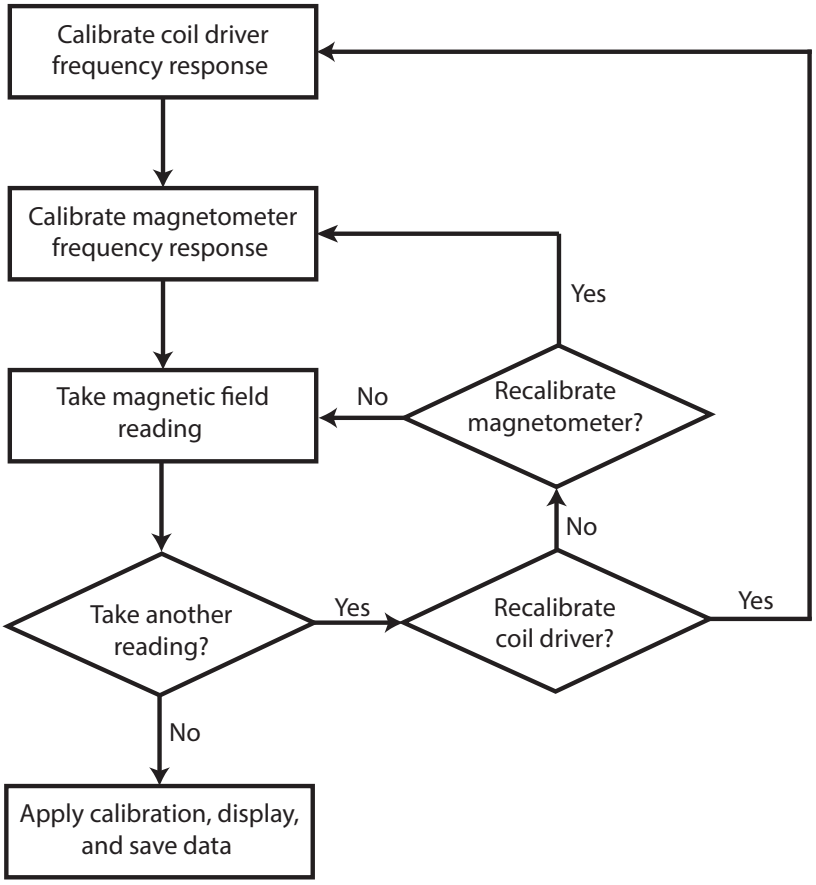



**Figure 4.2:** A simple schematic of the calibration and data collection procedure. Once a magnetic field measurement has been made, recalibration of the coil driver and magnetometer response is possible, though not always necessary. The coil driver response must only be re-measured if the output resistor or other circuit components are changed. The magnetometer should be recalibrated if any parameter of the apparatus has been varied (laser detunings/intensities, atomic density, etc.) or if the ambient fields have changed enough to warrant a re-nulling of the magnetometers.

large value of $\omega_{\text{max}}$ was chosen to ensure the entire spectrum of the magnetometer's response is captured: for frequencies above $\sim$350 Hz, the magnetometer's response to the chirp signal is usually not discernible from the noise. The amplitude of the chirp signal, $A_0(t)$, is generally a constant (an example chirp of this type is shown in Figure 4.3). To improve the calibration signal-to-noise ratio when the response is small (at higher frequencies on the magnetometer response, for example), it can modified to be a monotonically increasing function of time. This strategy was implemented in our gradiometry studies (Chapter 5) where large SNR at all frequencies of the calibration signal was particularly important.

The signal is digitally synthesized by the FPGA and copies are sent through the ADCs of four of the analog outputs, which in turn deliver the voltage to the coil drivers' AC input terminals.

### 4.4.2 Coil driver circuit calibration

The chirp signal is sent to the AC input of the coil driver circuit connected to the magnetometer's shell-mounted $\hat{y}$ coils and added to the DC input voltage provided by the Arduino-controlled MAX5735. The resultant voltage is monitored by connecting a terminal placed across the pair of output resistors to the FPGA input as described in Appendix A. The chirp signal is applied to the circuit three times, and each time series $V_{\text{in,i}}(t)$ is recorded. The Power Spectral Density (PSD) of each time series is computed by the LabVIEW code. To help reduce the effect of white noise during the measurement, the three PSDs $S_{\text{c,out,i}}(\omega)$ are averaged together.

The transfer function describing the circuit's conditioning of the input signal is given by

$$C(\omega) = \frac{S_{\text{c,out}}(\omega)}{S_{\text{in}}(\omega)} \tag{4.13}$$

where $S_{\text{in}}$ is the PSD of the signal produced by the FPGA system and $S_{\text{c,out}}$ is the PSD of the circuit's response. Because there are no filters or other circuit components with frequency-dependent responses between the circuit's AC input and monitor output, $C(\omega)$ is independent frequency across our calibration range. If the applied voltage $V_{\text{in}}$ is sufficiently small, narrow spikes due to electronic line noise are observed in $S_{\text{c,out}}(\omega)$ at $f = 60$, 120, and 180 Hz, so the circuit has the

*appearance* of having a large response at these frequencies, the effect of which can be seen in the final magnetometer calibration $\zeta(\omega)$ (see Section 4.4.3 and Figure 4.3b, for example).

Using the output resistor $R_{\text{out}}$ and the measured coil calibration $\chi(\omega)$, the magnetic field applied to the atoms during the chirp can be calculated:

$$B_{\text{app}}(\omega) = \frac{C(\omega)S_{\text{in}}(\omega)}{R_{\text{out}}}\chi(\omega) \equiv \eta(\omega)S_{\text{in}} \tag{4.14}$$

The function $\chi(\omega)$ is assumed to be a constant for the shell-mounted coils. The coils have low inductance ($\sim 1\,\mu\text{H}$) and at our frequencies of interest, have very low inductive reactance. Furthermore, they are well-separated from any mu-metal which would induce a frequency-dependent coil calibration. As demonstrated in Section 3.4.1, this does *not* hold true for the room coils (which are wrapped in close proximity to the mu-metal shield), so $\chi(\omega)$ must be appropriately modified (also discussed in Section 3.4.1) if the large room coils are used for calibration, which they are in our gradiometry studies.

Once the coil driver calibration is determined, the coil driver must only be recalibrated if $\eta(\omega)$ is changed, which is most often due to changing $R_{\text{out}}$.

### 4.4.3 Magnetometer calibration

Once the applied field $B_{\text{app}}(\omega)$ has been calculated, the magnetometer's response to that field can be recorded by once again applying the frequency chirp to the coil driver circuit but this monitoring the magnetometer signal output from the transimpedance amplifier. Once again, the time series $V_{\text{m,out}}(t)$ is recorded, and the PSD of the time series $S_{\text{m,out}}(\omega)$ is calculated. The magnetometer's response to the applied field is then given by

$$\zeta(\omega) = \frac{S_{\text{m,out}}(\omega)}{B_{\text{app}}(\omega)}. \tag{4.15}$$

$\zeta(\omega)$, which has units of volts per field, serves as the calibration for any measurement made under the calibration conditions. A sample chirp measurement and resultant $\zeta(\omega)$ are shown in Figure 4.3. A continuous function $\zeta^{\star}(\omega)$ is attainable either by curve-fitting the data to the theoretical frequency response of the magnetometer given by Equation (2.49) or by creating an interpolation

of the discrete points of $\zeta$. We have settled on the latter approach as it is not too processor intensive for our size of data sets, and it does a better job of reflecting the exact properties of the calibration.

Unlike the coil driver calibration, the magnetometer calibration changes regularly. Sometimes, this is due to a controlled modification to the sensor, such as changing the detuning or power of the pump or probe laser. Often, though, the magnetometer's sensitivity varies due to drifts in the ambient magnetic field or changes in the atomic density due to cell temperature fluctuations. Thus, it is good practice to recalibrate the magnetometer regularly to ensure that measurements accurately reflect the magnetic field that is being detected.

### 4.4.4 Environmental magnetic field noise measurements

With a calibration in hand, the sensors are now ready to accurately measure the magnetic field. 6-10 seconds of magnetometer signal $V_m(t)$ is acquired and digitized by the FPGA. The signal is Fast Fourier Transformed and the continuous calibration function $\zeta^\star(\omega)$ is applied to convert it to a real magnetic field measurement.

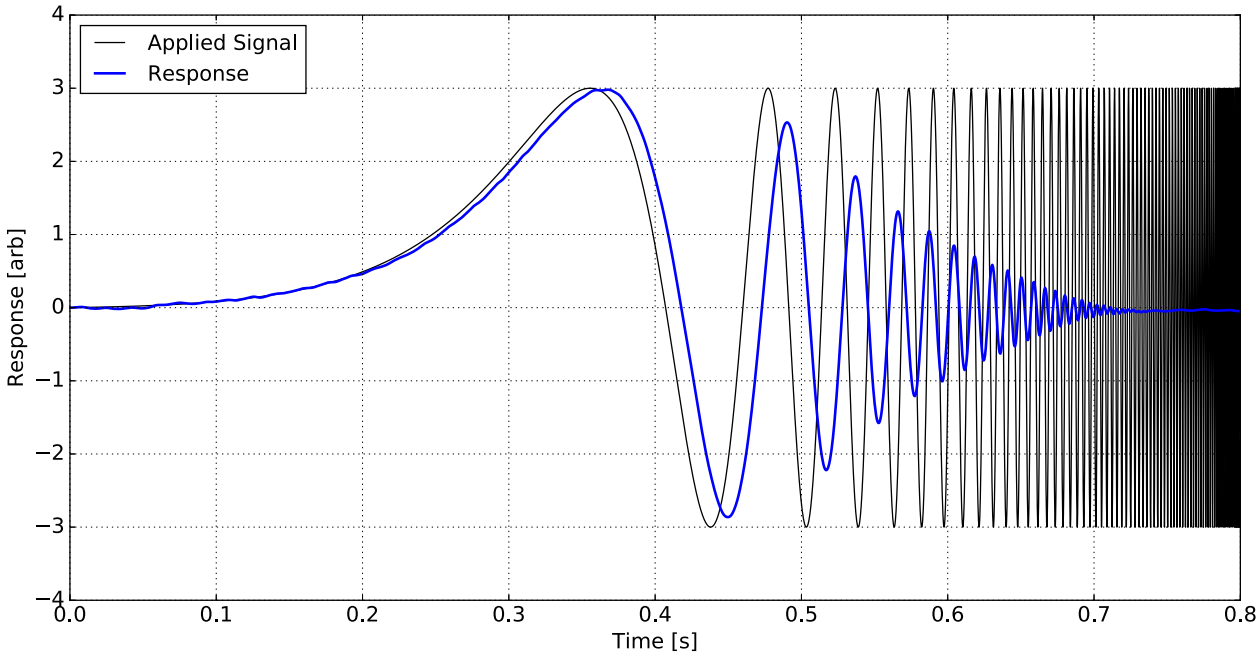
$$\begin{aligned} \tilde{V}_m(\omega) &= \text{FFT}\left[V_m(t)\right] \\ \tilde{B}(\omega) &= \frac{\tilde{V}_m}{\zeta^\star(\omega)} \\ B(t) &= \text{IFFT}\left[\tilde{B}(\omega)\right] \end{aligned} \tag{4.16}$$

Simultaneously, the PSD of the magnetometer signal $V_m(t)$ is calculated, and the magnetic field spectral density ($\sqrt{\text{PSD}}$) is calculated, which provides insight into the specific characteristics of the magnetic field of the room.

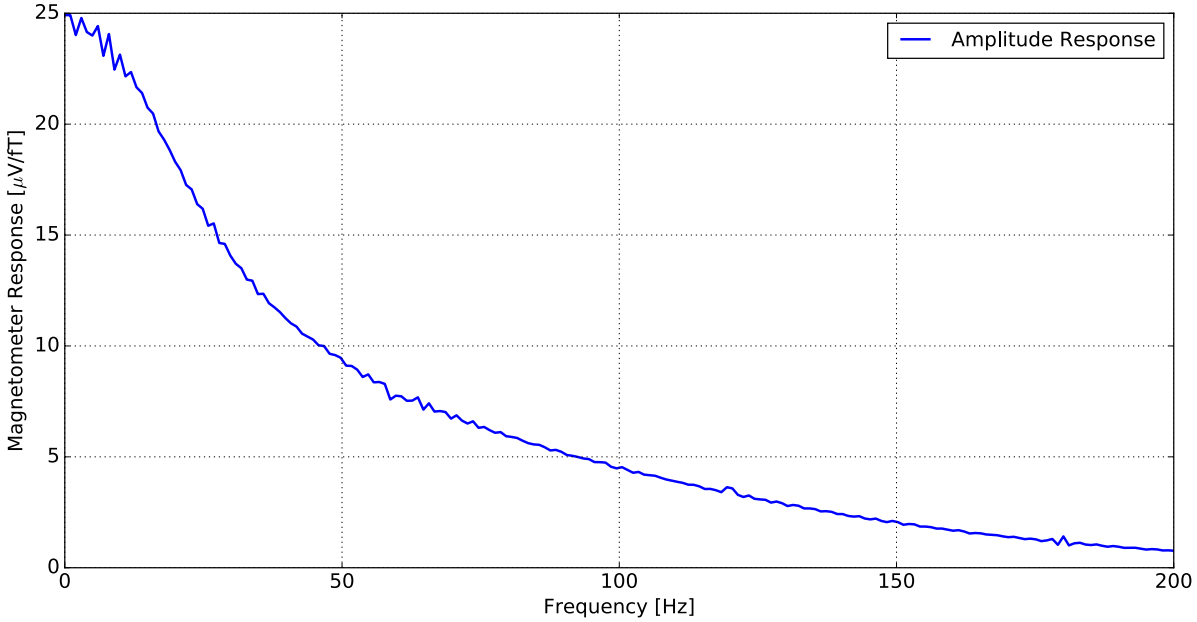
$$\frac{\delta B}{\sqrt{\text{Hz}}} = \frac{\sqrt{\text{PSD}\left[V_m\right]}}{\zeta'(\omega)} \tag{4.17}$$

### 4.4.5 Additional noise measurements

The signal obtained in Section 4.4.4 provides the user with a measure of the characteristics of the magnetic field in the room, but there are some subtleties that must be addressed. The measurement assumes that the signal acquired is generated solely by the magnetic field in the




**(a)** Application of synthesized chirp and times-space magnetometer response




**(b)** Magnetometer amplitude response to applied chirp

**Figure 4.3:** (a) Application of the chirp signal and the DC-mode magnetometer response to the signal. Note that for frequencies $f$ >350 Hz, the magnetometer response is not discernable from the noise. Larger calibration signals can be used to mitigate this effect, but data at these frequencies are not particularly interesting for a fMCG measurement. (Chapter 6). (b) The magnetometer response, $\zeta(\omega)$ to the applied chirp, in frequency space. Although the transform used to obtain $\zeta$ is discrete, the points are interpolated to create a continuous function $\zeta^{\star}(\omega)$. It is this function which is applied to all successive collected data to transform into magnetic field units.

room, yet this is hardly ever the case. Because the acquired a signal is actually a *voltage* (which can only be converted to a field through calibration), any noise sources which contribute to that voltage signal will *appear* to be adding magnetic noise to the system. These non-magnetic sources can be quantum (spin-projection noise and photon shot noise) or technical (probe polarization noise, electronic pickup, *e.g.*) in nature. It is necessary to evaluate the contributions to the total noise from these other sources to make sure they are not limiting our ability to measure the true magnetic field signal.

Recall Equation (4.9), the steady-state solution to the atomic polarization of the atoms of the magnetometer.

$$\begin{aligned} P_x &= P_z \left( \frac{\Gamma' \Omega_y + \Omega_x \Omega_z}{\Gamma'^2 + \Omega_z^2} \right) \\ &\approx P_z \frac{\Omega_y}{\Gamma'} \end{aligned} \tag{4.18}$$

Since we are sensitive to $P_x$ it would seem that making $P_z = 0$ would leave us insensitive to magnetic fields. This can be accomplished by simply blocking the pump beam (or in our setup, turning off the pump tapered amplifier). Now the detected magnetometer signal is only a result of the quantum and technical noise sources mentioned above. This signal can be scaled by $\zeta(\omega)$ to estimate the equivalent magnetic noise that these sources create.

There are two caveats to this insensitivity assumption though. The first is that any noise injected into the signal due to the pump laser itself will not manifest itself in this measurement. Fluctuations in the pump intensity or frequency will cause changes in $\langle P_z \rangle$ (which, by Equation (4.18) will cause detectable fluctuations in $P_x$). In addition, the same pump fluctuations will lead to variations in the field caused by the AC Stark shift along the $\hat{z}$ direction. This can couple to fields in the probe direction ($\hat{x}$), increasing the noise in the magnetometer signal. A third method for pump-induced noise is the scattering of pump photons directly to the detector. We have mostly discounted this effect by blocking the probe light and measuring the detector signal with the pump light blocked and unblocked. No discernible difference is seen in the noise spectra between these two measurements, leading us to believe there are very few scattered pump photons seen by the detector.

| Measurement | Method | Interpretation |
|---|---|---|
| Magnetic Noise | Nothing blocked | Higher than probe: If channels are ∼same, magnetic<br>If channels are different, pump<br>Same as probe: limited by probe noise |
| Prone Noise | Block pump beam | Noise from probe laser and fiber, plus fundamental (quantum) noise and electronics |
| Electronic noise | Pump and probe blocked | Noise from electronics (pickup) or DAQ (aliasing) |

**Table 4.1:** The three noise measurements commonly made to characterize the magnetometers. By making these three separate measurements, we can deduce what, if anything, is limiting the magnetometers' performances.

The second caveat to the technical noise assumption is that any small amount of circular polarization of the probe beam can optically pump the atoms, inducing a $P_x$. This, in turn, will make the magnetometer sensitive to magnetic fields in the $\hat{y}$ and $\hat{z}$ directions even with the pump light turned off. We try to limit this by implementing 100,000:1 extinction ratio linear polarizers just before the probe light enters the cell. However, birefringence in the glass cell wall can cause the linear light to pick up some ellipticity before interacting with the alkali vapor. This effect is magnified when the detuning from resonance of the probe beam is small and the probe's optical pumping rate is increased. The lower buffer gas pressures in the newest batch of magnetometer cells has led us to tune our probe laser closer to resonance to increase response, and this probe pumping effect has become apparent. During the balancing of the detectors (which is done with the pump light off and $P_z = 0$), the detector can sense magnetized object carried by the operator (as seen by corresponding variations in the detector output), so metallic objects must be left outside of the shielded room while the detectors are being balanced.

These two effects actually have opposite effects on the estimation of non-magnetic noise. Removing the pump-induced technical noise would lead one to believe the technical noise is lower than it actually is. On the other hand, the probe pumping makes the device sensitive to magnetic noise, which *adds* non-technical noise to a supposed "technical noise" measurement. Without attempting to calibrate the magnetometer with the pump off, it is difficult to know exactly how much noise each contributes to the technical noise measurement, though estimations are possible [Wyllie, 2012].

With these effects in mind, we take this noise measurement (with the pump off and probe on) to be a measure of the total non-magnetic noise in our system. Because this is almost always dominated by probe polarization and intensity fluctuations, we refer to this as simply the "probe noise". A third and final measurement taken with both pump and probe lasers turned off gives us an insight into what we call "electronic noise". That is, noise sourced from electronic pickup or aliasing effects in the data acquisition. Table 4.1 provides a summary of these three noise measurements and their interpretations with regard to the total noise of the system. A plot of these noise levels for all four channels during a typical run inside the MSR is shown in Figure 4.4.

Detector imbalance has been shown to be a large contributor to measured technical noise. This effect is made apparent in that the probe noise measurement will often yield higher noise levels than a "total noise" measurement. Realistically, this is impossible, but the predicament arises because during a measurement of the total noise (with the pump on), the magnetometer signal can be held at zero (whether by active feedback or the user adjusting the $\hat{y}$ fields). The magnetometer is insensitive to fields during a probe noise measurement, however, so this is impossible, and the detector balance is simply given by its orientation relative to the probe polarization. Sequential measurements made with the same magnetometer calibration but with varying imbalances of the detector show clear evidence of increasing probe noise with increasing imbalance. In conclusion, care should always be taken to balance the detector before measurement, but the distressingly-high noise captured during the probe noise measurement is often not indicative of its actual contribution during a real magnetic field measurement, where the balance is much better. Section 4.8.2.1 will explore this type of noise in more detail.

## 4.5 Magnetometer optimization

With the three noise measurements in hand, we have a good understanding of how the magnetometers are performing and can begin to optimize them if necessary. If the channels all show the roughly the same amount of magnetic noise signal and we are confident those levels are not limited by "probe noise", we will say that we are indeed measuring the true ambient magnetic noise at the location of the detectors. Significant variations in the total measured noise without obvious

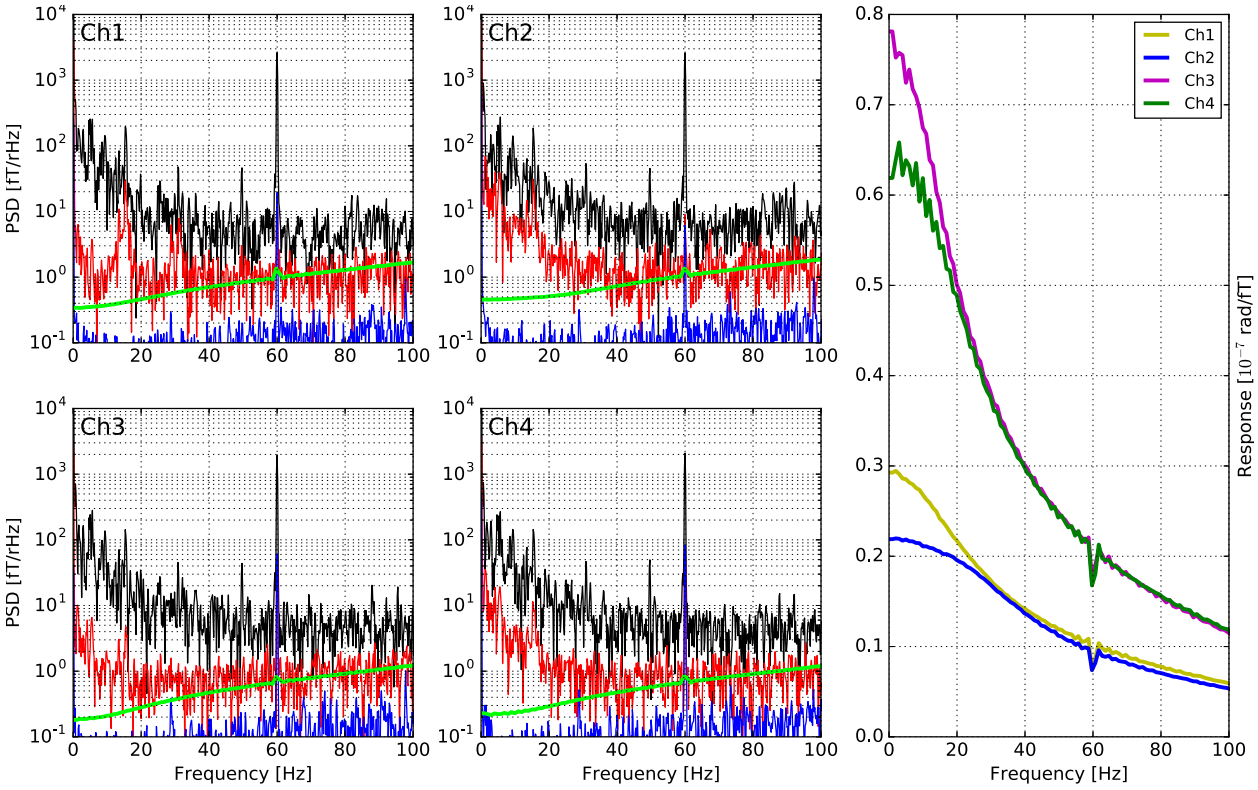



**Figure 4.4:** Magnetometer performance in DC-mode. The total noise (black), probe noise (red), and electronic noise (blue) for each magnetometer channel are shown in the first four panels.. The bright green line is the magnetic-equivalent photon shot noise (Section 4.8.3.2). All show a noise floor of $\sim$5 fT/$\sqrt{\text{Hz}}$ outside of the $1/f$ low-frequency noise contributions. The total noise is well-separated from the technical noise on all channels, implying that the total noise measurement is indeed a measurement of the ambient magnetic field, not limited by technical noise. The exact probe noise levels depend on the balance of the detectors, but all are $<$ 2 fT/$\sqrt{\text{Hz}}$ for $f$ $>$20 Hz. The probe noise spikes at $\sim$15 Hz are a result of vibrational room resonances which couple into the fibers and detectors. The pane on the right shows the magnetometer calibration $\zeta(\omega)$ used to convert the acquired voltages to magnetic fields.

indications of being probe-noise limited suggest the sort of pump-induced noise discussed in the last section. If the total noise level appears limited by technical noise sources, though, there are several avenues we can take to remove the limitations.

Electronic noise is perhaps the easiest source of noise to reduce. Pickup from the heater drive voltage can be seen clearly on the outputs of the transimpedance amplifiers if low-pass filters are not implemented. We have worked on reducing this noise by using grounded shielded cables for carrying the heater and photodiode signals. However, this noise (at $\sim$150 kHz) is far outside the measurement bandwidth of interest and does not pose much of an problem beyond a visual annoyance. Applying one of the low-pass filters on the transimpedance amplifiers removes this feature. In addition, the selection of too small a transimpedance gain can lead to a high amount of equivalent electronic noise. To avoid this, we push this gain to the point where the signal from the largest applied or detected magnetic field (typically the $\sim$200 pT calibration chirp) fills the range of the amplifier but does not saturate it.

If, instead, the total noise is dominated by probe noise, the equivalent probe noise can be decreased by increasing the intrinsic atomic response to magnetic fields. This does not indicate a real decrease in the amount of voltage noise generated by technical noise sources; rather, it is due to the fact that a larger $\chi(\omega)$ (measured in volts per magnetic field) will decrease the technical noise in equivalent magnetic field units.

The easiest way to attempt to increase the intrinsic magnetometer response is to change the intensity and detuning of the pump laser. To optimize these parameters, a magnetic field is applied in the $\hat{y}$ direction to all of the magnetometer channels simultaneously, either implementing the room coils or the array-mounted feedback coils. The same LabVIEW program used in the absorption measurement (LaserScan v1.1.vi) can be used to sweep the detuning of the pump and record the peak-to-peak response of the magnetometer outputs as the frequency is changed. A plot of the magnetometer response versus frequency is generated (see Figure 4.5), and the pump detuning which maximizes atomic response can be identified. The pump intensity can be varied by manually adjusting the drive current on the pump tapered amplifier while observing magnetometer response (at present, there is no way to automate this process). It is important to re-null the magnetometers

(especially the $\hat{z}$ field) after any change in the pump parameters, as a change in the AC Stark fields will inevitably follow.

The probe laser can also be adjusted to increase response. Figure 2.13 suggests an optimum detuning point exists on each side of the D2 probe resonance, so this detuning (or somewhere near it) should be utilized. The probe detuning can be swept in the same way as the pump; a sample set of data is shown in Figure 4.5. The optimum intensity of the probe beam can be somewhat more difficult to determine. Higher probe intensity provides for a larger signal size (which aids in reducing equivalent electronic and photon shot noise), but also increases the effects of probe-induced optical pumping and may lead to non-linear behavior of the photodiodes. Repeated technical noise measurements after varying probe intensities are required to ascertain an optimal setting.

Once the electronics and pump/probe lasers have been optimized, along with the re-nulling of the AC Stark fields which arise from these optimizations, the magnetometers should be recalibrated and the three noise spectra (total, probe, and electronic) should be taken again. The sensors will be considered "ready for action" when the following criteria are met:

1. The electronic noise is not the limiting factor in making a probe noise measurement.

2. The probe noise is not the limiting factor in making a total noise measurement.

3. The total measured noise is uniform across all magnetometer channels, suggesting the noise measured by each channel is the ambient noise of the MSR and not pump induced noise, which varies channel-to-channel.

## 4.6 Z-mode measurements

Before detailing the procedure to make a cardiographic measurement, we explain some of the differences in the calibration and measurement procedure if the array is operated in Z-mode, the theoretical basis for which is given in Section 2.5. This section assumes all of the DC-mode preparation up to Section 4.3 has been completed.

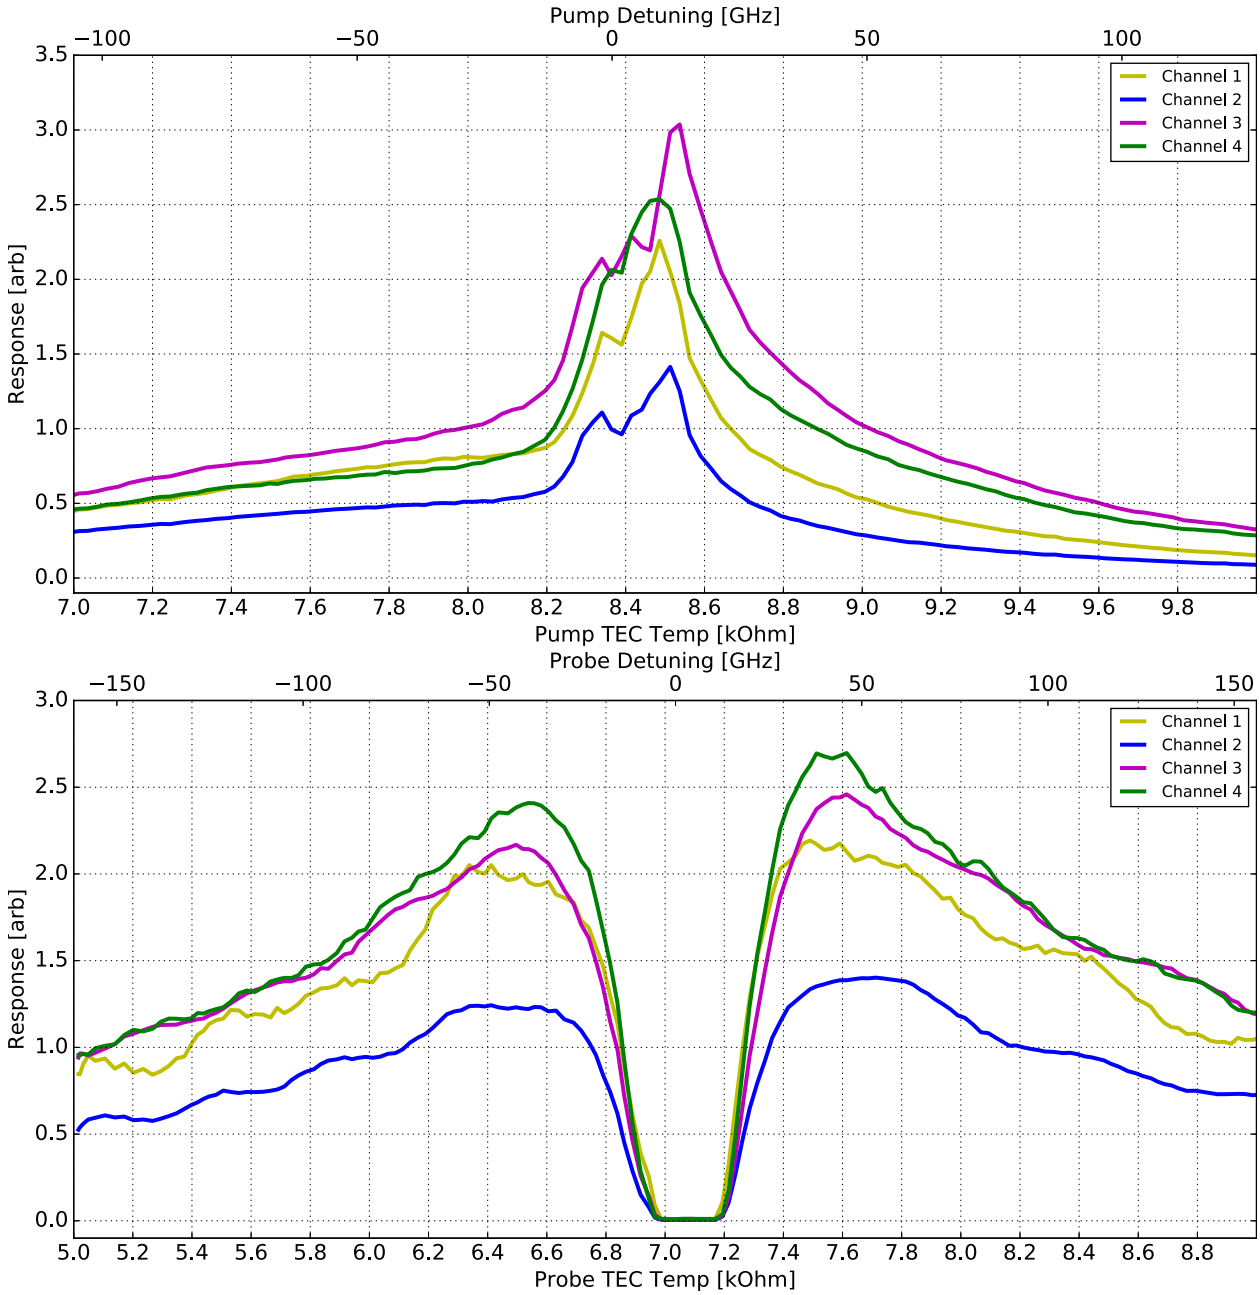



**Figure 4.5:** Results from sweeping the detunings of the pump laser (top) and probe laser (bottom) across the D1 and D2 resonances to find the point that maximizes response to magnetic fields. The sweeps are done sequentially. A uniform magnetic field $B(t) = B_0 \sin(\omega t)$ (with $\omega = 2\pi \times 15$ Hz) is applied to all channels while the detuning is swept (via ramping the laser diode temperature) and the peak-peak amplitude of the magnetometer response is recorded. The laser detunings are set to a point which gives a near-optimal response on all four channels. These data resulted in the pump TEC being adjusted to 8.52 kΩ and the probe TEC to 7.60 kΩ, corresponding to detunings of 16 GHz and 45 GHz, respectively.

### 4.6.1 Field nulling and modulation optimization

In Z-mode, a large oscillating field is applied along the pump ($\hat{z}$) axis of the magnetometer. Field-nulling when operating in Z-mode is somewhat more complicated than DC-mode. Ideally, we want the response to $\hat{y}$ fields to be as large as possible, but also to simultaneously optimize the response to $\hat{x}$ fields, which are detected at the first harmonic of the applied z-modulation frequency.

The same magnetometer program used in DC-mode for signal synthesis, acquisition, and analysis can be used for Z-mode as well. For initial field nulling and optimization, though, a stripped-down program (XY Modulation+Z-mode v1.5.vi) is implemented. The program generates the required field-modulation

$$\Omega_z(t) = \Omega_{z0}\sin(\omega_z t), \tag{4.19}$$

with $\omega_z \approx 2\pi \times 1$ kHz) signals for each magnetometer channel, acquires the magnetometer signal from the transimpedance amplifiers, and performs demodulation and low-pass filtering on each channel's signal to recover the $\hat{x}$ field information. The program simultaneously low pass filters the input signal to recover the $\hat{y}$ field information. The four $\hat{x}$ or $\hat{y}$ signals (one from each magnetometer channel) can be output to an oscilloscope, which greatly aids in the field nulling procedure.

The Z-mode nulling procedure utilizes two additional "test" modulation fields,

$$\Omega_x(t) = a\sin(\omega_1 t) \tag{4.20}$$

$$\Omega_y(t) = b\sin(\omega_1 t). \tag{4.21}$$

These fields have similar field strength ($a, b \sim \gamma \times 50$ pT) and frequency ($\omega_1 \sim 2\pi\times$ 20 Hz) to the ones used for DC-mode Section 4.3. Adjustments are made to the DC offset fields $B_{0x}$, $B_{0y}$, and $B_{0z}$ in the following manner (to reiterate: during a field measurement, the demodulated signal will be used to measure $B_x$ and the low-passed signal to measure $B_y$):

1. Bring the fields to the DC-mode null point as described in Section 4.3. This provides a good starting point for the optimization for Z-mode.

2. Apply the modulation signal $\Omega_z$ to the coil driver connected to the $\hat{z}$ shell-mounted coils. While watching the demodulated signal output, adjust $\Omega_{0z}$ to maximize response to an applied $\Omega_x$. The demodulation phase should also be adjusted to maximize the "in phase" signal size.

3. While still applying $\Omega_x$, adjust $B_{0x}$ to minimize any $\sin^2(\omega_1)$ behavior in the demodulated signal, and adjust $B_{0z}$ to maximize the response to $\Omega_x$. Verify that this also minimizes the response to $\Omega_x$ on the low-passed signal.

4. Adjust $B_{0y}$ to remove any remaining $\sin^2(\omega_1)$ behavior from the demodulated signal. In theory, this will also bring the low-passed signal to the center of its dispersion curve (*i.e.*, zero on the low-passed output). Readjust $B_{0z}$ to once again maximize response to $\Omega_x$ in the demodulated signal and minimize response to that field in the low-passed signal.

5. Repeat step 2 to ensure the demodulated response to $\Omega_x$ is still optimized. Note that adjusting $\Omega_0$ to maximize response to $\Omega_x$ actually goes beyond the "ideal" point where the $J_0^2$ and $2J_0J_1$ functions cross in Figure 2.15. In practice, though, this "ideal" point is somewhat difficult to find and we have found we achieve good Z-mode performance for detecting both $\hat{x}$ and $\hat{y}$ fields by setting $\Omega_0$ to the point which maximizes response to $\Omega_x$.

Again, each channel must be nulled separately using this procedure. Because of the close proximity of the magnetometers, and the fact that the magnetic field decoupling scheme of Section 3.5.4 only applies to the applied DC nulling fields, the optimum size of $\Omega_0$ for a particular channel will vary depending if the modulation signal $\Omega_z$ is being applied to the other nearby channels or not. Assuming the user desires to use all four channels to make a Z-mode measurement, it is a good idea to apply a roughly correct-amplitude $\Omega_z(t)$ (with $\Omega_{z0}$ $\sim$200 pT) to all four channels and then begin the nulling and optimization procedure channel-by-channel.

### 4.6.2 Z-mode data acquisition

Acquiring magnetic field information in Z-mode is very similar to DC-mode, though with two components of magnetic field to detect, two rounds of calibration must occur. We first calibrate

the coil driver circuit connected to the $\hat{y}$ shell-mounted coils, followed by the coil drive circuit connected to the $\hat{x}$ coils. Then both axes of the magnetometer are calibrated ($\hat{y}$ followed by $\hat{x}$) to obtain the calibration functions $\zeta_y(\omega)$ and $\zeta_x(\omega)$, respectively. The same three noise measurements (total noise, probe nosie, and electronic noise) are taken to evaluate the performance of the Z-mode magnetometers. Results from one dataset are shown in Figure 4.6.

There are several important aspects of operating in Z-mode:

1. The low-passed ($\hat{y}$) response is reduced compared to DC-mode, according to the $J_0^2$ term in Equation (2.83). This inevitably means that technical noise (most notably, the $1/f$ noise from the probe laser fluctuations) becomes more of an issue at low frequencies. Care must be taken to reduce this probe noise so as not to have it limit the measurement of $B_y$.

2. At the same time, the demodulated ($\hat{x}$) response of the Z-mode magnetometer is completely insensitive to this low-frequency probe noise, as it is not true field noise and not modulated by $\Omega_z$. The difference in probe noise contributions to the total noise level is evident in Figure 4.6.

3. In a notable improvement over Wyllie [2012], we were able to apply the large modulation fields without saturating the transimpedance amplifiers so the increased electronic noise associated with a reduction in transimpedance gain was not seen in our measurements. We were still able to work in a regime of baseline Z-mode noise between 5-15 fT/$\sqrt{\text{Hz}}$ on all eight measurement channels.

4. Feedback is somewhat difficult to implement in Z-mode given the large phase delays induced by the low-passing of the $\hat{x}$ and $\hat{y}$ signals. The phase delays cause instability in the loop after only a small amount of gain is applied. A careful design of a feedback circuit using either loop shaping (Appendix C) or a digital feedback scheme may allow us to work around these issues.

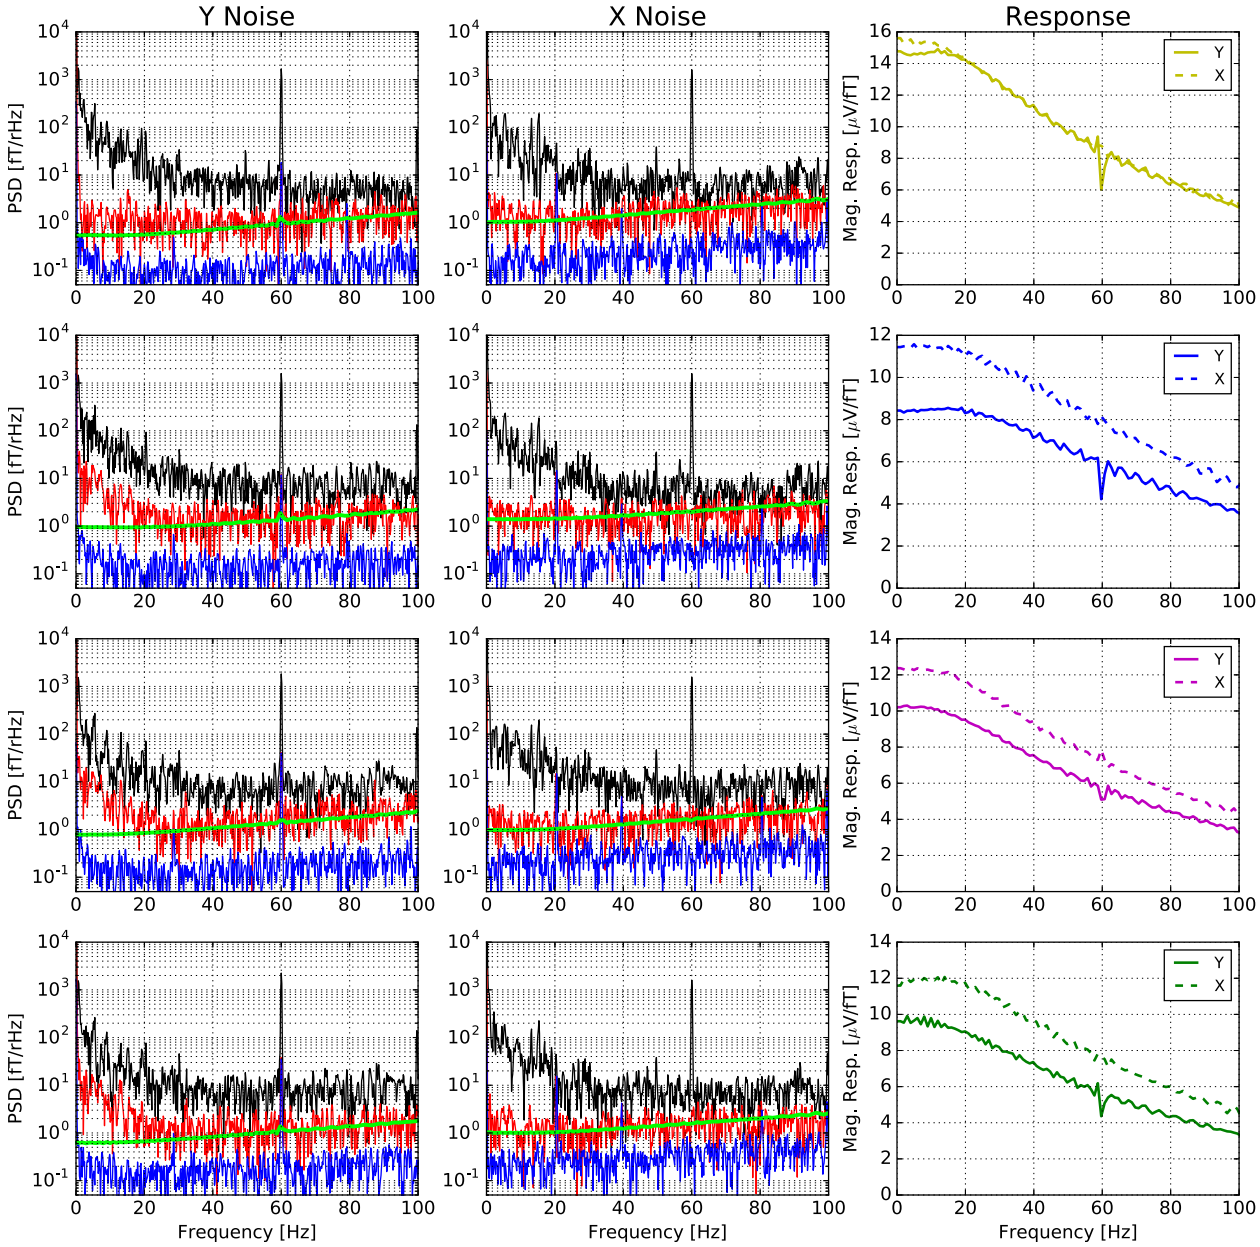



**Figure 4.6:** Response and measured noise with the magnetometers operating in Z-mode. From the top, results from channel 1, 2, 3, and 4. In the "noise" columns, total noise (black), probe noise (red), and electronic noise (blue) are shown. The responses $\zeta_i(\omega)$ for each channel are shown in the right column. Note that our operating procedure (Section 4.6.1) adjusts the applied modulation amplitude to optimize response to $\hat{x}$ fields, which necessarily sacrifices some $\hat{y}$ response. The modulation amplitude on channel 1 was adjusted to make the responses roughly equal. All channels show a total noise floor of $<$10 ft/$\sqrt{\text{Hz}}$ with the technical noise well-separated from the total noise level. Of note, the non-magnetic probe noise is suppressed to the detector photon shot noise level (bright green) at all frequencies for the demodulated ($\hat{x}$) signal.

