

Approaches to Continuous Excitation in a Transverse NMR Gyroscope

by

Susan Stoffer Sorensen

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

Doctor of Philosophy

(Physics)

at the

UNIVERSITY OF WISCONSIN–MADISON

2023

Date of final oral examination: 05/11/2023

The dissertation is approved by the following members of the Final Oral Committee:

Thad G. Walker, Professor, Physics

Deniz Yavuz, Professor, Physics

Mark Saffman, Professor, Physics

Jennifer Choy, Assistant Professor, Electrical and Computer Engineering

© Copyright by Susan Stoffer Sorensen 2023

All Rights Reserved

"Only the curious have something to find." ~Nickel Creek

Acknowledgments

I would like to thank, first and foremost, my advisor Thad Walker, for his patience, guidance, flexibility, and compassion. I cannot imagine making it through grad school with anyone else as my mentor. Thank you to Josh Weber, who trained me, and Daniel Thrasher, who was my collaborator for most of the work of this thesis. Thanks to Michael Bulatowicz for his electrical engineering know-how, especially when it came to the design of the PDM pulsing circuit. Thank you to Mike Larson for enlightening conversations and consultation. Thank you to Zaynab Yardim for your friendship and bringing me back to basics, and to Mike Vinje for checking in and keeping my lab clean.

Certainly, the support I have received during my time in grad school has extended far beyond academics. My husband, Eric Sorensen, has been a grounding influence during the most turbulent times. He has been a shoulder to cry on, a voice reassuring me that I have not, in fact, screwed up everything, when panic set in and it felt like the world was crashing in around me. I could write another thesis just about you. Thank you to my family for supporting me in uncountable ways. My parents raised me with a belief that I could do anything. They instilled in me a curiosity for the world around me, and a love of sharing what I know with others. They have provided invaluable counseling, guiding and supporting me to where I am now. I am thankful for the friendship of Linea and JD Kemsley and Shelby Shoaf, old friends who moved across the country so we could live near each other. Thank you to everyone from trivia night and my D&D groups, who have kept me (relatively) sane and provided much-needed escapes. I thank and love you all.

Contents

Contents iii

List of Figures vi

Abstract ix

1 Introduction 1

1.1 Device Overview 2

1.2 A Brief History of NMR Comagnetometers 5

1.3 Novel Features of This Device 7

1.4 Outline of This Thesis 9

2 Theory 13

2.1 Geometry and Systematics 14

2.2 The Bloch Equation 15

2.3 Comagnetometry 21

2.4 Performance Metrics 22

2.5 Chapter Summary 25

3 Apparatus 26

3.1 Vapor Cell 26

3.2 Magnetic Fields and Pulsing Circuits 31

3.3	<i>Optics</i>	39
3.4	<i>Digital Control</i>	44
4	Polarization Modulation	47
4.1	<i>Modulation and Drive Scheme</i>	47
4.2	<i>Demodulation</i>	49
4.3	<i>NMR Excitation</i>	51
4.4	<i>Comagnetometry</i>	53
4.5	<i>Cross Talk</i>	57
4.6	<i>Discussion</i>	59
5	Pulse Density Modulation	60
5.1	<i>Modulation and Drive Scheme</i>	60
5.2	<i>Gated Detection</i>	61
5.3	<i>Demodulation</i>	66
5.4	<i>Xe Excitation</i>	68
5.5	<i>Feedback</i>	70
5.6	<i>SNR</i>	74
5.7	<i>Transverse Field Sensitivity and Stability</i>	78
5.8	<i>Chapter Summary</i>	82
6	Hybrid Modulation	84
6.1	<i>Modulation and Drive Scheme</i>	84
6.2	<i>Demodulation</i>	88
6.3	<i>Xe Excitation</i>	93
6.4	<i>Cross talk</i>	95
6.5	<i>Comagnetometry</i>	97
6.6	<i>Chapter Summary</i>	106

7	Outlook	108
7.1	<i>Comparing PM, PDM, and Hybrid Operation</i>	108
7.2	<i>Revisiting PM</i>	115
7.3	<i>Second Apparatus</i>	118
A	Acronyms, Symbols, and Abbreviations	121
A.1	<i>Acronyms and Abbreviations</i>	121
A.2	<i>Symbols</i>	122
	References	126

List of Figures

2.1	Simplified Apparatus.	14
3.1	Temperature stabilization. RTD noise with and without feedback.	28
3.2	Vapor cell heater and cartridge mount photographs.	29
3.3	Inside of the magnetic shields.	30
3.4	T_1 Measurement.	31
3.5	Rb relaxation measurement.	32
3.6	3D printed rig for cell cartridge and magnetic field coils.	33
3.7	Three layer μ -magnetic shield.	34
3.8	PM pulsing circuit.	35
3.9	Voltage for 2π pulse vs pulsing repetition frequency.	36
3.10	H-Bridge pulsing circuit.	36
3.11	Pulse shape comparison.	37
3.12	Rb transient area with and without π pulse.	38
3.13	Experimental setup of optical elements	40
3.14	Probe, electronic, and photon shot noise.	43
3.15	Block diagram of digital system.	45
4.1	Square wave PM signals	49
4.2	NMR of each species	51
4.3	Phase noise	52

4.4	Transfer function of ^{129}Xe	54
4.5	Comagnetometry noise and stability	56
4.6	Measurement of cross talk	57
5.1	PDM product of Fourier coefficients vs modulation depth	62
5.2	Closed loop drift for PDM	63
5.3	Example of gated detection	65
5.4	S_z signals obtained using PDM.	69
5.5	Xe linewidth measurements for PDM.	70
5.6	Cross talk measurement for PDM	71
5.7	Inverted zero with lead compensation gain.	72
5.8	Inverted zero gain.	74
5.9	Field Suppression Factor measurement for PDM	75
5.10	SNR measurements for PDM	77
5.11	Transfer function with and without higher-order transverse field effects.	80
5.12	Lissajous figures of transverse field nulling procedure.	81
5.13	Rotation MDEVs before and after nulling transverse fields.	82
5.14	Comparison of previous best and new best PDM MDEVs.	83
6.1	Applied modulation waveforms for hybrid operation.	86
6.2	Fourier amplitude j vs PDM modulation index.	88
6.3	Measured and theoretical S_z measurements.	90
6.4	Xe linewidth measurements using hybrid operation.	94
6.5	Measurement of cross talk.	96
6.6	MDEV of comagnetometry calculation using open loop Xe measurements.	98
6.7	Hybrid FSF measurement.	100
6.8	Rotation MDEVs for transverse field scans.	102
6.9	MDEVs of long measurement cut into 5 ksec segments.	103

6.10	Best-to-date MDEVs for PDM and hybrid operation.	104
6.11	An example of magnitude drift when applying DSDO feedback.	107
7.1	Table summarizing the major results of this thesis.	109
7.2	Comparison of the measured and calculated Xe spin-exchange fields.	111

Abstract

In this thesis, I describe three approaches to the continuous operation of a transverse spin-exchange optically pumped NMR gyroscope. I demonstrate the simultaneous, continuous excitation of ^{131}Xe and ^{129}Xe using first modulations of the optical pumping (polarization modulation, PM), then using modulations of the applied bias field (pulse density modulation, PDM), and then using a combined approach that includes both polarization and field modulations (hybrid operation). I describe the theory behind the operation of the device, and the apparatus used to perform the experiments presented in this thesis. I then present results for each of the three operation modes - PM, PDM, and hybrid operation. For each, I give several key performance metrics, including linewidth, signal-to-noise ratio, field suppression factor, cross talk, angle-random walk, and bias instability. The field suppression factor and cross talk are unique metrics which have generally been ignored in this field of study. This thesis includes the best bias instability ever measured with this device, at ~ 420 nHz. I close with a comparison of the three methods and a discussion of proposed future works.

Chapter 1

Introduction

Spin-exchange (SE) pumped Nuclear Magnetic Resonance (NMR) comagnetometers measure non-magnetic spin-dependent interactions by monitoring deviations in the precession frequency of spin-polarized nuclei about an applied bias magnetic field. A vapor cell contains two isotopes of noble gas atoms and an alkali-metal vapor. The alkali-metal vapor is optically pumped using circularly polarized near-infrared laser light. Spin-exchange collisions between the polarized alkali-metal and noble gas atoms polarize the noble gas nuclei. When subject to a magnetic (or “bias”) field, the polarized species undergo Larmor precession. By simultaneously measuring the Larmor precession of each noble gas species, magnetic field fluctuations can be removed, thereby increasing the signal to noise ratio for measuring non-magnetic spin-dependent phenomena, such as inertial rotation. When used to measure inertial rotation [Walker and Larsen, 2016; Kornack et al., 2005; Karwacki, 1980], such devices are referred to as NMR gyroscopes. In this thesis, I will present an NMR gyroscope which utilizes a unique transverse geometry to avoid systematic errors. I will present three different methods for continuously exciting the NMR precession of the spin-polarized nuclei.

In this chapter, I will begin by giving a basic overview of the operation of the transverse NMR gyroscope presented in this work. I will then touch on some of the prior work done in this field, and discuss separately our device from others currently being studied. Finally, I

will give an outline of the work presented in the rest of this thesis.

1.1 Device Overview

The goal of the gyroscope is, as one might guess, to measure rotation. For an NMR gyroscope, this is accomplished by measuring deviations in the rate at which a polarized spin precesses about an applied magnetic field. The total precession frequency (Ω) of a spin in the presence of a magnetic field (B) is

$$\mathbf{s}\Omega = \mathbf{s}\gamma B + X \quad (1.1)$$

where γ is the gyromagnetic ratio, \mathbf{s} is the sign of the gyromagnetic ratio such that γ is positive, and X represents any non-magnetic, spin-dependent effect. Since $\gamma B \gg X$, it is convenient to define the precession frequency Ω to have the same sign as γB . Since $\mathbf{s}^2 = 1$, we can re-write Eq. 1.1

$$\Omega = \gamma B + \mathbf{s}X. \quad (1.2)$$

In the case of a gyroscope, X is the inertial rotation frequency. In order to obtain a measurement of X which is not limited by the magnetic field noise, two species (a and b) precessing about the same field are used to create a comagnetometer. In this configuration, two precession rates are measured,

$$\Omega^a = \gamma^a B + \mathbf{s}^a X^a \quad (1.3a)$$

$$\Omega^b = \gamma^b B + \mathbf{s}^b X^b \quad (1.3b)$$

and for simplicity, we can re-write this system of equations as a single, species-generic equation, written simply as Eq. 1.1. I will condense systems of equations in this way

periodically throughout this thesis, when the reader might otherwise be overwhelmed by a sea of superscripts and equations.

Comagnetometry is the use of two species to eliminate common magnetic field fluctuations. We do this by solving the system of equations 1.3 to obtain

$$\mathbf{s}^a X^a - \frac{\gamma^a}{\gamma^b} \mathbf{s}^b X^b = \Omega^a - \frac{\gamma^a}{\gamma^b} \Omega^b \quad (1.4)$$

or, in the case of rotation ($\Omega^R = X^a = X^b$)

$$\Omega^R = \frac{\Omega^a - \rho \Omega^b}{\mathbf{s}^a - \rho \mathbf{s}^b} \quad (1.5)$$

where $\rho = \gamma^a/\gamma^b$. Thus, one can obtain a measure of rotation which is independent of magnetic field. In our system, species a is ^{129}Xe and species b is ^{131}Xe , such that $\mathbf{s}^a = -1$ and $\mathbf{s}^b = 1$, meaning that the two species precess in opposite directions in response to a magnetic field.