## 4.7 Cardiographic measurements

The goal of this work is to implement our SERF magnetometers for magnetocardiography so it is useful to periodically measure patients (both adult and fetal) and use the results of these measurements to evaluate the effectiveness of our device. To perform these measurements, the magnetometer array is connected to the end of an out-of-use SQUID detector inside the MSR. The SQUID is maneuverable via a ceiling-mounted gantry which allows for easy positioning.

To reduce the time spent with setup while the subject is in the room, a few steps are taken before the patient arrives. First, the array is mounted and positioned near where the heart signal will be when the patient lies down on the bed. The cells are heated, the lasers are aligned, and the fields at each sensor are nulled using the room and local coils as previously described. Noise measurements are performed with the array mounted to ensure good performance.

For adult measurements, the patient lays down on the bed and the device is placed directly over the center of the chest. Because the adult signal is relatively large (the largest feature of the adult cardiographic signal can be anywhere from 50-100 pT if the array is positioned appropriately) precise alignment is generally not required to achieve sufficient signal to noise ratio for MCG. Once the array is positioned, the room is magnetically sealed, the magnetometers are re-nulled and re-calibrated, and an extended measurement (usually 60-300 s) is initiated. Subsequent measurements of the same type (DC-mode or Z-mode) may require recalibration depending on the length of time since the last recalibration; switching between modes necessitates a recalibration.

For fetal measurements, the signal is much smaller, and positioning of the array to maximize the size of the fetal signal becomes more important. To locate the fetal heart, an ultrasound is performed with the patient lying comfortably on the bed. The device is placed to attempt to center the fetal heart between the four magnetometers, as depicted in Figure 4.7. This often involves tilting and/or rotating the array (easily done with the gantry and SQUID mount); if this is done, all future magnetic field shimming should be done with the shim coils *only*, as the $(x', y', z')$ axes of the magnetometer channels no longer align with the $(x, y, z)$ axes of the room. Additional layers of thermal insulation on top of the magnetometers allow them to be placed very close to the mother's

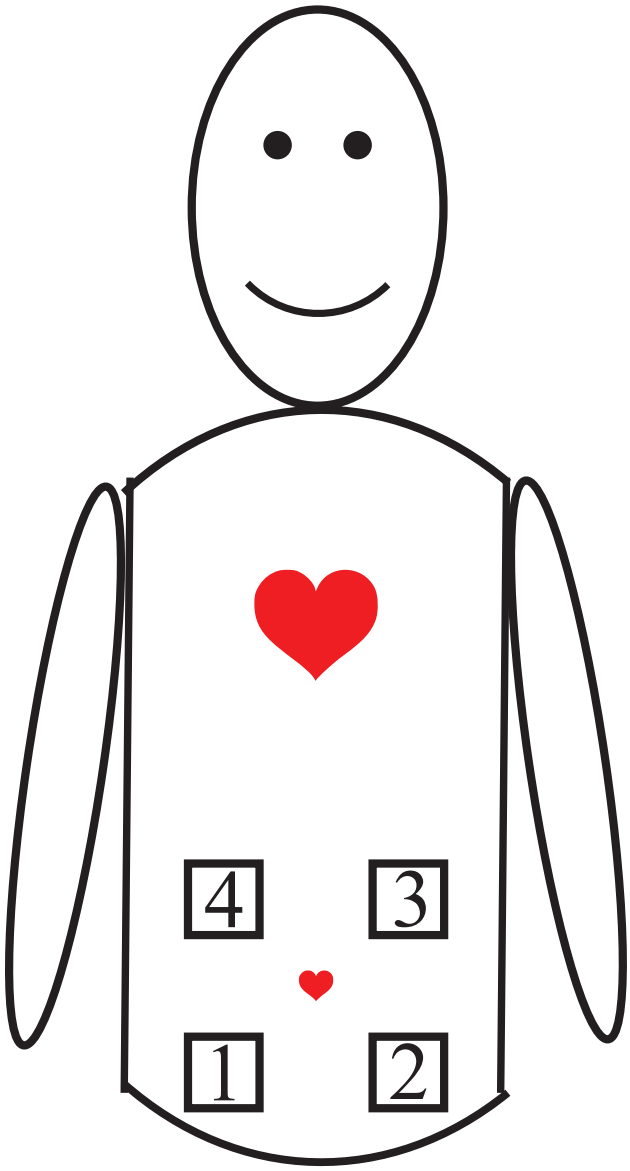



**Figure 4.7:** (Adapted from Wyllie [2012]) A rendition of the relative locations of the maternal and fetal hearts, along with the ideal arrangement of the magnetometers for an fMCG measurement. In this orientation, we expect to see a larger maternal signal on channels 3 and 4 and smaller signals on channels 1 and 2. The phase of the maternal signal should be the same across all channels. The fetal signal is roughly the same size on all channels, though features of the cardiogram will likely have opposite phases depending on the heart's orientation relative to the array. Note that while the subject in this figure is not wearing laser safety goggles, we require that all subjects do so to protect them from stray laser light.

skin without discomfort. We have even experimented with gently pressing against the mother's skin (a technique utilized by the SQUID fMCG group) to bring the detectors even closer to the fetus.

During the measurement, the fetal signal can be visually overwhelmed by noise sources, including the maternal heart signal. This unwanted noise is usually dominated by narrowband 60 Hz line noise along with its higher harmonics (120 Hz and 180 Hz). With LabVIEW, we have implemented a real-time preview of the signal that removes much of this noise and gives us a clearer picture of our signal of interest. The LabVIEW preview takes the raw magnetometer signal and bandpass filters between 1-80 Hz (the motivation for this band is discussed in Chapter 6), notch filters at 60 Hz, and outputs the signal to an oscilloscope. The filtered signal can be monitored for evidence of a fetal signal, and the array can be repositioned until the fetal signal is clearly visible.

## 4.8 Noise Analysis

The fetal magnetocardiographic signals we desire to detect with our atomic magnetometers are small, with the largest features (the so-called $QRS$ complexes, see Chapter 6) having magnitudes on the order of 1 pT when our device is aligned properly. Other cardiographic features' magnitudes can be ten times smaller. To maintain an SNR of $>$10 for these .1 pT features, a noise floor of 10 fT/$\sqrt{\text{Hz}}$ is required. Because of the strict noise requirements for this measurement, it is useful to do an analysis of the various types of noise sources that are present during a magnetic field measurement. We can then use this analysis to pinpoint what exactly could be limiting us when we make fMCG measurements.

The noise sources generally fall into one of three categories: magnetic noise, technical noise, and fundamental noise. Magnetic noise consists of noise sources that truly vary the magnetic field that we are measuring. These include magnetic Johnson noise from conductors and magnetic shields and current noise from electronics and circuits, both of which becomes magnetic field noise through the Biot-Savart Law, and time-varying fields generated by sources both inside and outside the magnetically-shielded enclosure. We define technical noise as noise that is not actually sourced from variations in the magnetic field, but is often indistinguishable from that type of noise during

a measurement. Recall that the raw signal we record is actually a voltage, which is converted to a field measurement by calibration procedure discussed in Section 4.4. Thus, effects like polarization noise in the probe beam or signal aliasing during acquisition can appear to add magnetic noise to the measurement, though they are non-magnetic in nature. Fundamental noise, most notably the spin-projection noise and photon shot noise of the probe laser measurement, give us an absolute limit on the sensitivity of our magnetic field measurement.

The other property of the noise that important for signal processing is whether the noise is correlated between channels or uncorrelated. Much like the magnetic signal we are trying to detect, correlated noise (noise from a magnetic field source outside of the room, *e.g.*) is observable on all magnetometer channels with some degree of amplitude and phase agreement between channels. Using signal-processing or gradiometry techniques, this type of noise can often be subtracted out from the signal, increasing the SNR of the measurement. Uncorrelated noise, which often arises from technical noise sources such as the probe laser, affects each channel differently and cannot to be suppressed using these techniques.

An excellent analysis of the noise present in our system is provided in Chapter 5 of Wyllie [2012]. While some of the specifics of the experiment have changed since that publication, the theory and most of the values for these sources' noise levels are maintained. Here, I will highlight some of the noise sources that are discussed in this work. Noise measurements are taken every time the magnetometer is operated (see Sections 4.4.4 and 4.4.5) to ensure that our experimental values match the claims we present here.

### 4.8.1 Magnetic noise

In this section, we will discuss three contributors to the magnetic noise that we encounter during a measurement: circuit-generated magnetic noise, Johnson noise from nearby conductors (including the MSR itself), and magnetic noise originating from outside of the room.

It should be noted that during a fMCG measurement, the most prominent source of "noise" (*i.e.*, unwanted recorded data) is actually the maternal MCG signal, whose $QRS$ complex can have an amplitude of 10 pT or more, roughly ten times the size of the fetal signal. Of course, this is an

*inescapable* noise source, as the beating heart of the mother cannot be simply "turned off" during a measurement! Strategies for extracting the fetal signal in the presence of a large maternal signal are discussed in Chapters 5 and 6.

#### 4.8.1.1 Circuit-induced magnetic noise

The four electronic components which comprise the magnetometer channel all have currents that potentially contribute to the magnetic noise seen by the magnetometer. The largest current by far is the one flowing through the heater panels that warm the cell ($\sim$200 mA). However, the operation of the heaters with a sinusoidal drive function at 150 kHz or higher (far outside our bandwidth of interest) and the careful design of the heater panels (counter-wound traces to reduce generated fields) mitigates the effect of these large currents on the total magnetic noise. The resistance temperature detectors (RTDs) that monitor the temperature of the cells only carry currents when the temperature is actually being measured by an ohm-meter, and they are always disconnected during a magnetic field measurement (though even leaving it connected shows no noticeable increase in detected noise level). The photodiode currents are usually very small ($<$100 $\mu$A) and the shielded twisted pairs of wires leading from the detectors produce negligible magnetic fields at the location of the magnetometers.

This leaves the coil driver circuits themselves. Unlike the other electronics and circuits mentioned above, the job of these drivers and coils is to *produce magnetic fields*. Thus, design considerations in the current-carrying wires (such as using shielded twisted pairs or counter-wound traces) to reduce current-created magnetic fields are irrelevant here. We have done our best to build low-noise circuits, but despite those efforts, they represent the dominant source of magnetic noise in our experiment. We measure the current output noise on the tri-axial "Sulai" current supplies to be $\sim$10-500 pA/$\sqrt{\text{Hz}}$ (most of the variation comes from the choice of output resistor), which corresponds to a noise spectral density of $\delta\tilde{B}_{\text{coils}}$ $\approx$1-10 fT/$\sqrt{\text{Hz}}$ . This conversion comes from the current-to-field calibration values for the shell-mounted coils shown in Table 3.2, as shown in Figure 3.12.

The upper limit on this noise is realized in the noise generated by the $\hat{z}$ coils, due to the need to use a larger output resistor to source the current required for Z-mode modulation. For this reason, though, this noise is primarily along the $\hat{z}$-axis of the magnetometer, and it only couples into the sensitive directions ($\hat{x}$ and $\hat{y}$) through imperfect nulling of those fields. For the $\hat{y}$ direction, we estimate a total white noise level of closer to 4 fT/$\sqrt{\text{Hz}}$ for the shell coils. The room-mounted coils (which use a similar current supply, but have a factor of $\sim$3 smaller current-to-field calibration value) add roughly 1.4 fT/$\sqrt{\text{Hz}}$ of magnetic noise. Thus, we can estimate our total magnetic noise due to electronics as

$$\delta\tilde{B}_{\text{coils}} = \sqrt{(\delta\tilde{B}_{\text{shell}})^2 + (\delta\tilde{B}_{\text{room}})^2} = 4.2 \text{ fT}/\sqrt{\text{Hz}}. \tag{4.22}$$

(Note: for the rest of this section, the notation $\delta\tilde{B}_{\text{j}}$ will be used to denote the noise spectral density (in units of field per root Hertz) from noise source j.) The above number represents the estimated *white*-noise level of the electronics. Many of the circuit components exhibit some amount of $1/f$ behavior at frequencies below 10 Hz. For example, examining Figure 3.12 shows the coil driver noise on the $\hat{y}$ coils (green) to be near 15 ft/$\sqrt{\text{Hz}}$ at a frequency of 2 Hz, but falling off to the predicted value of $\sim$4 ft/$\sqrt{\text{Hz}}$ by about 8 Hz.

#### 4.8.1.2 Johnson noise

Johnson noise arises from the random movement of thermal electrons inside of electrical conductors. The voltage noise (per root hertz) of a conductor of resistance $R$ and temperature $T$ is given by the well-known formula

$$\delta\tilde{V}_{\text{John}} = \sqrt{4k_B T R}, \tag{4.23}$$

where $k_B$ is the Boltzmann constant, $1.38\times10^{-23}$ J/K. The equivalent current noise is then simply

$$\delta\tilde{i}_{\text{John}} = \frac{\delta\tilde{V}_{\text{John}}}{R} = \sqrt{\frac{4k_B T}{R}}. \tag{4.24}$$

The magnetic field-equivalent Johnson noise levels for several different conductor geometries are provided in [Griffith et al., 2010, Lee and Romalis, 2008, Wyllie, 2012]. The most relevant ones are listed below.

1. *The innermost walls of the magnetically shielded room:* We model the room as six infinite planes of μ-metal at a distance of $a = 1.5$ m from the magnetometer. The noise contribution from one such plate is

$$\delta\tilde{B}_{\text{wall}} = \frac{1}{\sqrt{6\pi}}\frac{\mu_0\sqrt{k_B T\sigma d}}{a}, \tag{4.25}$$

where $T$ is the wall temperature, $\sigma$ is the electrical conductivity of the μ-metal, and $d$ is the wall thickness. Using appropriate values for these variables gives $\delta\tilde{B}_{\text{wall}} = .52$ fT/$\sqrt{\text{Hz}}$ . The noise contribution for each direction is just the uncorrelated (quadrature) sum of two such walls, so $\delta\tilde{B}_{\text{MSR}} = .73$ fT/$\sqrt{\text{Hz}}$ for each axis.

2. *Solid $^{87}$Rb coating the walls of the cell:* Additionally, in some cells there are droplets of excess rubidium that tend to condense at the coldest part of the cell, which is often near the windows. For simplicity, we will model all of this rubidium as a disc of thickness $d = 10$ nm, radius $r = 1$ cm at a distance of $a = 5$ mm from the measurement volume. Then the noise contribution is

$$\delta\tilde{B}_{\text{Rb}} = \frac{1}{\sqrt{8\pi}}\frac{\mu_0\sqrt{k_B T\sigma_{\text{Rb}} d}}{a}\frac{1}{1+a^2/r^2} = 0.9 \text{ fT}/\sqrt{\text{Hz}}. \tag{4.26}$$

3. *Wires forming the shell-mounted nulling coils:* We can model a single set of shell-mounted coils as ten wire segments (four that comprise the coil at one end of the cell, four at the other, and two that connect the coils together). Three such coil pairs gives a total of 30 wire segments. We will approximate each of these 28 gauge wires (with radius $y_0 = .160$ mm) having an approximate mean distance of 2.5 cm to the measurement volume. For a single wire, for $a \gg y_0$, the Johnson noise contribution is

$$\delta\tilde{B}_{\text{wire}} = \sqrt{\frac{3}{128}}\mu_0\sqrt{k_B T\sigma_{\text{Cu}}}y_0^2 a^{-5/2} \tag{4.27}$$

Summing 30 such wires in quadrature (ignoring the order-one factors which account for the direction of each magnetic field contribution and the correlation of noise between wires of the same loop) gives a noise contribution of $\delta\tilde{B}_{\text{coils}} = 0.13$ fT/$\sqrt{\text{Hz}}$ per axis.

4. *Heater panels:* The heater panels present a much more complex conduction path than the shell-mounted coils. Instead of attempting an analytical expression of their associated Johnson noise, a *Mathematica* simulation was performed using the real values of the platinum trace thickness and resistivity. We found that the four panels of oven combined contribute .15 fT/$\sqrt{\text{Hz}}$ of magnetic Johnson noise per axis.

#### 4.8.1.3 External noise sources

Finally, we consider magnetic noise sources that originate from outside of the MSR. These sources vary day-to-day and even minute-to-minute depending on what is going on around the laboratory building. The two largest sources of external noise during the work of Wyllie [2012] (the HVAC fan and construction near the building) were no longer present during this work–the fan was replaced with a non-magnetic version and the construction was largely completed. As a result of the fan being replaced, the large $\sim$6 Hz spike observed in the power spectral densities in Wyllie's thesis were non-existent.

However, there is still traffic on a road within 100 m of the shielded room, as well as nearby elevators which are frequently used. Additionally, we suspect that being part of the medical physics wing of the building, we are quite close to the large magnetic fields produced by MRI machines and other imaging equipment, though their exact locations are unknown. The fields we measure from these sources can be quite large (nT-sized), even inside the shielded room which offers a factor of $\sim$250 suppression of these fields. Often, the Faraday rotation of the probe beam caused by these fields will overwhelm the sensor, taking it out of the middle (*i.e.*, most responsive part) of the dispersive curve seen in Figure 2.10.

Despite the large magnitude of these noise fields, they are often less consequential because their frequency components generally lie in the $<$1 Hz band, which does not overlap with the MCG signals we are trying to detect. We can often use narrow bandwidth feedback to counteract these slowly drifting fields and hold the magnetometer signal at zero.

With all of this being said, we have noticed an increase (on the order of 50%) in the total noise in the DC-20 Hz band over the past couple of years. That is, the noise in this band appears to not

fall off as quickly to the noise baseline as was observed in measurements taken by Wyllie several years ago. By making total, technical, and electronic noise measurements (such as the ones in Figure 4.4), we are confident that this is real magnetic noise, not a result of technical noise limiting the measurement. After ruling out coil driver noise and other magnetic sources inside the MSR, we have come to the conclusion that this noise must be from a source outside the MSR, which is not being shielded sufficiently. We believe that a de-gaussing of the room (which as of this writing, has not been completed in over six years) may be required to suppress this noise and restore the low-frequency performance that we have seen in the past.

### 4.8.2 Technical noise

The non-magnetic technical noise in our experiment comes primarily from two sources: laser light fluctuations (from the laser diodes themselves, acquired in the optical fibers, or from air currents and other refraction index during free space propagation) and acquisition noise (due mostly to aliasing at the ADC). Here we will focus on what we believe to be the primary culprit: probe light intensity fluctuations and how the sensitivity to these fluctuations increases with increasing imbalance of the photodiode detector. We will also discuss how we limit the acquisition noise through anti-aliasing filters.

#### 4.8.2.1 Probe and detector imbalance noise

The need for modularity in our experiment necessitates the use of optical fibers to couple light generated outside of the magnetically shielded room to the apparatus. With polarization-*maintaining* fibers, significant noise can be picked up if the the linearly-polarized light is not precisely coupled down one of the axes (normally the slow axis) of the fiber. *Polarizing* fibers, which are used in this experiment, do not suffer from this issue, as any component of the light *not* aligned with the axis of the polarizing fiber is quickly attenuated in the fiber, (rather than sent down the fiber's fast axis, as it the case with PM fibers).

However, the fibers are specified at a 30 dB extinction ratio, meaning there is still opportunity for noise to be picked up in the fiber (at the $\sim$1/1000 level of the laser intensity). Additionally, the fiber splitter that divides the probe light into two separate beams is *not* polarizing, only polarization-maintaining. The result is that at the output of the pump tubes, some amount of polarization noise exists. With the arrangement of lasers and optics shown in Figure 3.5, any remaining polarization rotation noise is converted to intensity noise by high-extinction linear polarizers just before the probe light enters the cell. This noise has the same $1/f$ character as is common in fiber noise, caused by low-frequency vibrations and air currents disturbing the fiber. In particular, we have noticed that the MSR has vibrational resonances in the 14-16 Hz band. (These vibrational resonances were verified by the author re-creating the "Jump Around" tradition at Wisconsin football games inside the MSR while another researcher observed the detector signal). This noise is clearly visible in the probe noise measurements in Figure 4.4 (the measurements in this figure were made without the author attempting to recreate the "Jump Around" tradition the student section enjoys on Badger football Saturdays).

Our subtractive detector is specifically designed to mitigate the effects of this intensity noise. Any intensity noise should be present on both arms of the detector and subtracted. However, the *degree* to which this suppression takes place depends on the balance of the detector. Recall from Equation (3.3) that the difference current output from the detector is

$$\Delta I = I_2 - I_1 = -I_0 \sin(2\alpha) \approx -2\alpha I_0. \tag{4.28}$$

Therefore, any intensity noise $\delta I_0$ will result in noise on that difference signal which scales in a similar way:

$$\delta(\Delta I) \approx -2\alpha \delta I_0. \tag{4.29}$$

Notice that this noise is zero if the detector is perfectly balanced ($\alpha = 0$), and the noise increases linearly (for small $\alpha$) as the detector becomes more and more imbalanced. The exact level of noise is difficult to quantify, as the balance is one of many factors that influence its magnitude. However, we have seen equivalent probe noise levels at DC range from $<$1 fT/$\sqrt{\text{Hz}}$ to more than

100 fT/$\sqrt{\text{Hz}}$ depending on the balance of the detector. Clearly, suppression of this intensity noise through detector balancing is critical to maintaining an effective magnetometer.

#### 4.8.2.2 Signal aliasing and electronic noise

Signal aliasing occurs in any system where a continuous signal is discretely sampled. Nyquist's Theorem states that for a sampling rate of $f_{\text{samp}}$, the maximum frequency that can be uniquely represented by that sampling frequency is $f_{\text{Nyq}} = f_{\text{samp}}/2$. Any signal with frequency components above $f_{\text{max}}$ will have those components "aliased" back to the range DC-$f_{\text{Nyq}}$. For example, for a 2 kHz sampling rate ($f_{\text{Nyq}} = 1$ kHz), a sinusoidal function with $f = 1010$ Hz will appear to be a 10 Hz signal. This can be quite problematic for devices like ours where low-frequency noise in the 1-100 Hz range is particularly troublesome. Some of our noise-reduction strategies (such as operating the high-current heater signals at $f \approx 150$ kHz) can be foiled if the noise at that frequency is simply aliased down to our band of interest.

The solution to this issue is to apply low-pass filters to the signals before sampling, which can be accomplished using the built-in filters on the SRS DS345 transimpedance amplifiers. The problem with this strategy is the large phase shifts which accompany any low-pass filtering of a signal can degrade the performance of an active feedback circuit that is using that signal as an input. Thus, we generally leave the filters on the amplifiers disabled. Instead, we utilize a two-stage sampling procedure with our DAQ, where the data are sampled by the ADC at 500 kHz and immediately downsampled (via averaging) to 20 kHz while applying a digital anti-aliasing filter. These data are then demodulated (to recover $B_x$ if a Z-mode measurement is taking place) and low-passed again at 100 Hz to remove all traces of higher-frequency components. With this filtering chain in place, we generally see the white electronic noise at a level on the order of 1 pA/$\sqrt{\text{Hz}}$ , roughly equivalent to .1 fT/$\sqrt{\text{Hz}}$ in field noise (at DC).

### 4.8.3 Fundamental noise

The fundamental sensitivity of an atomic magnetometer is given by [Budker and Romalis, 2007, Kominis et al., 2003] and was given in Equation (2.3). A modified version is reprinted here

for reference.

$$\frac{\delta B}{\sqrt{\text{Hz}}} = \frac{1}{\gamma\sqrt{nVT_2}} \tag{4.30}$$

In a SERF magnetometer, this value can be theoretically as low as 0.01 fT/$\sqrt{\text{Hz}}$ [Kominis et al., 2003]. However, this value is never realized in our experiment as a result of the noise sources presented in the previous two sections. It is still useful to explore the two fundamental limits on noise: the quantum spin-projection noise of the polarized atoms and the photon shot noise limit of the detection. The photon shot noise level is particularly useful for comparing to the measured probe noise: a probe noise level far above the equivalent photon shot noise points to a severe problem with the probe beam or detector. Ledbetter et al. [2008] gives derivations for both of these quantities; we will follow their calculations for the spin projection noise, and then present a more experimentally-driven derivation of the shot noise.

#### 4.8.3.1 Spin projection noise

Spin projection noise arises from a measurement of $P_x$ which is perpendicular to the quantization axis $\hat{z}$, as defined by the pump laser. We begin by writing the uncertainty in the $x$ component of the total angular momentum ($F_x$) as $\delta F_x = \sqrt{\langle F_x^2\rangle - \langle F_x\rangle^2}$. Assuming we're working in a zero magnetic field spin-temperature regime ($\rho \propto \exp(\beta F_z)$, where $\rho$ is the population in a given $F_z$ sublevel and $\beta$ is the spin-temperature parameter defined in Equation (2.8)) with all of the atomic polarization along $\hat{z}$, $\langle F_x\rangle = 0$ by symmetry. We can evaluate the first term by noting that $\langle F_x^2\rangle = \text{Tr}(\rho F_x^2)$, and the trace evaluates to $q(P)/4$ per atom when the trace is normalized so that $\text{Tr}(\rho) = 1$ (the slowing-down factor $q$ varies with polarization $P$ according to Equation (2.11)).

By Poisson statistics, measuring a collection of $N = nV$ uncorrelated atoms will reduce this uncertainty by a factor of $\sqrt{nV}$, so

$$\delta F_x(P) = \sqrt{\frac{q(P)}{4nV}}. \tag{4.31}$$

After a long ($t \gg T_2$) measurement, the uncertainty is [Kornack, 2005]

$$\langle \delta F_x \rangle_t = \delta F_x(P) \sqrt{\frac{2q(P)}{\Gamma' t}} \tag{4.32}$$

$$= \frac{q(P)}{\sqrt{2\Gamma' n V t}} \tag{4.33}$$

where $\Gamma'$ is the total atomic relaxation rate defined in Equation (2.37). In spin-temperature, the total atomic spin is related to the electronic spin through the slowing down factor ($\mathbf{F} = q(P)\mathbf{S}$). Using $P_x = 2\langle S_x \rangle$, we can write

$$\delta P_x = \frac{2}{q(P)} \langle \delta F_x \rangle_t = \sqrt{\frac{2}{\Gamma' n V t}}, \tag{4.34}$$

which has no dependence on the slowing-down factor $q(P)$. Recall in DC-mode (Section 2.3), the transverse polarization $P_x$ is induced by a magnetic field $\Omega_y = \gamma B_y$. If we require that the fields are well nulled so $P_z = R/\Gamma'$ and $P_x = P_z(\Omega_y/\Gamma') = (R/\Gamma'^2)(\gamma B_y)$, then

$$\delta B_{y,\text{spn}} = \frac{\Gamma'}{\gamma R} \sqrt{\frac{2\Gamma'}{nVt}} \tag{4.35}$$

$$\delta \tilde{B}_{\text{y,spn}} = \frac{\Gamma'}{\gamma R} \sqrt{\frac{2\Gamma'}{nV}} \tag{4.36}$$

If we assume that we are working under optimal conditions for magnetometry, (i.e. $R = \Gamma = \Gamma'/2 = 500$ s$^{-1}$), assume the volume of atoms participating in the measurement is $\pi w_{pr}^2 l$ (where $w_{pr} = .3$ cm is the probe waist and $l = 1$cm is the length of the cell), and operate at a temperature of 155°C where the atomic density is $n = 1 \times 10^{14}$ cm$^{-3}$, then Equation (4.36) evaluates to a spin-projection noise level of 0.4 fT/$\sqrt{\text{Hz}}$ .

### 4.8.3.2 Photon shot noise

Ledbetter et al. [2008] also include a theoretical derivation of the photon shot noise given the initial probe photon intensity, and assuming a far-detuned probe laser. In this regime, the Faraday rotation scales as $1/\Delta_{\text{D2}}$, and the absorption as $1/\Delta_{\text{D2}}^2$, making it an advantageous regime to perform measurements. Normally, we operate in this regime, but at times we have desired to tune the probe laser closer to the D2 transition to make measurements, in which case the derivation

in [Ledbetter et al., 2008] does not necessarily hold true. Thus, in this work, we simply define the photon shot noise as shot noise in the photodiode currents resulting from detected photons.

Photons striking the photodiodes produce electrons at a rate commensurate with the photodiode's quantum efficiency. At 780 nm, the efficiency $\eta$ of our photodiodes is approximately .53 A/W, equivalent to .8427 emitted electrons per photon, or in other words, one electron emission per 1.18 absorbed photons. The shot noise associated with this photocurrent is [Hobbs, 1997]

$$\delta i_{\text{shot}} = \sqrt{2ei\Delta f} \tag{4.37}$$

where $e$ is the fundamental electron charge, $i$ is the generated photocurrent, $\Delta f$ is the measurement bandwidth. Our experiment utilizes two such photodiodes to make the Faraday rotation measurement, and the shot noise associated with the generated photocurrents is uncorrelated. Thus the total shot noise associated with our measurement is

$$\delta i_{\text{shot,total}} = \sqrt{\left(\delta i_{\text{shot,1}}\right)^2 + \left(\delta i_{\text{shot,2}}\right)^2} \tag{4.38}$$

Assuming that the detector is well-balanced (i.e., $i_1 \approx i_2 \equiv i_0$), then the contributions from each diode to the total shot noise are equal as well. Thus,

$$\delta i_{\text{shot,total}} = \sqrt{2\left(2ei_0\Delta f\right)} \tag{4.39}$$

$$= \sqrt{2ei_{\text{sum}}\Delta f} \tag{4.40}$$

$$= \sqrt{2e\left(\eta\mathcal{P}'_{\text{pr}}\right)\Delta f}, \tag{4.41}$$

with $i_{\text{sum}}\,(= i_1 + i_2 = 2i_0)$ the total current generated by the photodiodes as a result of a probe beam with transmitted power of $\mathcal{P}'_{\text{pr}}$. At first glance, it might seem that using a low-power probe would be the best way to reduce the shot noise. If all that mattered was the current noise, this would be true. However, what we are interested in is the equivalent magnetic field noise contribution from this current noise. Recall from Section 2.3.7 that the Faraday rotation signal is directly proportional to the incident probe power $\mathcal{P}_{\text{pr}}$ (Equation (2.58)), while the noise is only proportional to $\sqrt{\mathcal{P}'_{\text{pr}}}$. Assuming no non-linear absorption effects, the transmitted power should be the same fraction of the incident power regardless of the incident power level, so

$$\text{SNR} \propto \sqrt{\mathcal{P}_{\text{pr}}}. \tag{4.42}$$

Thus, the performance of the magnetometer (with respect to the shot noise) actually benefits from a *larger* probe intensity. On the other hand, too high a probe power can lead to $\Gamma_{\text{pr}} \approx \Phi_{\text{pr}}\sigma_{\text{pr}}(\Delta\nu)$, Equation (2.23)) becoming the dominant relaxation mechanism, so care must be taken to balance this effect with the desire for lower shot noise.

To calculate this noise in equivalent magnetic field units, one must simply divide this value by product of the current-to-voltage gain $g_{\text{IV}}$ from the transimpedance amplifier and the voltage-to-magnetic field calibration $\zeta^(\omega)$. The details of obtaining this calibration can be found in Section 4.4. In other words,

$$\delta B_{\text{shot}}(\omega) = \frac{\sqrt{2ei_{\text{sum}}\Delta f}}{g_{\text{IV}}\zeta(\omega)} \tag{4.43}$$

$$\delta \tilde{B}_{\text{shot}}(\omega) = \frac{\sqrt{2ei_{\text{sum}}}}{g_{\text{IV}}\zeta(\omega)}, \tag{4.44}$$

where in the last line, we divided by $\Delta f$ to express the noise in its spectral density form. Experimentally, we obtain an estimate for the photon shot noise by measuring $i_{\text{sum}}$ (using the detector in its "sum" mode), and calculating $\delta\tilde{B}_{\text{shot}}$ after calibration is complete. We tend to plot the photon shot noise every time we make a probe noise measurement. A probe noise measurement significantly above the shot noise level is often indicative of a significant imbalance of the detector (Section 4.8.2.1), which needs to be corrected

Using common experimental parameters of $i_{\text{sum}} = 250\ \mu A$, $g_{\text{IV}} = 1$ μA/V, and $\zeta(\text{DC}) = 30$ μV/fT (for DC-mode), we estimate the shot noise contribution to be $\delta\tilde{B}_{\text{shot}} = 0.3$ fT/$\sqrt{\text{Hz}}$ at DC, rising at higher frequencies where the magnetometer response $\zeta(\omega)$ is smaller. The exact behavior of the magnetic-equivalent shot noise as a function of frequency is shown in Figure 4.4.

# Chapter 5

# Gradiometry

In the previous chapters we have discussed atomic magnetometry and how we make measurements with our SERF magnetometers, with the end goal being the measurement of a MCG or fMCG signal. In this chapter, we explore setting up two of these magnetometers in a gradiometer setup: that is, subtracting the signals from two nearby SERF magnetometers with the goal of removing noise that is correlated between channels. One concern for a gradiometric approach in measuring fMCG is the inevitable loss of signal that accompanies any measurement which subtracts signals. However, we expect many of the noise sources (including the maternal signal) to be much more uniform than the fetal signal given the positioning of the array (Figure 4.7) so a gradiometric measurement of the magnetic field should increase the signal-to-noise ratio of the fetal signal.

This chapter has been adapted from a soon-to-be-published journal article.

## 5.1 Introduction

Magnetic gradiometers can exhibit higher signal to noise ratios than magnetometers for detecting weak local magnetic field sources. This is true in environments where the primary noise sources are far from the detectors. In this paper, we report studies of a gradiometer constructed from two independent SERF [Kominis et al., 2003] atomic magnetometers. We achieve common-mode rejection comparable to a commercial superconducting quantum interference device (SQUID) gradiometer, and explore potential benefits of using atomic gradiometers for fetal magnetocardiography (fMCG).

Presently, the preeminent method for measuring fMCG employs the use of low Tc SQUID gradiometers with many ($> 20$) channels [Quartero et al., 2002, Strasburger et al., 2008]. Commercially available SQUID gradiometer arrays have demonstrated robust, low-noise performance in clinical settings. However, fMCG is currently conducted inside magnetically shielded rooms of sufficient size to accommodate the relatively large liquid helium dewars required for operation of the low Tc SQUIDs. The large size of the shielded rooms, and the necessary cryogenics, make SQUID gradiometers expensive to set up and operate. Atomic gradiometers, by contrast, require no such cooling. They are relatively inexpensive to operate and have significantly lower setup costs due to a smaller shielding volume requirement.

Previous work [Wyllie et al., 2012a] has demonstrated the suitability of a SERF atomic magnetometer for making fMCG measurements. Here we investigate the potential advantages of using the atomic magnetometers in a gradiometric configuration. In addition to suppressing correlated noise natively present in the measurement environment, a gradiometer is also more insensitive to the magnetic field produced by the maternal heart, which given its greater distance from the gradiometer, has a more uniform spatial distribution than the nearby fetal signal. To determine what level of maternal signal suppression was desired for a fMCG measurement, maternal and fetal MCG waveforms were created using the ECGSYN [McSharry and Clifford, 2015] program. The waveforms were scaled appropriately to the field strengths generally observed during a fMCG measurement [Wyllie, 2012]. These waveforms, along with their Power Spectral Densities, are shown in Figure 5.1. Note at DC, the maternal signal is roughly ten times the size of the fetal signal. If we assume the maternal heart signal, which originates $\sim$60 cm from the device (compared to the $\sim$5 cm distance to the fetal heart), is uniform across the two channels of the gradiometer, a common-mode signal suppression of at least 100 is required to achieve a 10:1 fMCG:mMCG signal ratio.

Additionally, we endeavored to match or exceed the common-mode noise suppression performance of a commercial SQUID gradiometer array [Vrba, 1996], such as a device in use at the University of Wisconsin-Madison (Magnes BTI, 4D Neuroimaging, Inc., San Diego, CA). The

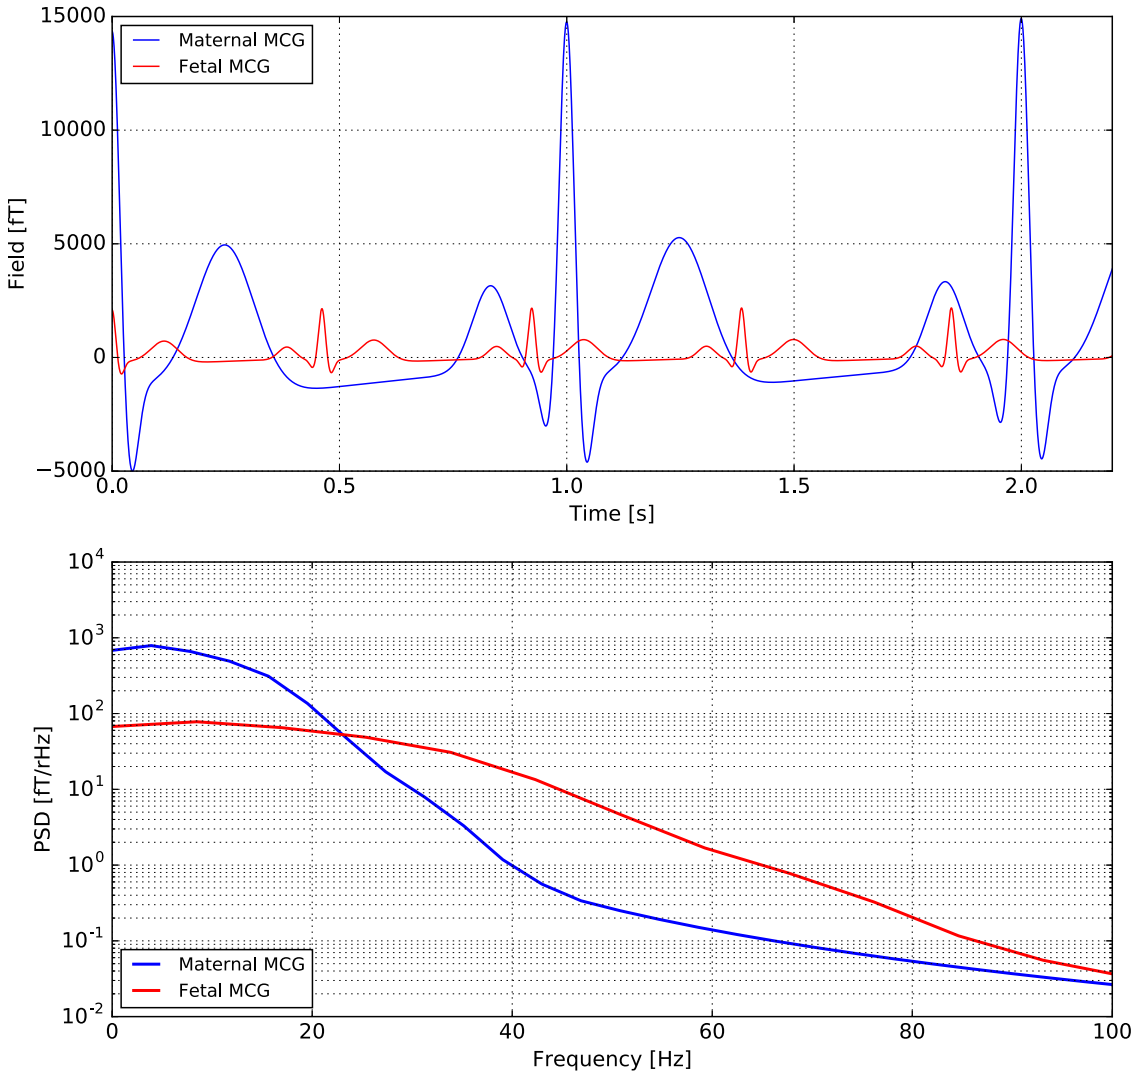



**Figure 5.1:** Time series generated using the ECGSYN program, along with their power spectral densities. To achieve a 10:1 fMCG:mMCG signal ratio, suppression of the maternal signal by at least a factor of 100 is required.

SQUIDs are coupled inductively, and measure the difference in detected flux between two counter-wound pickup coils. The array's ability to suppress common-mode noise is limited by matching of the areas and alignments of these coils, which are often wound by hand for many commercial devices. The areas of these hand-wound loops [Clarke and Braginski, 2005, **?**] are accurate to one part in 100 – setting the upper bound of the common-mode suppression at that level.

We demonstrate that our high-sensitivity atomic magnetometer surpasses the performance of a commercial SQUID in terms of its ability to suppress common-mode "correlated" noise and furthermore, can maintain that suppression over the timescale necessary to complete a typical fMCG measurement. We focus here on the performance of a single gradiometer channel – noting that the method can be easily extended to multiple channels by adding additional SERF magnetometers. With comparable sensitivity to a commercial SQUID used for fMCG, we believe a four-channel device measuring two components of the field (using parametric modulation of the sensors) provides a compelling alternative gradiometric device for fMCG.

## 5.2 Gradiometer theory and figure of merit

Consider a configuration as shown in Figure D.1. A magnetic field source is located a distance $r_s$ from a magnetometer, with a second magnetometer positioned at an additional distance of $L$ along the same axis. The source is assumed to generate a magnetic field with a power law $B(r_s) \propto (1/r)^p$. For example, a point source magnetic dipole in a uniform medium would have a profile

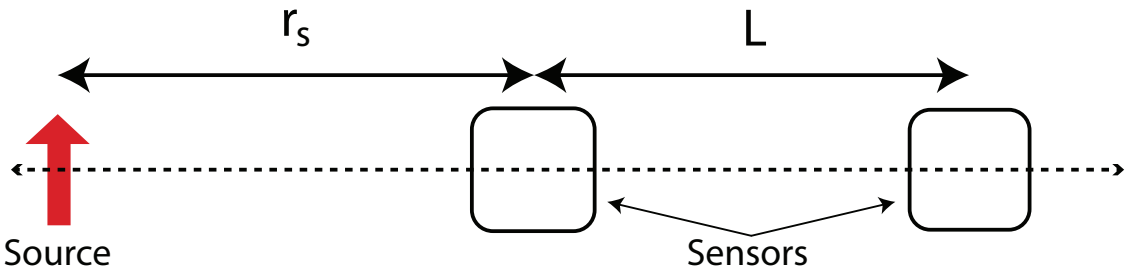



**Figure 5.2:** Configuration of source and magnetometers. The source, a fetal heart is modeled to be a current dipole $Id\vec{l}$. In the cartoon, the dipole is oriented perpendicular to the axis of the magnetometers, but it should be noted the intra-heart currents that generate the detected fields flow along all three axes.

with $p = 3$. On the other hand, the field from a heart – relevant for MCG – is often modeled as a current dipole [Cohen and Hosaka, 1976], with $p \sim 1.8 - 2.2$ .

We express the output signal $S_i$ of each magnetometer as $S_i = M_i B(r_i)$ where $M$ is a response function and $B(r_i)$ is the field magnitude at a distance $r_i$ from the source. The gradiometer signal $S_G$ is given by the difference of the two magnetometer signals:

$$S_G = S_1 - S_2 = M_1 B_1 - M_2 B_2 \tag{5.1}$$

The signal from a well-balanced (i.e. $M_1 \approx M_2 = M$) gradiometer in a magnetic field $B_i = b_0/r_i^p$ is

$$\begin{aligned} S_G \approx M(B_1 - B_2) &= M\left(\frac{B_0}{r_s^p} - \frac{B_0}{(r_s + L)^p}\right) \\ &= \frac{-r_s^p + (r_s + L)^p}{(r_s + L)^p} M B_1 \\ &= \frac{g^p - 1}{g^p} M B_1 \end{aligned} \tag{5.2}$$

where $g = 1 + L/r_s$, and $B_1$ is the magnetic field strength at the magnetometer closest to the source.

In practice, the magnetometers can have different amplitude and phase responses to identical magnetic field inputs. The resulting gradiometer consequently will acquire a residual sensitivity to correlated magnetic fields. The undesired residual sensitivity is inversely proportional to the common-mode rejection ratio $\xi$, where

$$\begin{aligned} \xi &= \frac{\frac{1}{2}(S_{1c} + S_{2c})}{S_{1c} - S_{2c}} \\ &= \frac{\frac{1}{2}(M_1 + M_2)}{M_1 - M_2}. \end{aligned} \tag{5.3}$$

$S_{ic}$ is the measured signal from magnetometer $i$ due to an applied correlated field. In SQUID gradiometry [Vrba, 1996], $\xi$ is limited by the matching pickup coil areas or angular alignment of the two coils. In the case of atomic magnetometers, it is limited by the calibration procedure.

In general, a measurement is made in a magnetically noisy environment with total noise $\delta B = \sqrt{(\delta B_u)^2 + (\delta B_c)^2}$. $\delta B_u$ and $\delta B_c$ are the uncorrelated and correlated sources of magnetic noise,

respectively. The SNR for the two-channel gradiometer is then

$$\begin{aligned} SNR_G &= \frac{\frac{g^p-1}{g^p} M B_1}{\sqrt{(M\delta B_{u1})^2 + (M\delta B_{u2})^2 + 2\left(M\delta B_c/\xi\right)^2}} \\ &= \frac{g^p-1}{g^p} \frac{B_1}{\sqrt{2\left(\delta B_u\right)^2 + 2\left(\delta B_c/\xi\right)^2}}. \end{aligned} \tag{5.4}$$

Here we have made the assumption that $\delta B_{u1} \approx \delta B_{u2}$ for the two magnetometers.

If, instead, the two magnetometers (again, with $M_1 \approx M_2 = M$) were not operated as a gradiometer, then the total signal is

$$\begin{aligned} S_M &= S_1 + S_2 = M\left(B_1 + B_2\right) \\ &= M\left(\frac{B_0}{r_s^p} + \frac{B_0}{(r_s+L)^p}\right) \\ &= \frac{g^p+1}{g^p} M B_1 \end{aligned} \tag{5.5}$$

Thus the signal-to-noise ratio for this two-channel magnetometer is

$$SNR_M = \frac{g^p+1}{g^p} \frac{B_1}{\sqrt{2\left(\delta B_u\right)^2 + 2\left(\delta B_c\right)^2}}. \tag{5.6}$$

Comparison of equations D.4 and D.6 illustrates the potential advantage of using a gradiometer, which we define as the gradiometer figure of merit $\mathcal{F}$.

$$\begin{aligned} \mathcal{F} = \frac{SNR_G}{SNR_M} &= \frac{g^p-1}{g^p+1} \frac{\sqrt{2\left(\delta B_u\right)^2 + 2\left(\delta B_c\right)^2}}{\sqrt{2\left(\delta B_u\right)^2 + 2\left(\delta B_c/\xi\right)^2}} \\ &= \frac{g^p-1}{g^p+1} \frac{\sqrt{\left(\frac{\delta B_u}{\delta B_c}\right)^2 + 1}}{\sqrt{\left(\frac{\delta B_u}{\delta B_c}\right)^2 + \xi^{-2}}} \end{aligned} \tag{5.7}$$

Gradiometers are only useful if the total noise is dominated by correlated noise between the two channels. Under the assumption that $\delta B_u/\delta B_c \ll 1$,

$$\mathcal{F} \approx \frac{g^p-1}{g^p+1} \frac{1}{\sqrt{\left(\frac{\delta B_u}{\delta B_c}\right)^2 + \xi^{-2}}} \tag{5.8}$$

This clearly demonstrates the the figure of merit depends on not only the noise ratio of $\delta B_u/\delta B_c$, but also how that ratio compares to $\xi$. If $(\delta B_u/\delta B_c)^2 \ll \xi^{-2}$, immediate and potentially substantial benefits to the figure of merit (and therefore gradiometer performance) will be gained through an increase of $\xi$. If, instead, $(\delta B_u/\delta B_c)^2 \gg \xi^{-2}$, further experimental efforts would be best focused elsewhere.

It is important to note here that the ratio $(g^p - 1)/(g^p + 1)$ is somewhat geometry-sensitive and will also depend on the power law characterizing the field, as shown in Figure D.2. In setting the appropriate baseline, a common practice is to match the baseline with the depth of the signal of interest, ie $L \approx r_s$, though simulations [Garachtchenko et al., 1999] have shown the optimal baseline is a function of balance, signal depth, and ambient noise source distance.

Taking $p = 2$ and $L = r_s$, for example,

$$\mathcal{F} \approx \frac{3}{5} \frac{1}{\sqrt{\left(\frac{\delta B_u}{\delta B_c}\right)^2 + \xi^{-2}}}. \tag{5.9}$$

A gradiometric measurement of a magnetic field is only beneficial if the signal-to-noise ratio in gradiometric mode is superior to the ratio for two separate magnetometers. Thus, we define the figure of merit $\mathcal{F}$ as the ratio of these two SNRs. A complete derivation of $\mathcal{F}$ can be found in the appendix.

$$\mathcal{F} = \frac{SNR_G}{SNR_M} = \frac{g^p - 1}{g^p + 1} \frac{\sqrt{(\delta B_u)^2 + (\delta B_c)^2}}{\sqrt{(\delta B_u)^2 + \left(\frac{\delta B_c}{\xi}\right)^2}}. \tag{5.10}$$

Here, again, $\delta B_u$ and $\delta B_c$ are the magnitudes of uncorrelated and correlated noise, respectively, present in the measurement environment, and $\xi$ is the common-mode rejection ratio (CMRR). A related parameter often found in the gradiometer literature is the balance $b$, which is simply the inverse of $\xi$.

Gradiometers are only useful if the total noise is dominated by correlated noise between the two channels. Under the assumption that $\delta B_u/\delta B_c \ll 1$,

$$\mathcal{F} \approx \frac{g^p - 1}{g^p + 1} \frac{1}{\sqrt{\left(\frac{\delta B_u}{\delta B_c}\right)^2 + \xi^{-2}}} \tag{5.11}$$

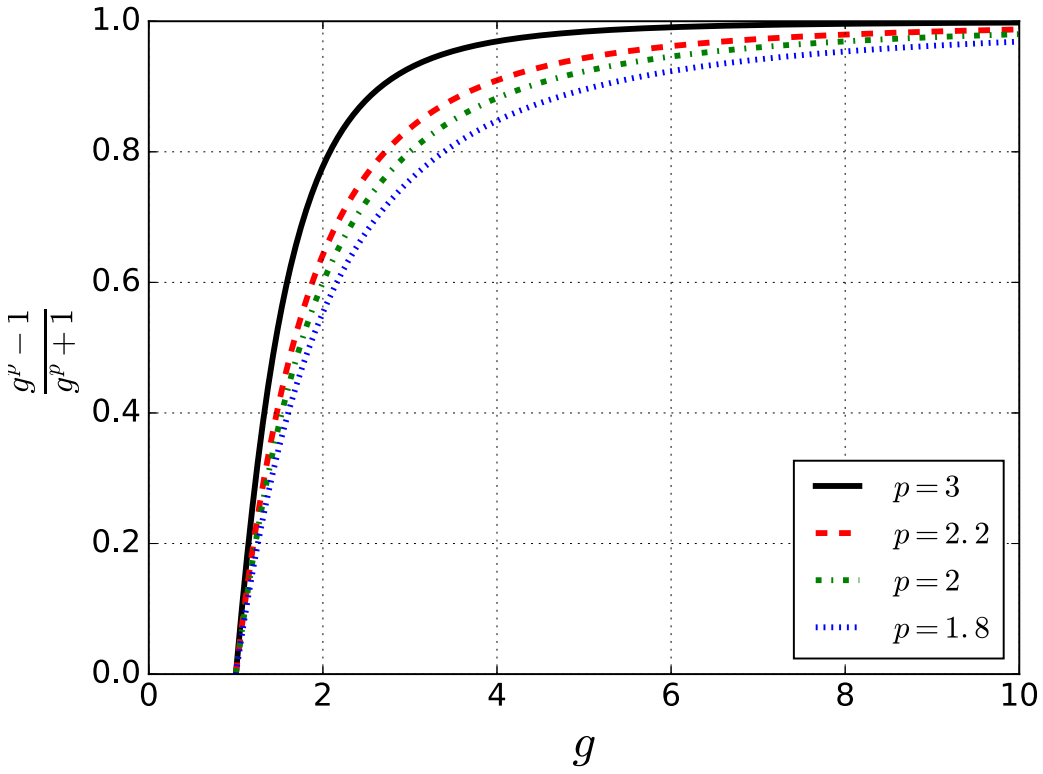



**Figure 5.3:** $(g^p - 1) \,/\, (g^p + 1)$ as a function of the geometric parameter $g$ for various values of $p$. $p = 3$ is appropriate for a purely magnetic dipole source, while $p = 2$ is appropriate for a purely current dipole source. $g = 2$ for the case $L = r_s$.

This clearly demonstrates the the figure of merit depends on not only the noise ratio of $\delta B_u/\delta B_c$, but also how this ratio compares to $\xi$. In Table 5.1, we make the assumption $(\delta B_u/\delta B_c)^2 \ll \xi^{-2}$, in which case

$$\mathcal{F} = \frac{g^p - 1}{g^p + 1}\xi \tag{5.12}$$

## 5.3 Previous Work

A number of groups have reported implementations of atomic gradiometers. Table 5.1 summarizes the performance of the different devices if they were to be used in an fMCG application. It must be noted here that the figures quoted for $\xi$ are based on the published results from those devices, and may not represent the true or absolute maximum attainable $\xi$ for that device. In our experiments, we deliberately added excess common-mode noise to measure the maximum value of $\xi$.

| AM Type | Baseline (cm) | $\xi$ | $\mathcal{F}$ | Sensitivity ($\mathrm{fT}/\sqrt{\mathrm{Hz}}$) | Reference |
|---|---|---|---|---|---|
| CPT | 2 | $> 800$ | $> 260$ | 26000[a] | [Affolderbach et al., 2002] |
| SERF | 0.5 | $> 5$[b] | $> 0.48$[b] | .16 | [Dang et al., 2010] |
| SERF | $0.3 < L < 0.9$[d] | $> 13$[b] | $> 0.8 - 2$[b d] | 0.54 | [Kominis et al., 2003] |
| SERF | $0.75 < L < 12$[d] | $> 17$[b] | $> 2.4 - 14$[b d] | 4 | [Kim et al., 2014] |
| Rb Scalar | 1.5 | $> 60$[b] | $> 15$[b] | 0.54 | [Sheng et al., 2013] |
| SERF | 0.5 | $> 3$[b] | $> 0.29$[b] | 5 | [Johnson et al., 2013] |
| NMOR | 2.5 | 20 | 7.7 | 80 | [Xu et al., 2006] |
| CsOPM | 10 | $> 5$[b] | $> 4$[b] | 54 | [Koch et al., 2015] |
| SERF | 3 | 6 | 2.6 | 9.3 | [Kamada et al., 2006] |
| SERF | $> 2.5$ | $> 6$[b] | $\sim 2.3$ | 5 | [Fang et al., 2014] |
| Mx | 5 | 1000 | 600 | 300[c] | [Bison et al., 2009] |
| SERF | 4 | 400 | 210 | 10 | This Work |
| SQUID | 5 | 100 | 60 | 5 | |

[a] May not meet the $\left(\frac{\delta Bu}{\delta Bc}\right)^2 \ll \xi^{-2}$ criterion required to achieve the calculated figure of merit in a clinical setting

[b] Estimated from published data: may under-represent the actual value

[c] May require an unshielded clinical setting to meet the $\left(\frac{\delta Bu}{\delta Bc}\right)^2 \ll \xi^{-2}$ criterion

[d] Variable baseline: array of sensors. Figure of merit calculated for first-order gradient measurement only.

**Table 5.1:** Summary of published results of atomic gradiometers. The figure of merit is calculated for $r_s =$ 5 cm, and $p = 2$, typical values of $p$ and $r_s$ for fetal magnetocardiography applications **?**. The values of $\xi$ for previously-published work should be considered lower bounds, and the resulting figure of merit is calculated for an ambient environment sufficiently noisy such that $\left(\frac{\delta Bu}{\delta Bc}\right)^2 \ll \xi^{-2}$. Specifications for a commercial SQUID gradiometer in use at University of Wisconsin-Madison are included for comparison.

## 5.4 Experimental methods

### 5.4.1 Open-loop gradiometry

We have developed a diffusion-mode SERF magnetometer array for use in fMCG [Sulai et al., 2013, Wyllie et al., 2012b], two channels of which are shown in Figure 5.4. The magnetometers operate with a baseline of 4 cm, a typical depth of a fetal heart source below the mother's skin. We typically operate the instrument in either of two operating modes, each with its own advantages and disadvantages.

In the standard DC operation mode, the component ($\hat{y}$) of the magnetic field perpendicular to the plane of our pump ($\hat{z}$) and probe ($\hat{x}$) lasers is detected [Allred et al., 2002]. In this mode, we can take full advantage of each channel's intrinsic gain [Wyllie, 2012]. However, the noise floor is often limited by low-frequency ($<$20 Hz) polarization fluctuations in the probe beam which are indistinguishable from real magnetic noise during measurements. Given that the polarization fluctuation is uncorrelated between the two channels, that source of noise cannot be suppressed in gradiometry.

To suppress the effects of these polarization fluctuations, we can operate the magnetometers with parametric modulation of the magnetic field along the pump direction at 1 KHz [Li et al., 2006], an operating mode which we call "Z-mode." In this mode, the first harmonic of the demodulated signal is proportional to the magnetic field along the probe direction ($\hat{x}$). Since the polarization noise in the frequency band around 1 kHz is negligible, this "lock-in" detection technique results in a higher signal to noise ratio for the $\hat{x}$-magnetic field component.

In either mode, there is a need for local magnetic field compensation. Working in a magnetically-shielded room (MSR), we measure residual DC fields (from the earth and nearby magnetized objects) on the order of 50 nT. In our apparatus, these fields are compensated using both a large ($D \approx$ 3 m) set of 3-axis "common" coils wrapped on the perimeter of the MSR and sets of small ($D \approx$ 4 cm) "local" coils wrapped around each individual magnetometer. One potential source of magnetic noise is the current noise from the current supplies used to generate the compensation fields.

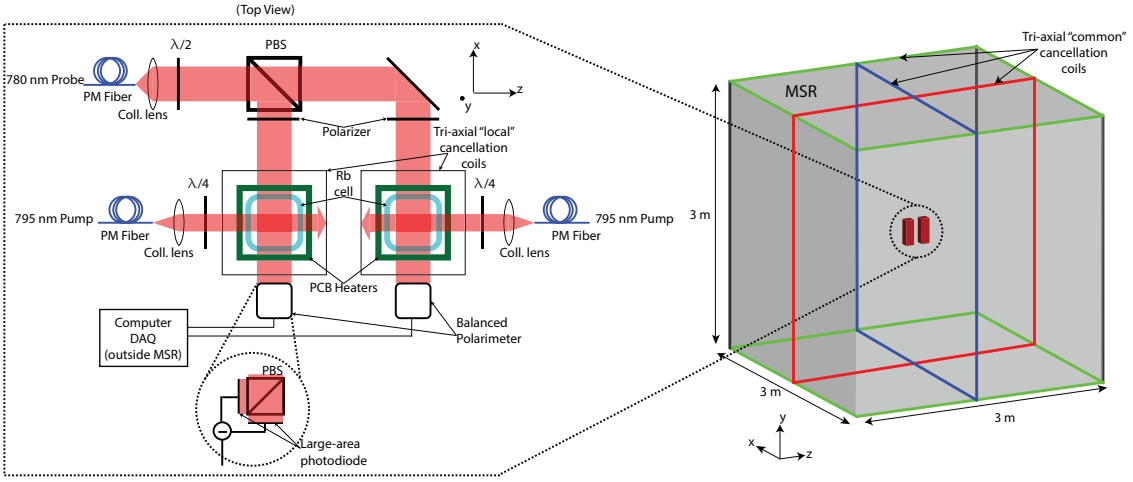



**Figure 5.4:** A schematic of the apparatus. Two magnetometers are separated by a distance of $L = 4$ cm. The magnetometers are operated in the Spin-Exchange Relaxation Free (SERF) regime. A circularly polarized pump beam slightly detuned from the Rb D1 line spin polarizes a rubidium vapor along the $\hat{z}$ axis. A linearly polarized probe beam detuned from the D2 line is directed perpendicularly ($\hat{x}$) to the pump beam. Magnetic fields in the $\hat{y}$ direction cause precession of the atoms's polarization, the direction of polarization of the transmitted probe light via Faraday rotation. This rotation is detected by a balanced polarimeter on the far side of the cell and recorded on a computer. In an alternative operation mode ("Z-mode"), parametric modulation is applied via the "local" $z$-coils. In this mode, the magnetometer is also sensitive to fields in the $\hat{x}$ direction. The two channels of the gradiometer (shown in dark red) are placed at the center of the magnetically shielded room ($right$). "Common" cancellation coils wrapped around the perimeter of the innermost shield are used to null the $x$ (red), $y$ (green), and $z$ (blue) fields at one of the magnetometers, as well as apply calibration and common-mode noise signals.

Although our supplies are very quiet, a 50 nT DC field requires a 150 dB$\sqrt{\text{Hz}}$ signal to noise ratio to reach 1 fT/$\sqrt{\text{Hz}}$. To maintain the gradiometry requirement that uncorrelated noise is small compared to correlated noise, the majority of this DC field must be shimmed with the common coil, such that noise associated with these shimming currents is common-mode. The common-mode noise produced by these large sets of coils was measured to be approximately 5 fT$/\sqrt{\text{Hz}}$ at the location of the magnetometers. In practice, we null the field of one of the channels using the common set of coils, and then null the residual fields in the other channel ($\leq$ 5 nT) using its set of local coils. This reduces the current noise requirement by 20 dB. Because the use of local coils inherently introduces uncorrelated noise, it is paramount to minimize the local sources of magnetic field (necessitating the use of local coils) in the construction of the sensors.

Due to differences in a variety of magnetometer properties including cell contents, cell temperature, pump and probe laser characteristics, etc, the two channels of the gradiometer must be calibrated against each other. The ability to suppress common-mode noise in gradiometric configuration is limited by the signal to noise ratio of the calibration so great care must be taken in this step. To calibrate, we apply a digitally synthesized frequency chirp to the common coils along the sensitive direction. The magnetometers' responses to this chirp, as well as the chirp itself, are digitized and recorded. The magnetometers' intrinsic bandwidth ($\sim$ 30 Hz) is narrower than the bandwidth of interest for fMCG (100 Hz), and for the sake of rapid calibration the chirp increases exponentially in frequency as a function of time, thereby reducing available amplitude and phase response information for calibration curve-fitting as a function of frequency. As such, a constant-amplitude chirp would result in a phase response signal to noise ratio that decays substantially at frequencies outside of the magnetometers' bandwidth. Thus, the chirp we apply has the form $B(t) = (B_0 + B_1 t)\sin\left(Ate^{Ct}\right)$. The increased calibration chirp response SNR at higher frequencies both improves calibration accuracy and improves $\xi$ across a frequency band wider than the magnetometers' intrinsic bandwidth. The "ramped" calibration signal applied to the magnetometers, as well as the magnetometers' responses, are shown in Figure 5.5.

In order to determine the common-mode rejection ratio $\xi$, we apply white magnetic noise in our bandwidth of interest (1-100 Hz) [Comani et al., 2005] to the large room coils. Applying the

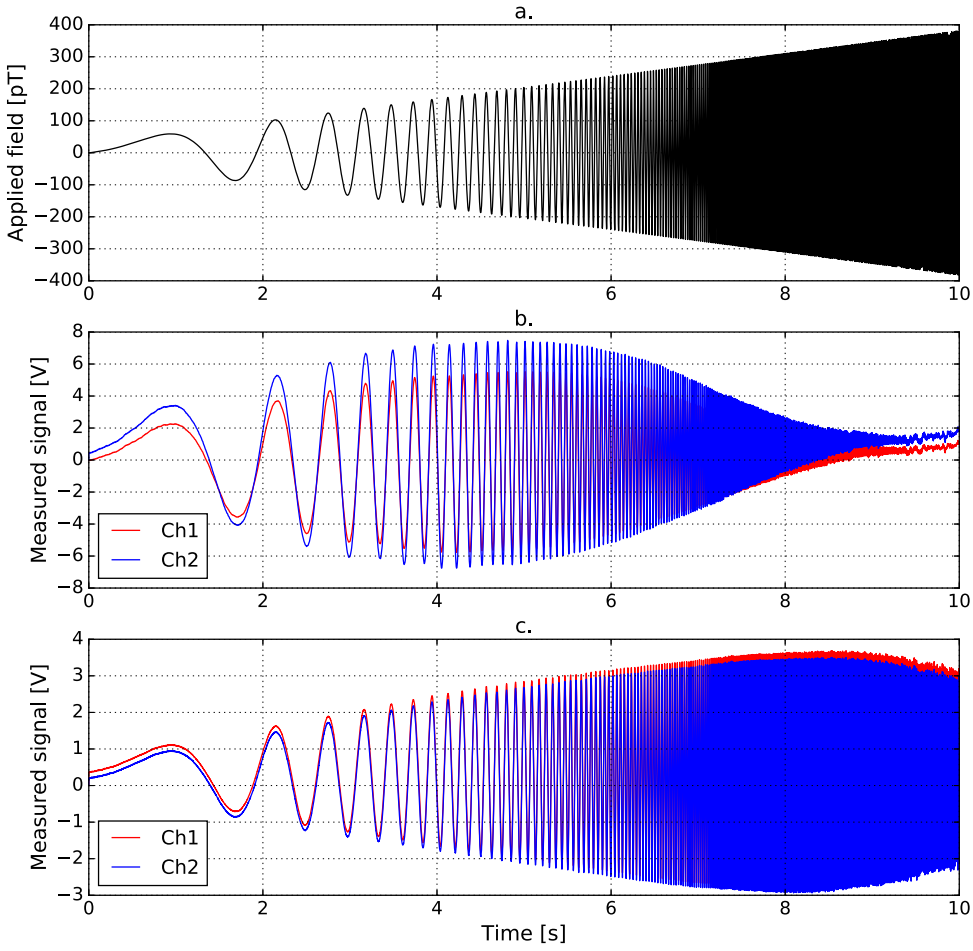



**Figure 5.5:** Calibrating the gradiometer. During calibration, the two channels of the were calibrated simultaneously by applying a time dependent field of the form $B(t) = (B_0 + B_1 t)\sin\left(Ate^{Ct}\right)$ (a.) to a common coil. The responses to the calibration field are recorded. Panel (b.) shows the responses of the two magnetometers in open-loop Z-mode operation. Panel (c.) shows the responses in closed-loop DC-mode operation.

calibration to the raw time series data, we obtain $B_1(t)$, $B_2(t)$, as well as the sum and difference of the two channels. $\xi(f)$ is then calculated by comparing the power spectra of those sum and difference channels. In both DC and Z-modes, we measured $\xi \gtrsim 150$ across the critical band, as shown in Figure 5.8.

### 5.4.2 Closed-Loop Gradiometry

To improve the performance of our gradiometer, we can also run in "closed-loop" operation mode with active feedback stabilizing the field (i.e. driving the measured field to zero) in the sensitive direction. This operation mode offers several benefits. First, the closed-loop operation mode exhibits substantially reduced sensitivity to changes in operational parameters. The amplitude and phase response of each individual channel becomes dominated by the feedback circuitry rather than the operational parameters of the cell and lasers. Thus, if identical components are used in the feedback circuit of each magnetometer channel, the magnetometers exhibit substantially equal amplitude and phase responses to a common-mode input signal. This reduces calibration drifts (and corresponding decay of $\xi$) associated with differences in sensitivity to potentially common-mode changes in operational parameters of the individual channels (cell temperature, laser power, etc).

Second, this mode eliminates observed issues with low-frequency magnetic field drifts in the shielded room along the sensitive axis direction, which often cause the detection system to "rail" during an open-loop measurement. Third, keeping the polarimeter signal "balanced" (i.e. equal light on each photodiode) reduces the uncorrelated probe noise during measurement. Fourth, the feedback mode allows us to "flatten" the response of the device through design of the feedback circuit. Although the native bandwidth of the magnetometers is $\sim 30$ Hz, we are able to broaden this to $\sim 120$ Hz by selecting appropriate circuit components for the feedback circuitry. While this does not improve the fundamental SNR of the detection system, the flattened response exhibits substantial improvements in calibration signal to noise ratio and corresponding improvements in $\xi$. Of particular interest are the substantial reduction in the relative phase response uncertainty of the calibration curve fit, most apparent at higher frequencies. Applying active feedback affords us excellent calibration SNR across our frequency range of interest, even at 100 Hz and above. With

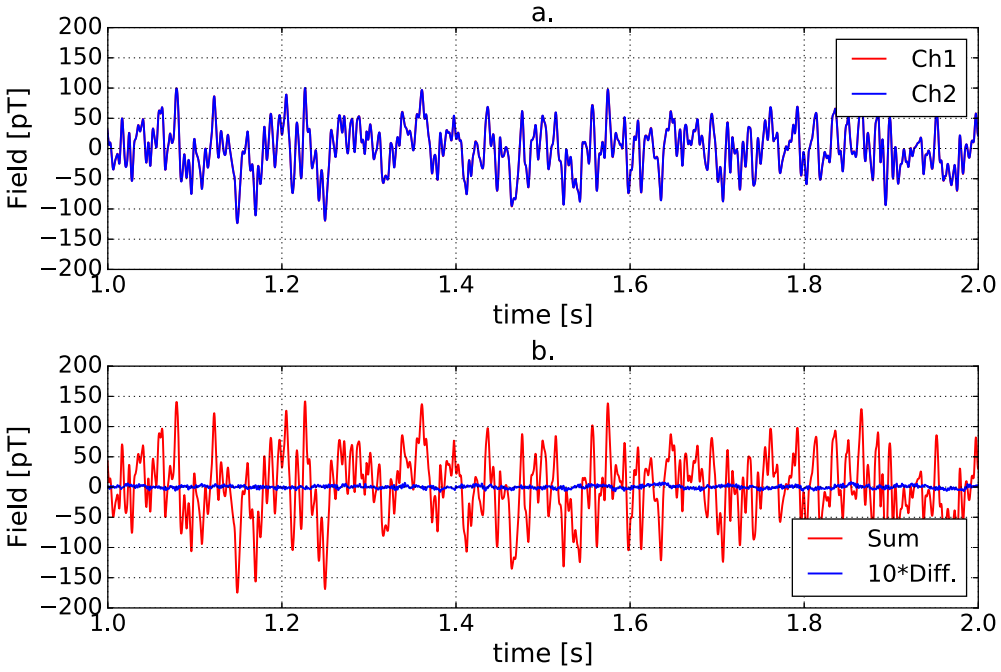



**Figure 5.6:** Determination of $\xi$. White noise is added to the magnetically shielded room via the large coils mounted on the inside surface of the room itself. The individual channels operating in closed-loop DC mode are shown in the top panel, (a.) In panel (b.) the difference channel is shown multiplied by a factor of 10 for clarity.

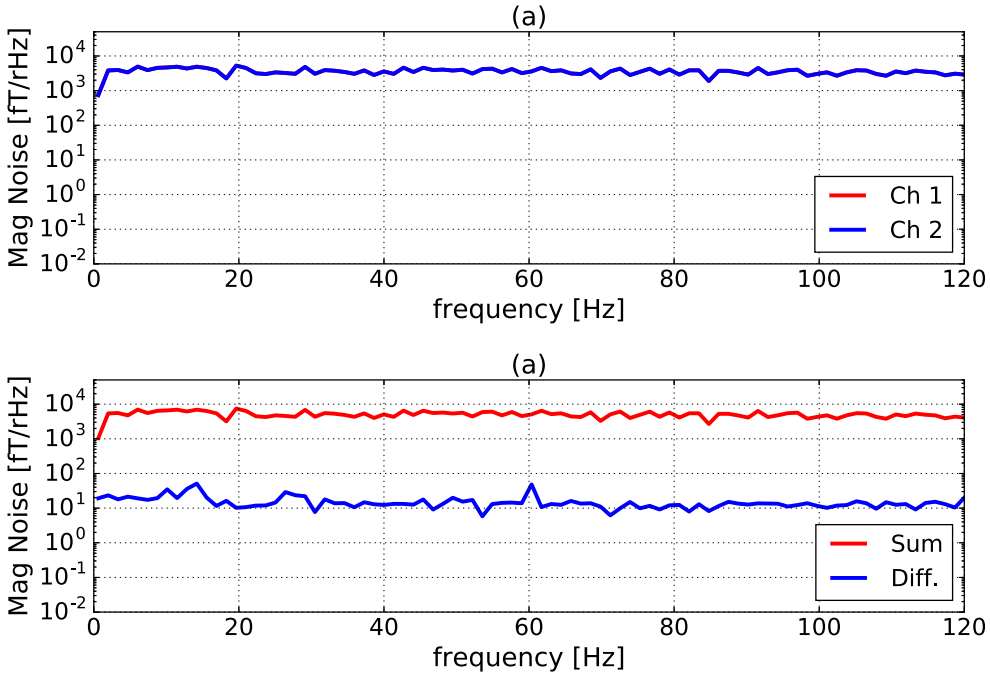



**Figure 5.7:** Power spectra of the data shown in Figure 5.6. (a.) Individiual channels. (b.) Sum and difference channels. Note the baseline noise on the difference signal of 10 fT/$\sqrt{\text{Hz}}$.

active feedback, we obtain measurements of $\xi \gtrsim 400$ across a wider frequency band, a factor of two improvement over the open-loop operation mode.

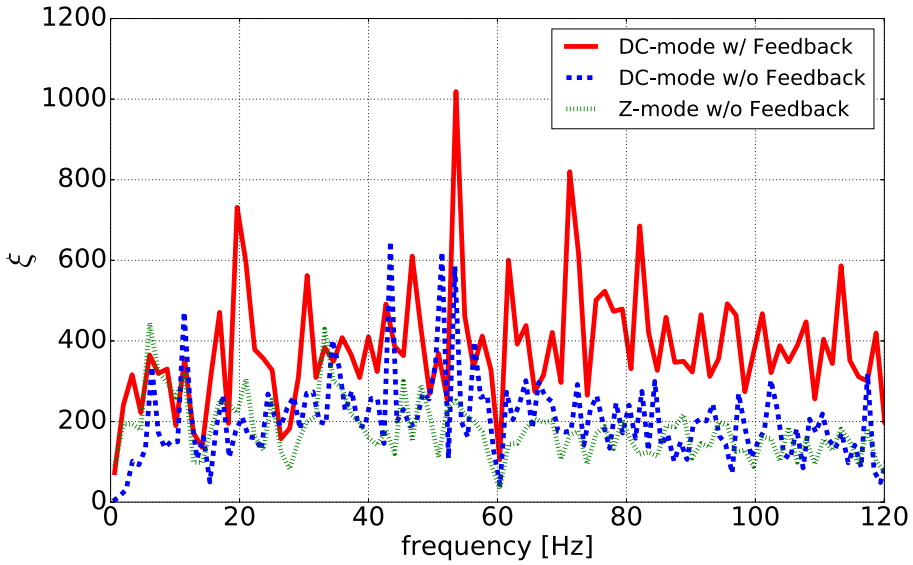



**Figure 5.8:** Comparison of the common-mode rejection ratio for various operating modes. $\xi$ was determined by taking the ratio of the sum and difference power spectra, for example the ones shown in Figure 5.7. The use of active feedback shows superior performance in suppressing common-mode noise. The large magnitude of uncorrelated 60 Hz line noise limits the performance of the gradiometer at that frequency.

## 5.5 CMRR Stability

The most important feature of the active feedback is the increased stability of the calibration over time. The length of a standard fMCG measurement is approximately 3 minutes. In our experience, after this amount of time has elapsed, the fetus or mother has often shifted position sufficiently to require recalibration of the magnetometers. Occasionally, the shift in position is significant enough to require readjustment of the position of the array to maximize signal, followed by a recalibration. Thus, for our gradiometer to be feasible, we require that the measured value of $\xi$ be above the target threshold of 100 [Clarke and Braginski, 2005, Hamalainen et al., 1993] over at least a 3 minute timescale following initial calibration. To test our device's robustness over time, we calibrated the device and then applied common-mode white noise over a period of 5 minutes, as shown in Figure 5.9. The average value of $\xi$ over the band of 2-100 Hz was recorded at various times after calibration in both "open-loop" and "closed-loop" modes.

We found that as in previous experiments, $\xi$ was superior when using the active feedback "closed-loop" mode, but more importantly, remained significantly above the target threshold for the full 5 minute measurement time frame. In this mode, we observed an unexpected and measurable degradation, as the amplitude and phase response to $\hat{y}$ fields (set by the feedback circuits) should remain unchanged over time. However, the magnetometers retain some residual sensitivity to fields in the other two directions [Wyllie, 2012]. This residual sensitivity can vary in both magnitude and phase from channel to channel. Therefore, field variations in the $\hat{x}$ and $\hat{z}$ directions will affect the relative calibration of the individual magnetometers depending on each channel's response to these transverse fields; it is hypothesized that the observed degradation is dominated by these transverse field effects. In contrast, the "open-loop" experiment showed poorer performance both initially and also after extended time periods. The value of $\xi$ in "open-loop" mode eventually dropped below the target threshold of 100, again indicating a substantially improved likelihood of a successful gradiometer measurement when using active feedback.

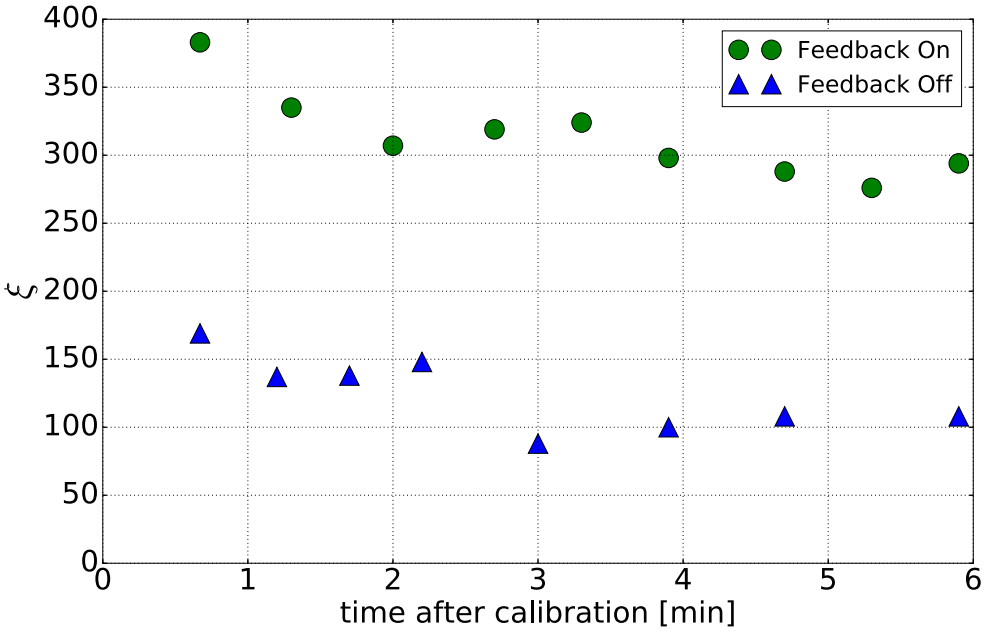



**Figure 5.9:** The stability of $\xi$ over time was measured in DC-mode with and without active feedback. With active feedback ("closed-loop" operation), the device was able to maintain an acceptably high level of common-mode noise suppression over a time period comparable to a fMCG measurement. When operating without the feedback, however, the device suffers both in the absolute value of $\xi$ as well as its stability. Most notably, $\xi$ value drops below the critical value of $\xi \approx 100$ after approximately 3 minutes. This evidence strongly supports using active feedback in future work.

In the future, we wish to combine the low-technical-noise benefits of Z-mode with active feedback. At present, the electronics and coils we use only allow for active feedback in the $(\hat{y})$ direction, though with a few modifications we could implement feedback in the $(\hat{x})$ or $(\hat{z})$ direction, or even all three directions simultaneously [Seltzer and Romalis, 2004]. Additionally, the filter settings on our lock-in amplifiers while acquiring data in Z-mode were inappropriate for active feedback. In an effort to suppress high-frequency noise, we selected a low pass filter of $\tau = 1$ ms. However, the additional phase shift introduced in this filtering made the overall phase shift of the device cross the critical $\pi$ threshold at a relatively low frequency of 70 Hz, and quickly led to feedback oscillation in early experiments. We are exploring creation of digital filters to help "compensate" this phase shift.

## 5.6 Conclusion and outlook

We have demonstrated a gradiometer constructed from two atomic magnetometers with baseline and common-mode rejection ratio that positions it favorably to other published atomic devices, and has performance in line with current commercially-available SQUIDs. Future work will include efforts to combine the benefits of parametric-modulation with active feedback to create a superior gradiometer device.

# Chapter 6

# Magnetocardiography

Over the past several chapters, we have discussed the theory behind our SERF atomic magnetometers, the way we characterize the devices and collect data, and an evaluation of our new gradiometer scheme. We finally arrive at the crux of this project: using our devices to detect magnetocardiographic signals from adults (MCG) and fetuses (fMCG). We will begin this chapter with a brief physiological description of the signals we are attempting to detect. A simulation of the expected magnetic signal will then be presented, alongside magnetic field power spectra collected in the magnetically-shielded room (MSR) with and without magnetocardiograpic signals present. This will provide an estimate of the relative size of all of these signals. Finally, we will present actual 8-channel MCG and fMCG data taken using Z-mode, which is the first time such data has been published. We will also discuss the signal processing techniques we use on our data and compare our data to that collected with a SQUID gradiometer, the device we hope to one day supplant with our atomic magnetometers.

## 6.1 Cardiac electrophysiology

The electrophysiology of the heart is based on a series of electrical signals traveling through a network of conductive cells surrounding its four chambers. These electrical signals can be detected quite easily in adults by placing conductive leads at spatially-separated locations on the subject's chest and measuring the potential difference between a them. This measurement, known as an *electrocardiogram* (ECG or EKG) is a common tool used by physicians to measure heart rate,

arrhythmias of the heart beat, and the strengths and timings of the electrical impulses as they pass through each part of the heart [NIH, 2016].

### 6.1.1 Heart cycle

Figure 6.1 depicts a normal heartbeat cycle [NIH, 2011]. The cycle begins in *diastole*, the "resting period" between heartbeats during which time the electrical activity of the heart is at a minimum and blood begins to fill the heart's chambers. At the onset of *systole*, an electrical impulse is spontaneously generated in the *sinoatrial* (SA) node, the heart's natural "pacemaker", and it begins to travel across the *atria*, or upper chambers of the heart. The signal passing around these chambers causes them to contract, expelling the blood into the *ventricles*, the lower chambers of the heart. The current that is associated with this initial step of systole generates the first measurable potential difference in the EKG. It is known as the *P* wave.

After passing around the atria, the signal arrives at the *atrioventricular* (AV) node, at which point there exists a brief pause in electrical activity (the so-called *PR* segment) while the left and right ventricles finish filling with blood. The AV node then releases an impulse that travels down bundled conducting branches between the ventricles. On the EKG, the conduction of this charge is represented by a small voltage spike called the *Q* wave. The impulse then begins to spread out around the lower chambers of the heart, which begin to contract, much like the atria did during the *P* wave. However, the two lower chambers do not contract simultaneously. The left ventricle, whose contraction expels oxygenated blood into the aorta and to the rest of the body, contracts an instant before the right ventricle, which sends de-oxygenated blood to the lungs. These back-to-back signals are known as the *R* and *S* waves, respectively, and their relative sizes on the EKG (the *R* being several times the size of the *S*) is a result of the strengths of the contractions required to send the blood from each chamber to its destination.