It is important to note that only magnetic fields common to both noble gas species will be eliminated through comagnetometry. Consider, for instance, a field which includes a spin-exchange field $b_S^{a,b} S$ due to alkali-metal electron spin polarization S . If b_S^a is not equal to b_S^b , the two noble gas isotopes will precess about slightly different fields and comagnetometry will not completely eliminate the spin-exchange field. In the transverse NMR gyroscope, the alkali-metal electron spin polarization is perpendicular to the applied bias magnetic field, suppressing the influence of spin-exchange fields on the precession measured about the bias field.

The fundamental operation of the device is as follows. Circularly polarized pump light passes through a cell containing alkali-metal atoms (^{85}Rb for this work) and two noble gas isotopes. This light optically pumps the alkali-metal atoms, polarizing the alkali-metal electrons. These polarized electrons undergo spin-exchange collisions with the two noble gas

species, polarizing the noble gas nuclei. When a bias magnetic field is applied, the polarized noble gas spins precess about the field. These are the precession frequencies Ω^a , Ω^b that we wish to measure.

Readout of the precession frequencies is done using the alkali-metal atoms as an embedded magnetometer. This is advantageous over using an external pick-up coil because, during atomic collisions, a Fermi-contact interaction enhances the field experienced by the alkali-metal atoms due to the the polarized noble gas nuclei. The Fermi-contact enhancement factor between the alkali-metal (^{85}Rb) and the noble gas (Xe) is approximately $\kappa = 500$ [Nahlawi et al., 2019]. Thus, the embedded Rb magnetometer is $500\times$ more sensitive to Xe fields than an external pick-up coil.

Due to optical pumping, the alkali-metal atoms remain polarized predominantly along the pump light axis while the noble gas nuclei are precessing. The noble gas then exerts a torque on the alkali-metal atoms, tipping the electron spin into the direction of the bias magnetic field. The amount of longitudinally-polarized alkali-metal spin then tracks with the precession angle of the noble gas. We can measure the amount of longitudinally-polarized alkali-metal spin using a linearly-polarized probe laser along the longitudinal axis and a balanced Faraday-rotation detector [Zapasskii, 1982].

One complication that arises in a transversely pumped NMR gyroscope is that the bias field is applied transverse to both the noble gas and alkali-metal polarizations. This means that the alkali-metal polarization would also precess about the bias magnetic field. If a constant DC bias field were applied, the rapid precession of the alkali-metal atoms would make it practically impossible to build up meaningful noble gas polarization through spin-exchange collisions, or to use the alkali-metal atoms as an embedded magnetometer. To avoid this complication, we apply our bias magnetic field as a series of short (low-duty-cycle) pulses, with $\int B dt$ (which we refer to as the pulse area) of each pulse set to result in a 2π radian rotation of the alkali-metal polarization. The noble gas polarization precesses much less than 2π radians during each pulse due to having a much smaller gyromagnetic ratio.

Therefore, by applying the bias field as a series of short pulses, the alkali-metal atoms spend most of the time experiencing zero precession, and we can approximate the precession of the noble gas nuclei with an effective bias field

$$B_p(t) = \omega_p(t)/\gamma^s \quad (1.6)$$

where ω_p is the repetition rate of the bias pulses and γ^s is the gyromagnetic ratio of the alkali-metal atoms. The co-precession of the noble gas and alkali-metal spins is referred to as "synchronous" or "Bell-Bloom" pumping [Bell and Bloom, 1957].

In order for a measurement of rotation to be useful in navigation applications, the device must have low noise (characterized by high Signal-to-Noise Ratio (SNR) or low Angle Random Walk (ARW) in this thesis) and high stability (characterized by low bias instability) over long times. These quantities are therefore of special import in this thesis, and will be discussed in more detail in chapter 2.

In our device, the two noble gas species are ^{129}Xe and ^{131}Xe , and the alkali-metal atoms are ^{85}Rb . More details about the vapor cell can be found in chapter 3. In order to measure for long times despite the short Xe transverse relaxation times (T_2), we wish to excite Xe continuously, rather than measure free induction decays (FIDs). In order to continuously excite Xe, we must apply modulations to the pump light polarization and/or to the bias field. These modulations will be discussed further in chapter 2.

1.2 A Brief History of NMR Comagnetometers

Spin-exchange pumped NMR devices [Walker and Larsen, 2016; Terrano and Romalis, 2021; Karwacki, 1980] have been studied for several decades. Spin-exchange optical pumping of He-3 was first demonstrated in 1960 [Bouchiat et al., 1960]. Such devices remain of interest in the precision measurement community for a variety of applications studying non-magnetic,

spin-dependent interactions [Bulatowicz et al., 2013; Feng et al., 2022; Thrasher et al., 2019a; Jiang et al., 2018; Lee et al., 2018; Allmendinger et al., 2014; Romalis et al., 2014; Smiciklas et al., 2011; Brown et al., 2010; Glenday et al., 2008; Rosenberry and Chupp, 2001; Allmendinger et al., 2019; Sachdeva et al., 2019]. Some of this interest has been driven by the potential for miniaturization of NMR devices, especially when compared to the size scaling of similar devices such as ring laser gyroscopes [Donley, 2010]. Currently studied applications for NMR comagnetometers fall primarily into two categories: studies of exotic physics, and inertial navigation.

NMR comagnetometers have been used to search for exotic physics interactions which are beyond the standard model. In order to search for axion-like interactions and place upper bounds on spin-mass couplings [Bulatowicz et al., 2013; Lee et al., 2018], a comagnetometer’s signal is monitored while a non-magnetic mass is placed at varying distances from the device. Searches for Lorentz invariance violations [Allmendinger et al., 2014; Romalis et al., 2014; Smiciklas et al., 2011; Brown et al., 2010; Glenday et al., 2008] are performed by monitoring a comagnetometer signal when the device is placed at different orientations in the lab frame. Searches for atomic electric dipole moments [Rosenberry and Chupp, 2001; Allmendinger et al., 2019; Sachdeva et al., 2019] involve measuring the comagnetometer signal in the presence of applied electric fields.

When coupled with an accelerometer, an NMR gyroscope can be used for inertial navigation applications. This application was first studied at Litton Industries [Grover et al., 1979], and has continued to find industrial research support. This work, for instance, was supported in part by Northrop Grumman Corporation. Interest in NMR gyroscopes [Walker and Larsen, 2016; Kornack et al., 2005; Jiang et al., 2018; Karwacki, 1980; Zhang et al., 2020] has been driven in part by demand for navigation systems which can operate in environments where communication with GPS satellites is limited or unavailable. The low power consumption and miniaturize-ability of these devices also contributes to their appeal in industrial applications.

Continuous operation is preferable for inertial navigation applications, and so the transverse NMR gyroscope presented in this work was developed to allow for continuous operation. The pulsed free-induction decay (FID) comagnetometers presented in [Bulatowicz et al., 2013; Limes et al., 2018] use a pulsed transverse field to excite noble gas precession and then allow the noble gas polarization to decay. FID comagnetometers do not run continuously, and are therefore more suited to exotic physics searches than gyroscope applications.

Another approach to continuous drive are NMR devices which operate about a compensation point [Kornack and Romalis, 2002; Kornack et al., 2005; Brown et al., 2010; Terrano and Romalis, 2021], which have demonstrated excellent short-term noise performance. However, the scale factor (aka, the conversion factor between between the measured signals and the desired output) of these self-compensating devices depends on experimental factors, and therefore must be calibrated. In comparison, the NMR gyroscope presented in this work has a scale factor which depends only on known physics constants, as shown in Eq. 1.5.

1.3 Novel Features of This Device

The NMR gyroscope presented in this work is unique in the field primarily because of its transverse geometry. This refers to the fact that the optical pumping is transverse to the bias field direction. NMR devices typically utilize longitudinal geometry, wherein the optical pumping and bias field are along the same axis. In this section, I will discuss several longitudinal approaches to NMR gyroscopes, and thereby highlight some of the benefits of the transverse NMR gyroscope presented in this work: namely, the suppression of systematic errors from longitudinal polarization and quadrupole shifts.

For comparison, consider the Northrop Grumman style continuous longitudinal NMR gyroscope in [Walker and Larsen, 2016], which uses a transverse RF field to drive noble gas precession. This device, like other longitudinal NMR devices, suffers from systematic errors due to longitudinal polarization. In longitudinal geometry, the alkali-metal electron

polarization and noble gas nuclear polarizations are predominantly along the bias field direction. This means that the spin-exchange fields present in the system contribute significantly to the bias field, with the spin-exchange field experienced by each species being slightly different (and therefore not completely removed by comagnetometry). In this way, longitudinal polarization causes systematic errors in longitudinal NMR comagnetometers. In the transverse NMR gyroscope, the alkali-metal electron polarization and noble gas nuclear polarizations are predominantly transverse to the bias field direction. Therefore, spin-exchange fields do not contribute significantly to the bias field, and systematic errors from longitudinal polarization are avoided.

As previously mentioned, our NMR gyroscope uses ^{131}Xe and ^{129}Xe as the two noble gas species. This choice is beneficial due to the very similar Fermi-contact enhancement factors for ^{85}Rb - ^{131}Xe and ^{85}Rb - ^{129}Xe . This means that the effective gain of our embedded Rb magnetometer is the same for both Xe species. We are certainly not the only group to have utilized these two Xe isotopes (for example, [Zhang et al., 2017; Walker and Larsen, 2016]), but it does present an additional complication that our device is uniquely well-suited to address. ^{131}Xe is spin 3/2, meaning that it is subject to frequency shifts due to quadrupole interactions from electric field gradients on the cell walls. Since ^{129}Xe is spin 1/2 and therefore not subject to quadrupole shifts, the rotation obtained by comagnetometry will have systematic error from the quadrupole effect. Here the transverse geometry helps us once again. As discussed in [Korver, 2015], we have not seen evidence of first-order quadrupole shifts in our transverse NMR gyroscope. We postulate that this is a direct result of transverse pumping. Our hypothesis to explain this insensitivity is as follows. Without the quadrupole interaction, we have four evenly spaced m levels. Quadrupole interactions shift two levels up, and two down. This results in one larger transition, one that has stayed the same, and one that has gotten smaller. For longitudinally pumped NMR, the population of the four states is like an upside-down pyramid, and you end up with the larger transition being driven much more than the smaller transition, so you have a net shift up. For transverse pumping,

we have a symmetric population between upper and lower states. Then the smaller and larger transitions are driven equally, and on average there is zero shift. The suppression of systematic error from first-order quadrupole shifts is another benefit of our transverse NMR gyroscope.

1.4 Outline of This Thesis

Prior to the work of this thesis, Anna Korver presented work on this device demonstrating the suppression of systematic errors in a transverse NMR gyroscope with pulsed bias field, and outlining the theory behind simultaneous dual species excitation [Korver et al., 2015]. This thesis will present results from three methods for continuous dual-species excitation of Xe precession in a transverse NMR gyroscope. Chapter 4 presents polarization modulation [Sorensen et al., 2020], Ch. 5 presents modulations of the pulsed bias field [Thrasher et al., 2019b], and Ch. 6 presents a hybrid approach, combining both polarization modulation and pulse density modulation [Sorensen and Walker, 2023]. The work presented in Ch. 4 and part of Ch. 5 was done in collaboration with my coworker Daniel Thrasher, and readers may look to his thesis [Thrasher, 2020] for details I may not have included. The hybrid approach may allow us to take advantage of the benefits of both PM and PDM to potentially improve signal detection and/or systematic errors. In order to properly discuss the results obtained for PM, PDM, and hybrid operation, I must first discuss the theory and operation of our apparatus.