After the signals have passed through the heart muscle and the heart has expelled the blood from its lower chambers, the heart muscles around the ventricles begin to relax. This results in a depolarization current that gives rise to the final element of the electrical activity: the *T* wave. The complete depolarization of the muscles returns the heart to diastole, and the process begins

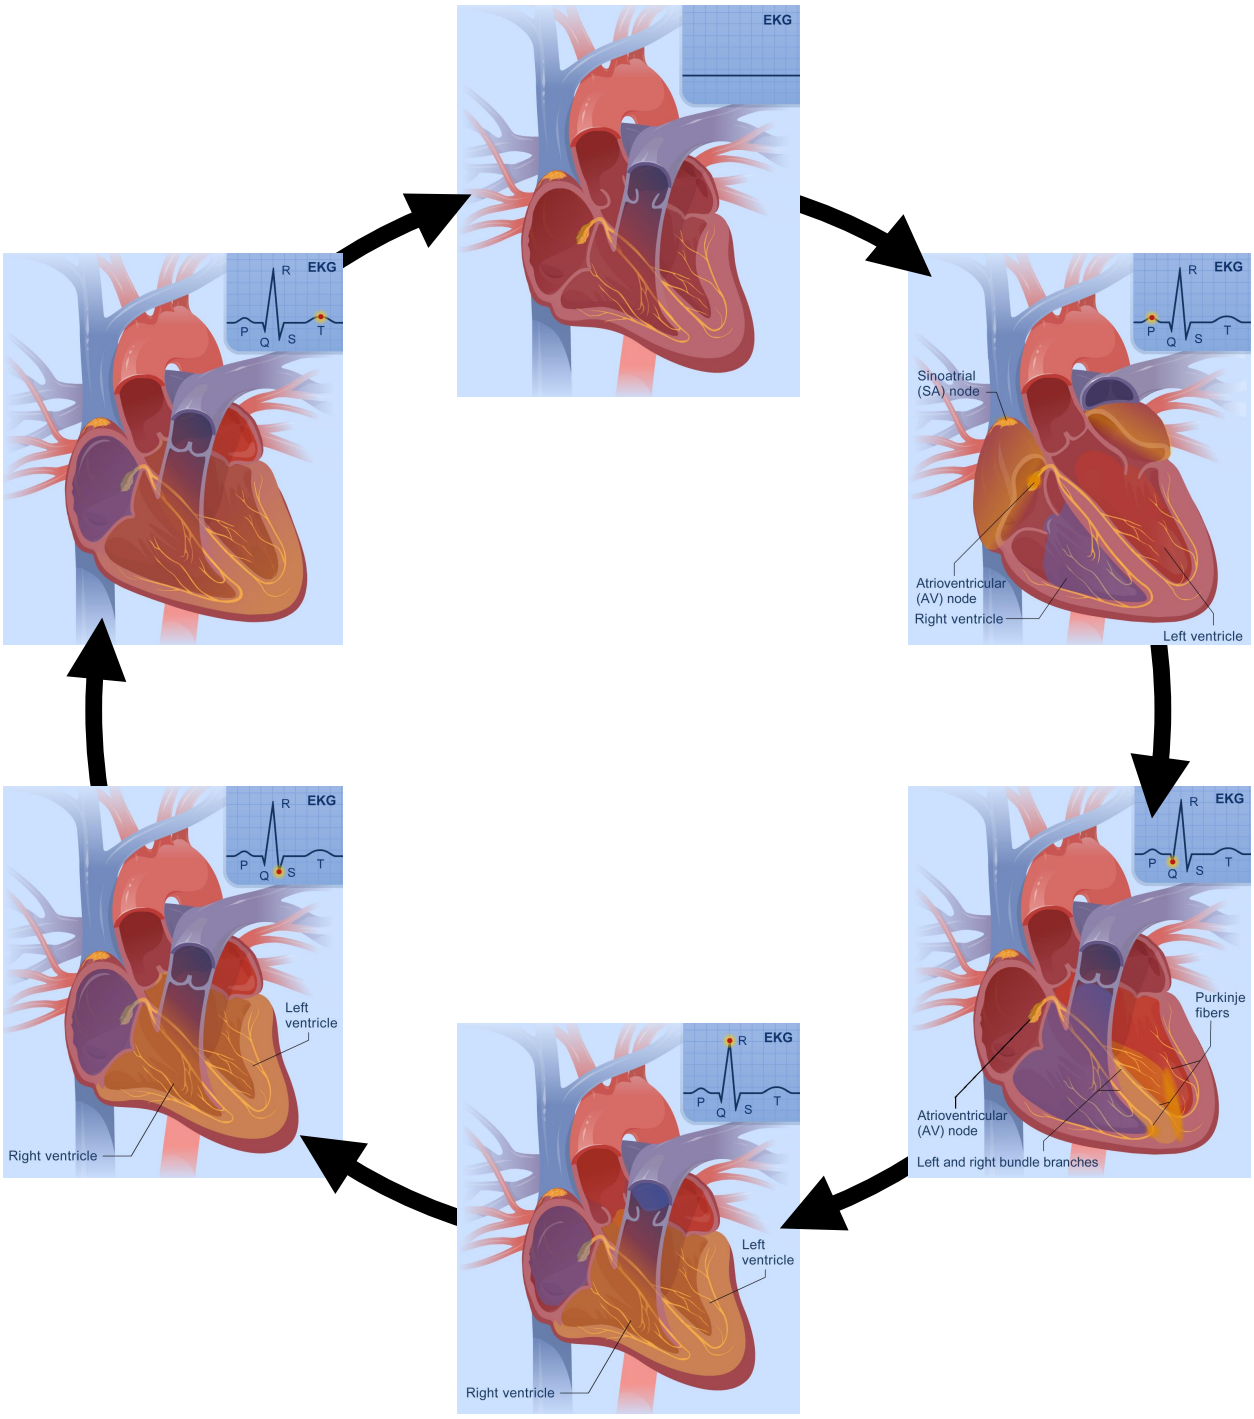



**Figure 6.1:** Electrophysiology of the human heart. The gold shading indicates the approximate location of the "pulse" that causes the contraction of the various chambers of the heart. These contractions pump blood between chambers as well as to the lungs and to the rest of the body. Each subfigure also contains an inset indicating part of the EKG corresponding to each stage of the cycle. The diastolic (resting) phase of the cycle is indicated at the top followed by the systolic sequence of (clockwise) *P, Q, R, S,* and *T* waves. Subfigures adapted from [NIH, 2011].

again. In a normal EKG, tens or hundreds of these cycles are recorded, which allow the physician to measure not only the size and timing of the individual waves, but also the intervals between the waves. These intervals are vital, as they are used to diagnose irregularities in the heartbeat, referred to as heartbeat arrhythmias. Some of the most-commonly measured intervals are indicated in Figure 6.2.

### 6.1.2 Need for MCG

Ampère's Law states that any electrical current will produce a magnetic field proportional to the current's magnitude. The cardiac currents described in the previous section thus produce magnetic fields that can be measured with sensitive magnetometers. It may seem redundant to measure the magnetocardiographic signals given the well-established nature of electrocardiography. For adult patients, magnetocardiography provides essentially information as electrocardiography. However, the non-contact nature of MCG could provide a way of measuring a cardiogram of a patient where the skin contact required by ECG is not advised (a severe burn victim, *e.g.*).

The benefits of MCG become most readily apparent when considering the measurement of a fetal heart. Placement of conductive leads directly on the fetus is an invasive procedure, which is generally inadvisable. Placing the leads close to the fetus on the mother's skin would appear to be the next-best way to acquire an EKG signal. However, it has been found that amniotic fluid in the uterus and the *vernix caseosa* (a waxy substance coating the skin of the fetus) severely attenuate the electric field, making the measurement of EKG difficult [Wakai et al., 2000]. The measurement of the fetal cardiogram is essential for diagnosing fetal heartbeat abnormalities *in utero*, including long-$QT$ syndrome [Zhao et al., 2006] and complete AV block [Das et al., 2008]. These conditions sadly often result in the death of the fetus, though there is evidence that these deaths could be preventable with a timely diagnosis and treatment plan [Strasburger et al., 2008].

The magnetic fields generated by the heart, on the other hand, do not suffer the same type of attenuation as the electric fields. In fact, the magnetic fields permeate the air outside of the mother's skin, so it is not even necessary to come in contact with the mother's skin while making

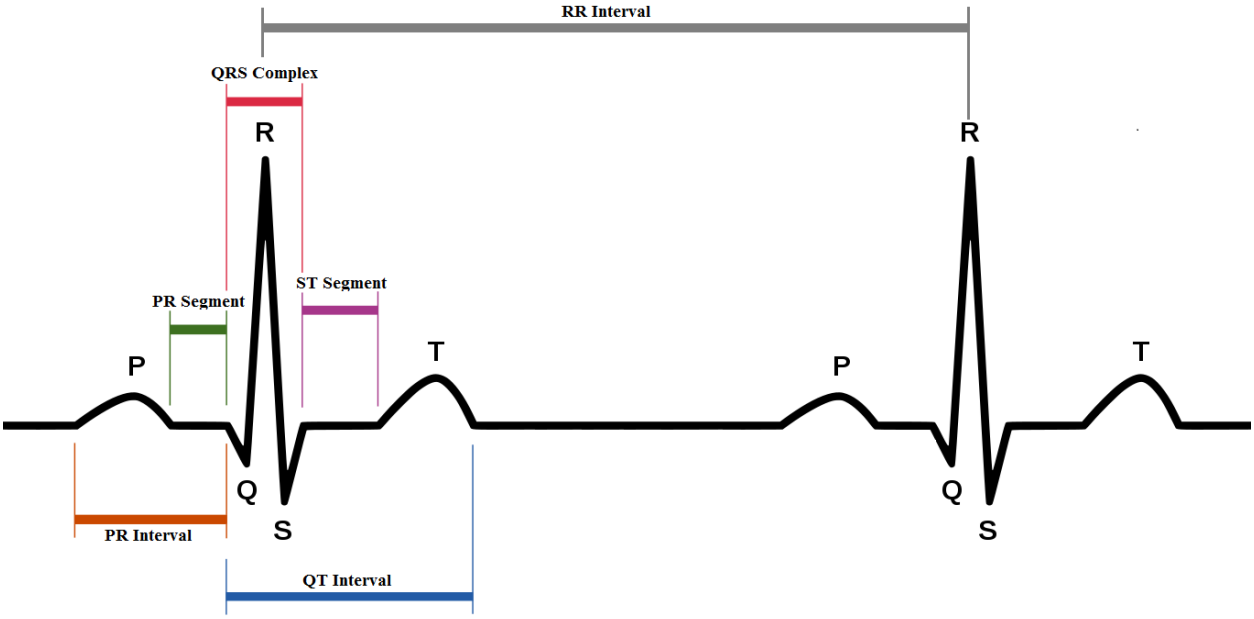



**Figure 6.2:** Two consecutive systolic cycles with a period of diastole in between [not to scale]. The intervals most useful for cardiologists in determining arrhythmias are indicated. The heart rate (in beats per minute) is equal to $60/T_{RR}$, where $T_{RR}$ is the $RR$ interval time in seconds. The $QT$ interval is of particular interest for fMCG.

a measurement. (We often do, however, simply to shorten the distance between the fetal heart and the detectors)

## 6.2 MCG and fMCG characteristics

We now turn to a more in-depth analysis of the MCG signals. To obtain a model of the signals we would be aquiring, we generated cardiographic waveforms using ECGSYN [McSharry and Clifford, 2015], a MATLAB program which generates cardiographic time series given a range of inputs. We created two time series of data: one to model a typical adult and one for a typical fetus. For an adult signal, we set ECGSYN to produce a 60-beats-per-minute signal; for the fetus, we chose a 140 bpm signal (which is consistent with studies of fetal heart rates at 25 weeks gestational age [Pildner von Steinburg et al., 2013]). We then scaled each signal to approximately match the sizes of the real field maternal and fetal *R* peaks measured by Wyllie et al. [2012a]: The maternal *R* peak was scaled to be 15 pT and the fetal *R* peak to be 2 pT. To better understand the frequency-space composition of these signals, we computer the power spectrum of each. These spectra, along with an actual power spectrum collected during a fMCG measurement, are shown in Figure 6.3. The maternal and fetal heart signals are apparent in the experimental power spectrum up until roughly 50 Hz, where the environmental noise inside of the MSR begins to dominate. It is clear that if we wish to capture the relevant cardiographic data (especially the *T* wave, which has a lower frequency makeup than the other structures), then minimizing experimental noise–especially in the DC-50 Hz range–is essential.

When performing a fMCG measurement, the maternal signal is generally regarded as "noise", as it is not the signal of interest. As seen in Figure 6.3, the maternal signal is roughly ten times the size of the fetal signal near DC. However, the maternal signal does not render a fMCG measurement impossible for the following reasons:

1. The diastolic portions of the maternal heartbeat offer a "quiet" period in which the fetal signal is more easily visible. The roughly 2:1 ratio of the heart rates mean that at least one fetal heart cycle occurs between the end of one maternal *T* wave and the beginning of the next maternal *P* wave.

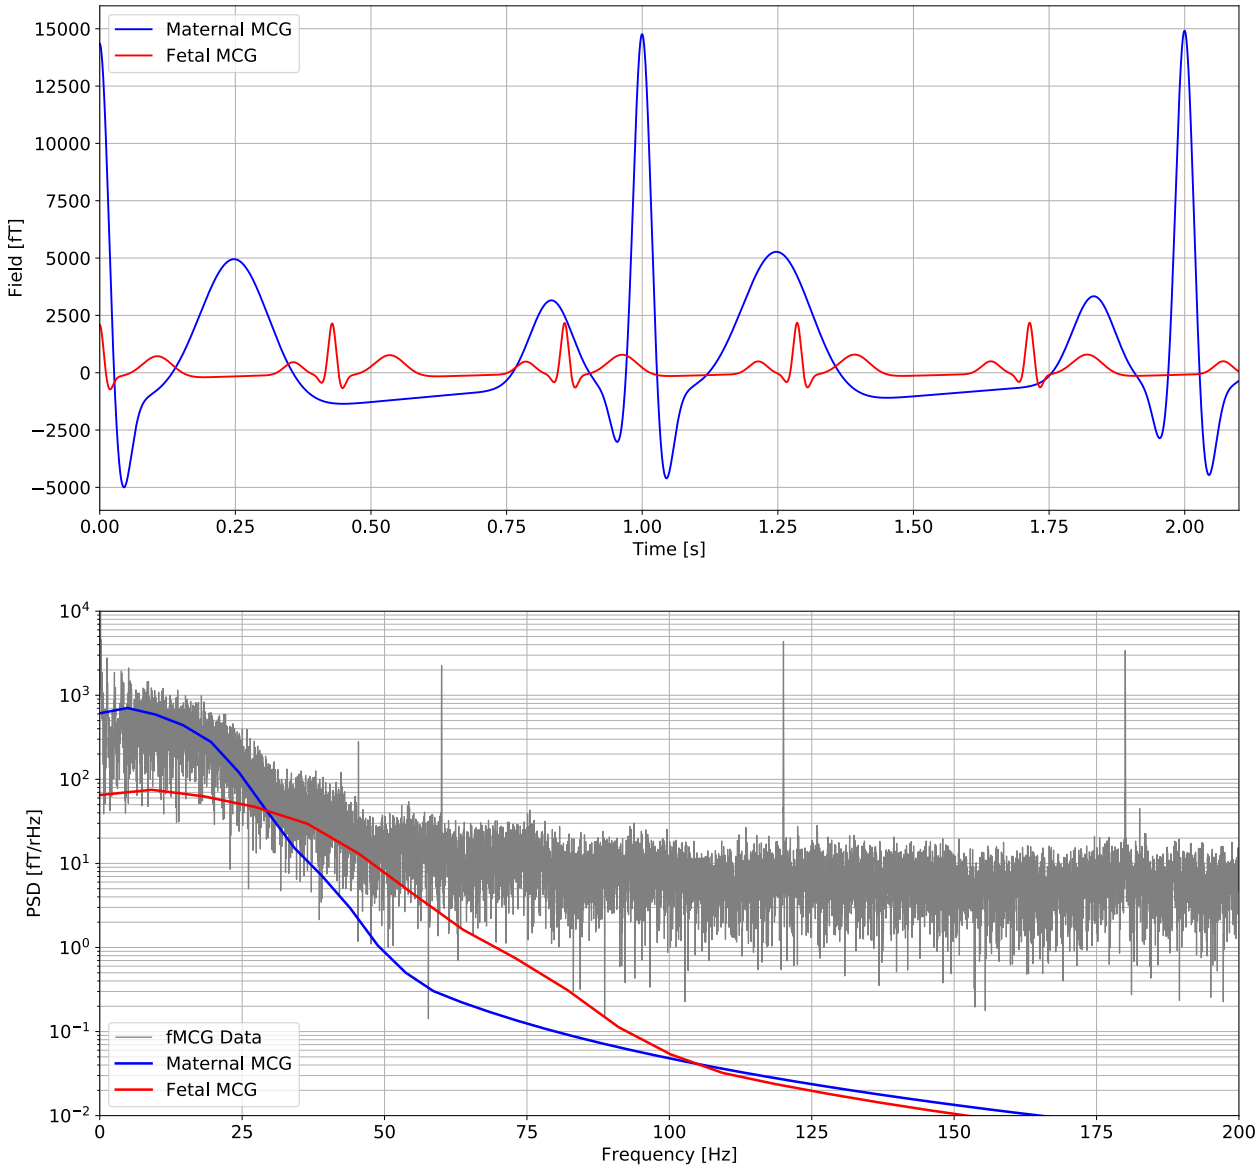



**Figure 6.3:** Top: Time series of the maternal and fetal heartbeats created using the ECGSYN program [McSharry and Clifford, 2015]. The magnitudes and heart rates for each were based on commonly-acquired maternal and fetal MCG data with our apparatus positioned appropriately for a fetal measurement. Bottom: The frequency-spectrum representation of the the maternal and fetal signals from the top panel. A power spectral density of a real fMCG measurement is shown in gray. Note the total signal is dominated by the maternal signal between DC and 30 Hz, by the fetal signal over the range 30-50 Hz, and then by environmental noise (which was at a level of roughly 10 fT/$\sqrt{\text{Hz}}$ during this measurement) above that. To capture as much of the cardiographic signal as possible, environmental noise must be kept to a minimum, especially at low frequencies where most of the cardiographic information exists.

2. Gradiometric techniques can be employed to suppress the maternal signal to some extent, which is more uniform than the fetal signal across the baseline of the gradiometer given its greater distance to the sensors. This effect is small though. As demonstrated by Wyllie [2012], the expected fractional difference in the magnetic field across the sensor separation distance is only $\Delta B/B \approx 2/3$. The common-mode suppression of $\sim$400 recorded in Chapter 5 was true for fields generated my sources much farther away, such that $\Delta B/B \approx 0$. Thus we cannot rely completely on gradiometry to suppress the maternal signal.

3. Fortunately, the maternal signal can be suppressed using nonlinear signal processing techniques, one of which, Splined Independent Component Subtraction, will be discussed in Section 6.6.3.

## 6.3 MCG and fMCG measurement procedure

### 6.3.1 Array setup and positioning

Our magnetometer array is located in a magnetically-shielded room shared with the Wakai group's SQUID magnetometers. The array is mounted on the end of the Wakai group's unused SQUID gradiometer, which itself is attached to a ceiling-mounted gantry. This allows for easy three-axis positioning of the magnetometer array, including the tilt and rotation needed to get the best possible positioning for the MCG or fMCG measurement. As mentioned in Chapter 4, the precision of the array positioning becomes more and more important as the signal size we are measuring gets smaller. For an adult, positioning the array 1-5 cm above the middle of the chest is generally sufficient to allow the acquisition of the MCG signal with a 10-100 SNR on all four magnetometer channels.

Acquiring the fetal signal requires more careful alignment, and we have found that translating the array by only a few cm or tilting it a few degrees can make the difference between an acceptable SNR or the complete inability to discern the fetal signal from the background noise. To aid in the placement of the array, a trained medical physicist performs an ultrasound on the mother's abdomen to locate the approximate position and depth of the fetal heart. We then attempt to center

the array such that the channels are roughly equidistant from the fetal heart. This may involve tilting the array to match the curvature of the mother's abdomen in order to get all of the sensors as close as possible to the heart. A photo of the array in position above a pregnant patient is shown in Figure 6.4.

Once the array is in position, the experimenters exit the room and re-null the magnetic fields at each magnetometer. By this time, the maternal heartbeat signal is usually large enough on all four channels to see on an oscilloscope, despite the additional distance between the sensors and the maternal heart (compared to a dedicated adult MCG measurement). Depending on the magnetic noise levels and positioning, the fetal heart signal may be visible at this point as well. If not, sometimes we will attempt to reposition the array in real-time to search for the fetal signal. To accomplish this, two experimenters are required. One experimenter will re-enter the room and make small adjustments to the positioning of the array while the second experimenter watches the magnetometers' signals on an oscilloscope. Because of the non-uniformities in the room's magnetic field, repositioning the array necessitates a constant re-nulling of the magnetic fields to keep the magnetometers in the SERF regime. Running the magnetometers in active feedback mode stabilizes the $\hat{y}$ fields to some extent, but without active feedback on the transverse $\hat{x}$ and $\hat{z}$ directions, some intervention is required for this positioning optimization procedure. In the future, 3-axis feedback will allow for this procedure to be done in a more efficient manner.

### 6.3.2 Magnetocardiogram measurement

Once we are satisfied with the position of the magnetometer array, we calibrate the magnetometers (Section 4.4), and then we collect a long time series of data (usually 1-5 minutes). During that time period, it is not uncommon to see the size of the fetal signal wax and wane as the distance from the fetal heart to the sensors changes, This is especially evident if the fetus is particularly active in the uterus. It is also common to detect the large, slow drifts associated with the movement of elevators, cars, and other large metal objects as they move in proximity to the MSR. If common or individual-channel active feedback is not employed, these fields must be countered in real time by

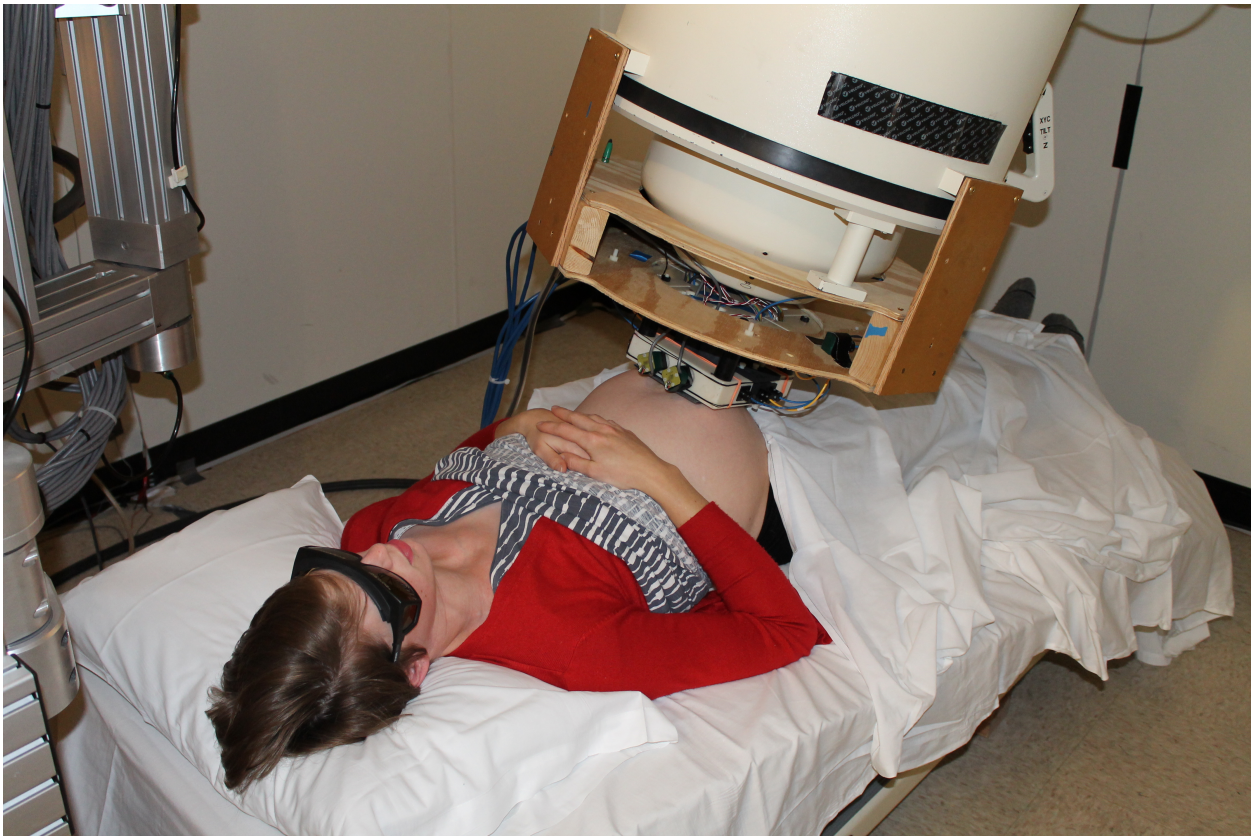


**Figure 6.4:** The array in position to make a fetal measurement. After the magnetometer array is mounted on the end of the out-of-use SQUID gradiometer (beige cylinder, top right), the mother lies down on the bed inside of the magnetically-shielded room. An ultrasound is performed to locate the fetal heart, and the array is translated and tilted so that the fetal heart lies roughly in the center of the four-unit array. Though laser light is completely contained within the apparatus, the patient wears infrared safety goggles to protect herself should unexpected stray laser light be present in the room.

re-nulling of the magnetic fields. Without user intervention, the photodiode signals can overload the current preamps rail and potentially valuable data may be lost.

A typical patient appointment with a patient generally lasts anywhere between 1-2 hours. During this time, roughly fifteen to twenty 1-5 minute measurements are performed, interspersed with breaks to accommodate the wishes of the patient, repositioning of the magnetometer, and recalibration. Normally, regular DC-SERF mode and Z-mode data are collected, both with and without active feedback. Analysis of the four types of datasets can reveal what was working well and what should be improved with a particular measurement technique in subsequent and measurements.

## 6.4 MCG signal processing

A high signal-to-noise ratio is critical for the realization of any measurement, and our apparatus has been designed to maximize this value. Its modularity and ease of positioning ensure that it can be positioned as close to the target heart as possible, maximizing signal. We have gone to great lengths (Chapters 3 and 4) to limit detected noise–both from external sources and from our apparatus. Despite these efforts, the raw data collected during a MCG or fMCG measurement is nearly always unsatisfactory for cardiographic analysis–especially for making the heartbeat interval measurements (Section 6.1.2) that are required to diagnose fetal heartbeat arrhythmias. A 10-second subset of raw data from a recent Z-mode fMCG measurement on a fetus at 28 weeks gestational age is shown in Figure 6.5. Without further signal processing, these data cannot be used for any sort of fetal cardiographic analysis.

To rectify this, all of the MCG data we collect undergoes rigorous signal processing to improve its SNR. The signal processing techniques range from simple single channel algorithms (such as notch-filtering the data at the power line frequency) to complex multiple channel procedures designed to isolate the individual sources of signal and noise in a multi-channel measurement. For the most part, the signal processing applied to the data in this thesis was done with *Saki*, a powerful software suite developed by the Wakai lab for use with their SQUID fMCG measurements. These tools, described in the following sections, have proved quite useful for our atomic magnetometry measurements as well.

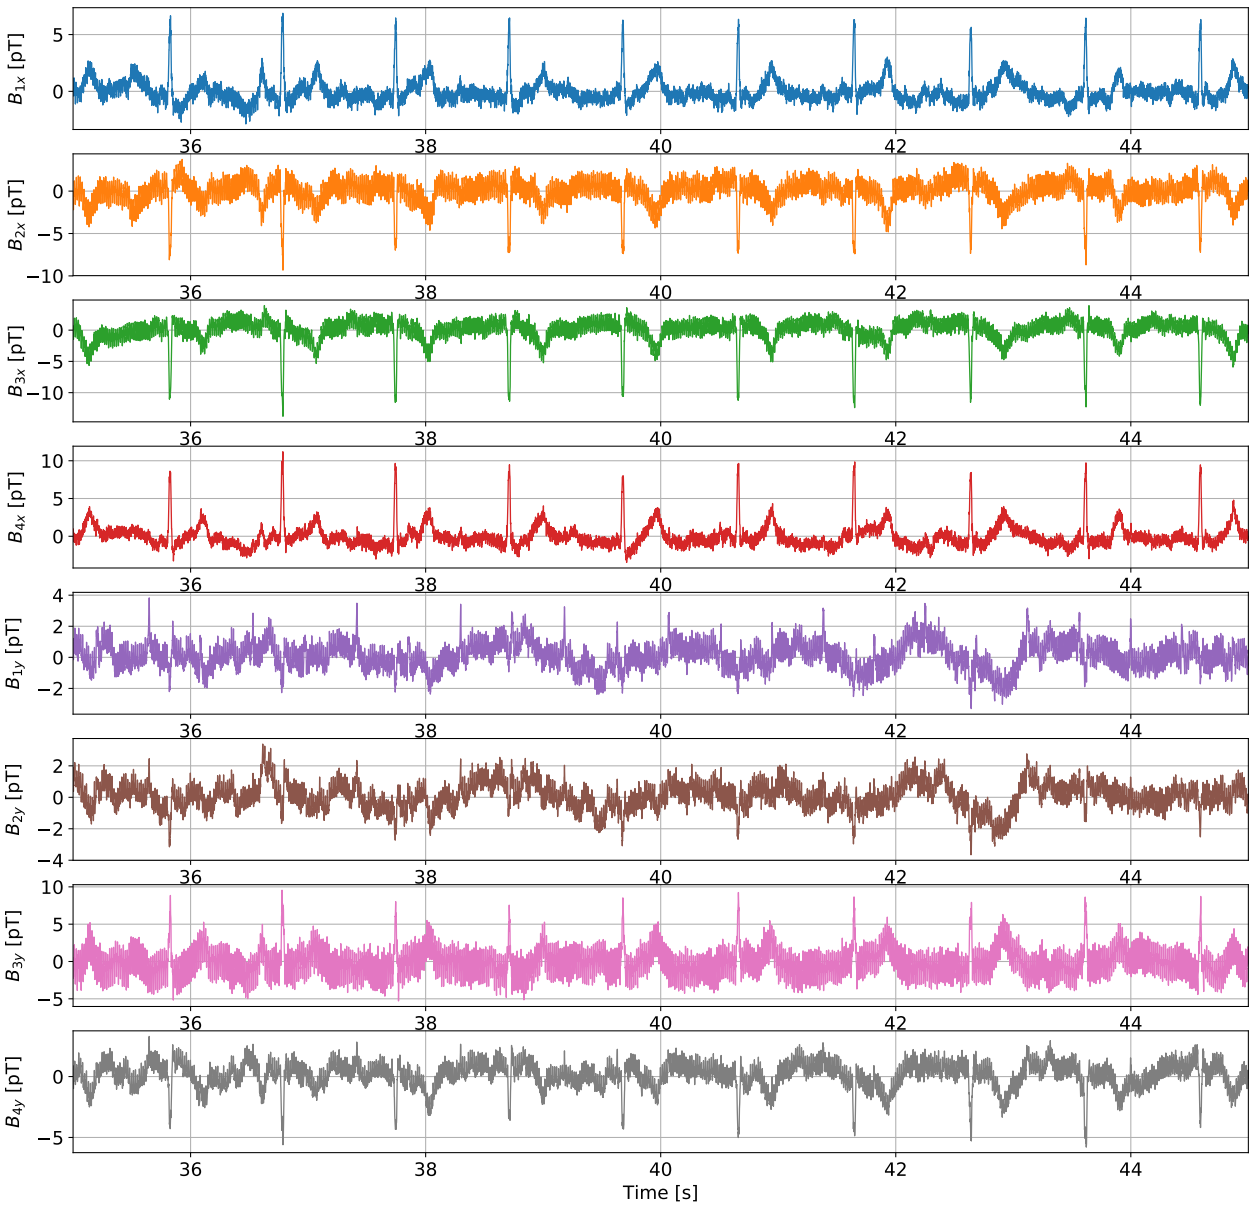



**Figure 6.5:** Calibrated z-mode data collected during a fMCG measurement with no additional signal processing applied. The top four traces are the recorded $B_x$ (Channels 1, 2, 3, and 4, respectively), while the bottom four are the measurement of $B_y$. The prominent peaks occurring every $\sim$1 second are the *maternal QRS*-complexes. The fetal data is mostly obscured by noise, though some indications of it are visible on traces 5 and 6 ($B_{1y}$ and $B_{2y}$, respectively). Channels 1 and 2 were farthest from the maternal heart, and during this segment of data, likely closest to the fetal heart. The quality of this raw data shows the need for further signal processing.

## 6.5 Single-channel processing

Before any real signal processing is done, the voltage time series $V_m(t)$ collected by the data acquisition system is converted into a magnetic field time series, as explained in Section 4.4.4: First the data are transformed to frequency space via FFT:

$$\tilde{V}_m(\omega) = \text{FFT}\left[V_m(t)\right]. \tag{6.1}$$

Next, the voltage to magnetic field calibration $\zeta^\star(\omega)$ is applied to the transformed voltage data.

$$\tilde{B}(\omega) = \frac{\tilde{V}_m}{\zeta^\star(\omega)}. \tag{6.2}$$

Finally, an inverse FFT is applied to recover the calibrated time series.

$$B(t) = \text{IFFT}\left[\tilde{B}(\omega)\right]. \tag{6.3}$$

This procedure is handled by the LabVIEW program which collects the data. While this procedure is essential for determining absolute signal sizes, as far as the data quality is concerned, this is simply a re-scaling of the $y$-axis and doesn't do anything to *improve* the quality of the data of and our ability to identify the various parts of the heart cycle.

After the data has been appropriately converted into magnetic field units, the true signal processing can begin. The most effective way to improve the fidelity of a single channel is by applying digital filters to the data to remove unwanted noise. Unfortunately, the environmental magnetic noise that we normally record inside the MSR (see Figures 4.4 or 4.6, for example) overlaps considerably with the signal of interest (Figure 6.3). Thus, we must take care in selecting the filters so as not to also eliminate large amounts signal along with the noise.

The first filter is aimed at removing the electric line noise at 60 Hz. This signal is problematic in that it is both large in magnitude and also sits squarely in our measurement band of interest. Much of the "fuzz" visible on the signals in Figure 6.5 is 60 Hz noise, which can be corroborated by noting size of the peak in the power spectral density plot in Figure 6.3. Fortunately, this noise is also extremely narrow-band and can be almost completely suppressed with the application of a steep (we use 4-pole) digital notch filter, which we apply using *Saki*. The higher harmonics (120

and 180 Hz) also figure prominently in the total observed noise. These peaks could be dealt with similarly, but since they lie outside of our frequency band of interest, we use a bandpass filter to limit the noise contributions at these frequencies.

Because of the limited bandwidth of the cardiographic signals (Figure 6.3), applying bandpass filters with parameters roughly matching the region of interest is an easy way to clean up the recorded signal for further processing. Several groups [Brisinda et al., 2005, Grimm et al., 2003, Stinstra et al., 2001] have experimented with different filters to improve signal quality for fMCG, but a more rigorous comparison of filters was provided in [Comani et al., 2005]. After testing multiple bandpass filters on sets of fMCG data (with fetal gestational ages ranging from 22-37 weeks), they determined that the optimal bandpass filter for preserving fetal signal while eliminating unwanted noise was one that spanned 1-100 Hz. We have followed their example and apply a 1-100 Hz 2-pole infinite impulse response (IIR) bandpass filter on all of our recorded channels. The application of this filter eliminates the slowly-varying ($\sim$0.1-0.5 Hz) fields produced by environmental sources such as cars, elevators, and MRI machines. At the same time, it removes much of the high-frequency white noise as well as the aforementioned 120 Hz and 180 Hz harmonics of the electric line noise.

Figure 6.6 shows the eight channels of data from Figure 6.5 after the notch and bandpass filters described in this section were applied. The fetal signals are now more prominent, especially on the fifth trace ($B_{1y}$), and evidence of those signals can be seen on other channels as well. While this is clearly a drastic improvement over the raw data, in order to fully realize the contents of our data, further signal processing is needed. This processing utilizes the power of all eight channels collectively, rather than individually, to obtain an even more refined set of data. This set will then be used to calculate the critical heart intervals that needed to diagnose fetal arrhythmias.

## 6.6 Multi-channel processing

In addition to the maternal signal, there are a multitude of noise sources whose bandwidth overlaps the 1-100 Hz band of the fetal signal. Excluding extremely narrowband noise sources like the 60 Hz line noise, it is impossible to simply apply bandpass filters to eliminate these noise sources

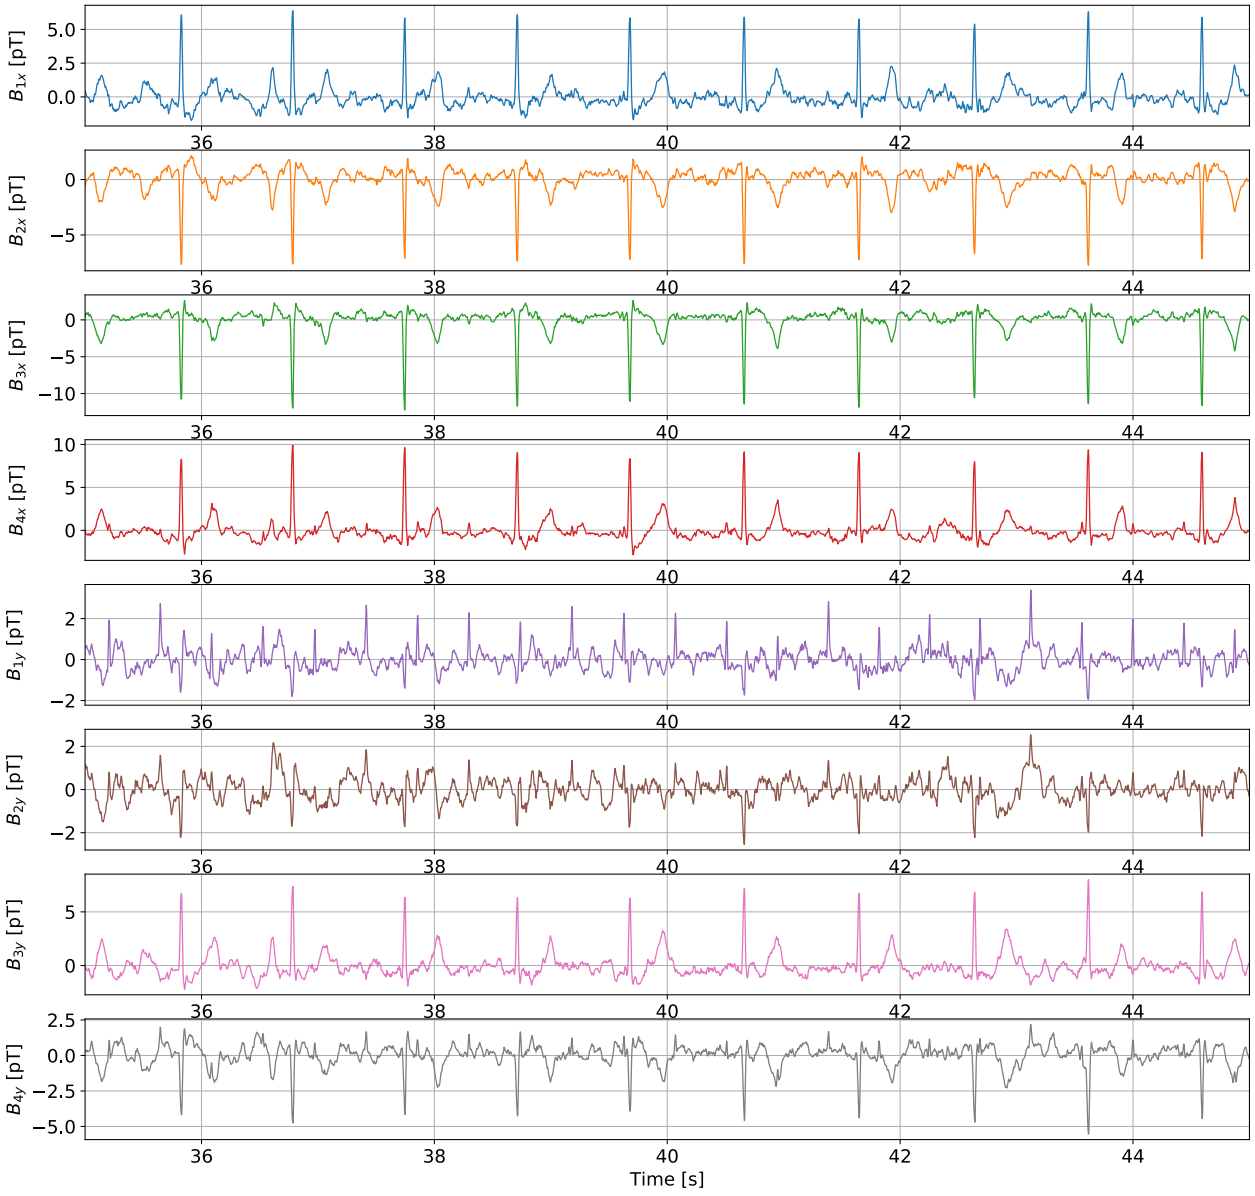



**Figure 6.6:** The data from Figure 6.5 after each channel is calibrated by its own calibration function $\zeta_i(\omega)$ and two digital filters are applied: a notch filter at 60 Hz to remove electronic line noise and a 1-100 Hz bandpass filter that removes all signal outside of our band of interest. The fetal *QRS*-peaks (occurring at roughly twice the rate of the maternal peaks) are now apparent, especially on Channel $1y$, and small indications of those peaks are visible on other channels as well. Note the $y$ axes scales for each channel are different.

without simultaneously compromising the interesting fetal data. We instead attempt a strategy known as *blind-source separation*. In blind source separation, the time series from a number of independent sensors (each of which contains information from several *sources* of signal) are used to attempt to reconstruct each of the individual source signals as accurately as possible. One of the methods for performing blind source separation is called Independent Component Analysis (ICA) [Comon, 1994], and its application has been successful in many previous experiments with fMCG [Brisinda et al., 2005, Comani et al., 2004, Mantini et al., 2006, Yu and Wakai, 2011].

(Note: From this point forward, we will focus solely on fMCG data. Adult MCG data does not usually go through the same rigorous multi-channel processing, as the SNR is already quite large, and further processing is not required. Some sample 8-channel Z-mode data acquired during an adult measurement is included in Appendix E.)

### 6.6.1 An overview of ICA

Consider a situation with $m$ independent magnetic field sources $s_1(t),\ s_2(t), ...,\ s_m(t)$ detected by $n$ magnetic field sensors. In order to independently discern the $m$ sources, ICA requires $n \geq m$. (In other words, a collection of $n$ sensors can discern *at most* $n$ independent sources.) The time series from each of the $n$ sensors is $x_i(t)$. Each $x_i(t)$ has contributions from each of the $m$ sources. The amplitude and phase of these contributions vary depending on their positions and orientations relative to the sources. That is,

$$\begin{aligned} x_1(t) &= a_{11}s_1(t) + a_{12}s_2(t) + ... + a_{1n}s_m(t) \\ x_2(t) &= a_{21}s_1(t) + a_{22}s_2(t) + ... + a_{2n}s_m(t) \\ &\vdots \qquad\qquad\qquad\qquad \vdots \\ x_n(t) &= a_{n1}s_1(t) + a_{n2}s_2(t) + ... + a_{nn}s_m(t), \end{aligned} \tag{6.4}$$

or, in more compact terms,

$$\mathbf{x} = \mathbf{A}\mathbf{s}, \tag{6.5}$$

where $\mathbf{x} = [x_1(t)\ \ x_2(t)\ \ ...\ \ x_n(t)]^T$, $\mathbf{s} = [s_1(t)\ \ s_2(t)\ \ ...\ \ s_m(t)]^T$, and $\mathbf{A}$ is a $n \times m$ *mixing matrix* that contains information pertaining to the amount of each source present in each recorded channel [Hyvärinen and Oja, 2000].

To reconstruct the sources $s_i(t)$, one could imagine simply solving Equation (6.5) for $\mathbf{s}$:

$$\mathbf{s} = \mathbf{W}\mathbf{x}, \tag{6.6}$$

where $\mathbf{W}$ is the $m \times n$ "unmixing" matrix. In fact, this is similar to the strategy we used to "decouple" the magnetic fields produced by the shell-mounted coils in Section 3.5.4. The key difference here is that *we do not know the coefficients* $a_{ij}$, so constructing the deconvolving matrix $\mathbf{W}$ directly is not possible.

Instead, we begin by "whitening" the components of $\mathbf{x}$ by Principal Component Analysis or Singular Value Decomposition. The whitening process linearly transforms $\mathbf{x}(t)$ to a vector $\mathbf{x}'(t)$, whose components are uncorrelated, each with a variance of unity. We then relate the whitened vector $\mathbf{x}'(t)$ to a source *estimate* vector $\mathbf{y}(t)$ by

$$\mathbf{y} = \mathbf{W}\mathbf{x}' \tag{6.7}$$

$\mathbf{y}$ is constructed by minimizing the Average Mutual Information (AMI) between the components; in other words, making the components as independent as possible. The Average Mutual Information (AMI) can be calculated by [Comani et al., 2004, Comon, 1994, Mantini et al., 2006]:

$$\text{AMI}(p_y) = \int p_y(u) \log \frac{p_y(u)}{\prod_{i=1}^{N} p_{yi}(u)} du, \tag{6.8}$$

Here, $p_y$ is the *joint* probability distribution function for all of the $y_i$, and the denominator represents a product of the *individual* probability distribution functions $p_{iy}$. If the vectors are truly independent, the numerator and denominator of Equation (6.8) would be equal and the AMI would be zero; in other words, the components would be completely independent. In general, though, this is not the case, so the minimum value of the AMI will identify the *most* independent case. There are several algorithms that can be used for this iterative minimization procedure, and studies have been performed to compare their effectiveness for fMCG [Mantini et al., 2006]. The Wakai Lab's

*Saki* program utilizes the CubICA algorithm, which is explained in detail by Yu and Wakai [2011] and Blaschke and Wiskott [2004].

When complete, the vector $\mathbf{y}(t)$ will be a very good estimate for the *actual* source vector $\mathbf{s}(t)$. After constructing $\mathbf{y}(t)$, that vector and $\mathbf{x}'(t)$ are used in conjunction to solve for the components $w_{ji}$ of the unmixing matrix. Constructing the matrix $\mathbf{W}$ and applying it to the original sensor time series $\mathbf{x}$ will yield the source vector $\mathbf{s}$, and the *independent source components* $s_i(t)$. With the matrix $\mathbf{W}$ complete, the construction of the original *mixing* matrix $\mathbf{A}$ is straightforward.