In Ch. 2, I will derive Bloch Equation solutions for the noble gas nuclei and alkali-metal electron polarizations (\mathbf{K} and \mathbf{S} , respectively) for our transverse NMR gyroscope. I will go through the the derivation generically such that the results can be applied to PM, PDM, or hybrid operation in their respective chapters. Over the course of the derivation, I will derive and define several variables of note throughout this thesis. These include the Rb z-polarization (S_z) used for our detection, the magnetometer phase shift (ϵ_z) which can

introduce errors in our measure of rotation, the steady state Xe amplitudes (K_{\perp}), the Xe phase shift (δ). I will derive an expression for open-loop comagnetometry, and discuss three assumptions in this expression that would lead to errors, and must therefore be addressed in our performance of comagnetometry for PM, PDM, and hybrid operation. Finally, I will discuss the performance metrics used to evaluate the performance of each mode of operation. Several of these metrics (signal-to-noise ratio (SNR), linewidth, bias instability, Angle Random Walk (ARW)) are standard to report in papers discussing NMR comagnetometers. I will make a case for widespread inclusion of another metric, the field suppression factor (FSF), which quantifies how well common fields are suppressed in measurements of rotation or other non-magnetic, spin-dependent effects.

I will present the experimental apparatus used for this work in Ch. 3. First, I will describe the vapor cell used in this work, as well as the heating elements and housing for the cell. I will then move on to discuss the aspects of our apparatus relating to magnetic fields. This includes the magnetic shielding, three-axis shimming coils, and pulsing coils along z to provide our bias pulses. I will present the development and application of a new circuit used to drive the pulsing coils which provides pulse areas which are 100-fold less dependent on pulsing frequency than the previous driving circuit. Next, I will discuss the optical elements and lasers used in the apparatus, and discuss how these optics were changed between PM, PDM, and hybrid operation. I will then present measurements of the noise of our system. Finally, I will briefly discuss the digital systems used for control and measurement.

After discussing the foundations of the experiment in chapters 2 and 3, I will move on to discussing the major results in chapters 4, 5, and 6. In each of these three chapters, I will discuss the performance metrics as outlined in Ch. 2. These chapters are ordered chronologically in terms of our experimental progress, starting with PM, then PDM, and finally hybrid operation. As such, there are some things presented in the earlier chapters which were later improved upon. For example, we developed a demodulation scheme for PDM that we believe would greatly improve our PM results if we were to go back and apply

it to PM.

Chapter 4 discusses Polarization Modulation operation, wherein the pulsed bias field is applied with constant frequency, but the handedness of circular polarization of the light used to optically pump the Rb is modulated as a phase-modulated square wave. We "rectify" the detected S_z by multiplying it by the optical modulation waveform, and then use lock-in amplifiers to obtain the Xe phase shifts. For simultaneous excitation of both Xe species, we obtain linewidths of 14 mHz for ^{129}Xe and 17 mHz for ^{131}Xe , and amplitudes of approximately 66 μG for ^{129}Xe and 78 μG for ^{131}Xe . The SNR for each isotope is $\sim 5000\sqrt{\text{Hz}}$. I will present a feedback method used to improve the comagnetometry. When we were using PM operation, we were not yet explicitly measuring the FSF, and so we can only give a lower-limit estimate of ~ 240 for the FSF based on open- and closed-loop measurements. We demonstrate an ARW of $\sim 16 \mu\text{Hz}/\sqrt{\text{Hz}}$ and a rotation bias instability of ~ 800 nHz. I will then show that our chosen demodulation results in significant signal mixing between our detection channels, which we refer to as "cross talk." This was the primary motivation for transitioning to PDM operation.

Chapter 5 discusses Pulse Density Modulation, wherein the optical pumping is constant in time, but the pulsed bias field is modulated. I will discuss how modulating the bias pulses necessitates changes to our detection scheme, and the implementation of a new "gated" detection. I will discuss a new demodulation method, matrix inversion least-squared fitting, which practically eliminated cross talk. For simultaneous, dual-species Xe excitation, we obtained ^{131}Xe and ^{129}Xe linewidths of 19 mHz and 11 mHz (respectively), and amplitudes of 16 μG and 14 μG (respectively). I will discuss a feature of the demodulation that initially limited the SNR. After addressing the issue, we obtained SNRs of about $2700 \sqrt{\text{Hz}}$ for both Xe species. I will describe the feedback scheme used for PDM comagnetometry, resulting in $\text{FSF} = 4700$. In order to contextualize the $\text{ARW} = 17 \mu\text{Hz}/\sqrt{\text{Hz}}$ and bias instability of 420 nHz, I will discuss the system's sensitivity to transverse fields, which was part of the motivation for the transition to hybrid operation.

Chapter 6 discusses hybrid modulation, wherein both the optical pumping and pulsed bias field are modulated. I moved to this mode of operation in hopes of gaining some of the benefits of both PDM (low-cross-talk demodulation, good bias instability) and PM (large Xe field sizes, AC Rb polarization). It was our hope that moving the Rb polarization back to AC would reduce the transverse field sensitivity. I will discuss the changes to our modulation, demodulation, and feedback necessary for hybrid operation. For simultaneous, dual-species excitation, I measure linewidths of 14 mHz and 16 mHz for ^{129}Xe and ^{131}Xe , respectively, with peak field sizes of 26 μG for ^{129}Xe and 11 μG for ^{131}Xe . Unfortunately, we still observe transverse field sensitivity. We obtain a field suppression factor of ~ 1400 , angle random walk of 21 $\mu\text{Hz}/\sqrt{\text{Hz}}$, and bias instability of ~ 480 nHz.

This thesis will conclude in Ch. 7 with a summary of the presented works. I will compare PM, PDM, and hybrid operation. Then I will discuss my recommendations for future experiments using a transverse NMR gyroscope apparatus.

Chapter 2

Theory

In this chapter, I define the geometry of a transversely-pumped NMR gyroscope, and discuss the effects of this geometry on sources of systematic error. I derive the behavior of the noble gas nuclear and Rb electron polarizations using the Bloch Equations. I show how the Rb electron longitudinal polarization can be used as a monitor of the Xe precession, and how stray magnetic fields along the longitudinal direction can cause a magnetometer phase shift in this readout.

Xe can be continuously excited using pump light polarization modulation (PM), bias field pulse density modulation (PDM), or a combination of the two (hybrid modulation). For the derivations in this chapter, the modulation scheme will be left undefined. Results specific to PM, PDM, and hybrid modulation will be shown in their respective chapters. I will present a modulation-generic implementation of comagnetometry to obtain a measure of rotation, and discuss some of the complications that must be addressed in the rest of this thesis to improve the measure of rotation.

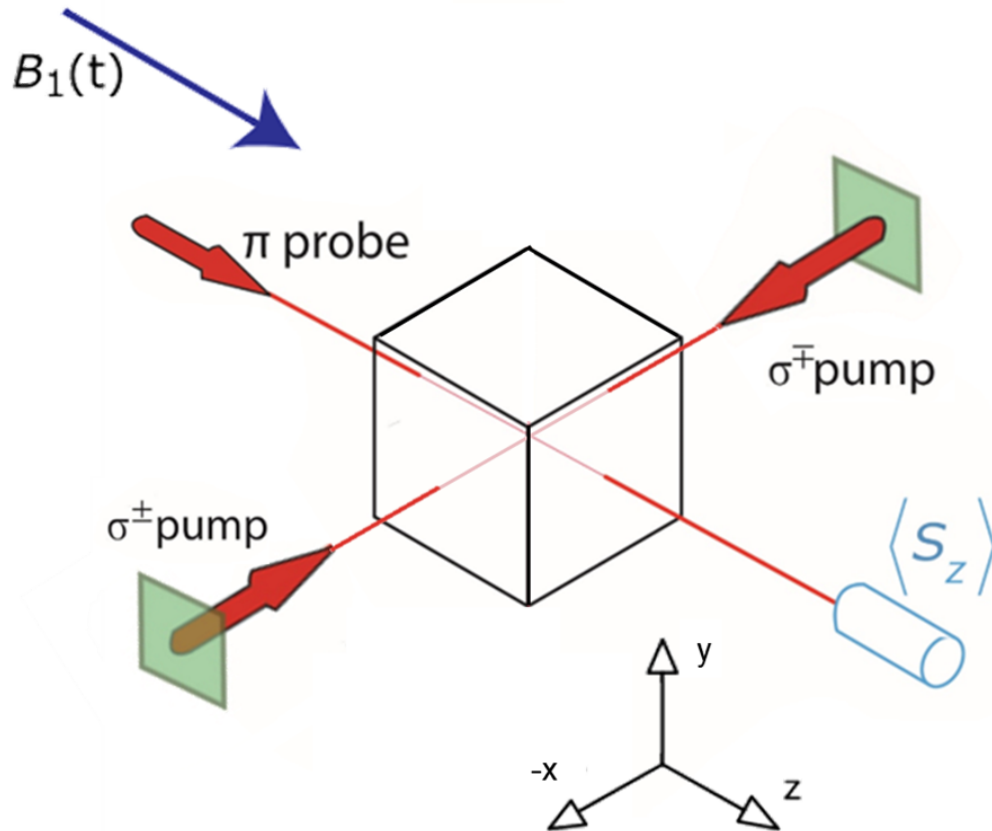


Figure 2.1: A simplified diagram of the transverse SEOP setup, showing the geometry of pump and probe lasers and bias field.

2.1 Geometry and Systematics

Throughout this thesis, I will make assumptions and coordinate assignments based on the geometry of our transversely pumped NMR gyroscope (see fig. 2.1). I assign the longitudinal direction to the z -axis. This is the axis along which our pulsed bias field is applied, and along which the probe laser propagates. The linearly polarized probe will therefore experience a Faraday rotation proportional to the z -polarized Rb (S_z). The pump light propagates along the x -axis, and therefore optical pumping polarizes the Rb electrons along \hat{x} . The Xe polarization then precesses primarily in the transverse (x - y) plane.

As a result of this transverse geometry, we suppress several systematic noise sources.

Since the Rb is polarized transverse to the bias field, and since we apply a "compensation" field to cancel the spin-exchange (SE) field produced by the polarized Rb [Korver, 2015], the Rb SE field does not contribute significantly to the bias field experienced by the Xe. While ^{131}Xe is subject to quadrupole interactions, we have found that transverse pumping suppresses quadrupole shifts to first order [Korver, 2015; Thrasher et al., 2019a].

2.2 The Bloch Equation

Because the Xe and Rb are co-located in our vapor cell, the Rb polarization (\mathbf{S}) will affect the Xe polarization (\mathbf{K}), and vice-versa, through spin-exchange magnetic fields and spin-exchange collisions. We assume the collisional spin polarization transfer from Rb to Xe is significant, but the back polarization from Xe to Rb is negligible [Limes et al., 2018]. We assume that \mathbf{K} precesses slowly enough that S_z has an adiabatic response. The two Xe isotopes, ^{129}Xe and ^{131}Xe , will be labeled species a and b, respectively.