### 6.6.2 ICA Implementation

Since we had $n = 8$ sensors (4 measuring $B_x$ and 4 measuring $B_y$), we could at most discern $m = 8$ independent magnetic field sources. Despite knowing in advance that we would only be interested in two sources (namely, the fetal and maternal heartbeats), we still opted to construct 8 ICA component time series. We knew from prior experience that the interesting fetal signal would likely show up on multiple channels, and others would contain pure noise. Performing ICA on our fetal dataset results in the components shown in Figure 6.7. It should be stressed here that the data shown in this and the preceding two figures shows a ten second sample of the data; the full dataset (60 seconds) was used when calculating the ICA components $s_i(t)$.

### 6.6.3 Splined Independent Component Subtraction

The value of separating out the components of the signal becomes clear when modified versions of the original time series $\mathbf{x}(t)$ are reconstructed from the components $\mathbf{s}(t)$. After identifying the channels with noticeable fetal content (in the case of Figure 6.7, components 4, 5, 6, and 7) and maternal content (components 2 and 3), the original time series can be reconstructed using a process called Splined Independent Component Subtraction (spICS) [Yu and Wakai, 2011]. spICS consists of the followings steps:

1. The components containing primarily the maternal signal are *splined*, or downsampled, which effectively removes the higher-frequency fetal signal in those components.

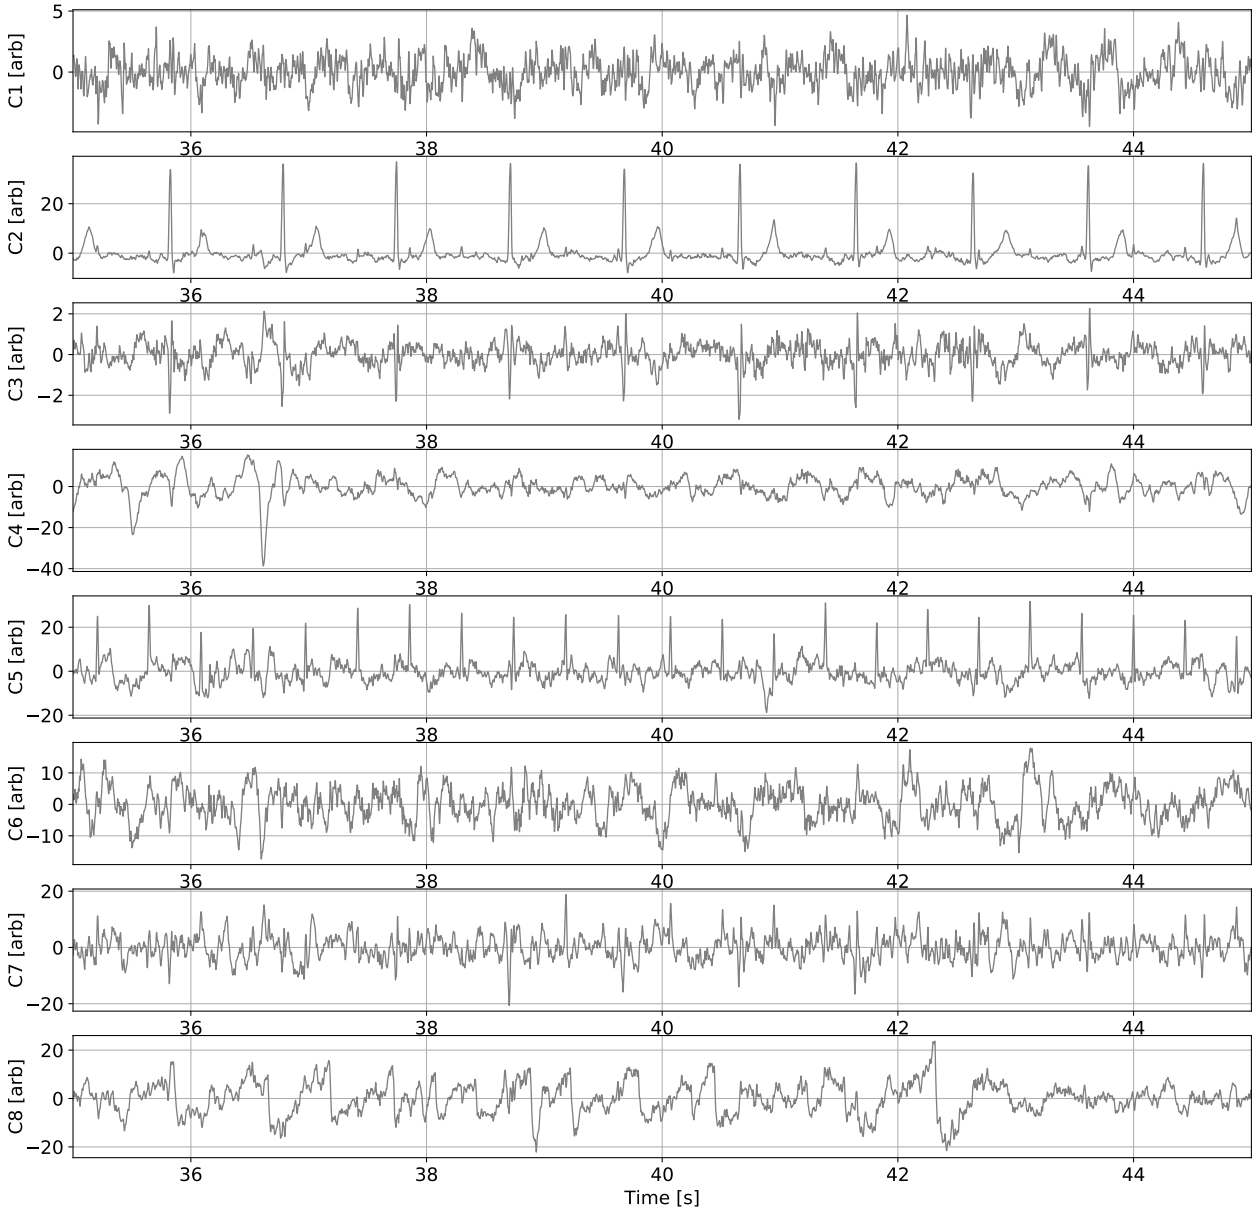



**Figure 6.7:** The eight ICA components after applying the CubICA algorithm to the filtered data in Figure 6.6. The maternal signal is primarily in component 2, while parts the *QRS* complex are also quite visible in component 3. Most of the fetal signal is confined to component 5, though evidence can also be seen in other components, including a small amount in the primary maternal component (2). The individual components were calculated using the *entire* 60 second time series of each magnetometer channel, not just the 10 second sample shown in Figures 6.5 and 6.6.

2. Next, the original time series $\mathbf{x}(t)$ is reconstructed using Equation (6.5), except the ICA components which that primarily noise (1 and 8) are removed (by setting that row of $\mathbf{A}$ identically equal to zero). The reconstructed signals should now contain primarily maternal and fetal content.

3. Finally, the splined versions of the maternal signal are projected back to magnetic field space and *subtracted* from each of the original signals. Because the splined versions of the maternal signal should contain little to no fetal signal, subtracting the reprojections of the splined components results in minimal loss of fetal signal.

(A somewhat simpler method, "normal" Independent Component Subtraction (ICS), treats the (unsplined) maternal components like noise components and simply "zeroes out" their contribution (using the procedure of step 2, above) during the reprojection step. However, this results in the complete loss in any fetal signal present in those channels. In Figure 6.7, some fetal signal is visible in component 2. This information would be lost with "normal" ICS.

The results of performing spICS on our dataset are shown in Figure 6.8. Notice the large features (*P* wave, *QRS* complex, and *T* wave) of the maternal signal are suppressed (for the most part). Some evidence of the maternal *QRS* peaks remains (see $B'_{2x}$ and $B'_{3x}$, for example), but those remaining features are more than ten times smaller than they were before the spICS process.

## 6.7 Averaging

Although the data in Figure 6.8 are of much higher quality than the original data (Figure 6.5), it would still be inadvisable to attempt to measure any heart intervals without first averaging each channel. Averaging can make the periodic heartbeat features more visible by eliminating some of the remaining non-periodic white noise. Our signal processing software has a helpful feature that allows the user to simply indicate on the time-axis location of each fetal *QRS* complex, and the software averages the time series around each of those peaks over a range encompassing three *QRS* complexes (one central complex and the next-nearest *QRS* peaks on each side) automatically.

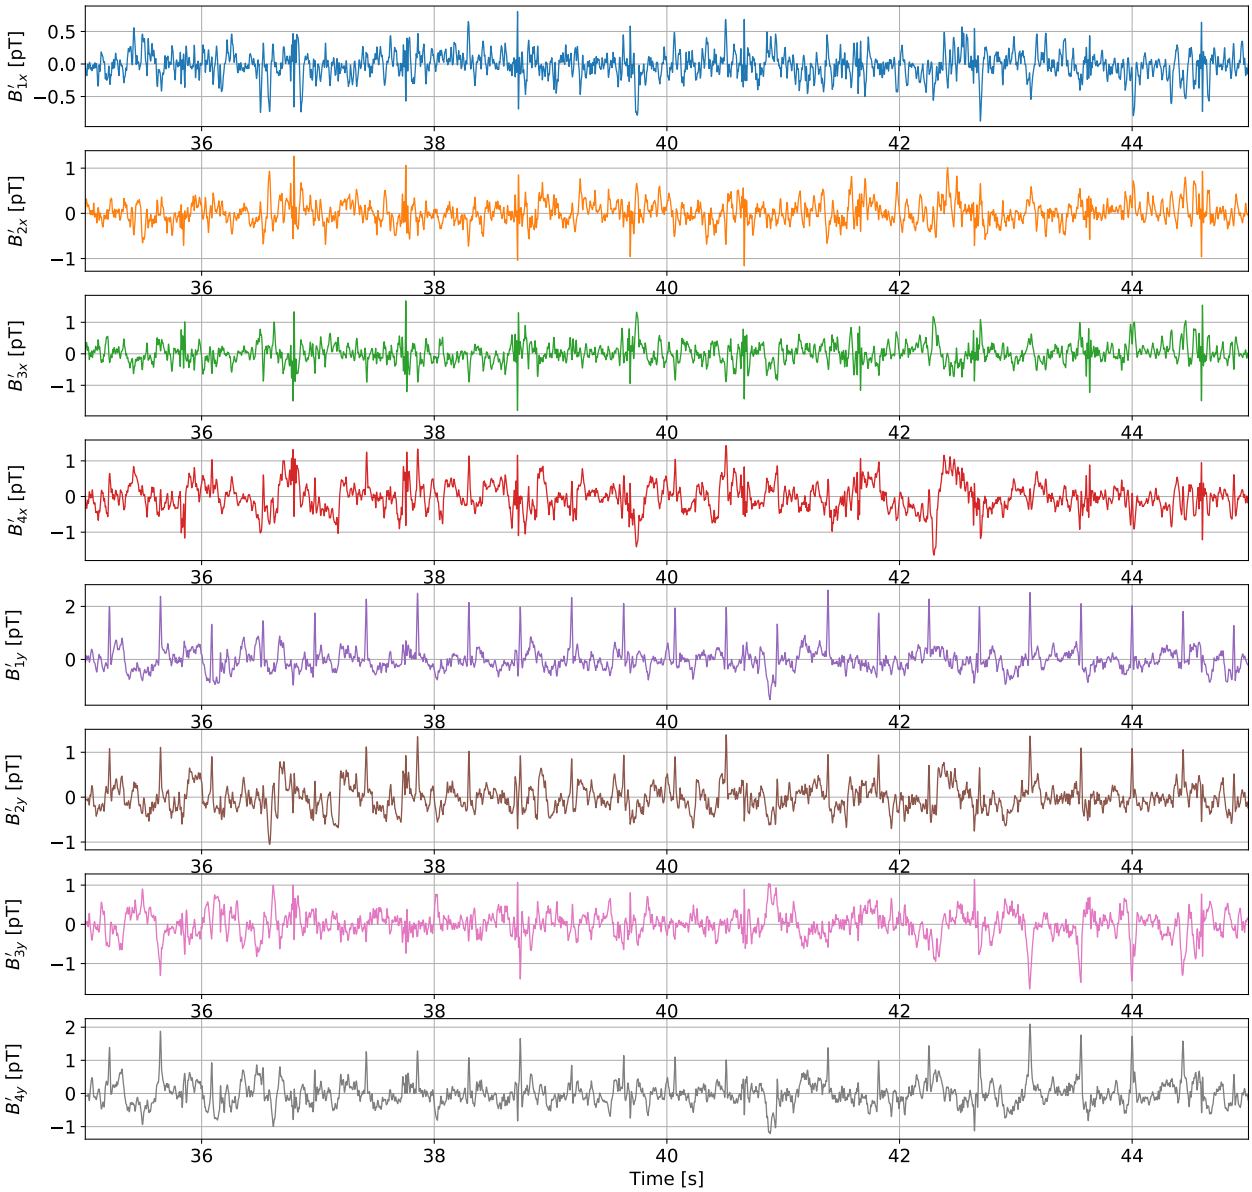



**Figure 6.8:** The eight magnetic field time series (the $\hat{x}$ and $\hat{y}$ components of the four channels of the array), after spICS reconstruction. The maternal features have been suppressed and some of the noise has been removed during the reprojection of the ICA components. The fetal *QRS* complexes are now clear on many channels, though the other structures in the cardiogram are still somewhat difficult to discern. The

Because of variations in the fetal heart rate over the measurement interval, the "peripheral" features (the nearest-neighbor *QRS* complexes, *e.g.*) are reduced in magnitude by the averaging process. Fortunately, those features are only required to calculate the *RR* interval, and with the *R* peaks being by far the largest feature of the cardiographic signal, even a reduced-amplitude *R* peak is clearly visible. The eight averaged traces are overlaid onto one plot, as shown in Figure 6.9. It is from these data that the final analysis of the important heartbeat intervals is extracted.

## 6.8 Discussion

Overall we are pleased with the quality of these data. This was the first time that our research simultaneously captured eight channels of data in a fMCG measurement. Furthermore, the gestational age of the fetus was 28 weeks, three weeks younger than the fetus measured by Wyllie [2012]. Using the *Saki* software, the final processed and averaged dataset (Figure 6.9 was judged by an fMCG specialist to be of sufficient quality to be used for a clinical evaluation.

With that being said, there is still room for improvement in signal quality. Although eight channels of data were captured during the measurement, a significant fetal signal is difficult to discern (even after averaging) on several of those channels. Part of issue is simply positioning; the pre-collection ultrasound gives us a best estimate of the position of the fetal heart, but fetuses do not always remain still during a measurement. For this reason, it is common to see little to no evidence of fetal signal on a particular sensor for some time, followed by a sudden appearance of the signal as the fetus shifts position inside the mother's uterus. A good example of this phenomenon can be observed near the 43-second mark in Figure 6.8 on the seventh and eighth channels.

The first four channels of data in Figure 6.8 also show little evidence of fetal signal. Recall that these are the channels measuring the $\hat{x}$ component of the magnetic field (one of the components in the plane of the mother's abdomen), and it appears from the lack of fetal signal that the magnetic field produced by the fetal heart is primarily perpendicular to the plane of the mother's abdomen (the $\hat{y}$ direction in our coordinate space). This is confirmed by the much stronger signal on our $\hat{y}$ magnetometer signals (channels 4-8). On the other hand, channels 1-4 show a strong and relatively clean *maternal* signal (see Figure 6.6). Conversations with the Wakai group (whose SQUID device

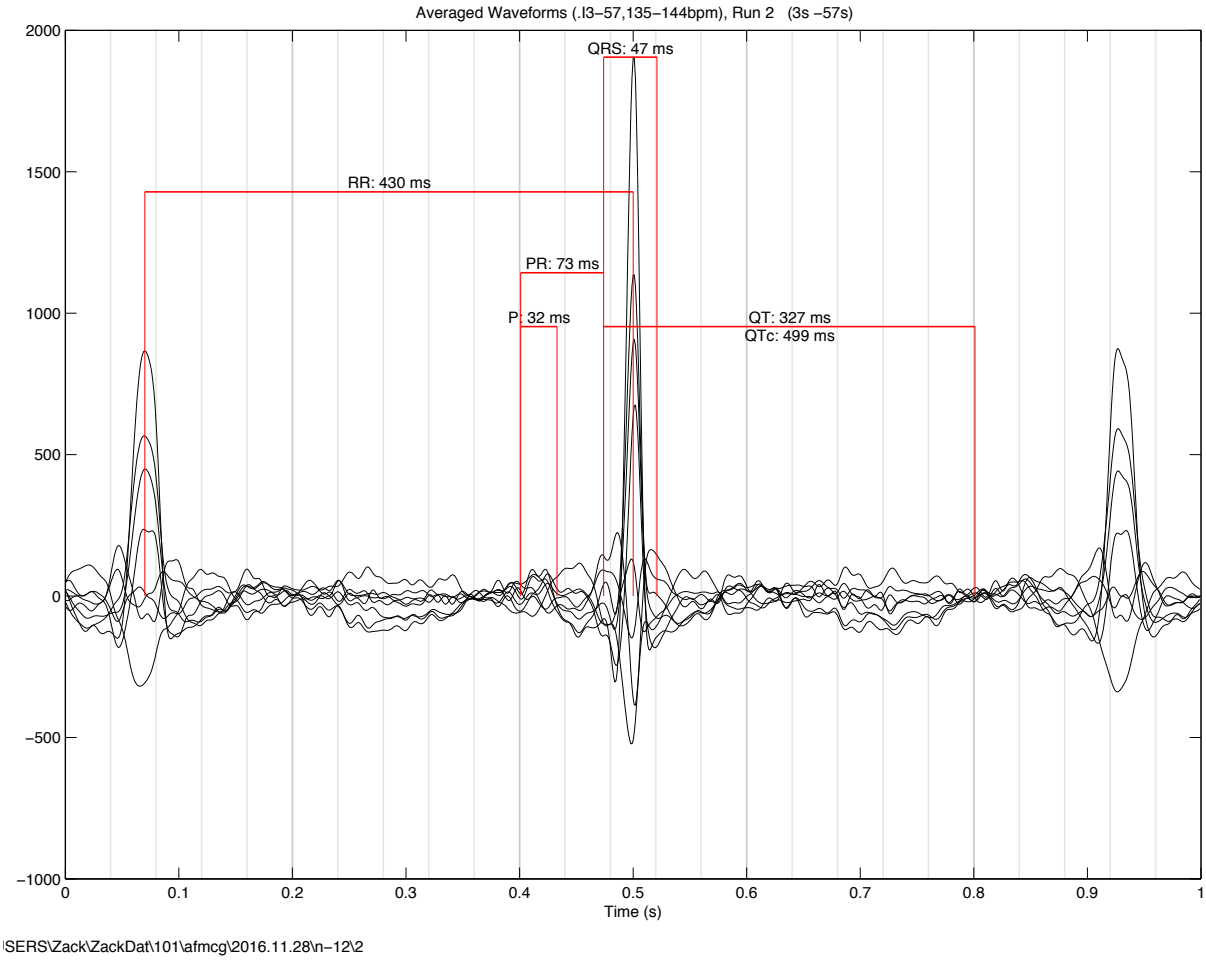



**Figure 6.9:** The averaged waveforms of the eight channels of processed fMCG data. The $y$ axis of this graph has units of femtotesla. Variations in the fetal heart rate mean that the averaging tends to "smear out" the features in the peripheral beats (such as the "QRS" complexes near .07 s and .93 s), but the central features are well-defined. The relevant cardiographic intervals are indicated by the red lines. The *QTc* intereval time is *not* a true measurement in milliseconds; rather, it is a value that combines the lengths of the *QT* and *RR* intervals to help fetal cardiologists determine if the fetus may be at risk for long-*QT* syndrome. This image was produced using the Wakai group's *Saki* software.

measures $B_x$, $B_y$, and $B_z$) corroborate these observations. We were indeed surprised at the relative lack of maternal signal in our $y$ channels, especially $B_{1y}$ and $B_{2y}$. These were the two channels positioned furthest from the maternal heart, and on those channels the fetal and maternal signals were of nearly the same magnitude! The exact reason that these maternal signals were so weak is still under investigation.

## 6.9 Comparison with SQUID

In the past, fMCG measurements using our group's magnetometer array were often followed by a measurement using the Wakai Group's 21-channel SQUID gradiometer device. This allowed us to compare the quality of our data to the established method of acquiring a fetal magnetocardiogram. In addition, after putting both datasets through identical signal processing, the heartbeat intervals could be extracted from both sets, and we could compare our intervals compare to the ones obtained from the SQUID measurement. The Wakai Group's SQUID, made by Tristan Technologies, is a FDA-approved device for the clinical measurement of patients deemed to be at risk of having a fetus with a heartbeat arrhythmia. Unfortunately, at the time of our measurement, the group's SQUID was undergoing a FDA-mandated review and was unavailable for use to measure our patient.

In general, the SQUID has a significant advantage over our device for two reasons. First, the significantly larger number of channels gives a better chance that a fetal signal will be present on more traces, which aids in the multi-channel processing described earlier in this chapter. Somewhat mitigating this advantage for the SQUID is the fact that those 21 channels are actually seven triaxial sets of gradiometers (measuring $B_x$, $B_y$, and $B_z$). Much like our findings, most of the fetal signal acquired by the SQUID exists on the $y$ channels of those triaxial sets. Secondly, the gradiometric configuration of *each* of the SQUID channels means a good deal of the ambient magnetic field noise interfering with our measurement is not present in a SQUID measurement. While we showed that we could significantly reduce the effects of uniform background noise by operating our magnetometers in a gradiometric configuration (Chapter 5), this would mean effectively reducing our channel count to 7 (or likely 6 if we wanted to make a gradiometric measurement of

both $x$ and $y$ fields). The need to measure as many channels of fetal signal as possible outweighed our desire to use gradiometry for this measurement. We are currently studying the use compact hardware atomic gradiometers (each gradiometer comprised of a single vapor cell with lasers configured to measure a field difference, rather than simply subtracting two nearby magnetometers) to mimic the configuration of the SQUID device. This potential advance will be discussed in the next chapter.

# Chapter 7

# Conclusion and Future Directions

In this thesis, we have presented a complete description of our magnetometer array for the measurement of adult and fetal magnetocardiograms. We began with a description of the theory of SERF magnetometry, which provides the basis for measuring very small magnetic fields. Next, the theory behind three important extensions to SERF-mode were explained: Z-mode, diffusion-mode (diff-SERF), and active feedback mode. These three extensions allow us to circumvent three problematic limitations to normal SERF operation: low-frequency technical noise, large AC-Stark shift fields, and drifts in the ambient magnetic field inside the magnetically shielded room, respectively.

We then presented an in-depth look at our apparatus, including all of the upgrades and modifications made to the device since the publishing of my predecessor's thesis [Wyllie, 2012]. In the next chapter, we explained our data-taking procedure and the data analysis that we use to characterize how well our magnetometer units are functioning. Their performance is dictated by the noise detected by the units in a non-patient setting, so we explained how we estimate the noise produced by fundamental, technical, and real magnetic noise sources. We ended this chapter with a discussion of the theory behind these noise sources and how it affects the measurements we make.

The final two chapters of this thesis were devoted to two major experimental results. First, spurred on by the successes of SQUID devices and atomic gradiometers, we desired to push the limit of common-mode noise suppression by operating two adjacent SERF magnetometer units in a subtractive gradiometric configuration. In this mode, we were able to suppress added uniform background noise by a factor of roughly 400 across our entire measurement bandwidth. This shows much promise for using atomic gradiometers to measure small non-uniform signals in uniformly

noisy environments. We concluded with a discussion of the biomagnetic signal of interest, the fetal magnetocardiogram (fMCG). We followed this by showcasing a recent 8-channel Z-mode measurement of the fMCG, along with the signal processing required to produce a clinically-viable set of data. (We again stress that the goal of this measurement was to evaluate the performance of the device. No clinical advice or recommendations were delivered as a result of our measurements).

To evaluate the success of this phase of our group's ongoing biomagnetism work, it is only natural to compare it against the work that immediately preceded mine [Wyllie, 2012]. We believe we have made significant progress from that work. Namely, we were able to successfully measure a younger fetus (with a weaker fMCG signal) and doubled the amount of magnetic field information collected by operating in Z-mode. Furthermore, we managed to shrink the overall footprint of the device and implement an option for individual-channel $y$-field feedback in addition to "common coil" feedback.

With the goal of making the device a viable replacement for a SQUID device for fMCG measurement, there is still much that can be done to improve our methods and apparatus. Taking a cue from the SQUID device, the implementation of hardware gradiometers would be a sensible upgrade to our apparatus. The gradiometry explored in this thesis was based on two nearby, but separate magnetometers. Because of the many differences between the two units (buffer gas pressure, pump and probe laser intensities, etc.), the magnetometer calibrations $\zeta_i(\omega)$ were often quite different. The normalization of the two $\zeta_i$s, and hence, the actual magnetic field subtraction, had to be completed after the fact in software. The suppression of common-mode fields is limited by our ability to reconcile these two calibrations; thus, a situation where the calibrations were identical would is desirable.

One possible implementation of a SERF gradiometer is a setup similar to our current device, except after the first pass of the probe beam through the cell, the beam is translated using two corner mirrors along the $\hat{z}$ direction and sent back through the optically-pumped vapor anti-parallel to the initial pass. Assuming the two regions of vapor are optically pumped equally, any Faraday rotation in the first pass would be canceled by an equal amount of rotation in the opposite direction in the second pass. That is, *unless* the field causing the rotation was not the same at the two paths through

the cell (*i.e.*, the baseline of the gradiometer). A uniform magnetic field would cause the probe to rotate one way in the first region, but an equal amount back the other way in the second region. Only a gradient field would give a non-zero final rotation. Another possible gradiometer could include a single probe beam but two *pump* beams with opposite handedness.

Another area in which improvements could be made is in active feedback for the nulling coils. We have shown that we can stabilize the $y$ field quite well, but the transverse fields ($x$ and $z$) generally operate in "open loop" mode. This can be problematic if the field drifts in these directions are enough to significantly reduce the sensitivity of the device (see Figure 2.10, for example). Seltzer and Romalis developed a scheme in which active feedback can be applied to all three axes of a SERF magnetometer simultaneously [Seltzer and Romalis, 2004], and we hope to implement a similar system. Robust stabilization of the magnetic fields may even allow us to do *in-situ* positioning of the array, where one experimenter adjusts the positioning and orientation of the device guided by a second experimenter who is the fetal signal in real time.

We must continue to recruit and measure test subjects on a regular basis. Several attempts were made in the course of this work to measure a fMCG with younger fetuses, as early as 20 weeks gestational age in some cases. Despite evidence that fetal signals should be detectable at this age [Golbach et al., 2000], we were unsuccessful in observing meaningful fetal signal in these cases, even after significant signal processing. Successful detection requires careful array positioning, not to mention a good amount of luck. The depth and position of the fetal heart inside of the mother's uterus, two things over which we have frustratingly little control, can be the difference between a successful or unsuccessful fMCG measurement. More regular measurement of subjects would not necessarily improve our success rate, but it should give us a larger amount of quality data to study.

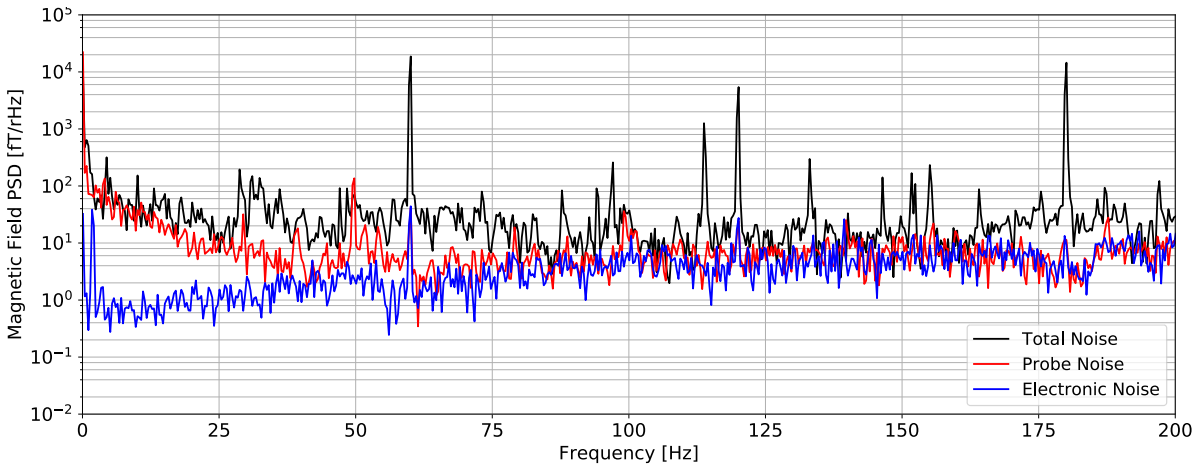
Finally, we would like to begin to take advantage of one of the main selling points of our apparatus: its small size compared to a SQUID measurement system. Most importantly, the smaller footprint of our device should allow us to work in a much smaller magnetic shield–a shield that would cost a fraction of the price of the SQUID system's magnetically shielded room where we currently operate. Our collaborators in Medical Physics have recently purchased such a shield, a 2.5 m length by 1.0 m diameter three-layer cylindrical shield with a patient bed that can slide in



**(a)** 3-layer cylindrical shield



**(b)** Cylindrical shield, end view




**(c)** Detected noise inside the cylindrical shield

**Figure 7.1:** Photographs and performance of the 3-layer cylindrical μ-metal shield. The outermost layer of the shield has a diameter of approximately 1 m and a length of 2.2 m. During measurement, the patient lies on the padded bed and is slid into the shield by the experimenters. The magnetometers are attached to the inside of the shield and lowered into place above the mother's abdomen. The far end of the shield is enclosed by three μ-metal caps. The near end can also be capped (the caps have small holes for fiber and cable access), though noise performance is mostly unaffected by taking that cap on and off. As seen in (c), the total magnetic noise observed inside the cylindrical shield (with the caps on one end off) is higher than the noise inside the MSR, but only by a factor of 2-4 in most frequency bands. The 60 Hz line noise is noticeably higher, but with the ability to notch filter that frequency, its detriment to a fetal measurement would be small. The elevated probe noise during this measurement was not due to the change in shields; more likely it was an improper balance of the detector.

and out. Some preliminary measurements have been made with our magnetometers inside of this shield (Figure 7.1), and apart from some significant noise increases at the line frequencies (integer multiples of 60 Hz), the noise levels inside of this shield are only a factor of two or so higher than the levels observed in the much larger and more expensive MSR. Attempts will be made to study a pregnant mother under these conditions.

With the advances made in this thesis and the interesting avenues for future work that they suggest, we hope to one day make our apparatus a reliable, relatively inexpensive alternative to SQUID magnetometers. By making the crucial fetal diagnostic benefits of SQUIDs available with a more accessible atomic device that could be used in more hospitals around the world, we believe our device can help improve the pregnancy outcomes of thousands of mothers in the years to come.

# Bibliography

C. Affolderbach, M. Stähler, S. Knappe, and R. Wynands. An all-optical, high-sensitivity magnetic gradiometer. *Appl. Phys. B Lasers Opt.*, 75(6-7):605–612, 2002. ISSN 09462171. doi: 10.1007/s00340-002-0959-8.

C. B. Alcock, V. P. Itkin, and M. K. Horrigan. Vapor Pressure of the Metallic Elements. *Can. Metall. Q.*, 23(3):309–313, 1984. ISSN 0008-4433. doi: 10.1179/cmq.1984.23.3.309.

E. Aleksandrov, M. Balabas, and A. Ivanov. Laser pumping in the scheme of an Mx-magnetometer. *Opt. Spectrosc.*, 78(2):292–298, 1995. ISSN ISSN 0030-400X.

J. C. Allred, R. N. Lyman, T. W. Kornack, and M. V. Romalis. High-Sensitivity Atomic Magnetometer Unaffected by Spin-Exchange Relaxation. *Phys. Rev. Lett.*, 89(13):130801, 2002. ISSN 0031-9007. doi: 10.1103/PhysRevLett.89.130801.

L. Anderson, F. Pipkin, and J. Baird. Hyperfine Structure of Hydrogen, Deuterium, and Tritium. *Phys. Rev.*, 120(4):1279–1289, 1960. ISSN 0031-899X. doi: 10.1103/PhysRev.120.1279.

S. Appelt, A. B. A. Baranga, A. R. Young, and W. Happer. Light narrowing of rubidium magnetic-resonance lines in high-pressure optical-pumping cells. *Phys. Rev. A - At. Mol. Opt. Phys.*, 59 (3):2078–2084, 1999. ISSN 10502947. doi: 10.1103/PhysRevA.59.2078.

J. Bechhoefer. Feedback for physicists: A tutorial essay on control. *Rev. Mod. Phys.*, 77(3): 783–836, 2005. ISSN 00346861. doi: 10.1103/RevModPhys.77.783.

W. Bell and A. L. Bloom. Optical Detection of Magnetic Resonance in Alkali Metal Vapor. *Phys. Rev.*, 107(6):1559–1565, 1957. ISSN 0031-899X. doi: 10.1103/PhysRev.107.1559.

G. Bison, N. Castagna, A. Hofer, P. Knowles, J. L. Schenker, M. Kasprzak, H. Saudan, and A. Weis. A room temperature 19-channel magnetic field mapping device for cardiac signals. *Appl. Phys. Lett.*, 95(17), 2009. ISSN 00036951. doi: 10.1063/1.3255041.

T. Blaschke and L. Wiskott. CuBICA: Independent component analysis by simultaneous third-and fourth-order cumulant diagonalization. *IEEE Trans. Signal Process.*, 52(5):1250–1256, 2004. ISSN 1053-587X. doi: 10.1109/TSP.2004.826173.

K. Bonin and V. Kresin. *Electric-dipole polarizabilities of atoms, molecules, and clusters*. World Scientific Pub Co, Inc, New York, 1997. ISBN 9810224931.

D. Brisinda, S. Comani, A. M. Meloni, G. Alleva, D. Mantini, and R. Fenici. Multichannel mapping of fetal magnetocardiogram in an unshielded hospital setting. *Prenat. Diagn.*, 25(5):376–382, 2005. ISSN 01973851. doi: 10.1002/pd.1160.

D. D. Budker and M. Romalis. Optical magnetometry. *Nat. Phys.*, 3(4):227–234, 2007. ISSN 1745-2473. doi: 10.1038/nphys566.

W. C. Chen, T. R. Gentile, T. G. Walker, and E. Babcock. Spin-exchange optical pumping of He3 with Rb-K mixtures and pure K. *Phys. Rev. A - At. Mol. Opt. Phys.*, 75(1):1–14, 2007. ISSN 10502947. doi: 10.1103/PhysRevA.75.013416.

J. Clarke and A. I. Braginski. *The SQUID Handbook*, volume 1. 2005. ISBN 9783527603640. doi: 10.1002/3527603646.

D. Cohen and H. Hosaka. Part II magnetic field produced by a current dipole. *J. Electrocardiol.*, 9 (4):409–417, 1976. ISSN 00220736. doi: 10.1016/S0022-0736(76)80041-6.

D. Cohen, E. A. Edelsack, and J. E. Zimmerman. Magnetocardiograms taken inside a shielded room with a superconducting point-contact magnetometer. *Appl. Phys. Lett.*, 16(7):278–280, 1970. ISSN 00036951. doi: 10.1063/1.1653195.

S. Comani, D. Mantini, P. Pennesi, A. Lagatta, and G. Cancellieri. Independent component analysis: Fetal signal reconstruction from magnetocardiographic recordings. *Comput. Methods Programs Biomed.*, 75(2):163–177, 2004. ISSN 01692607. doi: 10.1016/j.cmpb.2003.12.005.

S. Comani, D. Mantini, G. Alleva, S. Di Luzio, and G. L. Romani. Optimal filter design for shielded and unshielded ambient noise reduction in fetal magnetocardiography. *Phys. Med. Biol.*, 50(23): 5509–21, 2005. ISSN 0031-9155. doi: 10.1088/0031-9155/50/23/006.

P. Comon. Independent component analysis, A new concept? *Signal Processing*, 36(3):287–314, 1994. ISSN 01651684. doi: 10.1016/0165-1684(94)90029-9.

B. F. Cuneo, M. Ovadia, J. F. Strasburger, H. Zhao, T. Petropulos, J. Schneider, and R. T. Wakai. Prenatal diagnosis and In Utero treatment of Torsades de Pointes associated with congenital long QT syndrome. *Am. J. Cardiol.*, 91(11):1395–1398, 2003. ISSN 00029149. doi: 10.1016/S0002-9149(03)00343-6.

B. F. Cuneo, J. F. Strasburger, S. Yu, H. Horigome, T. Hosono, A. Kandori, and R. T. Wakai. In Utero Diagnosis of Long QT Syndrome by Magnetocardiography. *Circulation*, 128(20), 2013. ISSN 15378276. doi: 10.1038/jid.2014.371.

H. B. Dang, A. C. Maloof, and M. V. Romalis. Ultrahigh sensitivity magnetic field and magnetization measurements with an atomic magnetometer. *Appl. Phys. Lett.*, 97(15):28–30, 2010. ISSN 00036951. doi: 10.1063/1.3491215.

B. Das, B. F. Cuneo, M. Ovadia, J. F. Strasburger, C. Johnsrude, and R. T. Wakai. Magnetocardiography-guided management of an unusual case of isoimmune complete atrioventricular block complicated by ventricular tachycardia. *Fetal Diagn. Ther.*, 24(3):282–285, 2008. ISSN 10153837. doi: 10.1159/000151677.

J. Doyle, B. Francis, and A. Tannenbaum. Feedback Control Theory. *Design*, 134(6):219, 1990. ISSN 00223395. doi: 10.1016/0005-1098(86)90018-X.

J. Fang, T. Wang, H. Zhang, Y. Li, and S. Zou. Optimizations of spin-exchange relaxation-free magnetometer based on potassium and rubidium hybrid optical pumping. *Rev. Sci. Instrum.*, 85 (12), 2014. ISSN 10897623. doi: 10.1063/1.4902567.

A. Garachtchenko, A. Matlashov, R. H. Kraus Jr., and R. Cantor. Baseline Distance Optimization for SQUID Gradiometers. *IEEE Trans. Appl. Supercond.*, 9(2):3676–3679, jun 1999.

R. K. Ghosh. *Spin Exchange Optical Pumping of Neon and its Applications*. Phd, Princeton University, 2009.

H. M. Gibbs and R. J. Hull. Spin-exchange cross sections for Rb87-Rb87 and Rb87-Cs133 collisions. *Phys. Rev.*, 153(1):132–151, 1967. ISSN 0031899X. doi: 10.1103/PhysRev.153.132.

E. Golbach, J. Stinstra, P. Grot, and M. Peters. Reference values for fetal MCG/ECG recordings in uncomplicated pregnancies. *Proc. 12th Int. Conf. Biomagn.*, pages 595–98, 2000.

W. C. Griffith, S. Knappe, J. Kitching, I. M. P. M. Ledbetter, D. Savukov, V. Budker, S. Shah, J. Knappe, D. J. Kitching, S. Michalak, and A. Xu. Femtotesla atomic magnetometry in a microfabricated vapor cell. *Opt. Express*, 18(26):27167–27172, 2010. ISSN 1094-4087. doi: 10.1364/OE.18.027167.

B. Grimm, J. Haueisen, M. Huotilainen, S. Lange, P. Van Leeuwen, T. Menendez, M. J. Peters, E. Schleussner, and U. Schneider. Recommended Standards for Fetal Magnetocardiography. *PACE - Pacing Clin. Electrophysiol.*, 26(11):2121–2126, 2003. ISSN 01478389. doi: 10.1046/j.1460-9592.2003.00330.x.

S. Groeger, G. Bison, J. L. Schenker, R. Wynands, and A. Weis. A high-sensitivity laser-pumped Mx magnetometer. *Eur. Phys. J. D*, 38(2):239–247, 2006. ISSN 14346060. doi: 10.1140/epjd/e2006-00037-y.

M. Hamalainen, R. Hari, R. J. Ilmoniemi, J. Knuutila, and O. V. Lounasmaa. Magnetoencephalography theory, instrumentation, and applications to noninvasive studies of the working human brain. *Rev. Mod. Phys.*, 65(2):413–497, 1993. ISSN 00346861. doi: 10.1103/RevModPhys.65.413.

W. Happer. Optical Pumping. *Rev. Mod. Phys.*, 44(2):169–250, apr 1972. ISSN 0034-6861. doi: 10.1103/RevModPhys.44.169.

W. Happer and A. C. Tam. Effect of rapid spin exchange on the magnetic-resonance spectrum of alkali vapors. *Phys. Rev. A*, 16(5):1877–1891, 1977. ISSN 10502947. doi: 10.1103/PhysRevA.16.1877.

W. Happer and W. A. Van Wijngaarden. An optical pumping primer. *Hyperfine Interact.*, 38(1-4): 435–470, 1987. ISSN 03043843. doi: 10.1007/BF02394855.

W. Happer, Y.-Y. Jau, and T. Walker. *Optically Pumped Atoms*. Wiley, Weinheim, Germany, 1st edition, 2010. ISBN 9783527629503. doi: 10.1002/9783527629503.

P. C. Hobbs. Ultrasensitive laser measurements without tears. *Appl. Opt.*, 36(4):903–920, 1997. ISSN 0003-6935. doi: 10.1364/AO.36.000903.

P. C. Hobbs. *Building electro-optic systems*. Wiley Interscience, Hoboken, NJ, 1st edition, 2000. ISBN 978-0471246817.

A. Hyvärinen and E. Oja. Independent component analysis: algorithms and applications. *Neural networks*, 13(4-5):411–430, 2000. ISSN 0893-6080. doi: 10.1016/S0893-6080(00)00026-5.

K. Ishikawa and T. Yabuzaki. Diffusion coefficient and sublevel coherence of Rb atoms in N2 buffer gas. *Phys. Rev. A - At. Mol. Opt. Phys.*, 62(6):065401–065401, 2000. ISSN 10502947. doi: 10.1103/PhysRevA.62.065401.

Y. Ito, H. Ohnishi, K. Kamada, and T. Kobayashi. Sensitivity Improvement of Spin-Exchange Relaxation Free Atomic Magnetometers by Hybrid Optical Pumping of Potassium and Rubidium. *IEEE Trans. Magn.*, 47(10):3550–3553, oct 2011. ISSN 0018-9464. doi: 10.1109/TMAG.2011.2148191.

C. N. Johnson, P. D. D. Schwindt, and M. Weisend. Multi-sensor magnetoencephalography with atomic magnetometers. *Phys. Med. Biol.*, 58:6065–77, 2013. ISSN 1361-6560. doi: 10.1088/0031-9155/58/17/6065.

K. Kamada, Y. Ito, S. Ichihara, N. Mizutani, and T. Kobayashi. Noise reduction and signal-to-noise ratio improvement of atomic magnetometers with optical gradiometer configurations. *Nat. Phys*, 3(5):12668–12671, 2006. ISSN 1094-4087. doi: 10.1364/OE.23.006976.

K. Kim, S. Begus, H. Xia, S. K. Lee, V. Jazbinsek, Z. Trontelj, and M. V. Romalis. Multi-channel atomic magnetometer for magnetoencephalography: A configuration study. *Neuroimage*, 89: 143–151, 2014. ISSN 10538119. doi: 10.1016/j.neuroimage.2013.10.040.

J. Kitching, S. Knappe, and E. a. Donley. Atomic Sensors: A Review. *IEEE Sens. J.*, 11(9): 1749–1758, 2011. ISSN 1530-437X. doi: 10.1109/JSEN.2011.2157679.

S. Knappe, V. Gerginov, P. D. D. Schwindt, V. Shah, H. G. Robinson, L. Hollberg, and J. Kitching. Atomic vapor cells for chip-scale atomic clocks with improved long-term frequency stability. *Opt. Lett.*, 30(18):2351–2353, sep 2005. ISSN 0146-9592. doi: 10.1364/OL.30.002351.