There are many magnetic fields present in the system, but not all species experience the same fields. I use the phrase "experience" to emphasize when we can ignore fields that do not significantly influence the precession of a given species. For example, because the bias pulses cause 2π Rb rotations, the time-average Rb polarization is insensitive to the pulsed bias field. Spin-exchange fields are also not experienced equally by every species, where $b_K^S K$ is the spin-exchange field due to K experienced by S and $b_S^K S$ is the spin-exchange field due to S experienced by K . When assigning symbols for the various fields in the system, subscripts will generally be used to identify the source or direction of a field. Superscripts may be used to specify which species is experiencing a given field, but will often be omitted since it can generally be inferred from context. If I am deriving an equation for the Rb polarization, one can assume the fields (and the resulting precession frequencies, Ω) are from the fields experienced by the Rb. Similarly, when deriving expressions for the Xe precession, we will be concerned with the fields experienced by the Xe.

Rb Bloch Equation

The Rb electron polarization, \mathbf{S} , is utilized in two ways in our system. The transverse polarization, S_+ , is used to polarize the Xe nuclei through spin-exchange collisions. The longitudinal polarization, S_z , is used to monitor the precession of the Xe.

We begin with the time-average solution to the Bloch equation for the Rb polarization,

$$\mathbf{S} = \frac{\mathbf{R}\Gamma' + \boldsymbol{\Omega} \times \mathbf{R} + \boldsymbol{\Omega}(\boldsymbol{\Omega} \cdot \mathbf{R})/\Gamma'}{\Gamma'^2 + \boldsymbol{\Omega}^2} \quad (2.1)$$

where \mathbf{R} is the optical pumping rate, Γ' is the total relaxation rate (including optical pumping), and $\boldsymbol{\Omega}$ is the total precession frequency of the Rb. For $\boldsymbol{\Omega} \ll \Gamma'$, 2.1 can be expanded as

$$\mathbf{S} = \frac{\mathbf{R}}{\Gamma'} + \frac{\boldsymbol{\Omega} \times \mathbf{R}}{\Gamma'^2} + \frac{\boldsymbol{\Omega} \times (\boldsymbol{\Omega} \times \mathbf{R})}{\Gamma'^3} + \dots \quad (2.2)$$

In order to obtain S_+ and S_z from this expression, we take into consideration the geometry of our system. Our optical pumping is along \hat{x} such that $\mathbf{R} = R(t)\hat{x}$. The Rb spins precess as $\boldsymbol{\Omega} = \gamma\mathbf{B} + \mathbf{X}$, but we can assume $X \ll \gamma^S B$ due to the large gyromagnetic ratio of Rb compared to relatively small expected rotations. We expect fields of order tens of μG and γ^S is of order MHz/G . In comparison, if we were to measure Earth's rotation, it is about $10 \mu\text{Hz}$ in Madison. The precession of the Rb can then be written

$$\boldsymbol{\Omega} = \gamma^S(\mathbf{B}_{\text{ext}} + b_a^S \mathbf{K}^{\mathbf{a}} + b_b^S \mathbf{K}^{\mathbf{b}}) \quad (2.3)$$

where $b_K^S K$ is the spin-exchange field due to K experienced by S . Note that the Rb effectively does not precess about the pulsed field because each pulse results in a 2π radian precession, bringing it back to where it started. We can then perform the relevant cross products from

Eq. 2.2 to obtain

$$S_+ = \frac{R(t)}{\Gamma'} + i \frac{R(t)\Omega_z}{\Gamma'^2} = \frac{R(t)}{\Gamma'} e^{i\epsilon_z}, \quad (2.4a)$$

$$S_z = -\frac{R(t)}{\Gamma'^2} (\Omega_y - \frac{\Omega_z}{\Gamma'} \Omega_x) = \frac{-R(t)}{\Gamma'^2} \text{Im}[\gamma^S b_K^S K_+ e^{-i\epsilon_z}]. \quad (2.4b)$$

We define a magnetometer phase shift

$$\epsilon_z = \tan^{-1}\left(\frac{S_y}{S_x}\right) \equiv \tan^{-1}\left(\frac{B_z^S}{B_w}\right) \ll 1 \quad (2.5)$$

where B_z^S is the field experienced by the Rb and $B_w = \Gamma'/\gamma^S$ is the magnetic width of the magnetometer. When the Rb experiences a field along \hat{z} , the polarization tips from \hat{x} into the x-y plane. As shown in Eq. 2.4, this will affect both the polarization of the Xe (via S_+) and the detection of the Xe precession (via S_z). This change in the angle of the polarized Rb from a z-field is difficult to distinguish from an actual rotation of the device, and I will show in Sec. 5.5 and Sec. 6.5 how we limit the influence of ϵ_z in our rotation measurements.

Equation 2.4b shows that S_z is dependent on K_+ , hence S_z can be used to monitor the Xe precession in the transverse plane. The exact form of S_z depends on the modulation scheme used. Because of the Fermi contact enhancement factor between Rb and Xe, the Xe fields we measure by using the Rb as a magnetometer are about a factor of 500 larger than what we could detect using pick-up coils [Nahlawi et al., 2019].

I will now use the Rb results (Eq. 2.4) in the Xe Bloch Equations. In this chapter, the Xe Bloch Equation derivation will be left generic. When necessary, specific results for PM, PDM, and hybrid operation will be given in their respective chapters.

Xe Bloch Equation

The Xe polarization behavior is dominated by three terms: Larmor precession, relaxation, and spin-exchange collisions with the Rb. The Bloch equation for the Xe can be therefore

written

$$\frac{d\mathbf{K}}{dt} = -(\mathbf{s}\Omega) \times \mathbf{K} - \Gamma\mathbf{K} + \Gamma_S^K \mathbf{S}, \quad (2.6)$$

where Γ is the total relaxation rate, and Γ_S^K is the spin-exchange pumping rate from the Rb electrons to the Xe nuclei. The Xe precession is $\mathbf{s}\Omega = \mathbf{s}\gamma(\mathbf{B}_p(t) + \mathbf{B}_{\text{ext}} + b_S^K \mathbf{S} + b_{K'}^K \mathbf{K}') + \mathbf{X}$, where $b_{K'}^K \mathbf{K}'$ is the spin-exchange field experienced by one species of Xe (\mathbf{K}) due to the other Xe species (\mathbf{K}'). Though I mention it here, the spin-exchange fields between the two Xe species are small enough that they have not yet been relevant to our experiment. The term $b_{K'}^K \mathbf{K}'$ is therefore assumed to be negligible.

We separate Eq. 2.6 into transverse ($K_\perp = K_x + iK_y$) and longitudinal (K_z) components

$$\frac{dK_+}{dt} = -(si\Omega_z + \Gamma_2)K_+ + \Gamma_S^K S_+ + si\Omega_+ K_z, \quad (2.7a)$$

$$\frac{dK_z}{dt} = \mathbf{s}(\Omega_y K_x - \Omega_x K_y) - \Gamma_1 K_z, \quad (2.7b)$$

where Γ_1 and Γ_2 are the longitudinal and transverse relaxation rates, respectively. We use shim coils and our compensation field to null the transverse fields experienced by the Xe. This suppresses K_z such that K_+ is much more sensitive to Ω_z than Ω_+ .

It is useful to proceed with an eye toward how the Xe behavior will be measured in our real-life system. While details of our demodulations will be explored in the chapters for each operation mode, it is useful at this point to note that the demodulations will eventually result in one term proportional to the amplitude of the Xe signal, and one term proportional to the phase of the precessing Xe signal. Therefore, we will now make the substitution $K_+ = K_\perp e^{-si\phi}$ where K_\perp and ϕ are the amplitude and phase of the transverse Xe precession.

Equation 2.7a then becomes

$$\frac{dK_\perp}{dt} - si\frac{d\phi}{dt}K_\perp = -(si\Omega_z + \Gamma_2)K_\perp + \Gamma_S^K S_\perp e^{i\epsilon_z} e^{si\phi} \quad (2.8)$$

where $S_{\perp} = R(t)/\Gamma'$ (see Eq. 2.4a).

Up to this point, the derivations are generic for any field ($B_p(t)$ in Ω_z) or pumping ($R(t)$ in S_{\perp}) modulations. In order to proceed, we will make assumptions about the form of the applied modulations. For a given $R(t) = R * R_{mod}(t)$ (where R is constant and $R_{mod}(t)$ is the modulation), we assume our polarization modulations can be written as a Fourier series,

$$S_{\perp} = \frac{R}{\Gamma'} \sum_n \sum_p s_{p,n} e^{ip_n \omega_{PM,n} t} \quad (2.9)$$

where p and n are integers, s_p are Fourier coefficients, and $\omega_{PM,n}$ is the frequency of the n -th polarization modulation. In the case of no PM, $\omega_{PM,n} = 0$ and $s_{p,n} = 1$. For square wave PM (which is used in this thesis), p must also be odd. We assume any field modulations are applied as pulse density modulations, and that the pulsed field has the form

$$B_p(t) = B_{p0} * \left(1 + \sum_m b_m \cos(\omega_{PDM,m} t)\right) \quad (2.10)$$

where B_{p0} is the average field from the pulses (aka the center field around which the modulations are applied), m is an integer, $b_m = B_m/B_{p0}$ is the modulation depth (with B_m the field amplitude of the m -th modulation), and $\omega_{PDM,m}$ is the frequency of the m -th pulse density modulation. In the case of no PDM, $b_m = 0$. By defining our field and polarization modulations in this way, the following derivation will be valid for all of the modulation schemes presented in this thesis (PM, PDM, and hybrid).

Note that the field modulation will enter into Eq. 2.8 in the term Ω_z^K , but also in the Xe precession phase ϕ . To make this dependence more explicit, we will make the substitution $\phi = \delta + \alpha$, where δ is the Xe phase shift and α is the expected Xe precession phase based only on the applied modulations. We can write

$$\alpha = \int (\omega_d + \gamma B_{mod}(t)) dt, \quad (2.11)$$

where ω_d is the frequency at which we drive the Xe, and $B_{mod}(t) = B_p(t) - B_{p0}$ is the modulated part of the field. If the Xe were driven perfectly on resonance and there were no rotation or fields other than an ideal pulsed field, δ would be zero.

In order to find the transverse Xe polarization amplitude (K_{\perp}), we take the real part of Eq. 2.8 and solve for the steady state ($dK_{\perp}/dt = 0$). Using Eq. 2.2, the last term of Eq. 2.8 will have exponential factors

$$\sum_n \sum_p e^{ip_n \omega_{PM,n} t} e^{-si\alpha} e^{i(-s\delta + \epsilon z)}. \quad (2.12)$$

In order to find the steady state solution of the real part of these exponential factors, it is useful to consider the time-average of $\sum_n \sum_p e^{ip_n \omega_{PM,n} t} e^{-si\alpha}$. Given the form of our field modulation in Eq. 2.2, we can write $\alpha = \omega_d t + \gamma \sum_m \frac{B_m}{\omega_{PDM,m}} \sin(\omega_{PDM,m} t)$. Then we can use the Jacobi-Anger identity

$$e^{iz \sin(\theta)} = \sum_q J_q(z) e^{iq\theta} \quad (2.13)$$

where J_q is the q -th Bessel function of the first kind. We then keep only the terms that mix to DC (which eliminates the sums over p and q). This leaves us with a Xe resonance condition

$$-s\omega_d = \sum_n p_n \omega_{PM,n} + \sum_m q_m \omega_{PDM,m} \quad (2.14)$$

which must be satisfied for $\omega_d = \Omega^K$ in order to continuously excite Xe. We call our choices of p 's and q 's our "drive scheme," and there are many drive schemes that satisfy Eq. 2.2.