H.-C. Koch, G. Bison, Z. D. Grujić, W. Heil, M. Kasprzak, P. Knowles, A. Kraft, A. Pazgalev, A. Schnabel, J. Voigt, and A. Weis. Design and performance of an absolute 3He/Cs magnetometer. *Eur. Phys. J. D*, 69(8):202, 2015. ISSN 1434-6060. doi: 10.1140/epjd/e2015-60018-7.

I. K. Kominis, T. W. Kornack, J. C. Allred, and M. V. Romalis. A subfemtotesla multichannel atomic magnetometer. *Nature*, 422(6932):596–599, 2003. ISSN 0028-0836. doi: 10.1038/nature01484.

T. W. Kornack. A Test of CPT and Lorentz Symmetry Using a K-3He Co-magnetometer. *Dissertation*, 2005.

B. Lancor. *Studies of the Efficiency of Spin Exchange Optical Pumping*. Phd, University of Wisconsin-Madison, 2011.

M. P. Ledbetter, I. M. Savukov, V. M. Acosta, D. Budker, and M. V. Romalis. Spin-exchange-relaxation-free magnetometry with Cs vapor. *Phys. Rev. A - At. Mol. Opt. Phys.*, 77(3):1–7, 2008. ISSN 10502947. doi: 10.1103/PhysRevA.77.033408.

S. K. Lee and M. V. Romalis. Calculation of magnetic field noise from high-permeability magnetic shields and conducting objects with simple geometry. *J. Appl. Phys.*, 103(8):1–10, 2008. ISSN 00218979. doi: 10.1063/1.2885711.

Z. Li. *Development of a Parametrically Modulated SERF Magnetometer*. Phd, University of Wisconsin-Madison, 2006.

Z. Li, R. T. Wakai, and T. G. Walker. Parametric modulation of an atomic magnetometer. *Appl. Phys. Lett.*, 89(13):23–25, 2006. ISSN 00036951. doi: 10.1063/1.2357553.

L.-A. Liew, S. Knappe, J. Moreland, H. Robinson, L. Hollberg, and J. Kitching. Microfabricated alkali atom vapor cells. *Appl. Phys. Lett.*, 84(14), 2004.

D. Mantini, K. E. Hild, G. Alleva, and S. Comani. Performance comparison of independent component analysis algorithms for fetal cardiac signal reconstruction: a study on synthetic fMCG data. *Phys. Med. Biol.*, 51(4):1033–1046, 2006. ISSN 0031-9155. doi: 10.1088/0031-9155/51/4/018.

P. McSharry and G. Clifford. ECGSYN: A realistic ECG waveform generator, 2015. URL `https://www.physionet.org/physiotools/ecgsyn/`.

NIH. Your Heart's Electrical System, 2011. URL `https://www.nhlbi.nih.gov/health/health-topics/topics/hhw/electrical`.

NIH. Electrocardiogram, 2016. URL `https://www.nhlbi.nih.gov/health/health-topics/topics/ekg`.

S. Pildner von Steinburg, A.-L. Boulesteix, C. Lederer, S. Grunow, S. Schiermeier, W. Hatzmann, K.-T. M. Schneider, and M. Daumer. What is the "normal" fetal heart rate? *PeerJ*, 1:e82, 2013. ISSN 2167-8359. doi: 10.7717/peerj.82. URL `http://www.ncbi.nlm.nih.gov/pubmed/23761161`.

H. W. P. Quartero, J. G. Stinstra, E. G. M. Golbach, E. J. Meijboom, and M. J. Peters. Clinical implications of fetal magnetocardiography. *Ultrasound Obstet. Gynecol.*, 20(2):142–153, 2002. ISSN 09607692. doi: 10.1046/j.1469-0705.2002.00754.x.

M. V. Romalis. Hybrid Optical Pumping of Optically Dense Alkali-Metal Vapor without Quenching Gas. *Phys. Rev. Lett.*, 105(24):243001, dec 2010. doi: 10.1103/PhysRevLett.105.243001.

M. V. Romalis, E. Miron, and G. D. Cates. Pressure broadening of Rb D-1 and D-2 lines by He-3, He-4, N-2, and Xe: Line cores and near wings. *Phys. Rev. A*, 56(6):4569–4578, 1997.

I. M. Savukov and M. V. Romalis. Effects of spin-exchange collisions in a high-density alkali-metal vapor in low magnetic fields. *Phys. Rev. A - At. Mol. Opt. Phys.*, 71(2):2–9, 2005. ISSN 10502947. doi: 10.1103/PhysRevA.71.023405.

I. M. Savukov and S. J. Seltzer. SERF Magnetometers. In D. Budker and D. F. Jackson Kimball, editors, *Opt. Magnetometry*, chapter 5, pages 85–104. Cambridge University Press, Cambridge, 1st edition, 2013. ISBN 978-1-107-01035-2.

P. D. D. Schwindt, B. Lindseth, S. Knappe, V. Shah, J. Kitching, and L. A. Liew. Chip-scale atomic magnetometer with improved sensitivity by use of the Mx technique. *Appl. Phys. Lett.*, 90(8): 3–5, 2007. ISSN 00036951. doi: 10.1063/1.2709532.

S. J. Seltzer. *Developments in Alkali-Metal Atomic Magnetometry*. Phd thesis, Princeton University, 2008.

S. J. Seltzer and M. V. Romalis. Unshielded three-axis vector operation of a spin-exchange-relaxation-free atomic magnetometer. *Appl. Phys. Lett.*, 85(20):4804–4806, 2004. ISSN 00036951. doi: 10.1063/1.1814434.

V. Shah and M. V. Romalis. Spin-exchange relaxation-free magnetometry using elliptically polarized light. *Phys. Rev. A - At. Mol. Opt. Phys.*, 80(1):1–6, 2009. ISSN 10502947. doi: 10.1103/PhysRevA.80.013416.

V. K. Shah and R. T. Wakai. A compact, high performance atomic magnetometer for biomedical applications. *Phys. Med. Biol.*, 58(22):8153–61, 2013. ISSN 1361-6560. doi: 10.1088/0031-9155/58/22/8153.

D. Sheng, S. Li, N. Dural, and M. V. Romalis. Subfemtotesla scalar atomic magnetometry using multipass cells. *Phys. Rev. Lett.*, 110(16):1–5, 2013. ISSN 00319007. doi: 10.1103/PhysRevLett.110.160802.

D. A. Steck. Rubidium 87 D Line Data. 2001. URL `http://steck.us/alkalidata`.

J. G. Stinstra, E. G. M. Golbach, M. J. Peters, and H. W. P. Quarteo. Extracting reliable data from the fetal MCG. *Biomag2000 Proc. 12th Int Conf Biomagn.*, pages 591–594, 2001.

J. F. Strasburger, B. Cheulkar, and R. T. Wakai. Magnetocardiography for Fetal Arrhythmias. *Hear. Rhythm*, 5(7):1073–1076, 2008. doi: oi:10.1016/j.hrthm.2008.02.035.

I. Sulai, R. Wyllie, M. Kauer, G. Smetana, R. Wakai, and T. Walker. Diffusive suppression of AC-Stark shifts in atomic magnetometers. *Opt. Lett.*, 38(6):974–976, 2013. ISSN 1527-5418.

J. Vrba. Squids in Real Environments. pages 117–178, 1996.

M. E. Wagshul and T. E. Chupp. Laser optical pumping of high-density Rb in polarized He3 targets. *Phys. Rev. A*, 49(5):3854–3869, 1994. ISSN 10502947. doi: 10.1103/PhysRevA.49.3854.

R. T. Wakai, J. M. Lengle, A. C. Leuthold, and I. Address. Transmission of electric and magnetic foetal cardiac signals in a case of ectopia cordis: the dominant role of the vernix caseosa. *Phys. Med. Biol.*, 45(7):1989–1995, 2000. ISSN 0031-9155. doi: 10.1088/0031-9155/45/7/320.

T. G. Walker and W. Happer. Spin-exchange optical pumping of noble-gas nuclei. *Rev. Mod. Phys.*, 69(2):629–642, 1997. ISSN 0034-6861. doi: 10.1103/RevModPhys.69.629.

R. Wyllie. *The Development of a Multichannel Atomic Magnetometer Array for Fetal Magnetocardiography*. Doctoral dissertation, University of Wisconsin-Madison, 2012.

R. Wyllie, M. Kauer, G. Smetana, R. Wakai, and T. Walker. Magnetocardiography with modular spin-exchange relaxation free atomic magnetometer array. *Phys Med Biol.*, 57(9):2619–2632, 2012a.

R. Wyllie, M. Kauer, R. T. Wakai, and T. G. Walker. Optical magnetometer array for fetal magnetocardiography. *Opt. Lett.*, 37(12):2247–9, 2012b. ISSN 1539-4794. doi: 10.1364/OL.37.002247.

S. Xu, S. M. Rochester, V. V. Yashchuk, M. H. Donaldson, and D. Budker. Construction and applications of an atomic magnetic gradiometer based on nonlinear magneto-optical rotation. *Rev. Sci. Instrum.*, 77(8), 2006. ISSN 00346748. doi: 10.1063/1.2336087.

S. Yu and R. T. Wakai. Maternal MCG Interference Cancellation Using Splined Independent Component Subtraction. *IEEE Trans. Biomed. Eng.*, 58(10):2835–2843, 2011. ISSN 1946-6242. doi: doi:10.1109/TBME.2011.2160635.

H. Zhao, J. F. Strasburger, B. F. Cuneo, and R. T. Wakai. Fetal cardiac repolarization abnormalities. *Am. J. Cardiol.*, 98(4):491–6, 2006. ISSN 0002-9149. doi: 10.1016/j.amjcard.2006.03.026.

**DISCARD THIS PAGE**

# Appendix A: Circuits

## A.1 Platinum RTD temperature sensors

We use platinum resistance temperature detector (RTD) elements from Heraeus Sensor Technology (L220, manufacturer part number 32207400) to obtain a rough idea of the vapor cell's temperature. Monitoring the temperature prevents damaging the vapor cell through overheating. We also use the temperature as a proxy for the rubidium vapor pressure. The sensors are attached to one of the four heater panels on the oven with EpoTek epoxy. The element's resistance at a given temperature $R(T)$ is governed by the Callendar-Van Dusen equation:

$$R = R_0\left(1 + AT + BT^2\right) \qquad T > 0°\text{C} \tag{A.1}$$

where $R_0$ is the temperature at 0°C, and $A$ and $B$ are coefficients particular to the particular class of RTD. For our RTDs, $R_0 = 100\ \Omega$, $A = 3.9083 \times 10^{-3}\ °\text{C}^{-1}$, and $B = -5.775 \times 10^{-7}\ °\text{C}^{-2}$.

More useful to us is $T(R)$, or the temperature as far as the read-out resistance. This function is found simply by solving the quadratic equation for $T$:

$$T = \frac{-A + \sqrt{A^2 - 4B\left(1 - \frac{R}{R_0}\right)}}{2B}. \tag{A.2}$$

Because $B$ is small, $T(R)$ is nearly linear in our temperature range of interest (100-200°C) and can be approximated as such by simply setting $B = 0$ in Equation (A.1). Doing so gives

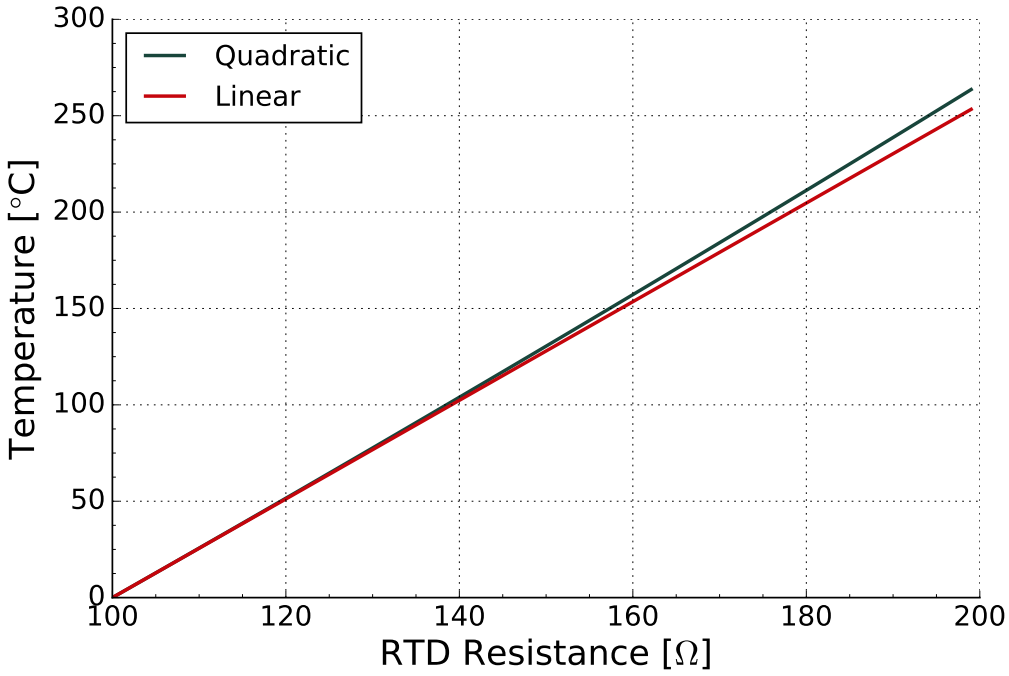
$$T \approx \frac{\frac{R}{R_0} - 1}{A}. \tag{A.3}$$

Figure A.1a shows both linear and exponential curves. The rubidium density (in $\text{cm}^{-3}$) is given by [**?**]:

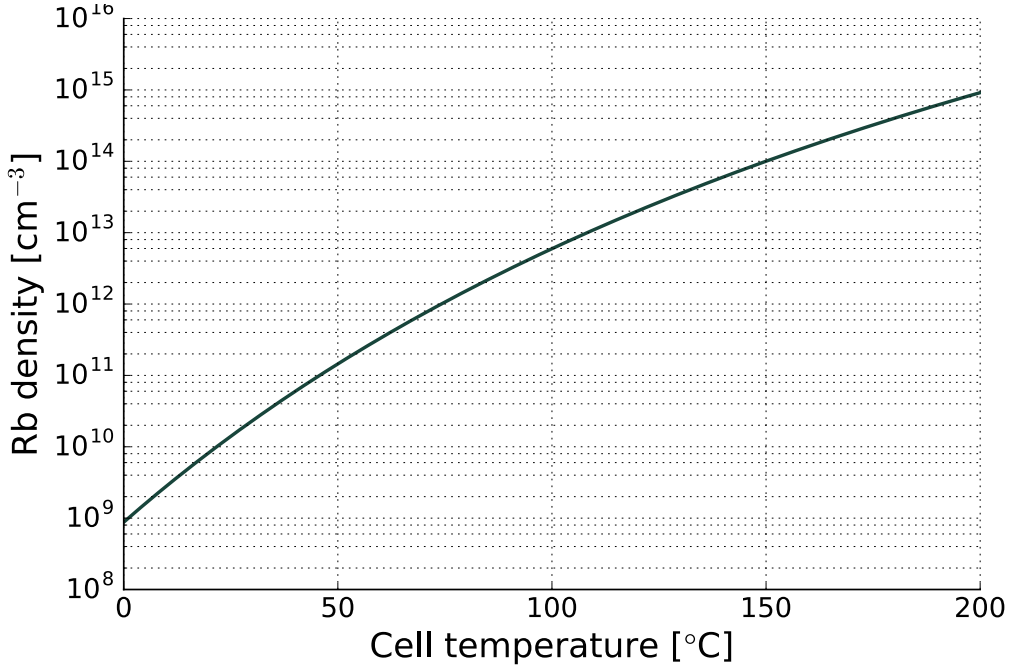
$$n_{\text{Rb}} = \frac{(7.336 \times 10^{21}\ \text{cm}^{-3}\ \text{K})}{T} 10^{4.312 - (4040\ \text{K})\ T^{-1}} \tag{A.4}$$

where $T$ is in kelvins.

Because of the relatively small nominal resistance of the RTD (compared to, say, a $R_0 = 10$ k$\Omega$ NTC thermistor), the extra resistance in the leads and Cat6A cables between the RTDs and




**(a)**




**(b)**

**Figure A.1:** (a) Temperature as a function of resistance for a Hereaus L220 RTD. (b) Saturated rubidium vapor pressure as a function of temperature.

the measurement location outside of the MSR can comprise a non-negligible additional resistance. Additionally, because the RTD is placed on one of the heater panels and not the cell itself, the measured temperature might be somewhat different than the glass temperature. Given these considerations, we do not use the measured temperature and Equation (A.4) as a precise measure of the density, only a ballpark estimate. A transmission curve fit (Section 4.2) provides a much better estimation of the density, though it is much more time-consuming to complete.

## A.2 Heater drivers

To heat the cells, we employ four ceramic pieces with printed-on counter-wound platinum traces (Section 3.1.2). Each panel has a resistance of $\sim$200 $\Omega$, so four panels connected in parallel gives the oven a resistance of approximately 50 $\Omega$. To reach the required $T \approx 145 - 165$°C temperatures required for SERF operation, we determined that roughly 5-8 W of heater power was required. Heating the cells at DC would be a relatively straightforward operation, but we found that despite our best efforts to create heater panels that self-cancel their created magnetic fields, applying DC voltages to the panels create magnetic fields and field *gradients* which degrade magnetometer performance. Thus, AC heating is necessary.

The frequency which we drive the heater panels should be much larger than the associated Larmor precession frequency of the atoms in the field the panels may create. Additionally, the frequency should be far outside of the measurement bandwidth of interest (1-100 Hz). A frequency of several hundred kHz was chosen to easily satisfy both of these requirements, and we usually run the heaters at roughly 150 kHz, though experimentation with running the heaters at even higher frequencies ($\sim$300-500 kHz) has shown slight improvements in magnetometer performance. Outputting the required voltage ($\sim$15-20 $V_{rms}$) requires a serious chip with the following requirements:

1. $\gtrsim 25$ V peak output voltage
2. $\gtrsim 200$ mA peak output current
3. $\gtrsim 100$ V/μs slew rate

In the end, we decided to take an "overkill" approach and used an op-amp that far surpassed these requirements: the Apex Electronics PA107DP. While somewhat costly, the reliability and ease of designing a circuit around this op-amp made it the winner in our search. The circuit is shown in Figure A.2.

The circuit runs off of $\pm 72$ V rails and also requires $\pm 15$ "auxiliary" voltages, which are produced by first stepping down the rail voltage and then using a 78L15/79L15 $\pm 15$V regulator pair. The circuit itself is a simple inverting amplifier circuit with a gain of 20. A 5 pF capacitor limits the bandwidth of the circuit to help stabilize it. A potentiometer (R1) forms a voltage divider with R2 which allows us to scale the input voltage appropriately. This allows us to use a single signal to drive all four circuits and scale the output voltage of each channel appropriately to reach and maintain the optimal temperature. The signal is AC-coupled ($C_{\mathrm{out}}$) before being sent to the oven.

In the future, we would like to employ a simple feedback circuit which monitors the temperatures of the channels and drives the to their optimum values. In this type of configuration, the input voltage could be controlled by the feedback circuit and the potentiometer could be left open.

## A.3 "Sulai" coil drivers

Bob Wyllie designed coil drivers for the previous iteration of the experiment [Wyllie, 2012]. The new circuits, designed primarily by Ibrahim Sulai, used those circuits as a foundation before making the following additions/modifications (see Figure A.3 for details):

1. To save space, the circuit blocks which attenuated the DC and AC input voltages by factors of10 and 100 were removed, as well as the BUF634 buffer circuits (which allowed the previous circuit's output current to reach 200 mA). All of the DIP components with surface-mount replacements, and what was once three 5"x5" boards were condensed to a single 6"x4" PCB.

2. Rather than using a pair of LM399 voltage references with an accompanying potentiometer to set the DC value of the applied voltage, we now use the voltage from a MAX5734 Digital

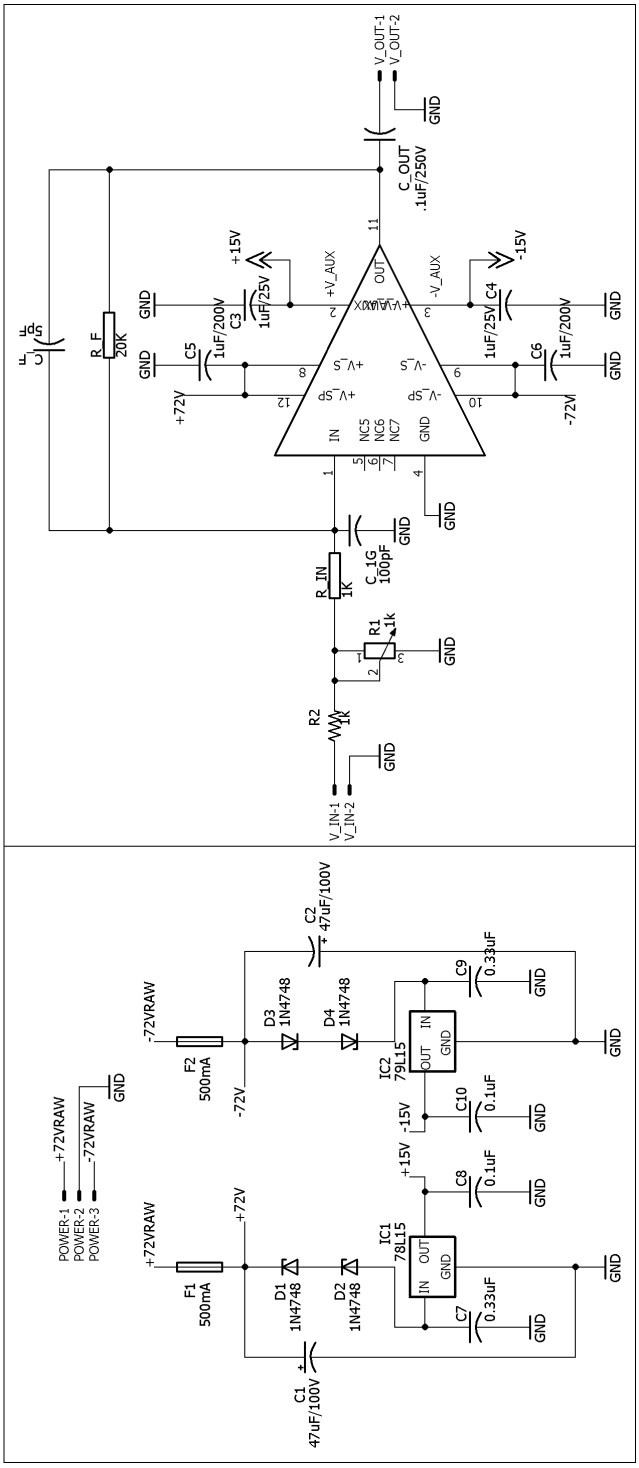



**Figure A.2:** PA107DP heater driving circuit with component values.

to Analog Converter chip to determine the input DC level. This was done to facilitate the decoupling scheme described in Section 3.5.4.

3. The push switches which determined which output resistor pair was used to drive the coils were replaced by two HI9P302-5 chips, each of which is a pair of DPST switches. The opening and closing of these switches is controlled with a front panel-mounted rotary switch which applied +5 V (the "high" logic level) to the desired DPST which in turn selected that pair of output resistors.

4. The same rotary switch also gives us a clever way of indicating which output resistor pair was selected. The +5 V is also applied to one of four resistors on the input of an inverting amplifier (bottom of Figure A.3). With the indicated resistor values (40k, 22k, 15k, and 10k), the voltage output at the GAIN MONITOR will be .375, .68, 1.0, or 1.5 V, respectively. This voltage can be read by a DMM to verify the output resistor selection. The original intent was to have this voltage read by the Arduino and adjust the coefficients in the decoupling matrix (Section 3.5.4) appropriately if two magnetic fields that "couple" to each other (Channel 1Y and Channel 2Y, e.g.) were being generated via different output resistors. (Recall that the Arduino+DAC combination is only able to output *voltages*, and the decoupling scheme only functions correctly if these voltages are first converted to currents and later to magnetic fields in a well-defined way. The Arduino code assumes that the output resistors are equal.) This idea never quite came to fruition, and was found to be largely unnecessary as the room coils allowed us to shim most of the large DC magnetic fields away and use the same large output resistors (5k for $\hat{x}$ and $\hat{y}$ directions, 1k for $\hat{z}$ direction) on every channel, all of the time.

The "Wyllie" supplies are still very much in use, notably driving the room and common-feedback coils and often connected to the shell coils on a separate test-bed setup we have at Chamberlin Hall. However, we have found the noise performance of the two to be comparable and primarily use the "Sulai" supplies for the shell coils on our magnetometer array.

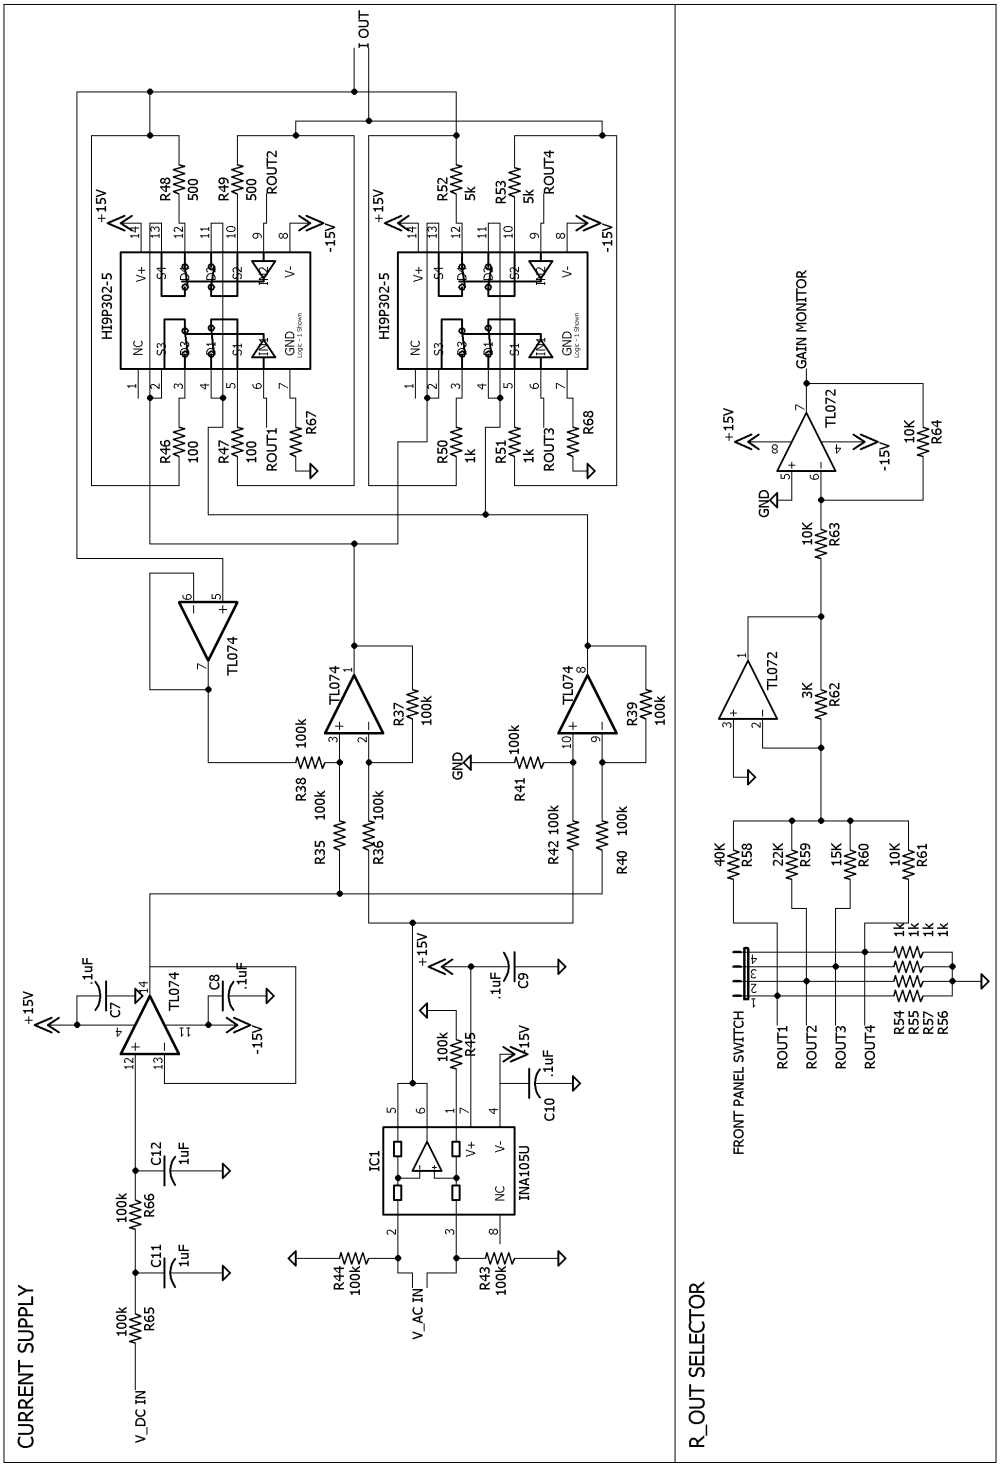



**Figure A.3:** One channel of the "Sulai" coil drivers with component values. The wires labeled ROUTx are connected on the printed circuit board, though they have been suppressed here for cleanliness.

## A.4 Feedback circuits

For this work, we employed two types of feedback circuits: one which controlled feedback to the shell-mounted coils and another which controlled feedback to the array-mounted "common" coils).

### A.4.1 Shell coil feedback circuits

Active feedback can be applied to the shell coils by simply plugging the feedback circuit into the AC voltage input of the appropriate coil driver (in our experiment, this has exclusively been $\hat{y}$, though in theory this could be extended to both $\hat{x}$ and $\hat{z}$ as well). The goal of this feedback is to uniformly hold the magnetometer signal at zero - which significantly reduces the unsubtracted intensity noise that we would otherwise see if the detector were not appropriately balanced (Section 4.8.2.1).

The feedback circuit is a basic Proportional+Integral scheme depicted in Figure A.4, which has been arranged in a "block diagram" format which allows for easier visualization of what each component of the circuit does. The magnetometer signal (ER_IN) is taken from the Stanford Research Systems preamplifiers and can be low-passed if desired (often, though, we leave C1 open as any low-pass filters add phase shifts which can degrade feedback performance). After being buffered, the signal is simultaneously sent to a proportional stage with adjustable gain and an integral stage which includes a compensation resistor (R6) in parallel with the capacitor to reduce the gain at DC. The proportional gain rolls off at $f = 1/(2\pi R_2 C_2) = 318$ Hz and can be enabled or disabled by flipping a front panel-mounted switch. The proportional and integral signals are summed with a simple voltage adder.

If the feedback is enabled, the signal is sent through another summing circuit which adds the calibration chirp signal generated by the FPGA card. After passing through a stage which allows the signal to be optionally inverted to ensure the phase of the feedback signal is *opposite* to that of the magnetometer signal which will drive that magnetometer signal to zero. Before being sent to the coil driver, the signal can be optionally low-passed again.

Because of the wide bandwidth of the feedback, the magnetometer signal should be held at zero, so recording this signal to analyze provides almost no information about the actual magnetic fields in the room. Instead, we record a scaled copy of the feedback signal, which, after all, should just be the negation of the signal produced by the magnetometer. The same P+I signal that is eventually sent back to the coil drivers is high-passed at 0.5 Hz (recall the biomagnetic signals contain frequency components primarily between 1 and 100 Hz), buffered, and then multiplied by an appropriate factor such that the signal approximately fills the voltage range of the FPGA's analog-to-digital converter (-10 V to +10 V). Note that this does not improve the SNR of the measurement, though it can help matters if the measurement is digitization-limited. (For the 16-bit ADC on the FPGA, the digitization limit is $[10\text{ V} - (-10\text{ V})]/2^{16} = .305$ mV.) After scaling, the signal is sent to the ADC.

Clever design of the circuit means that the magnetometers can also be run without active feedback without serious plugging and unplugging of signal cables. Flipping the DPDT "enable" switch (center of Figure A.4) to "off" sends the original magnetometer signal straight through to the ADC and grounds the feedback signal input to the coil driver signal. The chirp signal is still sent to the coil driver regardless of whether the feedback is engaged or not.

Admittedly, this feedback design is not perfect, and often we cannot apply much proportional gain before the magnetometer begins to oscillate. A more effective approach to feedback circuit design is "Loop Shaping" and is described in Appendix C. Feedback based on these loop-shaped circuits will be employed in future iterations of the apparatus.

### A.4.2 Large coil feedback circuits

We also use feedback to the feedback coil wrapped around the perimeter of the array (Section 3.3.5) to counter the large, slow drifts that are often observed during a measurement. Because we believe that these drifts are caused by far-away sources such as cars and elevators, we assume that their magnitudes are approximately uniform across the span of the four magnetometers. Thus these fields can be canceled with a set of common coils, similar to the feedback scheme that was used in [Wyllie, 2012]

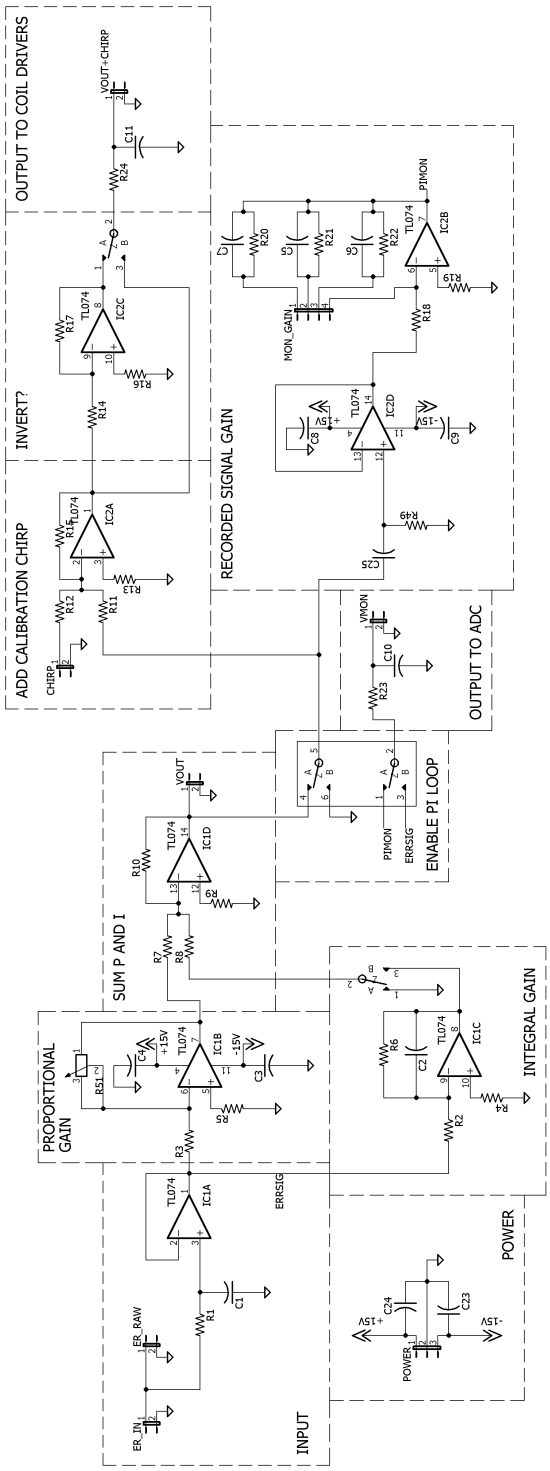



**Figure A.4:** PI circuit controlling feedback to the shell-mounted coils. Feedback can be enabled by flipping a front panel-mounted switch which simultaneously sends the P+I signal to the coil driver and a scaled version of the signal to the FPGA ADC.

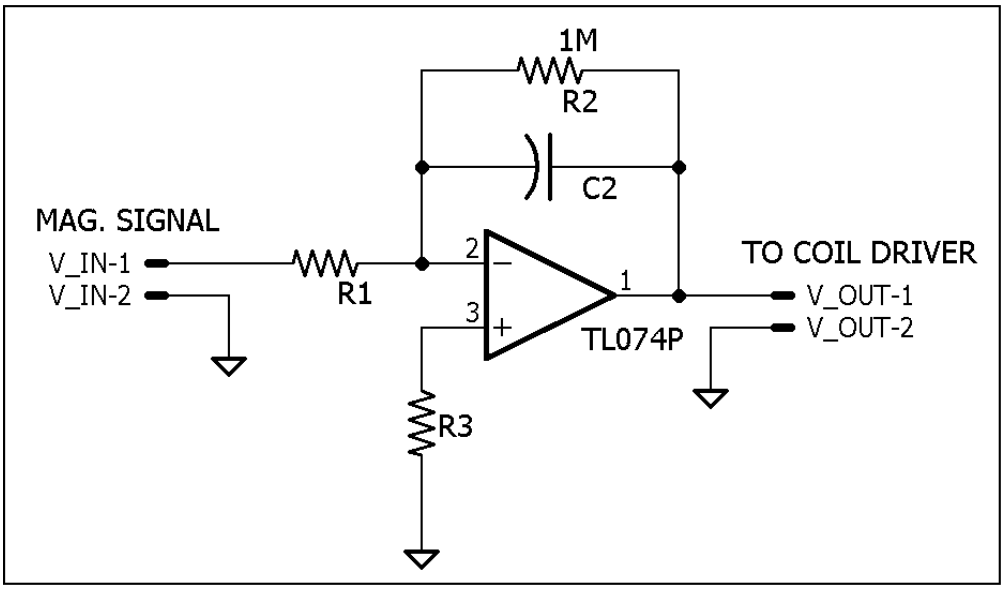



**Figure A.5:** Feedback circuit for common coil.

The feedback circuit is a simple one-stage integrator circuit bandwidth limited at 0.5 Hz; the large drifts generally have frequency components of 0.1 Hz or less. Much like the integrator circuit in the previous section, a large (1 MΩ) resistor is placed in parallel with the capacitor to limit the DC gain. We use the raw magnetometer signal from one of the magnetometers (usually channel 1) and send the feedback signal to a "Wyllie" coil driver connected to the feedback coil. Because the feedback bandwidth does not overlap with our 1-100 Hz band of interest, the raw magnetometer signal can still be recorded by the ADC: the biomagnetic signals will be unaffected by this feedback stage.

# Appendix B: Magnetometer Cell Fabrication

## B.1 Motivation

Production of rubidium vapor cells to use in our magnetometer array traditionally has been handled by third-party sources, namely Triad Technologies and Opthos Instruments, who produced all of the cells that were used by Wyllie [2012]. These cells performed adequately in the magnetometer array, but suffered from a few practical drawbacks. Namely, if a cell were to break or otherwise fail, we would be beholden to one of these companies to produce a new cell for us, possibly putting measurement on hold for weeks. The specialty nature and low quantity of purchased cells also means a high cost of acquisition.

Additionally, hybrid cells containing a mixture of rubidium and potassium atoms have shown promise for use in atomic magnetometers [Ito et al., 2011, Romalis, 2010]. The cell production method employed by Triad and Opthose (as well as other companies) involves heating a separate chamber containing an alkali metal, and then using heat to "chase" the vapor to the magnetometer cells, where it condenses after cooling. This procedure, while very effective in the transfer of alkali atoms, is not precise enough to allow the cell-makers to move exact amounts of alkali to each cell. It is therefore difficult to set the exact alkali density ratio for each magnetometer cell if hybrid cells are being made, an important parameter for cell performance.

With these limitations in mind, we endeavored to create a system whereby we could fabricate our own magnetometer cells, and circumvent the two pitfalls mentioned above. Our work on this front was inspired by the work of Kitching's group [Knappe et al., 2005, Liew et al., 2004], who used similar processes to create vapor cells for their microfabricated magnetometers. The alkali metal is produced by reacting a chlorinated salt of the alkali with barium azide:

$$2\,\mathrm{RbCl(s)} + \mathrm{Ba(N_3)_2(s)} \longrightarrow 2\,\mathrm{Rb\,(s)} + \mathrm{BaCl_2(s)} + 3\,\mathrm{N_2(g)}. \tag{B.1}$$

The reaction in Equation (B.1) is actually a two-step process with the two reactions taking place at different elevated temperatures. The details are explained later in this chapter.

With a four-channel magnetometer array, we desired the ability to produce batches of four cells. Variations in cells inevitably require different pump and probe laser characteristics to optimize the performance of each magnetometer channel. While the laser spatial profiles and powers can be set independently to some extent, the frequencies cannot be, as all channels share the same pump and probe laser diodes. Creating cells in batches of four would decrease the cell-to-cell variations, including in rubidium isotope (natural abundance 85/87 or isotopically pure 87), buffer gas pressure, and optical glass type and quality, and would allow us to optimize each cell at the same point in pump and probe detuning space.

## B.2 Fabrication Procedure

### B.2.1 Manifold Construction and Preliminary Chemistry

The Pyrex cells for our fabrication process roughly matched the 1cm x 1cm x ∼5cm dimensions of those used by Wyllie [2012]. The cells were purchased from Starna Cells, Inc. (part number 3-PX-10). Starna offered several types of less-expensive optical-quality glass, but these alternate materials were deemed unable to stand up to the high temperatures associated with the glassblowing required during fabrication. The cells were sent to Steve Betts of Beste Scientific Glass in Grafton, WI, who attached the cells to a long glass tube with a glass-to-metal seal at one end. The seal allowed for attachment of the manifold to a vacuum system. The completed manifold is shown in Figure B.2.

Rubidium metal immediately oxidizes upon contact with air and thus is usually packaged in its chlorinated form RbCl. The barium likewise is packaged as a compound; however an important restriction of this compound is the complementary piece in the compound must be removed prior to the reaction with rubidium chloride, else no pure rubidium metal will be produced. Barium azide, $Ba(N_3)_2$, was deemed suitable for this purpose, as the nitrogen gas produced during the barium azide decomposition can be quickly removed by a vacuum pump.

Aqueous solutions of rubidium chloride and barium azide were prepared. With the end goal of of ∼5-10 mg Rb/cell, stoichiometric calculations were performed to determine the total amount of each solution that should be added to the combined aqueous mixture. The calculations and

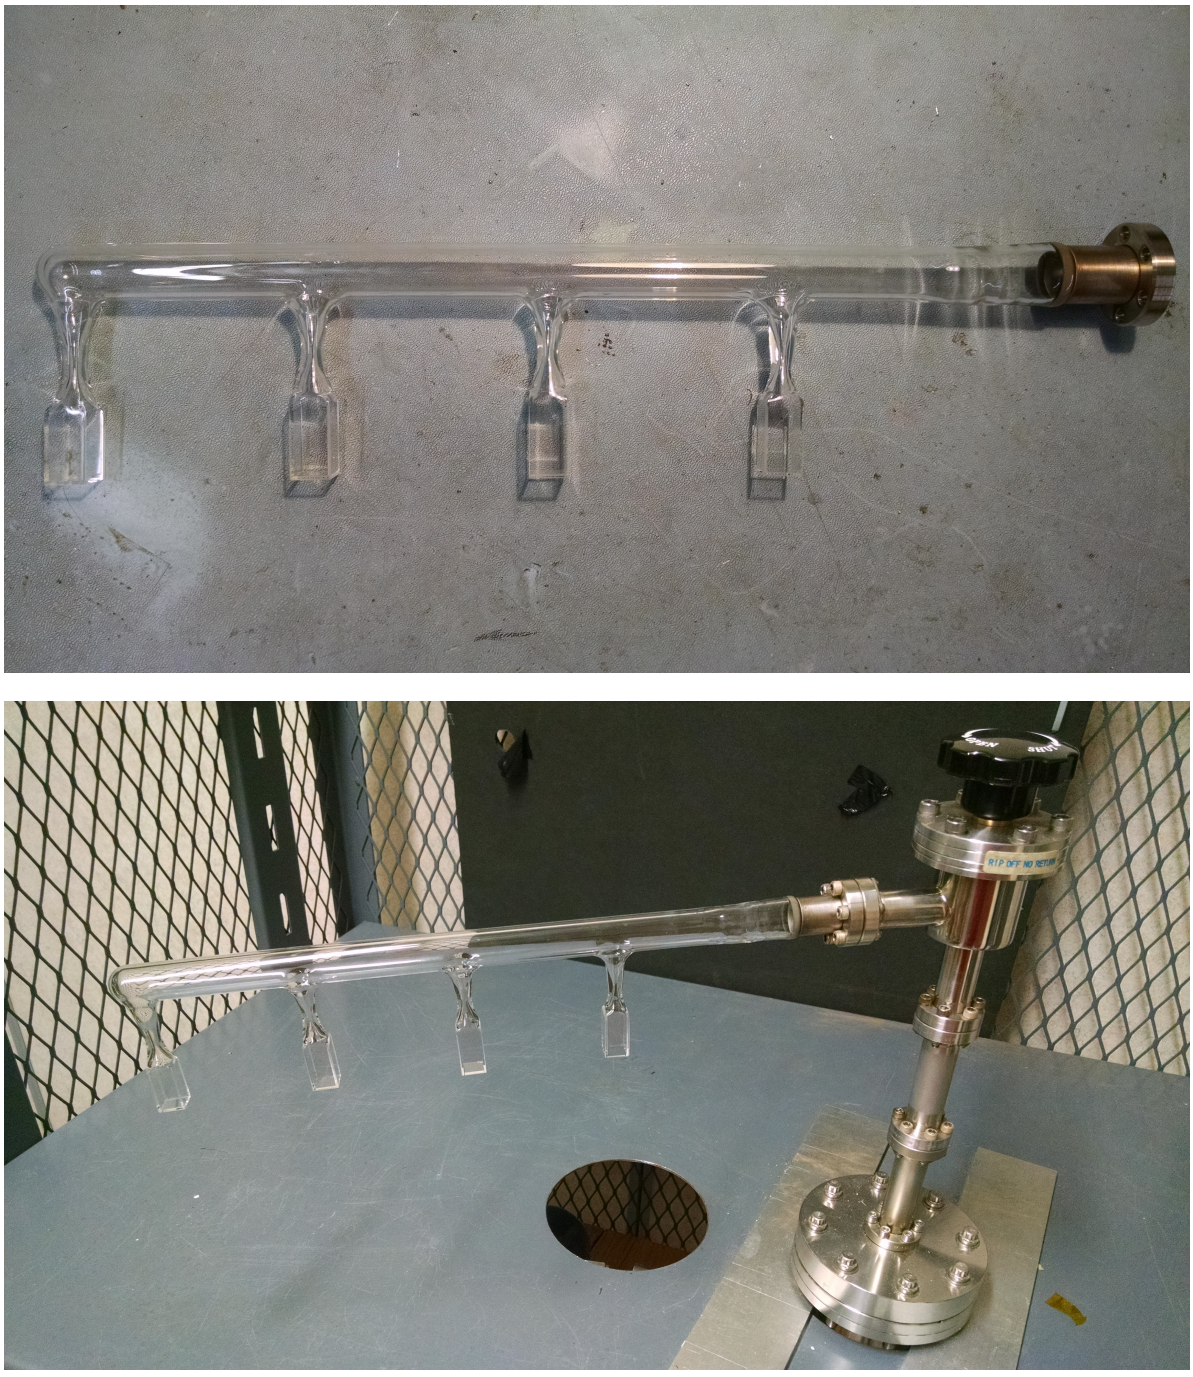


**Figure B.1:** (a) The manifold containing the four future magnetometer cells. The glass-to-metal seal (right) allows for attachment to a vacuum system. (b) the manifold attached to the vacuum system (see Section B.2.2).

results for the reaction performed 25 May 2015 are shown in Table B.1. The combined solution was agitated to ensure proper mixing of the reactants, and then the proper amount of solution (generally 200-300 μL) was delivered via auto-pipette to each cell. Because of the need for high-quality optical access when the cells are complete, care must be taken to avoid getting the solution on the walls of the cells.

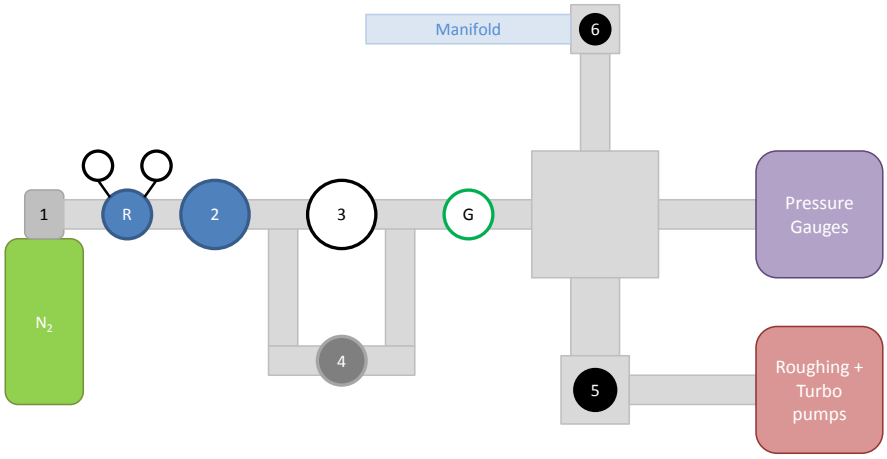
The reaction in Equation (B.1) must occur at high vacuum to prevent oxidation of the rubidium metal. However, before being attached to the vacuum system, it is first necessary to remove the water solvent that was used to make the solution. To do so, we place the manifold in a toaster oven set at a fairly low temperature ($T \approx 60 - 70$°C) to allow the water to *evaporate* (not boil) over the course of a couple of days. At this point, a well-mixed salt of rubidium chloride and barium azide is all that remains in the cell.

### B.2.2 Vacuum System

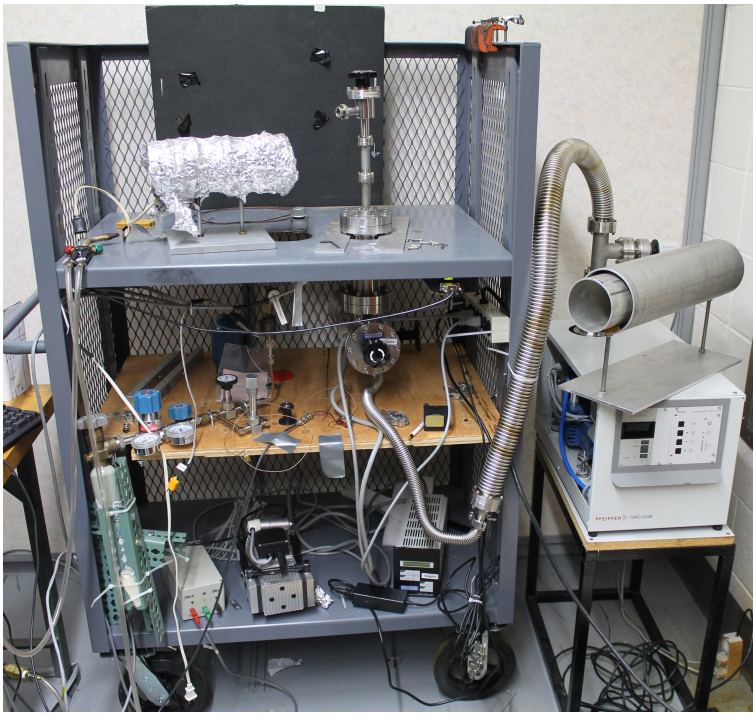
To reach the high-vacuum ($P \approx 10^{-7}$ torr) levels needed for the reaction, a custom vacuum system was designed and constructed, primarily by undergraduate student Greg Smetana. The system features dual vacuum pumps: a roughing pump to decrease the pressure down to the medium-vacuum level ($P \approx 10^{-4} - 10^{-5}$ torr) and a turbopump that can take the system down from medium-vacuum to the high-vacuum level. The system also includes a second "arm" that allows us to back-fill the cells with a desired amount of buffer gas. This arm also contains a getter used to remove oxygen and other impurities present in the vacuum. Valves and pressure gauges at various points in the system allow us to isolate the various parts of the vacuum system from each other and monitor the pressure. A complete schematic of the system is shown in Figure B.2a and a photo in Figure B.2b.

| Quantity | Value | Unit |
|---|---|---|
| Rb Isotope | 87 | |
| Rb molecular mass | 87.0000 | g/mol |
| RbCl molecular mass | 122.453 | g/mol |
| RbCl solubility @ 20C | 91 | mg/100 μL |
| RbCl mass | 69.7 | mg |
| RbCl amount | $5.69 \times 10^{-4}$ | mol |
| Rb amount | $5.69 \times 10^{-4}$ | mol |
| Rb mass | 49.5 | mg |
| $H_2O$ required | 77 | μL |
| $H_2O$ used | 100 | μL |
| Rb concentration | .495 | mg/μL |
| Ba required | $2.85 \times 10^{-4}$ | mol |
| Ba molecular mass | 137.327 | g/mol |
| $Ba(N_3)_2$ stock solution concentration | 53.3 | mg Ba/mL |
| $Ba(N_3)_2$ stock solution required | 733 | μL |
| $Ba(N_3)_2$ stock solution used | 773 | μL |
| Ba amount | $3.00 \times 10^{-4}$ | mol |
| Rb:Ba molar ratio | 2:1.05 | |
| $H_2O$ added to sol'n | 1000 | μL |
| Total volume of sol'n | 1873 | μL |
| Rb concentration | .0264 | mg/μL |
| Sol'n dispensed per cell | 200 | μL |
| **Rb per cell** | 5.28 | mg |

**Table B.1:** Stoichiometry for the chemical reactions in Equations (B.2a) and (B.2b) for a sample cell-making run. The $H_2O$ required indicates the *minimum* amount of water required to dissolve the given amount of rubdium chloride. Additional water was added to ensure the rubidium. Similarly, excess $Ba(N_3)_2$ solution was added to ensure that all of the rubidium reacted.




**(a)** Vacuum system schematic



**(b)** Vacuum system photograph

**Figure B.2:** (a) Cartoon schematic of the filling system. The valves are labeled 1-6 and G. When the manifold is being pumped down to vacuum level pressures and when the rubidium production reaction is occurring (Section B.2.3), valves 1, 2, 3, 4, and G are closed. During the nitrogen back-fill (Section B.2.4, valve 5 is closed to isolate the manifold from the vacuum pump. The other valves and the regulator (R) are used to precisely control the pressure of nitrogen (from the tank, far left) in the manifold. At all times, the pressure in the system is monitored by two guages, a hig-pressure gauge that is accurate from atmospheric pressure down to .1 torr, and low pressure that is accurate at pressures from .1 torr down to $1 \times 10^{-8}$ torr. (b) A photo of the (portable!) vacuum system. The manifold has been removed from the system. The oven used raise the manifold to the temperatures required for the rubidium production reaction is seen wrapped in aluminum foil on the top shelf of the cart. The combination roughing/turbo pump is at the far right.

### B.2.3 Cell Chemistry

As previously mentioned, Equation (B.1) is actually two separate processes whose reaction rates are small until elevated temperatures of $< 180°\text{C}$ [Liew et al., 2004] are reached:

$$\mathrm{Ba(N_3)_2(s)} \rightarrow \mathrm{Ba\,(s)} + 3\,\mathrm{N_2(g)} \qquad T \approx 120°\text{C} \qquad \text{(B.2a)}$$

$$2\,\mathrm{RbCl\,(s)} + \mathrm{Ba\,(s)} \rightarrow 2\,\mathrm{Rb\,(s)} + \mathrm{BaCl_2(s)} \qquad T \approx 200°\text{C} \qquad \text{(B.2b)}$$

To reach the required temperatures, an oven was constructed from a long brass tube wrapped with resistive tape. The tube is slipped over the manifold so that all cells are enclosed. A variable AC transformer is used to apply a large voltage to the tape which conductively heated the brass tube. The cells, still under vacuum, were heated inside the oven while temperature sensors placed on two of the cells' walls and in free space inside of the oven monitored the heating process. These temperatures, along with the pressure inside of the system, were recorded by a LabVIEW DAQ program.

As the temperature inside the oven slowly increases, small amounts of outgassing from the glass manifold, as well as the natural increase of pressure with elevated temperature, cause the total pressure inside the system to increase, even with the system under constant vacuum. However, this pressure increase is normally only a few μTorr. The signature of the azide decomposition, however, is a rapid increase of several hundred μTorr, which occurs at about 800 s in Figure B.3. Because of temperature gradients inside the oven, and to a lesser extent, because of temperature gradients across one individual cell, not all of the barium azide will decompose at the same time. This leads to a series of pressure "spikes" as the different pockets of barium azide decompose and release nitrogen gas, and there is a subsequent decrease in pressure as the released nitrogen is removed by the vacuum system. During the period of time in which the azide decompositions are occurring, the heat is kept on to ensure that even the cells near the coolest parts of the oven still reach the temperature required for the azide reaction to proceed.

Eventually, even with the cell temperatures still rising, there is a noticeable decrease in the system pressure due to a lack of nitrogen being released. At this point ($\sim$1200 s in Figure B.3), it

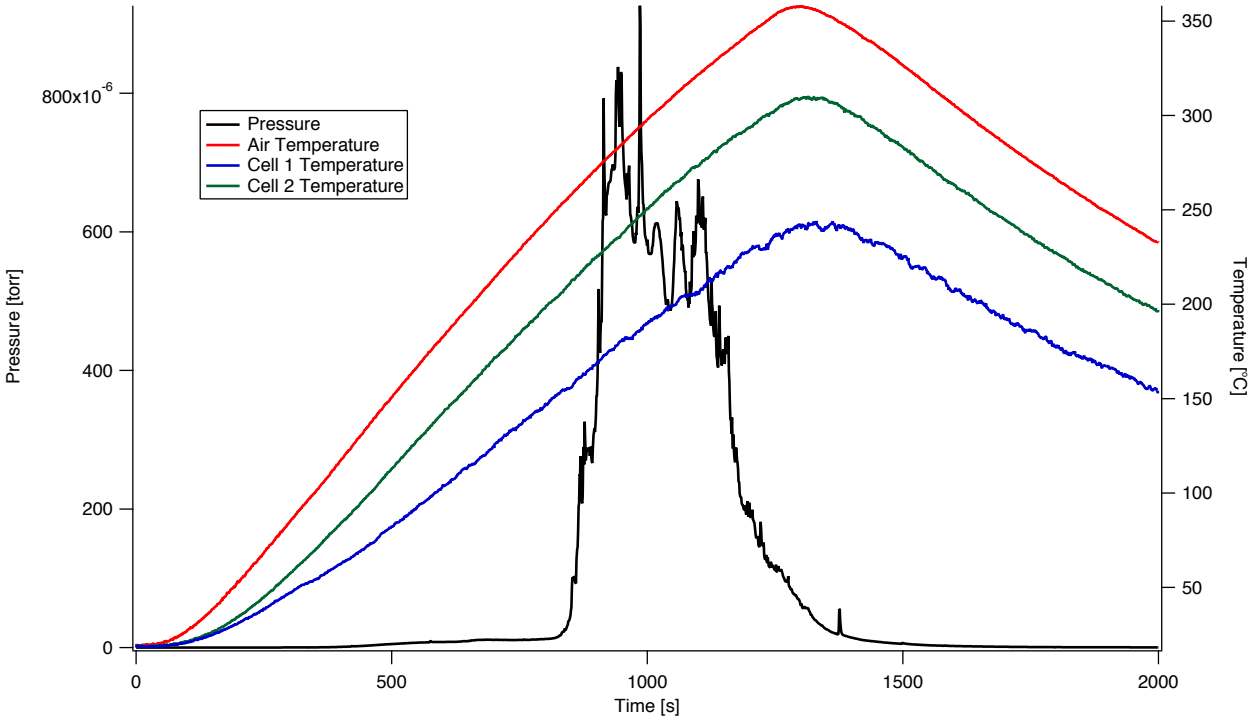



**Figure B.3:** Azide decomosition reaction. The temperature was slowly increased until we began to see evidence of the $N_2$ gas production (Equation (B.2a)). The many pressure spikes seen in this figure were a result of each cell (and even different parts of a single cell) being at different temperatures and thus undergoing the azide decomposition at different times. Note the difference in temperatures between Cell 1 (which as placed near the oven) and Cell 2 (which was closer to the middle). For this particular run, we used procedure 2 from Section B.2.3, turning off the heat to the oven as soon as we were confident that all of the azide had decomposed. The barium/rubidium reaction (Equation (B.2b)) took place at a later time.

is theorized that all available azide has been decomposed and the resulting nitrogen evacuated. At this point in the fabrication process, one of two experimental procedures can be followed:

1. The heat is left on to allow the temperatures inside the oven to reach the required level for the rubidium production reaction indicated in Equation (B.2b). After this reaction, only a rubidium film along with a barium chloride salt will remain in the cell.

2. The heat is turned off and the cells are allowed to cool. The cells will therefore contain rubidium chloride salt intermixed with barium metal.

Varying success with each of these methods have been realized; a full discussion will be presented in Section B.3.

### B.2.4 Buffer Gas Filling and Final Steps

Regardless of the procedure followed after the azide reaction, the next step in the fabrication procedure is to fill the cells with a background buffer gas to limit the diffusion of optically-pumped rubidium to the spin-destroying cell walls during magnetometry. The optimal pressure for filling the cells depends on several factors, most prominently the type of buffer gas being used and whether the magnetometers will be running in conventional SERF operation mode or in diff-SERF mode. In conventional SERF mode, where the pump and the probe beams are of comparable waists, a high background gas pressure ($\sim$1 atm) is desirable to confine the atoms to the region being pumped and probed. In diff-SERF mode, however, the pressure must be low enough to allow the polarized atoms to diffuse out of the narrowly-pumped center of the cell to the larger region being probed. Too low a gas pressure, however, will increase the spin-destruction rate to a level unsustainable for magnetometry. A complete calculation of optimum buffer gas pressures for each of these modes is presented in Section 2.4.3.

As we primarily worked with diff-SERF magnetometers, the cells were generally filled with 150-300 torr purified $N_2$ buffer gas. The backfilling is accomplished by first closing the valve between the vacuum pumps and the cell manifold, then opening the leak valve to the other arm of the system which allows a slow addition of the buffer gas. Pressure is constantly monitored inside

the system, and the manifold is sealed when the pressure reaches the desired value. At this time, a glass-blower is enlisted to "pull off" the cells from the manifold and seal them without exposing the contents to the oxygen, which would render them unusable.

If procedure 2 was followed during the cell chemistry process and the rubidium/barium reaction has not yet been pushed forward, that can be done now. An "unreacted" cell is placed on a slab of high-temperature calcium silicate and both are inserted back into the same brass oven that was used to heat the manifold. This time, however, we send a laser beam tuned near the rubidium D2 line axially down the oven, and the transmission of this beam (an indicator of rubidium vapor inside of the cell) is monitored by a photodiode at the opposite end. The laser is continuously swept across the absorption resonance as the cell is heated. The temperature of the air inside of the oven is monitored with a thermocouple. Generally, hints of rubidium absorption begin to be seen at temperatures around 180 °C, with full on-resonance absorption (indicating a significant rubidium density) occurring closer to 250 °C. To ensure that Equation (B.2b) is completed, the oven temperature is held at 275-300 °Cfor several minutes, until the transmission measurements indicate complete absorption over a $\sim$ 200 GHz range around the absorption line. At this point, the heat is turned off, and the oven and cell are allowed to return to room temperature. The cells can then be tested for rubidium by heating to a more modest temperature (150-170 °C, the operating regime of the SERF magnetometers) and again scanning the D2 laser across the resonance. If absorption curves similar to the ones in Figure 4.1 are observed, we know we have been successful.

## B.3 Results

Despite our best efforts and constantly refining our fabrication process, we have never been able to *consistently* produce high-quality magnetometer cells. Our goal was always to produce "batches" of four identical cells, yet after many trials, we were unsuccessful in doing so. Although several obstacles were overcome during our trial-and-error cell-making runs (see Table B.2 for details), frequently, the end result was cells that either showed no evidence of rubidium or cells that showed this evidence only transiently.

Our best hypothesis for the lack of rubidium is the fact that at elevated temperatures, the rubidium metal reacts with the glass walls of the cell. Therefore, although we had to raise the temperature of the cells to $> 200°$C to produce the rubidium through the reaction with barium, we had to be careful not to push the temperature *too* high for *too* long, or else we would risk losing the precious rubidium we had just produced.

The weakness of the solid brass heating tube is our inability to constantly probe the cells with the resonant laser to check for rubidium absorption during the initial heating process. The only indication we had that we were close to reaching an appropriate temperature was the large spike in pressure when the azide decomposition occurred. This was the primary impetus for moving to the alternate procedure mentioned in Section B.2.4 in which the rubidium/barium reaction was run at a later time. This setup allowed for a resonant laser to track rubidium absorption as the cells were heated a second time. Switching to this method allowed us to quickly cool the cells as soon as we had convincing evidence of rubidium in the cells, which limited the amount of rubidium that was lost to reactions with the cell walls.

In the cases where significant rubidium absorption was observed, we often observed a dark blue film on the interior of the cells after removing them from the oven. We suspect this in indicator of rubidium condensation. When these "blue cells" were used in magnetometers, they generally performed very well for several hours operating at modest temperatures (150-170°C), but performance degraded afterward, which coincided with a marked decrease in absorption. When these cells were later examined, the blue film had disappeared only to be replaced by a yellowish-brown hue. We hypothesize this is an indicator of a reaction between the rubidium and the glass, leaving us little unreacted rubidium in the cell.

Despite our troubles with reliability, some of these "homemade" cells gave us beautiful magnetometer results, in some cases surpassing the performance of the cells purchased from Triad and Opthos that were used by Wyllie [2012]. The promise of these results, coupled to the facts that our homemade cells have a turnaround time of less than a week and (when considering only the price of materials and labor) are considerably cheaper, have us considering new methods to improve the reliability of our procedure.

## B.4 Future

We have conversed with professional cell-makers and believe that a modification to our procedure could more reliably give us the type of high-quality cells we desire. Most cell-making companies implement a procedure whereby the chemical reaction in Equation (B.2) takes place in a heated volume separate from the cells to be used as magnetometers. The magnetometer cells are *not* heated (and in fact may be actively cooled) so that the rubidium produced will naturally condense in these cells, rather than where the reaction is happening. Rubidium that condenses elsewhere in the manifold can be actively "chased" to the magnetometer cells using a localized hot air gun. The advantages of such a strategy are two-fold:

1. With the magnetometer cells outside of the oven, evidence of rubidium condensation can be directly observed by eye throughout the process. Alternatively, moderate heat can be applied to each of the cells and rubidium can be observed via an absorption measurement.

2. The barium chlorideis never introduced to the magnetometer cells. At best, the barium chloride coats the bottom of the cell with residue and at worst can adhere to the sides of the cells obscuring optical access. By never allowing this material near the magnetometer cells, the optical quality of the windows is improved dramatically.

The above strategy can be implemented with our current setup with only a couple of modifications. First, a $\sim$1.5 cm slot can be cut in the bottom of the brass oven, wide enough to alllow the cells to protrude beneath the oven. When the manifold is constructed, the stems of three of the four cells should be extended such that when the oven is placed over the manifold, only one of the cells (the "reaction cell" sits inside the oven with the other 3 exposed below. The RbCl and $Ba(N_3)_2$ solutions can be mixed just as before, except only be dispensed into the reaction cell. An excess of solutions should be prepared in this cell such that any small amount of rubidium lost to wall reactions or condensation in other parts of the manifold will not leave the magnetometers rubidium-deficient. The buffer gas backfilling, as well as the pull-off, would be identical to the processes mentioned in Section B.2.4.

One potential disadvantage to such a scheme is the inability to set precise alkali ratios if we ever desired to make potassium-rubidium hybrid magnetometer cells. However, we believe that before those types of cells are even attempted, a reliable method for making rubidium cells must be perfected.

| Issue | Hypothesis | Implemented Solutions |
|---|---|---|
| Not all cells reached correct temperature for azide decomposition to occur when forced-air oven was used | Large cell-to-cell temperature gradients: Cells close to air inlet much hotter than far away cells | Replaced forced-air oven with brass tube wrapped with resistive tape. Configured in this way, the temperature of the tube is relatively uniform, and the cells are heated much more evenly |
| No rubidium detected inside cells at any point during fabrication process | Cell environment was contaminated and rubidium was immediately oxidizing | Discarded rubidium chloride solution; began a regimented procedure when making solution to avoid contamination during solution-making; installed a getter in vacuum system to aid in impurity removal; added more $RbCl/Ba(N_3)_2$ solution to each cell to ensure enough rubidium remained after the reaction |
| Rubidium observed during heating process, but no evidence after pull-off | Temperature was pushed too high during heating, causing rubidium to react with walls OR rubidium was escaping through neck into manifold after production | Cells were only heated to temperature required to run the azide decomposition reaction; rubidium/barium reaction was run after back-fill and pull-off in an environment where the presence of rubidium could be carefully monitored (procedure 2 from Section B.2.3) |
| Barium chloride residue left over from reaction obstructing optical access through cell windows | $RbCl/Ba(N_3)_2$ splashed on walls | Modified solution dispenser to carefully place the liquid only on the bottom of the cell; reduced temperature of evaporation oven to prevent any boiling |

**Table B.2:** Issues and resulting modifications to cell-making procedure

# Appendix C: Loop Shaping for Feedback

In Section 2.6, we discussed using active feedback with the magnetometer units. The circuitry which controls the feedback range simple PID loops similar to those found in myriad devices like temperature controllers, to much more sophisticated systems. Recall from Equation (2.88) the closed-loop behavior of a magnetometer in feedback mode:

$$y(s) = \frac{G(s)K(s)}{1 + G(s)K(s)} r(s) \equiv T(s)r(s). \tag{C.1}$$

where $G(s)$ is the "open loop" frequency response of the magnetometer without feedback, $K(s)$ is the frequency response of the feedback servo controller, and $r(s)$ is the control signal which we desire the magnetometer to follow. The degree to which we can accomplish this is determined through the careful design of the servo response $K(s)$. This design process, which is known as "loop shaping" and is described in detail below, follows three main steps:

1. Measure the transfer function $G(s)$ of the system. This is also known as the "loop kernel" of the system. The loop kernel takes into account all components (including electronics) that add gain or phase shifts.

2. Fit the transfer function using a product of basis functions, or "building blocks".

3. Choose the *complementary* building blocks to construct the servo $K(s)$.

We will describe these steps and how they can be used to create an effective feedback servo $K(s)$.

## C.1 Measuring the Loop Kernel

The loop kernel is determined by applying a frequency-varying calibration signal to the system and measuring the system's response $G(s)$. This can be accomplished either by applying a set of discrete frequency signals to the system and measuring the relative amplitudes and phases of the output, or by applying a frequency "chirp" and transforming it into the frequency and phase domain using either a Fourier or Laplace transformation function.

## C.2 Loop Building Blocks

Once the transfer function $G(s)$ (amplitude and phase) has been measured, the next step is to approximate the transfer function using a series of "building blocks". The four building blocks are the pole, the zero, the inverted pole, and the inverted zero. In the Laplace domain, these building blocks have the forms shown in Table C.1. The amplitude response (gain) of a given block is simply the modulus of the Laplace-domain function $\tilde{f}$

$$\sqrt{\tilde{f}\tilde{f}^*},$$

while the phase response is the argument of $\tilde{f}$

$$\arctan\left(\frac{\text{Im}\,\tilde{f}}{\text{Re}\,\tilde{f}}\right).$$

The amplitude and phase responses of each of the blocks are shown in Figure C.1. The goal is

| Building Block | Laplace Representation | Gain | Phase |
|---|---|---|---|
| Pole | $p\left(\frac{1}{1+\frac{s}{\omega_0}}\right)$ | $p\sqrt{\frac{1}{1+\left(\frac{\omega}{\omega_0}\right)^2}}$ | $\arctan\left(-\frac{\omega}{\omega_0}\right)$ |
| Zero | $p\left(1+\frac{s}{\omega_0}\right)$ | $p\sqrt{1+\left(\frac{\omega}{\omega_0}\right)^2}$ | $\arctan\left(\frac{\omega}{\omega_0}\right)$ |
| Inverted Pole | $p\left(\frac{1}{1+\frac{\omega_0}{s}}\right)$ | $p\sqrt{\frac{1}{1+\left(\frac{\omega_0}{\omega}\right)^2}}$ | $\arctan\left(\frac{\omega_0}{\omega}\right)$ |
| Inverted Zero | $p\left(1+\frac{\omega_0}{s}\right)$ | $p\sqrt{1+\left(\frac{\omega_0}{\omega}\right)^2}$ | $\arctan\left(-\frac{\omega_0}{\omega}\right)$ |

**Table C.1:** Laplace domain representations of the four loop-shaping building blocks. $s\,(=i\omega)$ is the independent variable related to the frequency and $\omega_0$ is the "center frequency" of the particular building block.

to fit the phase and gain response using a series of these building blocks $G_i$ centered at various frequencies $\omega_i$, such that the loop kernel can be decomposed as

$$G(s) = \prod_i G_i. \tag{C.2}$$

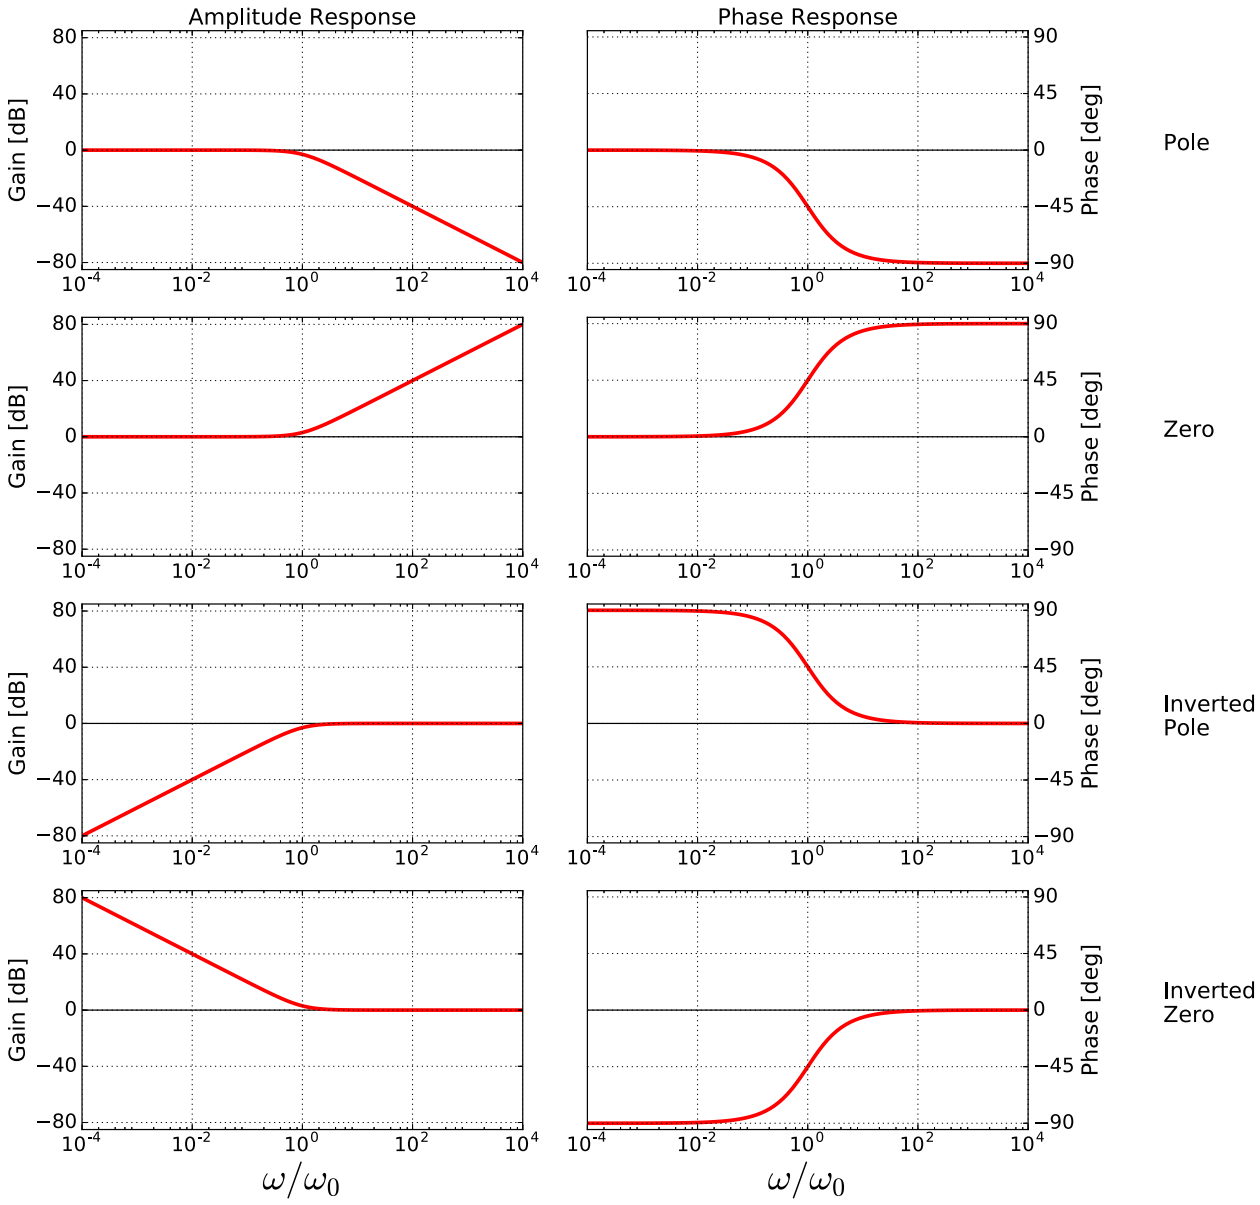



**Figure C.1:** The amplitude (left column) and phase (right column) response for the building blocks. From the top, pole, zero, inverted pole, inverted zero. Note that the frequency has been normalized to the "center frequency" $\omega_0$. For the amplitude response, the gain $p$ has been made unity.

## C.3 Determination of Servo Components

Once the transfer function $G(s)$ has been decomposed into its composite pieces $G_i$, the servo $K(s)$ can be constructed using "phase complementary" pieces $K_i$. Each $K_i$ "flattens" the phase contribution from the corresponding $G_i$. The phase complementary pieces are shown in Table C.2.

| Building Block | Phase Complementary Block |
|---|---|
| Pole | Inverted Zero |
| Zero | Pole |
| Inverted Pole | Inverted Zero |
| Inverted Zero | Pole |

**Table C.2:** Phase complementary blocks

The center frequencies $\omega_{0i}$ of each $K_i$ should match the center frequencies of the corresponding $G_i$. For example, if a particular $G(s)$ were decomposed into a pole at 100 Hz and a zero at 1 kHz, the phase complementary pieces that compose the $K(s)$ would be an inverted zero at 100 Hz and a pole at 1 kHz.

It is worth noting that apart from offsets, the phase responses for the pole and inverted pole, and for the zero and inverted zero, are identical. From a phase standpoint, a pole could be "phase-flattened" by either a zero or an inverted zero. However, the gain of the complementary piece should also be considered when making a selection. In most cases, high gain is desirable at low frequencies. Thus, the pole and inverted zero are the advantageous choices for complementary pieces.

## C.4 Construction of the Servo

Now that the servo function $K(s)$ has been modeled, it can be constructed digitally using a computer or FPGA or in analog using a combination of op-amps, resistors, capacitors, and inductors. Here we will demonstrate the analog approach, as that method was used for the work in this thesis.

The amplitude and phase response for each building block shown in Figure C.4 can be created using active operational amplifier filters, as shown in Figure C.2. For each of these circuits, the center frequency $\omega_0$ is set by varying the values of $R_2$, $C$, and $L$. Namely, for pole circuits and inverted zero circuits, the center frequency is $\omega_0 = \frac{1}{R_2 C}$. For zero circuits and inverted pole circuits,

the center frequency is $\omega_0 = \frac{L}{R_2}$. The gain $p$ for each stage can be set by choosing an appropriate value for $R_1$. In practice, this can be done by using a potentiometer, which makes the gain tunable.

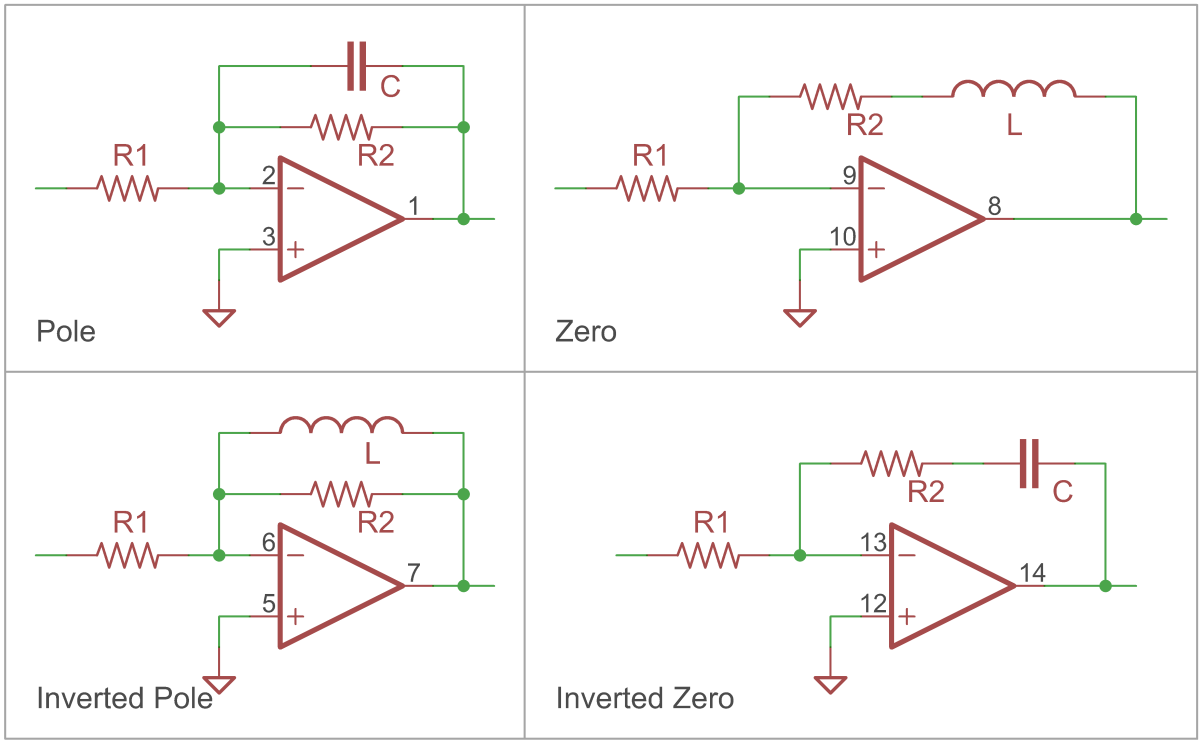



**Figure C.2:** The analog circuit representations for (clockwise from top left): pole, zero, inverted zero, inverted pole. For the pole and inverted zero, the center frequency is given by $\omega_0 = \frac{1}{R_2 C}$. For the zero and inverted pole, $\omega_0 = \frac{L}{R_2}$. The gain for the stage can be controlled by making $R_1$ a variable resistor, as the overall gain for the stage is $\frac{Z_{fb}}{R_1}$, where $Z_{fb}$ is the impedance in the "feedback" part of the stage (between the inverting input and output).

The phase-complement pieces are arranged in series to construct the feedback circuit. Placing op-amp followers between the various stages is beneficial for isolating the potentially mismatched impedances of each stage. For the simple 1-inverted-zero, 1-pole servo from above, a suitable circuit might look like Figure C.3. In this case, the circuit components would satisfy the relationships

$$R_2 C_1 = 2\pi(100\text{Hz})$$

and

$$R_4 C_2 = 2\pi(1000\text{Hz}).$$

The benefits of selecting the pole and inverted zero as the phase-complement pieces become even more apparent when examining their analog filter representations in Figure C.2. The pole and inverted zero can be constructed using only resistors and capacitors, which are readily available in a wide range of values in our lab and department electronics shop . This allows construction of circuits with precise values of $\omega_0$. The inverted pole and zero require the use of inductors, which are inferior to resistors and capacitors in both availability and accuracy.

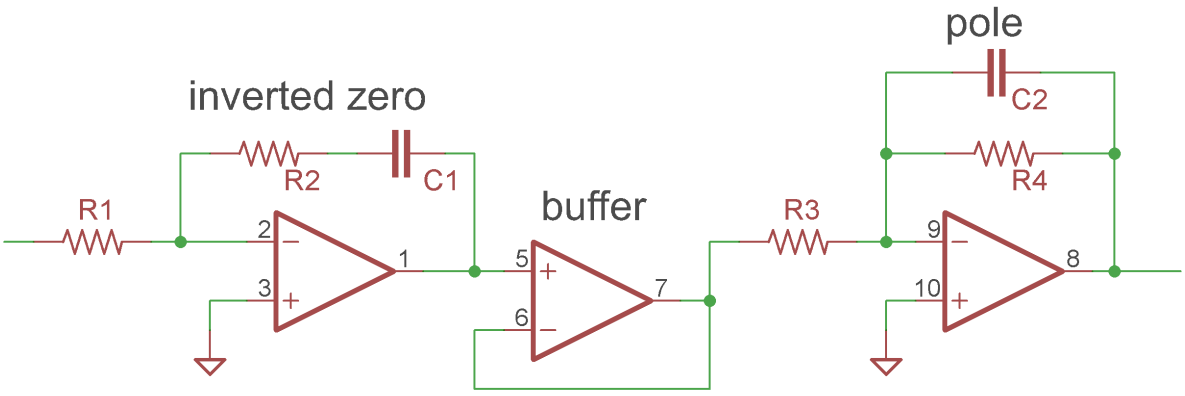



**Figure C.3:** A simple two-stage circuit designed to phase-complement a system with a pole at 100 Hz and zero at 1000 Hz. The

## C.5 Feedback Requirements

To avoid an oscillating when the system is operating in "closed-loop" mode, care must be taken when designing the phase response of the servo. Inspecting Equation C.2 reveals a potential catastrophe which can occur if the denominator of $1 + G(s)K(s)$ should go to zero. At this point, the gain is infinite and feedback causes oscillations in the system to occur. To avoid oscillation, we

*require* that the phase lag $|\text{Arg}\,[GK]\,|$ be smaller than 180ẇhen the amplitude product $|GK| = 1$, so that $1 + GK > 0$.

However, this is an absolute upper limit. To provide ourselves with some leeway (taking into account uncertainties in circuit component values, for example), it is wise to instead set the limit that

$$|\text{Arg}\,[KG]\,| < 135 \quad \text{when} \quad KG = 1.$$

The other consideration for good feedback is large gain across our frequency band of interest. For a fetal magnetocardiogram application, this region of interest is DC-100 Hz. Of course, this large gain consideration is subject to the *requirement* that KG ¡ 1 when the phase shift hits 180.

## C.6 Loop-shaping Example

We now present an example for implementing a loop-shaped servo into our magnetometer system. The goal of the servo is to maintain the "balance" of the Faraday rotation detector; in other words, to maintain at zero the y-component of the magnetic field at the sensor location. Thus for the "control signal", we have $r(t) \equiv 0$.

### C.6.1 Measuring Loop Kernel

The first step was to measure the loop kernel of the magnetometer and accompanying electronics. Notably, these electronics included the Stanford Research Systems SR570 current preamplifier with variable gain and filtering. For these measurements, the gain was set to 1 $\mu$A/V and the filters were disabled. Data were collected by the FPGA at a 20 kHz sampling rate and immediately downsampled to 1 kHz. After setting up the magnetometer and nulling the fields using the procedure described in Chapter 4, a constant-amplitude frequency chirp of amplitude .017 V and frequency band DC-20 kHz was applied to the y-coils of the magnetometer. This results in a nearly constant-amplitude 150 pT applied field over the frequency range. The recorded time series transformed to the frequency domain, and the result is shown in Figure C.4.