The resonance condition dictates that the remaining exponents will all sum to zero. If we then use Euler's formula to convert from exponentials to trig functions, the argument of the sine and cosine must be zero. Taking the resulting integral over the range $(0, T)$, we arrive at

$$\overline{e^{ip_n \omega_{PM,n} t} e^{-si\alpha}} = \sum_m \frac{J_{q_m} \left(\frac{\gamma B_m}{\omega_{PDM,m}} \right)}{T} (T + i * 0) \quad (2.15)$$

Taking the real and imaginary parts of each side of the equation, we arrive at the time averages $\overline{\cos(p_n\omega_{PM,n} + \mathbf{s}\alpha)} = \sum_m J_{q_m}(\frac{\gamma B_m}{\omega_{PDM,m}})$ and $\overline{\sin(p_n\omega_{PM,n} + \mathbf{s}\alpha)} = 0$.

We can now solve for the steady state transverse Xe amplitudes using the real part of Eq. 2.8. Taking $dK_{\perp}/dt = 0$ and assuming $\delta + \mathbf{s}\epsilon_z$ is small, we find the steady state amplitudes to be

$$K_{\perp} = \frac{\Gamma_S^K R \sum_n s_{p,n} \sum_m J_{q_m}(\frac{\gamma B_m}{\omega_{PDM,m}})}{\Gamma\Gamma_2}. \quad (2.16)$$

In subsequent chapters (Ch. 4, 5, 6), I will consider the product of the Fourier coefficients, $j = \sum_n s_{p,n} \sum_m J_{q_m}(\frac{\gamma B_m}{\omega_{PDM,m}})$, and compare the influence of different drive schemes on the expected steady-state Xe amplitudes (see Ch. 7).

We next find an expression for the Xe phase shift, δ , by taking the imaginary part of Eq. 2.8 and substituting $\phi = \delta + \alpha$. Using our steady state solution (Eq. 2.2) and our previous results for time-averages of sine and cosine, the time dependence of the Xe phase shift (for $\delta, \epsilon_z \ll 1$) simplifies to

$$\frac{d\delta}{dt} = -\Delta_0 - \Gamma_2(\delta + \mathbf{s}\epsilon_z), \quad (2.17)$$

where $\Delta_0 = \omega_d - \gamma B_{z,0} - \mathbf{s}X^K$ and $B_{z,0}$ is the total z-field minus the modulated part (B_{mod}). In the Fourier domain, the Xe phase shift noise is

$$\tilde{\delta} = -\frac{\tilde{\Delta}_0 + \mathbf{s}\Gamma_2\tilde{\epsilon}_z}{\Gamma_2 + i\omega}, \quad (2.18)$$

where the notation $\tilde{f} = f(\omega)$ is used.

2.3 Comagnetometry

We perform comagnetometry to obtain a measure of X , which in the context of a gyroscope is the rotation frequency $\tilde{\Omega}^R$. Given measurements of δ^a , δ^b , and Γ_2 , and assuming ϵ_z is

negligible, we can eliminate field noise and solve for $\tilde{\Omega}^R$,

$$\frac{\rho\tilde{\delta}^b(i\omega + \Gamma_2^b) - \tilde{\delta}^a(i\omega + \Gamma_2^a)}{(1 + \rho)} = \tilde{\Omega}^R. \quad (2.19)$$

Unfortunately, every one of those assumptions introduces errors in our measure of rotation. First off, we do not directly measure δ^a and δ^b . We must demodulate the measured Faraday rotation signal to obtain a phase shift measurement. In the chapters to follow, we will show that our demodulations do not output δ^a and δ^b exactly. Second, we assumed knowledge of Γ_2 for each species. While we can make independent measures of the average Γ_2 , we do not measure it during our rotation measurements. Therefore, any noise on Γ_2 will be mapped onto our calculated rotation. Third, we assumed ϵ_z is negligible. We have shown in ref. [Thrasher et al., 2019a] that ϵ_z is not, in fact, negligible, and can cause field noise to be mapped onto our calculated rotation.

In the chapters for dual PM, dual PDM, and hybrid modulation, we will show how we address each of these complications through demodulation methods and feedback [Bechhoefer, 2005]. Using more sophisticated demodulation methods allows us to be more confident about what exactly we are measuring, which we can take into account in our analysis. Performing feedback to the z-field minimizes the effect of ϵ_z . Using the measured phase shifts to feed back to one or more modulation frequency can allow us to bypass the need for Γ_2 entirely. For each of the three modulation schemes presented in this thesis, we will discuss how we address (or don't) each of these complications to arrive at a better measure of rotation noise.

2.4 Performance Metrics

Once we have obtained a measure of rotation, we must quantify the performance of our gyroscope in order to compare it to other NMR gyroscopes, or to compare our different modulation schemes to each other. The following are metrics of interest.

The Signal-to-Noise ratio (SNR) is obtained from a phase noise measurement, where $SNR = 1/\bar{\delta}$ (with units of \sqrt{Hz}). The average phase noise is measured at a high enough frequency so as to not be limited by 1/f magnetic field noise (generally over the range 0.2 - 1 Hz). The SNR sets a limit on how small of Xe precessions (and hence rotations) we are able to distinguish from noise.

The Xe linewidths (Γ_2) are obtained from Lorentzian fits to measurements of the Xe in-phase and quadrature signals as the drive frequency is scanned through the Xe resonance. A smaller linewidth corresponds to greater sensitivity to the Xe detuning (aka, how far from resonance the Xe is being driven). As shown in [Korver, 2015], applying an x-compensation field allows us to narrow the linewidths to the Γ_1 (longitudinal relaxation) limit.

Cross talk is a measure of how well the demodulation isolates independent signals for each Xe species. Ideally, demodulation accurately separates the contributions of the two Xe isotopes to our S_z signal. In practice, it is possible that some signal content from ^{131}Xe would show up on the demodulation channels for ^{129}Xe , and vice-versa. Such signal mixing is referred to as "cross talk." Suppose there exists cross talk in both channels such that $\omega^a = \gamma^a B_z + \beta\omega^b - \Omega^R$ and $\omega^b = \gamma^b B_z + \beta'\omega^a + \Omega^R$ where β and β' represent the fraction of one species that shows up on the detection channel for the other species. Solving for Ω^R we find

$$\Omega^R = \frac{\omega^b(\rho + \beta') - \omega^a(\rho\beta + 1)}{1 + \rho}, \quad (2.20)$$

where if $\beta = \beta' = 0$ we return to the expected expression for rotation without cross talk. Non-zero cross talk is undesirable because the accuracy with which it is known (or measured) limits the accuracy of conversion from measured precession frequencies to rotation (or any other non-magnetic spin-dependent interaction). A measurement of cross talk is vital since an important alleged feature of our comagnetometer is that it has a scale factor which is determined solely by physics constants. The methods for measuring the cross talk coefficients β and β' are discussed in the chapters for PM, PDM, and hybrid operation. The fundamental

idea behind these measurements is to drive only one Xe species on resonance while observing the detection channels for both species. The cross talk coefficients are then the ratio of the signal on the off-resonant species' detection channel to the signal of the on-resonant species' detection channel.

The field suppression factor (FSF) is a measure of how well comagnetometry (and field stabilization, when present) suppresses magnetic field noise on our measure of rotation. To obtain FSF, we apply an AC magnetic field of known size (B_{app}) at frequency ω_{app} along \hat{z} . If we were using a single species of Xe, we would expect the applied field to result in oscillations of the Xe Larmor resonance frequency as γB_{app} . We could use either Xe species for this calculation. We choose to use our less magnetically sensitive species, ^{131}Xe , as it would give a less-field-sensitive measure of rotation if used alone. We compare this to the oscillations present on a measure of rotation at the frequency of the applied field. We define $\text{FSF} = \gamma^b B_{app} / \Omega^R(\omega_{app})$ (unitless). Since this factor quantifies how much field noise can be expected on a measurement of rotation, we believe FSF is an important quantity to report when presenting results of an NMR gyroscope.

Our remaining metrics, Angle Random Walk (ARW) and bias instability, are obtained from Modified Allan Deviations (MDEVs) of the rotation frequency [Allan and Barnes, 1981]. We use MDEVs rather than Allan Deviations (ADEVs) because MDEVs can distinguish between white phase modulation and flicker phase modulation noise types, and averages away these noise types faster than a standard ADEV. Since we are not limited by white phase modulation or flicker phase modulation, an MDEV allows us to see our limiting noise sources at lower integration times (τ). ARW is found from the slope of a $\tau^{-1/2}$ trend multiplied by $\sqrt{2}$ (where $\sqrt{2}$ is a scaling factor that results from using MDEVs instead of ADEVs). Lower ARW is desirable, but the lowest we can expect is the SNR-limited $\text{ARW} = \Gamma_2 / \text{SNR}$. ARW is reported in units of $\text{Hz} / \sqrt{\text{Hz}} = 21600 \text{ deg/Hr}$. At some point, the MDEV stops following a $\tau^{-1/2}$ trend and levels out or turns up. The lowest point on the MDEV is referred to as the bias instability, reported in units of Hz. We also report the integration time at which we

reach our bias instability. Improving this metric would correspond to reaching a smaller bias instability, or reaching the same bias instability at a shorter integration time (often as a result of improving the ARW). Reaching our bias instability corresponds to the point where we are limited by a noise trend other than ARW.

2.5 Chapter Summary

In this chapter, we have defined the geometry of our transverse spin-exchange optically pumped NMR gyroscope, and discussed some of the effects of that geometry in suppressing systematic noise sources. We have started derivations for the time dependence of the Xe and Rb polarizations, which will be completed in chapters 4, 5, and 6. We have discussed comagnetometry, and why various assumptions would degrade a measure of rotation. We have defined the performance metrics that will be reported in the chapters for each modulation scheme.

Chapter 3

Apparatus

In this chapter, I will describe the experimental apparatus used to obtain the data presented in this work. I will discuss the vapor cell used, and how the vapor cell is heated and mounted into the system. I will then discuss the magnetic shields and the coils and circuits used to apply the pulsed bias field and shim fields. I will then discuss the pump and probe lasers, and the optical setup of our apparatus. Finally, I will touch on the digital system which is used to control the experiment. While PM, PDM, and hybrid operation all used fundamentally the same apparatus, some changes were made between the different modes of operation. Throughout this chapter, I will explicitly highlight differences in the apparatus between PM, PDM, and hybrid operation.

3.1 Vapor Cell

An 8 mm inner diameter cubic Pyrex cell filled with 40 Torr enriched Xe (about 90 percent ^{131}Xe and 10 percent ^{129}Xe), ^{85}Rb , 50 Torr N_2 , and 5 Torr H with a hydride coating was prepared and provided by Northrop Grumman Corp. The hydride coating enables longer T_2 times by limiting wall relaxation [Kwon et al., 1981; Kwon and Volk, 1984]. The cell is mounted in a ceramic housing with holes for laser light to enter the cell. The ceramic

housing has four symmetric faces which fit together like a jigsaw puzzle (and so we call these “jig” heaters). Conductive traces are printed on each face, through which AC current at ~ 150 kHz (produced by a function generator and amplified by an audio amplifier) is passed to heat the ceramic. The jig heaters were manufactured by Thick Film Technologies using high-solids, low-magnetic-impurity platinum thick film material on an aluminum oxide wafer, and were designed in accordance with U.S. Patent 8,138,760. The conductive traces are arranged to produce minimal stray magnetic fields (including gradient magnetic fields). An RTD is pressed against the glass of the cell to monitor the cell temperature. The RTD is pressure-held in place between the cell and a piece of insulation around the cell stem (see Fig. 3.2). The cell is heated in the range 108-119°C, corresponding to Rb vapor pressures of around $1.7\text{-}3.4 \times 10^{13}$ cm^{-3} [Lancor, 2011]. The procedure for measuring the Rb density can be found in [DeLand, 2017]. We chose the operating Rb density such that the optical pumping results in about 50% Rb polarization. For PM operation, the temperature was open-loop. For PDM and hybrid operation, temperature stabilization was performed as follows. Using the RTD as one of the resistors in a Wheatstone Bridge circuit [Horowitz and Hill, 2015], we can apply feedback [Bechhoefer, 2005] to the amplitude of the AC signal sent through the jig heaters, thereby stabilizing temperature drifts as shown in Fig. 3.1. The feedback gain is an inverted zero plus lead compensation of the form

$$G = P \left(1 + \frac{\omega_0}{i\omega} \right) \left(\frac{1 + A \frac{i\omega}{\omega_1}}{1 + \frac{i\omega}{\omega_1}} \right). \quad (3.1)$$

The ~ 1 mm gap between the vapor cell and ceramic heating jig is shimmed with a 1.5 mm thermally conductive and slightly compressible gap fill (model: TG 977, manufacturer: T-Global Technology, see Figure 3.2). The ceramic is wrapped with aerogel (a high temperature insulating material) and secured with Kapton tape and then fitted into a 3D printed (high temperature nylon) cartridge with holes to allow laser light to enter the cell. The compressible nature of the aerogel produces a friction fit keeping the ceramic jig

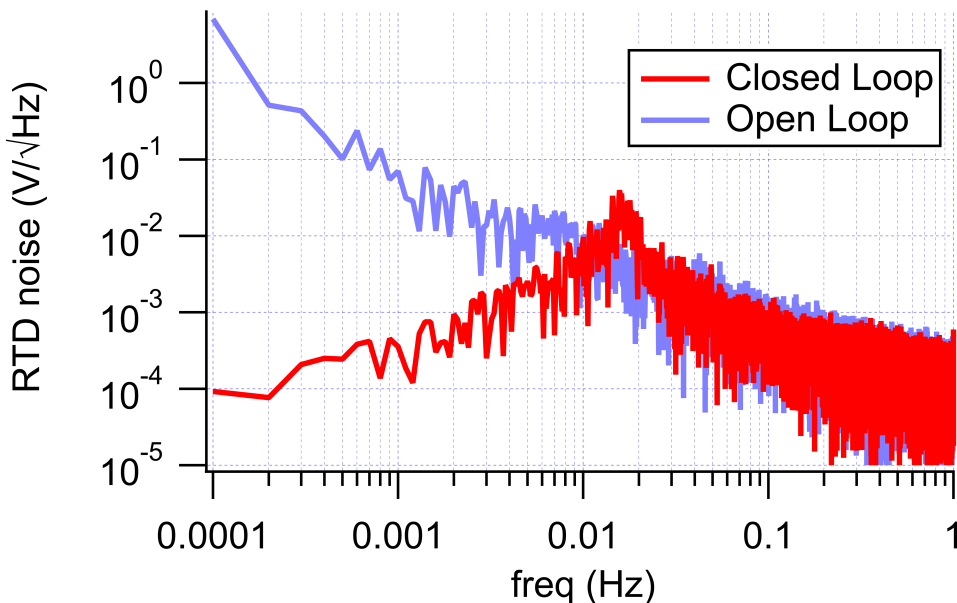


Figure 3.1: Temperature stabilization. RTD noise with and without feedback.

structure fixed within the cartridge. The cartridge itself is mounted in a 3D printed (ABS plastic) rig (Fig. 3.6). This rig has support arms which extend out three of the magnetic shield portholes described below (Fig. 3.3). These support arms are secured directly to an optical table on which the entire apparatus is mounted.

I will now discuss the measurements used to characterize the Xe and Rb relaxation rates. The longitudinal relaxation rate of the Xe, T_1 , is found using free-induction decay (FID) measurements. First, we perform DC pumping with no applied bias field until we reach maximum Xe polarization (times of order hundreds of seconds). A large "pinning" field is applied along the pumping direction during this time to ensure the polarization does not precess away from the x-axis. We then apply a Rabi flopping pulse to cause a π radian rotation of one of the Xe species. We then wait a variable amount of time, called the delay time, during which the polarization of the chosen Xe species will decay and eventually build back up in the pumping direction. After the delay time, the pinning field is turned off, the pumps are blocked, and a constant pulsed bias field is applied to start the FID. The Xe amplitude at the start of the FID will vary with delay time (t) as $\|A(1 - 2e^{-t/T_1})\|$, where




Figure 3.2: Vapor cell heater and cartridge mount photographs. **(a)** Four ceramic jig heater sides with aerogel pillows attached to outer faces. AC power enters via MMCX connectors at the bottom of each jig heater. **(b)** 1 cm^3 vapor cell with stem tucked into fiberglass insulation. A temperature sensor sits between the insulation and outer cell wall. **(c)** Cartridge with cell installed. **(d)** Gap fill shim on one jig heater face. **(e)** View looking into the jig heater with the cell installed. Note the four pieces of gap fill between each outer cell wall and the jig heater inner faces.

A is a constant. We can then fit amplitude vs delay time to obtain T_1 . Figure 3.4 shows measurements and fits for both Xe species. We find $T_1^{131} = 11 \text{ sec}$, and $T_1^{129} = 27 \text{ sec}$ at a Rb density of about $4 \times 10^{12} \text{ cm}^{-12}$.

The Xe transverse relaxation rate, T_2 , is measured by measuring the in-phase and quadrature Xe signals while varying the detuning between the drive and resonance frequencies. This measurement can be fit to a Lorentzian lineshape in order to obtain T_2 (or, equivalently, the linewidth Γ_2). The linewidth depends on how well compensated we are (aka. how well we have cancelled the Rb spin-exchange field) [Korver, 2015], and therefore measurements of Γ_2 are shown for PM, PDM, and hybrid operation in their respective chapters.




Figure 3.3: Inside of the magnetic shields. Left: view of probe laser entry direction. Right: view of probe laser exit direction.

In order to obtain the Rb relaxation rate, Γ' , we measure the magnetometer S_z response to an applied DC B_y of varying size. Assuming pumping along \hat{x} and negligible Ω_x , Ω_z , Eq. 2.1 gives a fit function of the form

$$S_z = \frac{\Omega_y R}{\Gamma'^2 + \Omega_y^2} + C \quad (3.2)$$

where C is an arbitrary experimental offset (eg. offset from photodiode mis-balance). Figure 3.5 shows a measurement of the relaxation rate at a Rb density of about $1 \times 10^{13} \text{ cm}^{-3}$, with the fit giving approximately $\Gamma' = 49000 \text{ rad/sec}$.

The Rb relaxation rate Γ' is related to the magnetic width of the magnetometer, B_w , according to $B_w = \Gamma'/\gamma^S$, and so the terms "magnetic width" and "Rb relaxation" can be used somewhat interchangeably. It is often useful to consider B_w , for example when

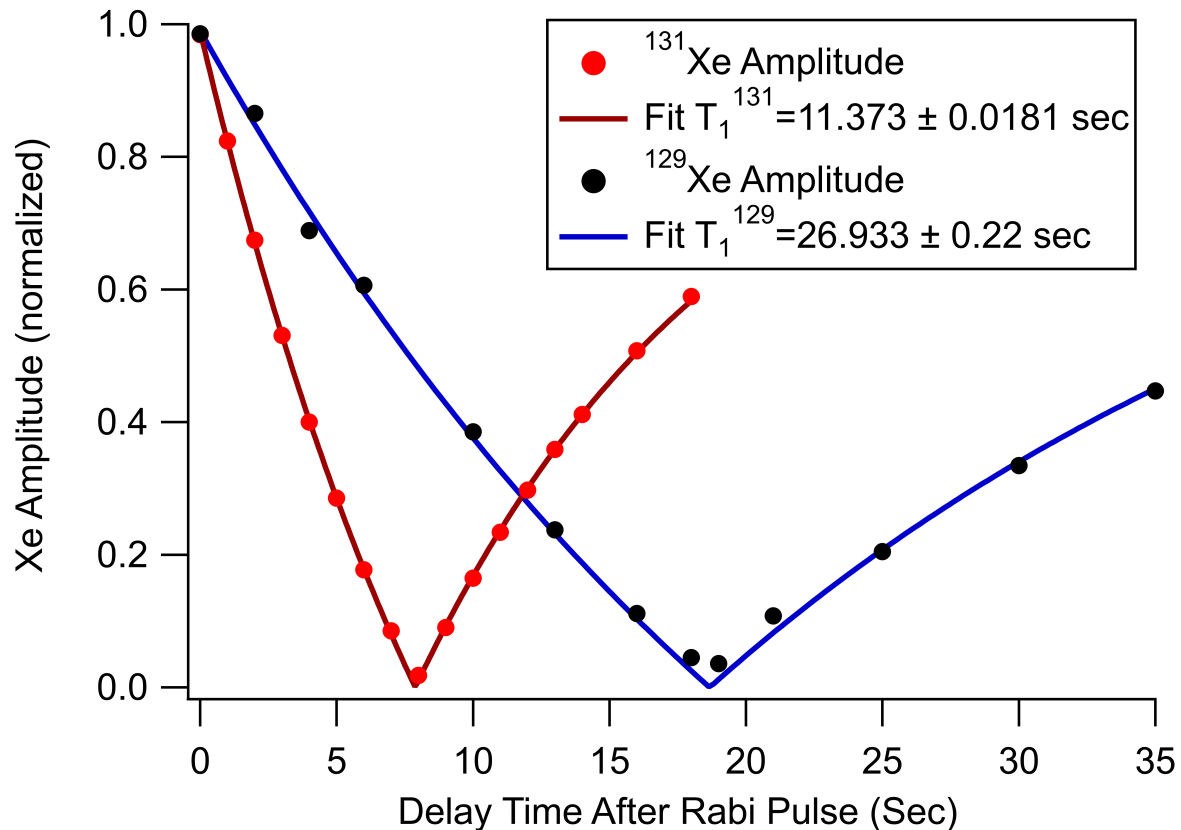


Figure 3.4: Measurements of normalized Xe amplitude vs delay time for ^{131}Xe (red) and ^{129}Xe (black). Dots are measured data, lines are fits for T_1 .

discussing the magnetometer phase shift (see Eq. 2.5). From our measure of B_w , we find the magnetic width to be about 3 mG.

3.2 Magnetic Fields and Pulsing Circuits

The rig used to hold the cell cartridge also has four sets of square coils, as shown in Fig. 3.6. Three of these are the x, y, and z "shim" coils used to cancel (or "null") external DC fields and to apply AC fields, as needed. For example, the x-compensation field is applied using the x shim coil. The circuit used to drive the shim coils was custom-made and is described in [Wyllie, 2012]. The final set of coils is used to apply the pulsed bias field, as described below.