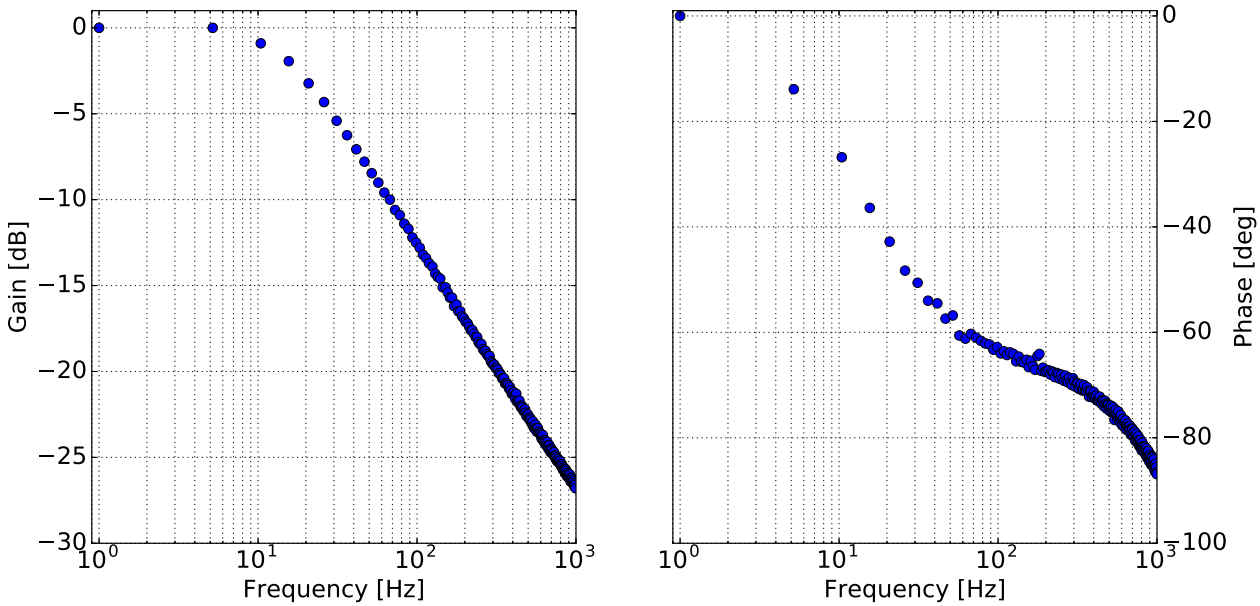



**Figure C.4:** Amplitude and phase response for the DC-SERF magnetometer.

### C.6.2 Fitting Loop Kernel

We turn our attention to the phase response of the system, which we will use to construct $G(s)$. We will demonstrate two methods for fitting this curve: one which is based more on the intuition of the human doing the curve-fitting, and one based on more of a "brute-force" method.

The intuition method relies on the user's ability to look at the phase response plot and "guess" the components that compromise the particular curve. For our data, the phase starts at zero near DC and then begins to lag, eventually reaching the 45 degree phase point near 20 Hz. Using the phase responses from Figure C.1 we would surmise there may be a pole at this frequency. From the pole phase response given in Table C.1, we might try initially fitting the curve of the form

$$C_\phi(f) = \arctan\left(-\frac{f}{20\,\text{Hz}}\right). \tag{C.3}$$

The result of this attempt is the red curve in the top plot of Figure C.5. This first approximation does indeed do a good job of fitting the data below 20 Hz, but does quite poorly after that. A function with a positive phase-frequency relationship is needed to fit the data as it begins to level

off at frequencies above 30 Hz. A zero satisfies these requirements, thus we might try modifying the fit to include the phase contribution from a zero centered at $f_0 = 150$ Hz.

$$C_\phi(f) = \arctan\left(-\frac{f}{20\,\text{Hz}}\right) + \arctan\left(\frac{f}{150\,\text{Hz}}\right). \tag{C.4}$$

We can continue adding more pieces to $C_\phi(f)$ until we get a close approximation of the curve. The fourth plot in Figure C.5 is the result of fitting the data with a 3-pole 1-zero curve fit of the form

$$\begin{aligned} C_\phi(f) = \arctan\left(-\frac{f}{20\,\text{Hz}}\right) + \arctan\left(\frac{f}{150\,\text{Hz}}\right) \\ + \arctan\left(-\frac{f}{300\,\text{Hz}}\right) + \arctan\left(-\frac{f}{10000\,\text{Hz}}\right). \end{aligned} \tag{C.5}$$

Once a suitable model (Equation C.5) is constructed, the exact frequencies can be fine-tuned using a least-squares curve-fitting regression program. Using the curve-fitting tool from the SciPy package in Python, the curve in the bottom plot of Figure C.5 was constructed with the parameters listed in Table C.3.

| Building Block | Center Frequency [Hz] |
|---|---|
| Pole | $f_{p_1} = 21.6$ |
| Zero | $f_{z_1} = 175$ |
| Pole | $f_{p_2} = 345$ |
| Pole | $f_{p_3} = 11000$ |

**Table C.3:** Result of curve fitting the phase-response to the 3-pole, 1-zero model first approximated in Equation C.5.

An alternate method for curve fitting, the so-called "brute-force" method, involves creating a fit with multiple components of each type and forcing the curve-fitting process to determine the frequencies. A fit of this type may have the form

$$C_\phi(f) = \sum_{j=1}^{a} \arctan\left(-\frac{f}{f_{p_j}}\right) + \sum_{j=1}^{b} \arctan\left(\frac{f}{f_{z_j}}\right) + \sum_{j=1}^{c} \arctan\left(\frac{f_{ip_j}}{f}\right) + \sum_{j=1}^{d} \arctan\left(-\frac{f_{iz_j}}{f}\right), \tag{C.6}$$

where the subscripts $p$, $z$, $ip$, $iz$ indicate the center frequencies for the poles, zeros, inverted poles, and inverted zeros, respectively. Thus $C_\phi(f)$ could fit a function with $a$ poles, $b$ zeros, $c$ inverted poles, and $d$ inverted zeros. In theory, this method is superior, as it requires no intuition on the user's part and can be expanded to fit very complicated phase responses. In practice, however, curve fitting data with complex fits such as the one in Equation C.6 can be difficult for the curve-fitting algorithm unless suitable initial "guesses" are provided by the user. The necessity of these guesses inevitably leads back to using the "intuition" method provided earlier.

The amplitude response of the loop kernel can be fit with the same set of pieces as the phase response. In fact, one could theoretically start by fitting the amplitude response, but in general the phase response is more indicative of the individual poles, zeros, inverted poles, and inverted zeros that compose the response. Put another way, the signal-to-noise ratio for curve fitting is often superior for the phase response. Additionally, phase shifts surpassing 180 with gain $\geq 1$ cause the loop to fail entirely, and this can be avoided by carefully fitting the phase response in advance of selecting the phase complementary components.

With the knowledge of the pieces required to model the phase response, the amplitude response can be modeled similarly. Because the amplitude response was measured in dB, the curve fit is simply a sum of the "Gain" functions from Table C.1, with each stage gain $p_i$ equal to unity.

$$C_A(f) = 20\log_{10}\left(\sqrt{\frac{1}{1+\left(\frac{f}{f_{p_1}}\right)^2}}\right) + 20\log_{10}\left(\sqrt{1+\left(\frac{f}{f_{z_1}}\right)^2}\right)$$
$$+ 20\log_{10}\left(\sqrt{\frac{1}{1+\left(\frac{f}{f_{p_2}}\right)^2}}\right) + 20\log_{10}\left(\sqrt{\frac{1}{1+\left(\frac{f}{f_{p_3}}\right)^2}}\right)$$
$$= 10\log_{10}\left[\frac{1+\left(\frac{f}{f_{z_1}}\right)^2}{\left(1+\left(\frac{f}{f_{p_1}}\right)^2\right)\left(1+\left(\frac{f}{f_{p_2}}\right)^2\right)\left(1+\left(\frac{f}{f_{p_3}}\right)^2\right)}\right]$$

where $f_{p_1}$, $f_{p_2}$, $f_{p_3}$, and $f_{z_1}$ are the frequencies determined from the frequency response curve fit, and whose values are listed in Table C.3. The resulting curve fit to the amplitude response is shown in Figure C.6.

### C.6.3 Servo Construction

Once the appropriate pieces and center frequencies of $G$ have been determined, all that is left is to select the appropriate phase-complementary pieces and determine the resistor/capacitor combinations for those frequencies. This process is illustrated in Table C.4. Because the individual stages are impedance-isolated by buffers, the absolute values of the components are not of high importance. For example, the 7.32 k$\Omega$, 1 $\mu$F $RC$ combination which creates the 21.6 Hz pole could have equally well been a 73.2 k$\Omega$, 0.1 $\mu$F pair.

The resulting circuit is shown in Figure C.7. Notice that the inverted zeroes have been combined in parallel on a single op-amp. The phase response of this stage will be dominated by the inverted zero whose center frequency is closest to the signal frequency. At low frequencies, the $RC$ combination for the 21.6 Hz inverted zero will shape the phase while the phase response of the 345 Hz and 11000 Hz inverted zeros will be essentially flat. Near 345 Hz, the $RC$ combination for that particular frequency will dominate, and the 21.6 and 11000 Hz will not contribute, and so on.

| $G$ element | Center Frequency [Hz] | $K$ element | $R$ [k$\Omega$] | $C$ [$\mu$F] | $(2\pi RC)^{-1}$ [Hz] |
|---|---|---|---|---|---|
| Pole | 21.6 | Inverted Zero | 7.32 | 1 | 21.6 |
| Zero | 175 | Pole | 9.09 | .1 | 175 |
| Pole | 345 | Inverted Zero | 4.64 | .1 | 343 |
| Pole | 11000 | Inverted Zero | 1.47 | .01 | 11000 |

**Table C.4:** Result of curve-fitting the phase-response to the 3-pole, 1-zero model first approximated in Equation C.5. The resistors and capacitors are off-the-shelf components avaliable at the department electronics shop. The resultant filter frequency using those components is shown in the last column.

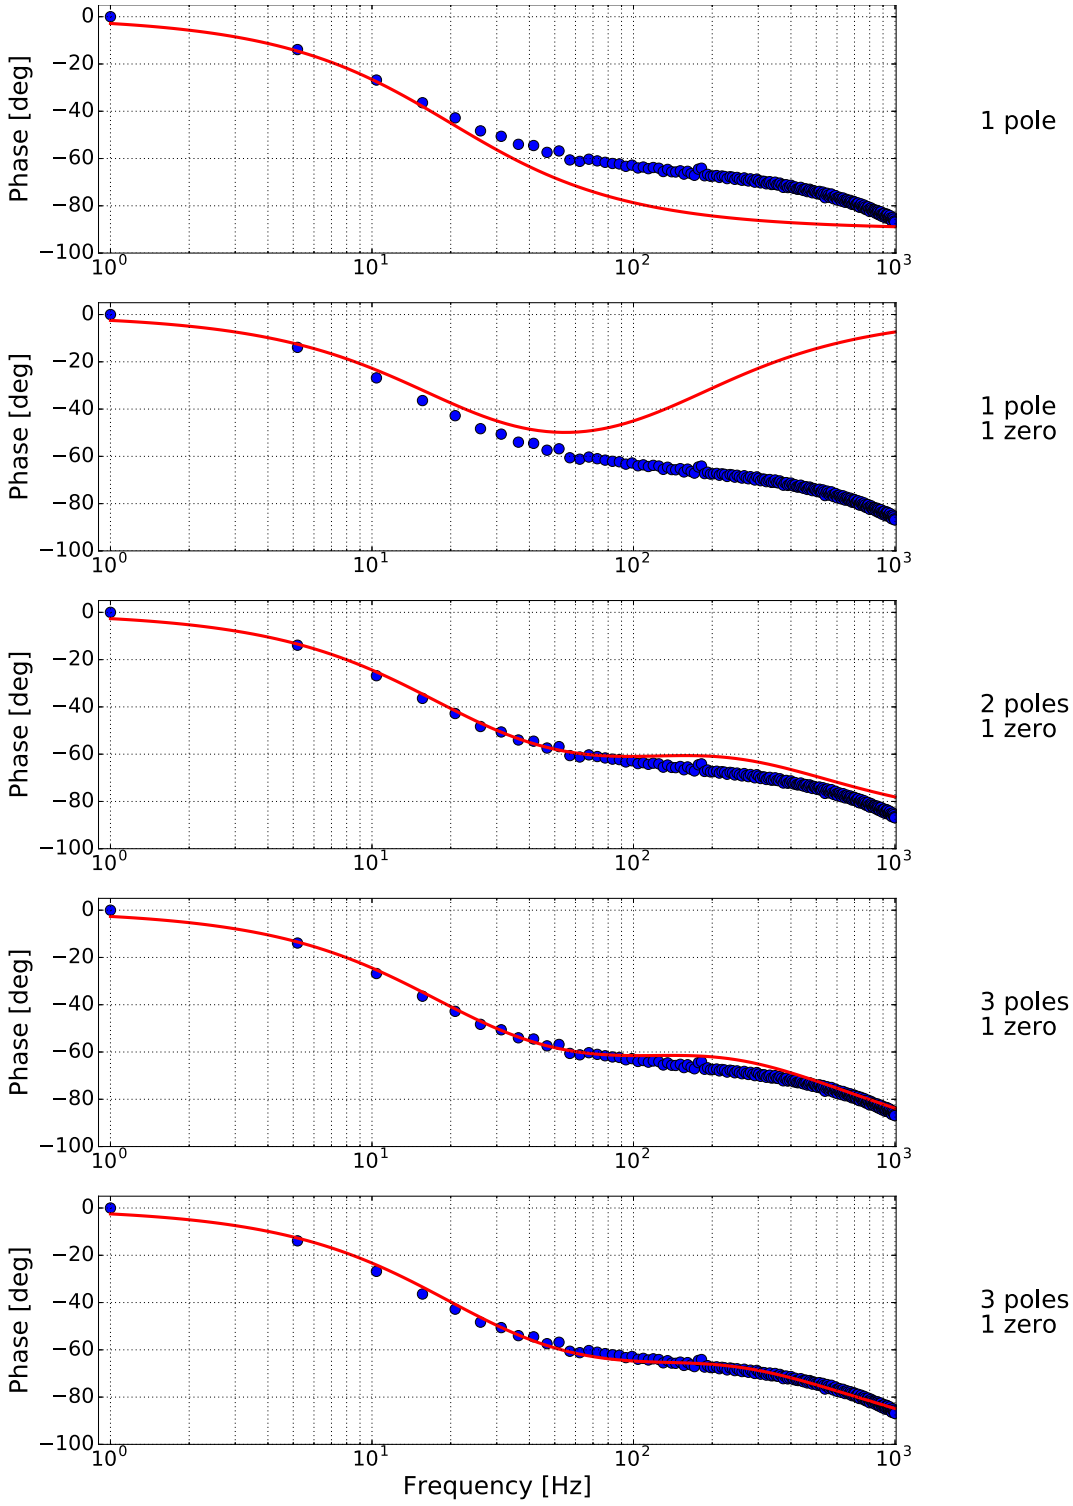



**Figure C.5:** Successive approximations to curve-fitting the phase response data. From top: 1-pole fit; 1-pole, 1-zero fit; 2-pole, 1-zero fit; 3-pole, 1-zero fit; 3-pole, 1-zero fit with center frequencies $f_{0_i}$ optimized by a least-squares algorithm.

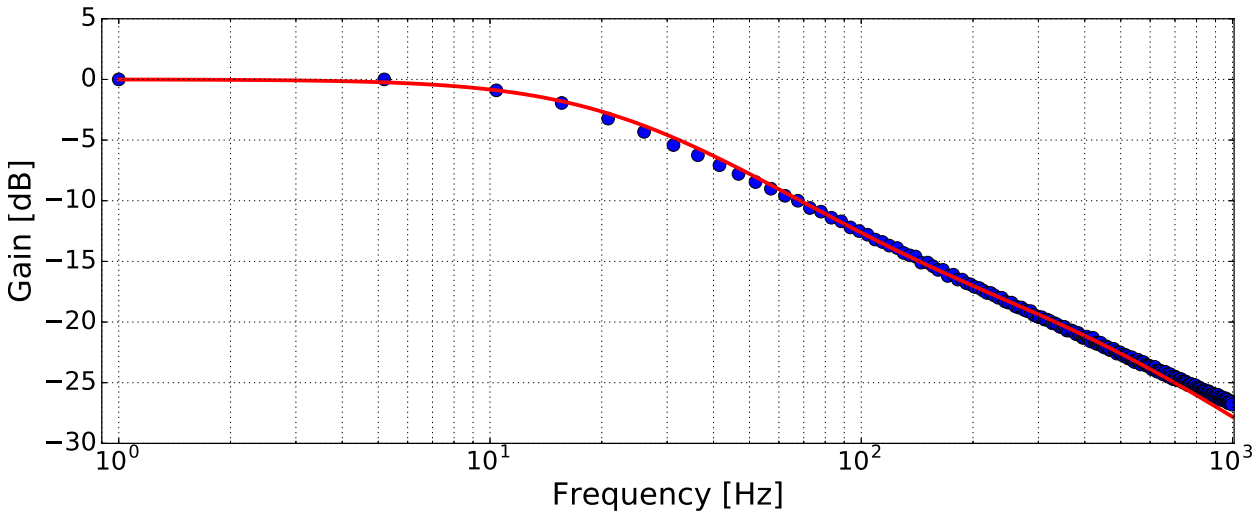



**Figure C.6:** Curve fit to the amplitude response using a 3-pole, 1-zero fit with the center frequencies from the phase response fit.

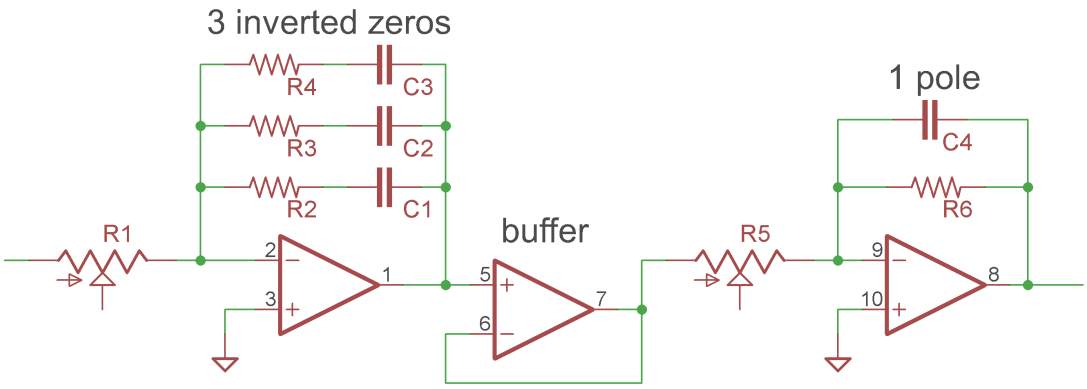



**Figure C.7:** Servo circuit created with the phase complementary pieces determined by the curve-fitting procedure. The component values can be found in Table C.4. The voltage follower between the poles and zero stages serves to impedance-isolate the stages from each other.

### C.6.4 Results

The servo $K(s)$ modeled in Section C.6.3 has an amplitude and phase response shown in Figure C.8. Note that the phase offset has been adjusted modulo [-180,180]. As desired, the phase response of the servo has the opposite behavior of the system. The offset of the amplitude response can be adjusted by an overall constant to suit the precise amplitude gain of $G$.

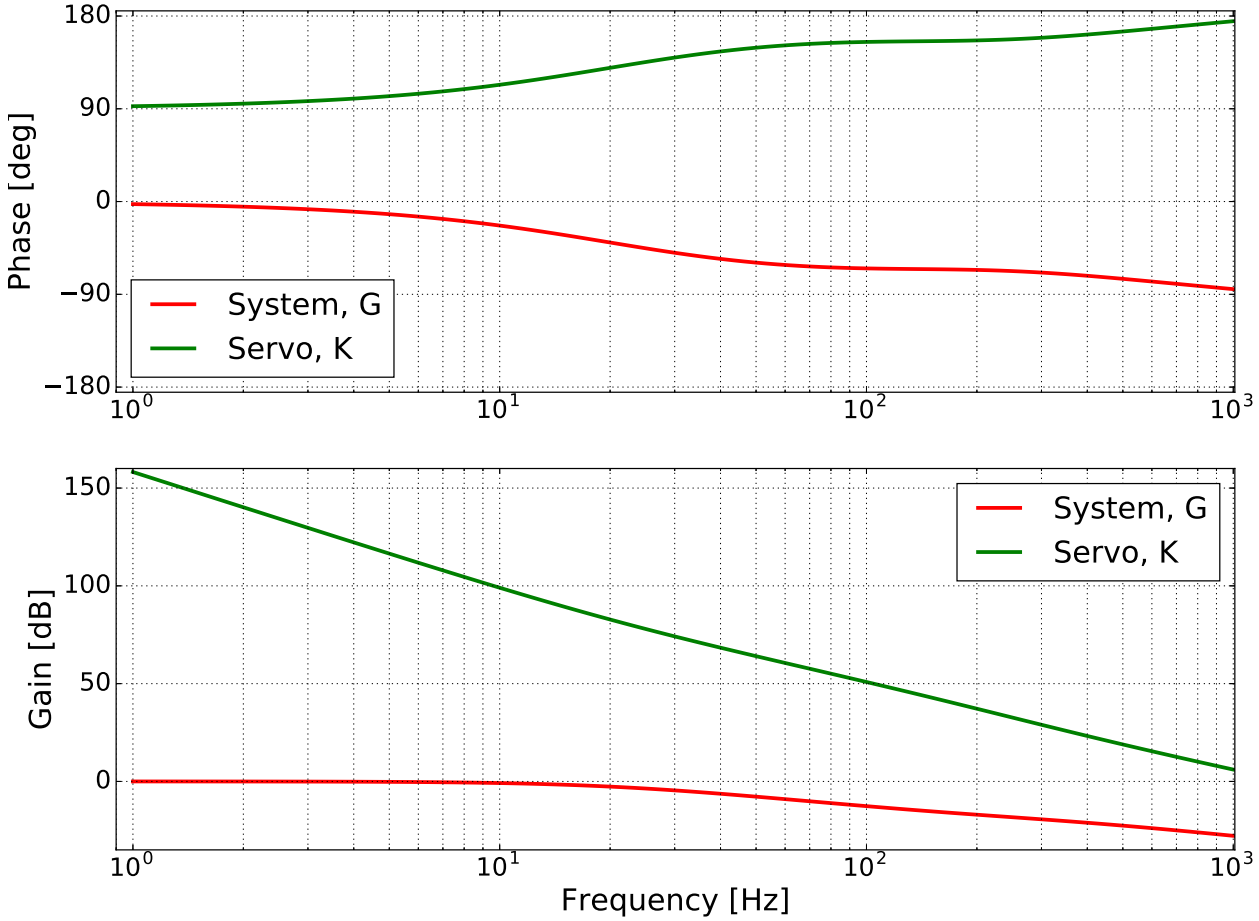



**Figure C.8:** Phase and amplitude response of the system $G$ and the servo $K$. Note the oppposite behaviors of the phase responses of $G$ and $K$.

The transfer function of the servo and the loop kernel is found by adding the phases of the two transfer functions as well as adding the gains (in dB). The resultant transfer function $GK$ is shown in Figure C.9. The desired flat phase response has been achieved over the range DC-1 kHz, though we expect this response to begin to tail off at higher frequencies, eventually crossing the oscillation

point of -180. With the additon of a frequency-independent adjustable gain stage, we can ensure that $KG$ is below 0 dB at this frequency.

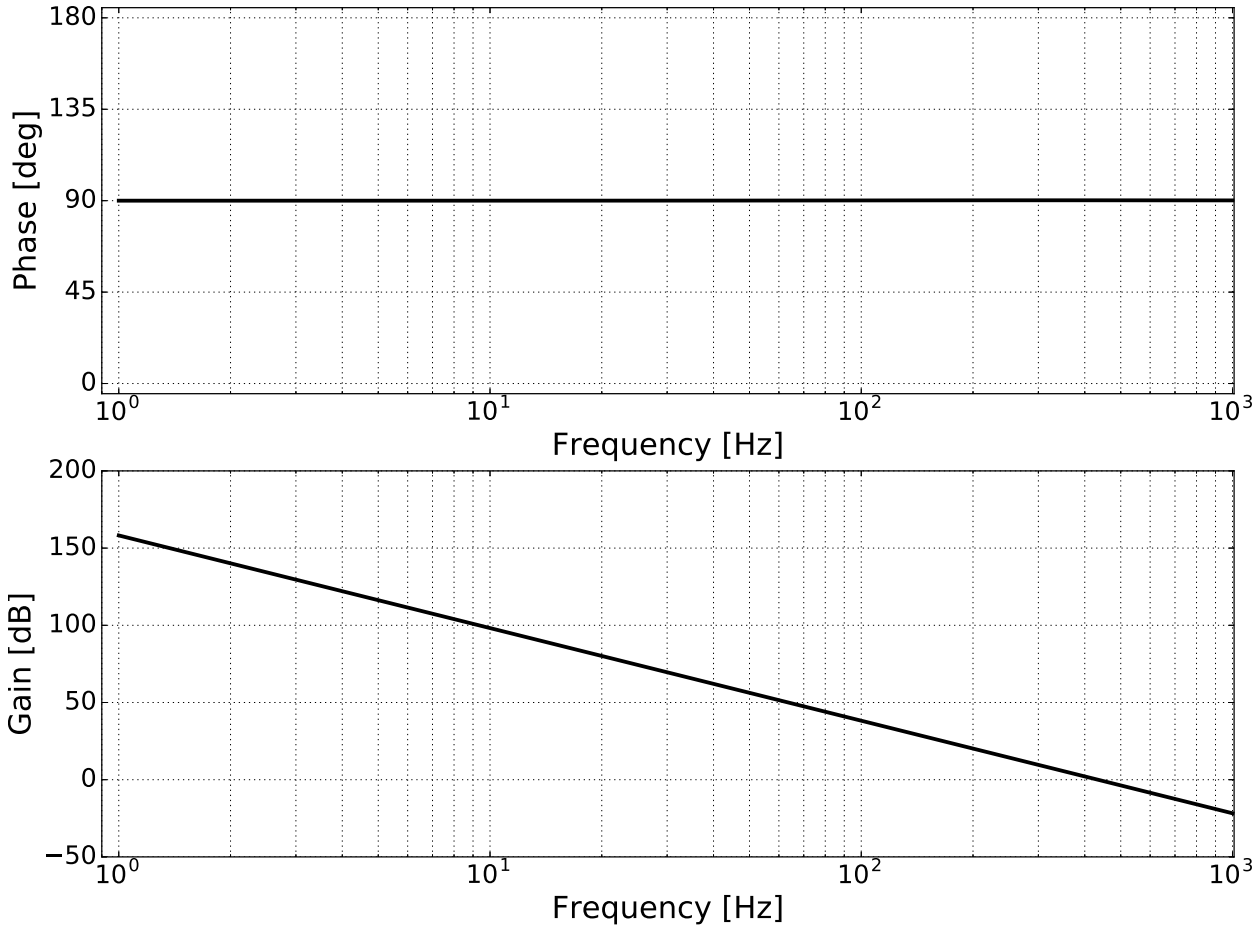



**Figure C.9:** Phase and amplitude response complete system $GK$.

# Appendix D: Gradiometry Figure of Merit

Here we provide a derivation of the gradiometry figure of merit $\mathcal{F}$ used in Chapter 5.

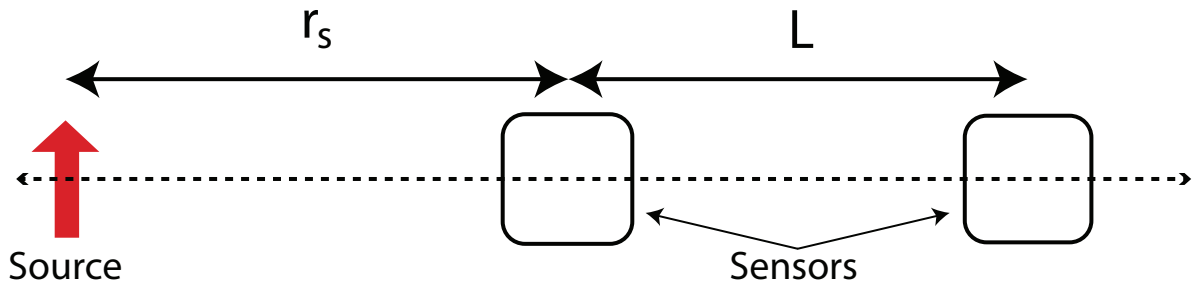



**Figure D.1:** Configuration of source and magnetometers. The source, a fetal heart is modeled to be a current dipole $Id\vec{l}$. In the cartoon, the dipole is oriented perpendicular to the axis of the magnetometers, but it should be noted the intra-heart currents that generate the detected fields flow along all three axes.

Consider a configuration as shown in figure D.1. A magnetic field source is located a distance $r_s$ from a magnetometer, with a second magnetometer positioned at an additional distance of $L$ along the same axis. The source is assumed to generate a magnetic field with a power law $B(r_s) \propto (1/r)^p$. For example, a point source magnetic dipole in a uniform medium would have a profile with $p = 3$. On the other hand, the field from a heart – relevant for MCG – is often modeled as a current dipole **?**, with $p \sim 1.8 - 2.2$ .

We express the output signal $S_i$ of each magnetometer as $S_i = M_i B(r_i)$ where $M$ is a response function and $B(r_i)$ is the field magnitude at a distance $r_i$ from the source. The gradiometer signal $S_G$ is given by the difference of the two magnetometer signals:

$$S_G = S_1 - S_2 = M_1 B_1 - M_2 B_2 \tag{D.1}$$

The signal from a well-balanced (i.e. $M_1 \approx M_2 = M$) gradiometer in a magnetic field $B_i = b_0/r_i^p$ is

$$\begin{aligned} S_G \approx M(B_1 - B_2) &= M\left(\frac{B_0}{r_s^p} - \frac{B_0}{(r_s+L)^p}\right) \\ &= \frac{-r_s^p + (r_s+L)^p}{(r_s+L)^p} M B_1 \\ &= \frac{g^p - 1}{g^p} M B_1 \end{aligned} \tag{D.2}$$

where $g = 1 + L/r_s$, and $B_1$ is the magnetic field strength at the magnetometer closest to the source.

In practice, the magnetometers can have different amplitude and phase responses to identical magnetic field inputs. The resulting gradiometer consequently will acquire a residual sensitivity to correlated magnetic fields. The undesired residual sensitivity is inversely proportional to the common-mode rejection ratio $\xi$, where

$$\begin{aligned} \xi &= \frac{\frac{1}{2}(S_{1c} + S_{2c})}{S_{1c} - S_{2c}} \\ &= \frac{\frac{1}{2}(M_1 + M_2)}{M_1 - M_2}. \end{aligned} \tag{D.3}$$

$S_{ic}$ is the measured signal from magnetometer $i$ due to an applied correlated field. In SQUID gradiometry **?**, $\xi$ is limited by the matching pickup coil areas or angular alignment of the two coils. In the case of atomic magnetometers, it is limited by the calibration procedure.

In general, a measurement is made in a magnetically noisy environment with total noise $\delta B = \sqrt{(\delta B_u)^2 + (\delta B_c)^2}$. $\delta B_u$ and $\delta B_c$ are the uncorrelated and correlated sources of magnetic noise, respectively. The SNR for the two-channel gradiometer is then

$$\begin{aligned} SNR_G &= \frac{\frac{g^p-1}{g^p} M B_1}{\sqrt{(M\delta B_{u1})^2 + (M\delta B_{u2})^2 + 2\left(M\delta B_c/\xi\right)^2}} \\ &= \frac{g^p - 1}{g^p} \frac{B_1}{\sqrt{2\left(\delta B_u\right)^2 + 2\left(\delta B_c/\xi\right)^2}}. \end{aligned} \tag{D.4}$$

Here we have made the assumption that $\delta B_{u1} \approx \delta B_{u2}$ for the two magnetometers.

If, instead, the two magnetometers (again, with $M_1 \approx M_2 = M$) were not operated as a gradiometer, then the total signal is

$$\begin{aligned} S_M &= S_1 + S_2 = M\left(B_1 + B_2\right) \\ &= M\left(\frac{B_0}{r_s^p} + \frac{B_0}{\left(r_s + L\right)^p}\right) \\ &= \frac{g^p + 1}{g^p} M B_1 \end{aligned} \tag{D.5}$$

Thus the signal-to-noise ratio for this two-channel magnetometer is

$$SNR_M = \frac{g^p + 1}{g^p} \frac{B_1}{\sqrt{2\left(\delta B_u\right)^2 + 2\left(\delta B_c\right)^2}}. \tag{D.6}$$

Comparison of equations D.4 and D.6 illustrates the potential advantage of using a gradiometer, which we define as the gradiometer figure of merit $\mathcal{F}$.

$$\begin{aligned} \mathcal{F} = \frac{SNR_G}{SNR_M} &= \frac{g^p - 1}{g^p + 1} \frac{\sqrt{2\left(\delta B_u\right)^2 + 2\left(\delta B_c\right)^2}}{\sqrt{2\left(\delta B_u\right)^2 + 2\left(\delta B_c/\xi\right)^2}} \\ &= \frac{g^p - 1}{g^p + 1} \frac{\sqrt{\left(\frac{\delta B_u}{\delta B_c}\right)^2 + 1}}{\sqrt{\left(\frac{\delta B_u}{\delta B_c}\right)^2 + \xi^{-2}}} \end{aligned} \tag{D.7}$$

Gradiometers are only useful if the total noise is dominated by correlated noise between the two channels. Under the assumption that $\delta B_u / \delta B_c \ll 1$,

$$\mathcal{F} \approx \frac{g^p - 1}{g^p + 1} \frac{1}{\sqrt{\left(\frac{\delta B_u}{\delta B_c}\right)^2 + \xi^{-2}}} \tag{D.8}$$

This clearly demonstrates the the figure of merit depends on not only the noise ratio of $\delta B_u / \delta B_c$, but also how that ratio compares to $\xi$. If $\left(\delta B_u / \delta B_c\right)^2 \ll \xi^{-2}$, immediate and potentially substantial benefits to the figure of merit (and therefore gradiometer performance) will be gained through an increase of $\xi$. If, instead, $\left(\delta B_u / \delta B_c\right)^2 \gg \xi^{-2}$, further experimental efforts would be best focused elsewhere.

It is important to note here that the ratio $(g^p - 1) / (g^p + 1)$ is somewhat geometry-sensitive and will also depend on the power law characterizing the field, as shown in figure D.2. In setting the appropriate baseline, a common practice is to match the baseline with the depth of the signal of interest, ie $L \approx r_s$, though simulations [Garachtchenko et al., 1999] have shown the optimal baseline is a function of balance, signal depth, and ambient noise source distance.

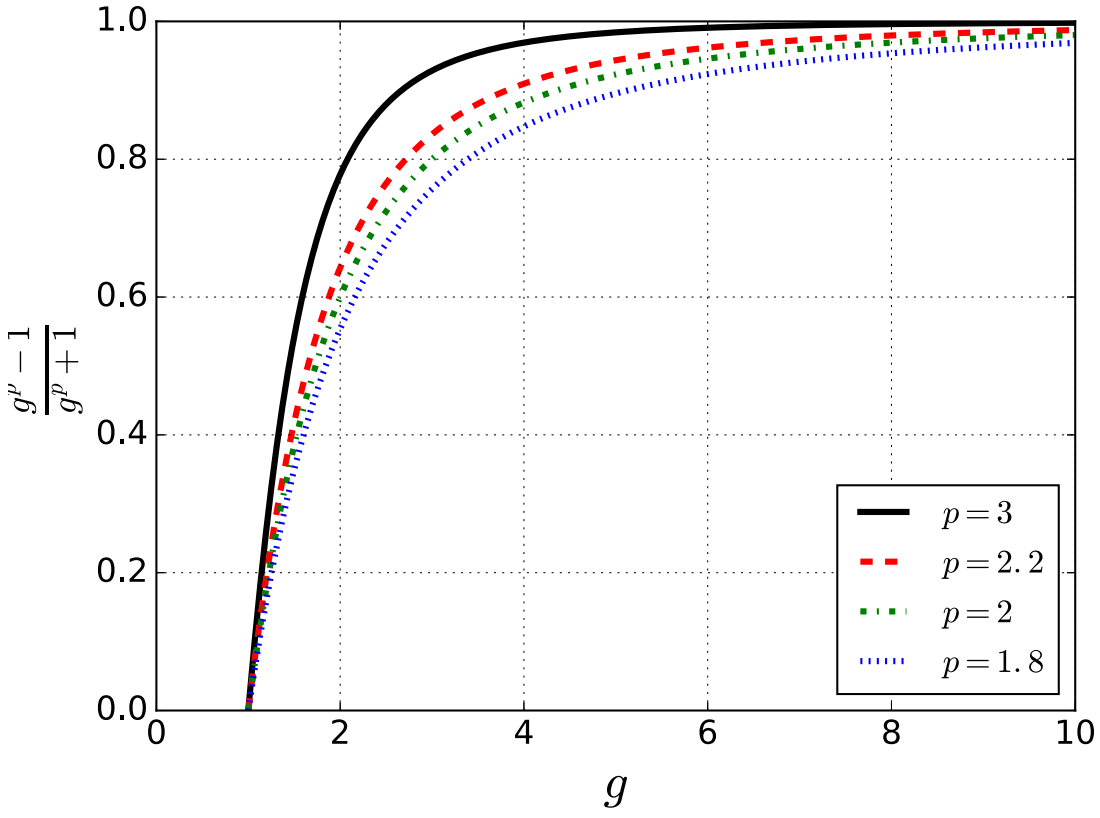



**Figure D.2:** $(g^p - 1) / (g^p + 1)$ as a function of the geometric parameter $g$ for various values of $p$. $p = 3$ is appropriate for a purely magnetic dipole source, while $p = 2$ is appropriate for a purely current dipole source. $g = 2$ for the case $L = r_s$.

Taking $p = 2$ and $L = r_s$, for example,

$$\mathcal{F} \approx \frac{3}{5} \frac{1}{\sqrt{\left(\frac{\delta B_u}{\delta B_c}\right)^2 + \xi^{-2}}}. \tag{D.9}$$

# Appendix E: Adult MCG results

The acquisition of the adult magnetocardiogram (MCG) was not the purpose of this thesis, as it has been captured using atomic magnetometers many times previously ([Wyllie et al., 2012a], for example). Still, testing our apparatus on adult subjects was often done as a peace of mind check or to experiment with new apparatus components or procedures. Figure E.1 shows some example 8-channel Z-mode data collected from an adult male patient. The magnetometer array was placed directly over the patient's chest when the measurement was made, thus the signal size is significantly larger than the maternal signal observed in the fetal measurement done in Chapter 6.

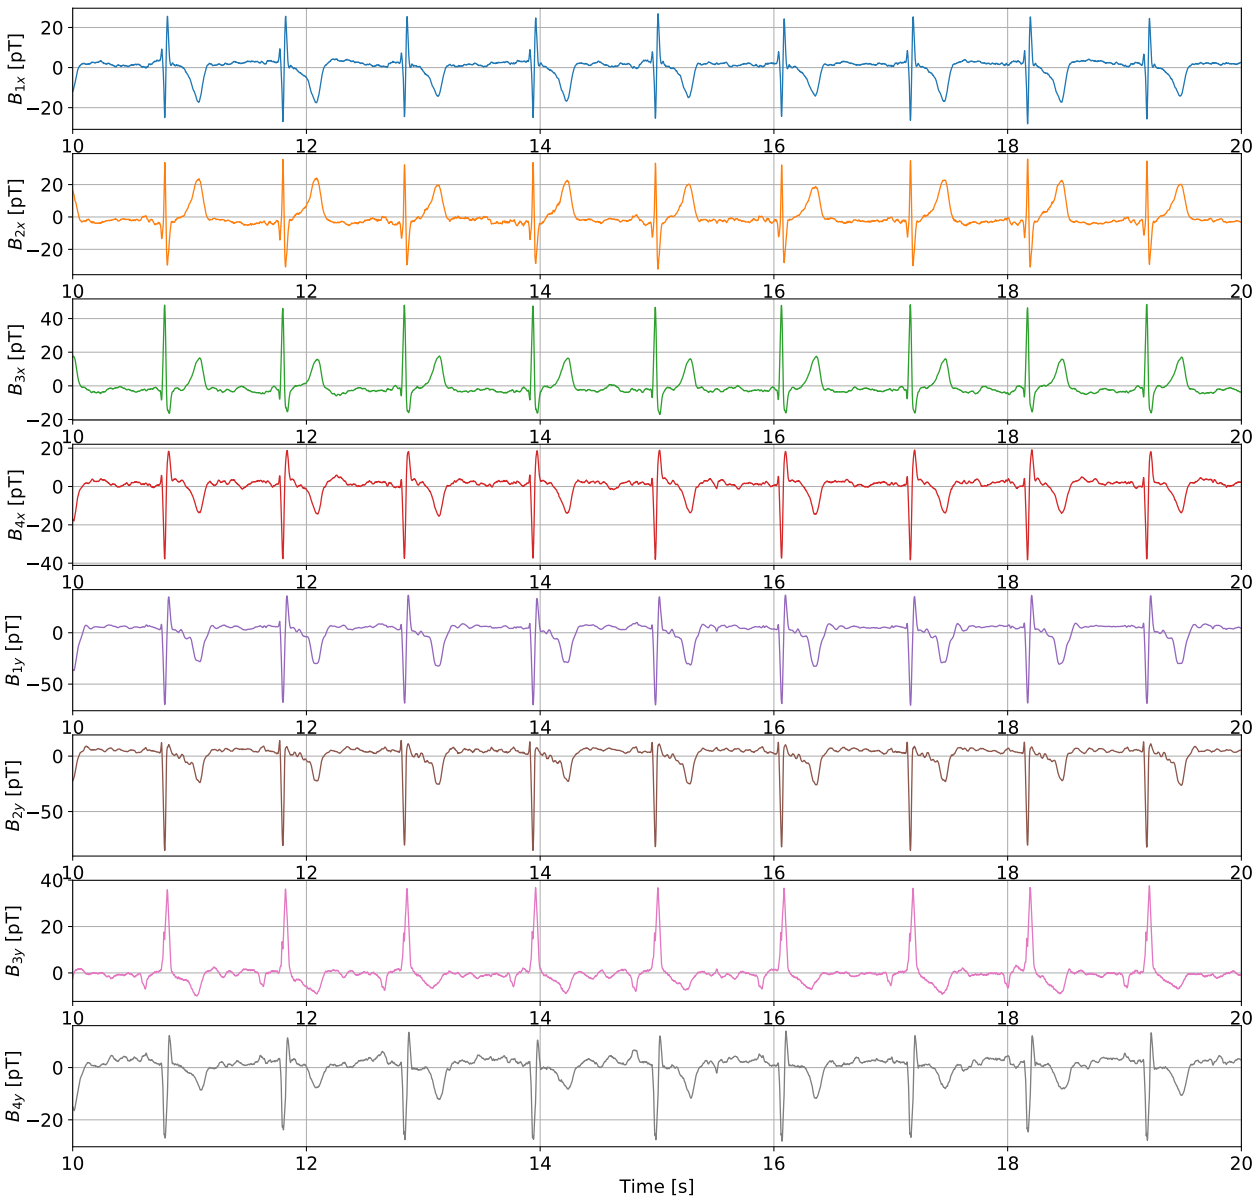



**Figure E.1:** An eight channel measurement of an adult MCG. The data presented here has undergone the single-channel noise processing techniques described in Section 6.5. None of the multi-channel processing techniques were employed here: it was obviously not desirable to remove the adult signal from these data as we did during the fMCG signal processing.