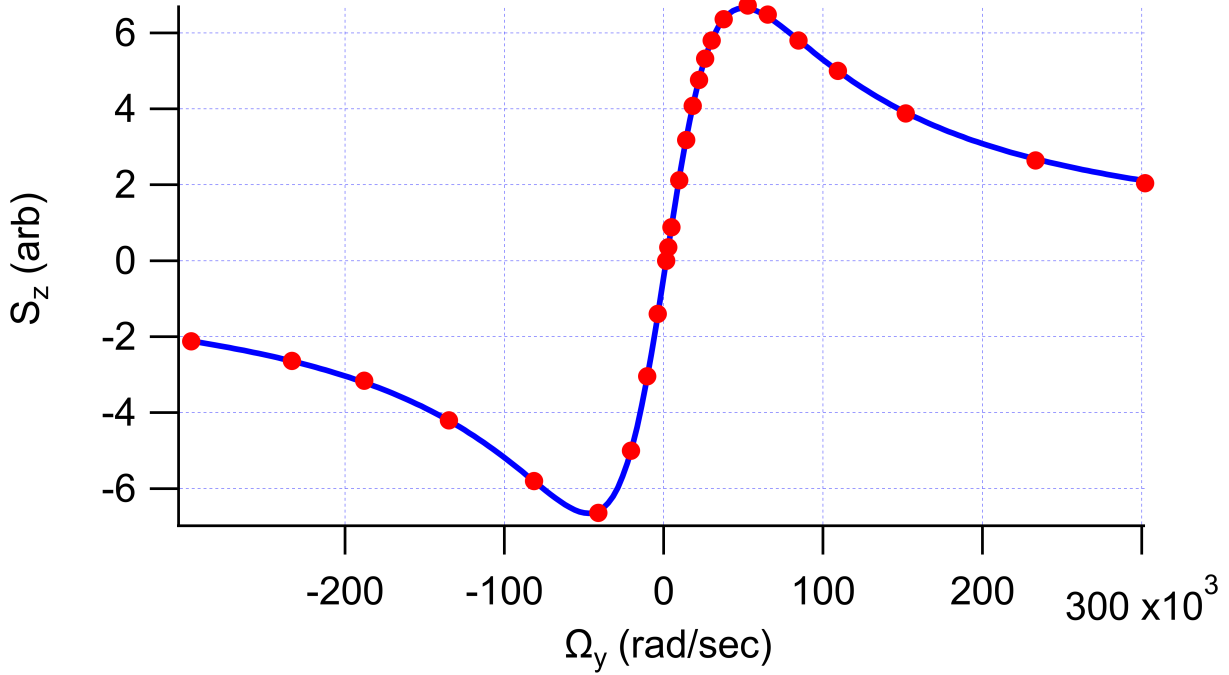


Figure 3.5: Rb relaxation measurement. Dots are measured data, line is the fit.

The three layer μ -magnetic shield we use is cylindrical with eight access ports (two along the axis of symmetry and six oriented tangentially). Since the tangential access ports are not located halfway between the two ends of the cylindrical shields, and since pump laser light must enter through the ports, we must place our cell such that it is not equidistant from the end caps. The asymmetry in distance to end caps informs the design of our bias pulsing coil set. To minimize coupling to the shield end caps and maximize uniformity across the volume of the cell, the pulsing coil set consists of two pairs of square Helmholtz coils with differing side lengths wound in series with opposite polarity. The purpose of the ancillary counter wound coils is to suppress the field produced by the coil set at the nearest end cap. See [Korver, 2015] for specific design details.

The bias field is comprised of short pulses ($< 5 \mu\text{sec}$) of ~ 1 Ampere peak current, with the exact pulse area tuned to obtain a 2π rotation of the Rb polarization. Custom pulsing circuits were used to drive the pulsing coils. I will now present the two pulsing circuits used for the results presented in this thesis. The first was used for PM operation, and the second




Figure 3.6: The 3D printed rig that holds the cell cartridge and the magnetic field coils. X,Y, and Z (along with their primed counterparts) are the shim coils. P1, P1', P2, and P2' are the bias pulsing coils.

for PDM and hybrid operation.

For PM (presented in Ch.4), we used the pulsing circuit shown in Fig. 3.8 which is described in detail in [Korver, 2015]. This circuit provided stable pulse areas for continuous pulsing at a single frequency, and was thus appropriate for use with PM operation (in which the pulses are constant).

When moving to PDM and hybrid operation (presented in Ch. 5 and Ch. 6), we found that the pulse area of the PM circuit was highly dependent on pulse repetition rate (see Fig. 3.9). Both PDM and hybrid operation utilize modulation of the pulse repetition rate. Therefore, if we tried to use the PM circuit for PDM or hybrid operation, our pulses would




Figure 3.7: Three layer μ -magnetic shield.

not consistently result in a 2π rotation of the Rb polarization.

Operating with PDM required us to re-design the circuit used to apply our bias field pulses. The new pulsing circuit (see Fig. 3.10) for PDM and hybrid operation uses an H-bridge design. I built this circuit in close collaboration with Michael Bulatowicz. Critically, all diodes used in this new circuit are "zero-recovery" Schottky diodes, known for their excellent high-speed performance. We modify the standard H-bridge to include "ringing suppression" gates (C_A and C_B). These gates are necessary to dampen current ringing across the inductor (aka, our magnetic field coils) after the pulse is applied.

The basic operation of the H-Bridge circuit is as follows. Two low-duty-cycle square pulses, produced by a function generator and triggered by a field-programmable gate array




Figure 3.8: PM pulsing circuit.

(FPGA), are used to drive the circuit. During the first pulse, the A gates are simultaneously opened, allowing current to flow through them across the pulsing coils. When the A gates are closed, the circuit discharges through the diodes to the lower left and upper right of the bias pulse coils. If we did nothing else, we would see a decaying oscillation of current over the coils. The second driving pulse then opens one of the ringing suppression gates, C_A , to damp this oscillation. In this way, we can apply a single z-field bias pulse. One of the benefits of the H-bridge design is that we can apply bi-directional pulses. Current can be sent through the inductor in the opposite direction by repeating the above procedure, but opening B and C_B gates instead of A and C_A .

Figure 3.11 shows a model of the current and voltage across the pulsing coils for the two different circuits. The effect of the ringing suppression gates can be seen by comparing the tail end of each pulse.

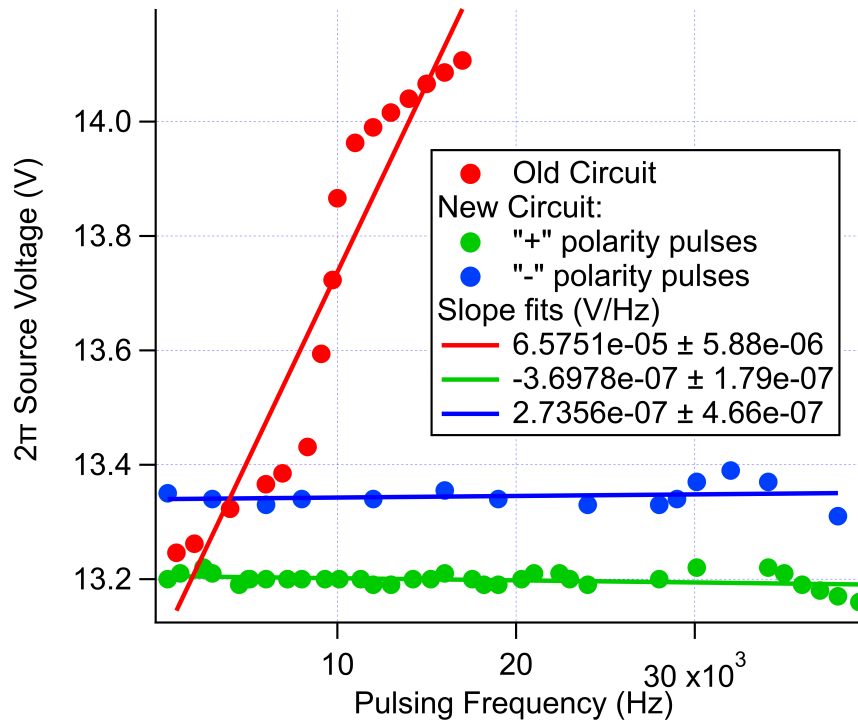


Figure 3.9: The voltage required to give a 2π pulse vs pulsing repetition frequency for the "old" PM circuit (red), and positive (green) and negative (blue) polarity pulses for the "new" H-Bridge circuit. Dots are data and lines are linear fits.

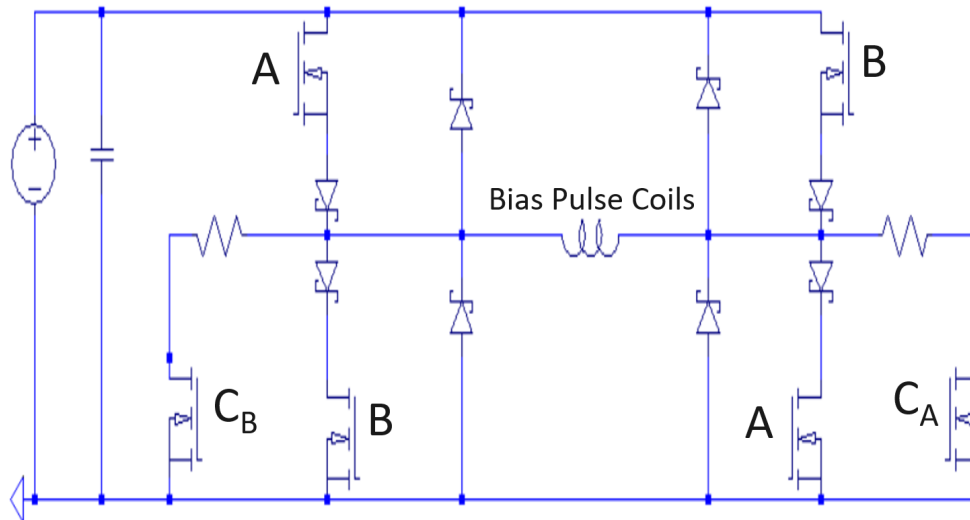


Figure 3.10: H-Bridge pulsing circuit.

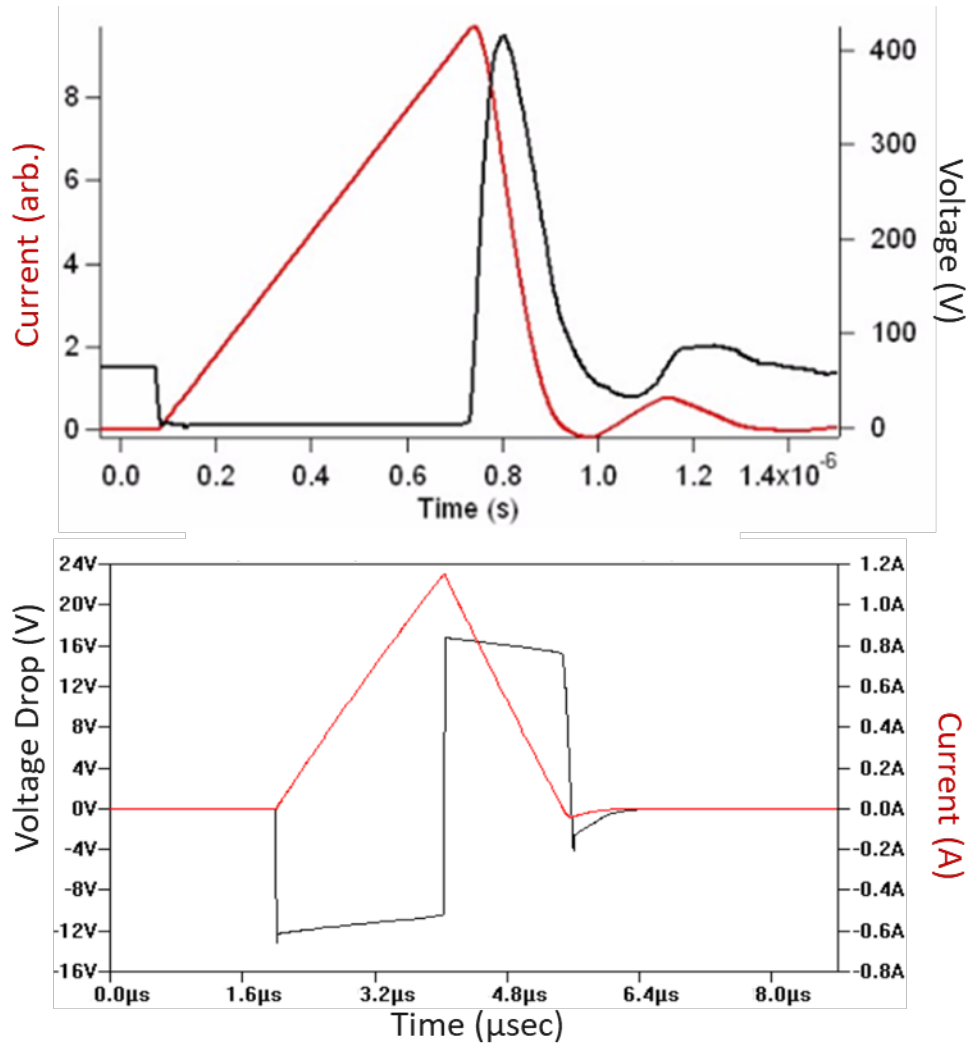


Figure 3.11: Pulse shape of the current (red) and voltage (black) applied across the pulsing coils by the PM pulsing circuit (top) and the H-Bridge circuit (bottom).

We find that the pulse areas applied using the H-Bridge circuit are 100-fold less dependent upon pulsing frequency than for the PM circuit. Figure 3.9 shows the source voltage required to provide a 2π Rb pulse versus pulsing repetition rate for the PM circuit and both polarities of the H-Bridge circuit. If we do not include the nonlinearity around 30 kHz, we find the slope for both polarities of the H-Bridge circuit to be consistent with zero.

While PDM operation does not include any polarization modulation, this H-Bridge circuit may also be useful in the future for reducing Rb transient times in PM operation.

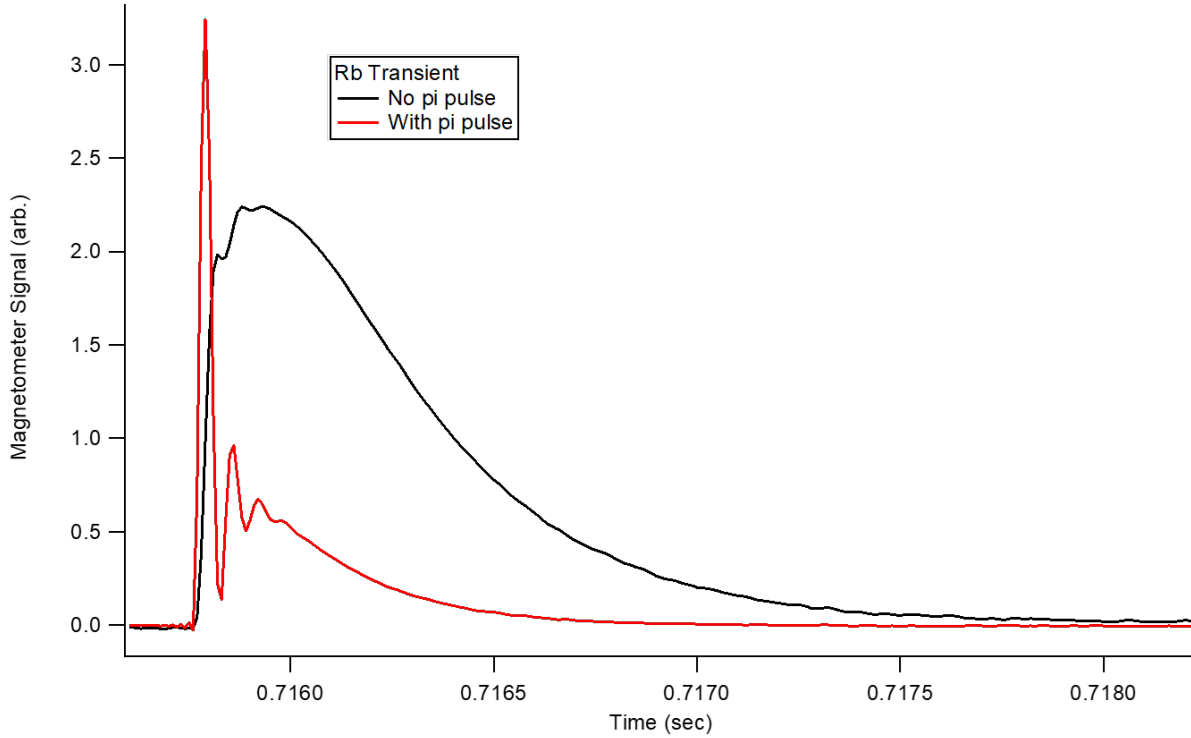


Figure 3.12: Rb transient area with (red) and without (black) a π pulse synchronized with the polarization flip.

Figure 3.12 shows how applying a π pulse in sync with the Rb polarization flip reduces the area of the transient by around 75 percent.

We null the x, y, and z fields by considering the form of the S_z signal as given in Eq. 2.4b. We null B_y using the $R(t)\Omega_y$ term, applying a square-wave $R(t)$ and adjusting the DC B_y shim until the square wave is no longer visible on S_z . We use the $\Omega_x\Omega_z$ term to null B_x , the AC Stark balance, the B_z shim, and to set the bias pulse area.

First, we apply an AC B_z and adjust both the DC B_x shim and the pump detunings (thus changing the AC Stark balance, see Sec. 3.3) to eliminate the AC response on S_z simultaneously for σ_+ and σ_- optical pumping. As an experimental note, I find it useful to adjust the frequency of only one pump laser (and use the same laser every day) to avoid slow wandering of the average detuning over time.

To null the DC B_z shim and set the bias pulse area, we apply an AC B_x and try to

eliminate the response on S_z . First, we null the DC B_z shim with the pulses off. Then, with the pulses on (at constant repetition frequency), we adjust the width of the applied bias pulse to again give no S_z response to AC B_x . The pulse width adjustments from day to day are generally of order 1 ns. The smallest increment that makes a noticeable change to the S_z response is around 0.2 ns, corresponding to a 0.006% change to the pulse length. The source voltage for the pulsing circuit is originally set in a similar way, though this is not changed from day to day. We apply AC B_x and slowly turn up the source voltage. The first "null" crossing seen in this way is a π rotation of the Rb, and the second crossing occurs at a 2π rotation. We are able to make finer adjustments of the pulse area using the pulse width than with the source voltage, so the source voltage is used to get us in the ballpark, and the pulse width is used for fine tuning.

3.3 Optics

A schematic of the optical elements of the apparatus is shown in Figure 3.13. The optical elements used vary slightly between PM, PDM, and Hybrid operation. The only optical change between the setups is in the inclusion or exclusion of the quarter wave plates (QWPs) labeled "A" and "B" on the schematic, as will be explained below.

To perform optical pumping of the Rb, we could theoretically use a single pump laser tuned on the Rb D1 resonance. However, due to the optical thickness of the vapor cell, this would require more laser power than our laser diode can provide. Instead, the outputs of two distributed feedback laser diodes (pump A and pump B) are detuned around 12 GHz to either side of the Rb D1 transition and overlapped. The detuning is chosen to optimize the magnetometer response to B_y . The exact detuning of each laser is adjusted to balance the AC Stark shift for the two lasers, thus canceling out the effective field experienced by the Rb due to the AC Stark effect. In practice, balancing the AC Stark shift often results in slightly mis-matched pump detunings (by a few GHz) due to slight differences in the power

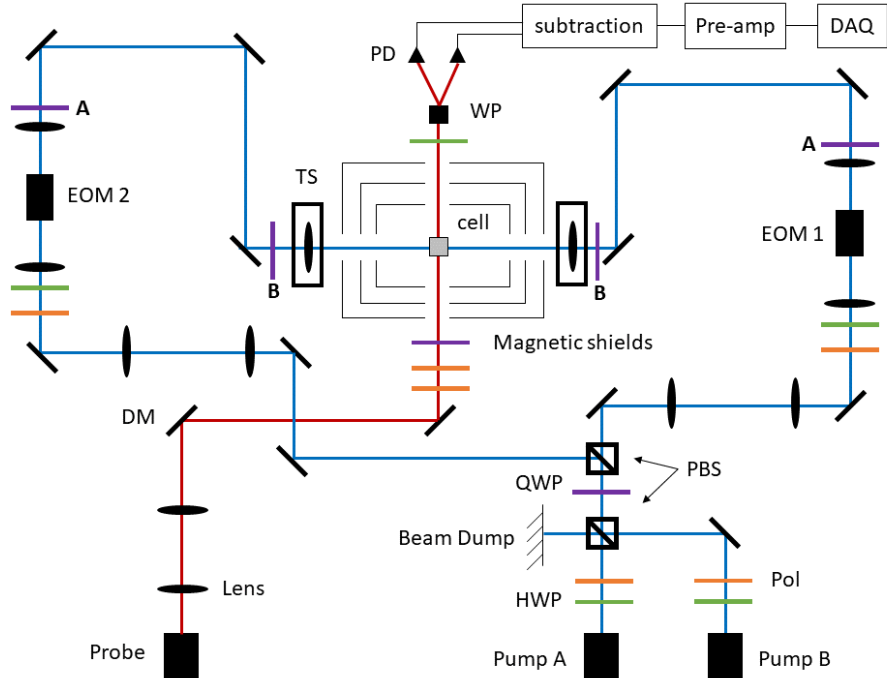


Figure 3.13: Experimental setup of optical elements (not to scale). DM: dichroic mirror, Pol: polarizer, HWP: half wave plate, QWP: quarter wave plate, PBS: polarizing beam splitter, WP: Wollaston prism, PD: photodiode, EOM: electro-optic modulator, TS: two-axis translation stage with lens. The three-axis magnetic shim and pulsing coils are not shown. The setup fits on a four foot square optical table. The QWPs labeled A and B are changed between PM, PDM, and hybrid operation.

and pointing of the two lasers. In the future, we plan to explore pumping with a single laser on resonance (using a Tapered Amplifier to provide enough power) to avoid the need for balancing the AC Stark shift. The overlap of the beams is accomplished by polarizing pump A so that it is mostly transmitted by a polarizing beam splitter (PBS) and pump B so that it is mostly reflected by the same PBS. The combined beam is then sent through a quarter wave plate and then separated into two beams using a PBS once again. The orientation of the quarter wave plate is chosen such that both pump A and pump B have half their power in each output beam. The overall beam resizing is set by telescopes to somewhat overfill the 7 mm diameter aperture of the ceramic heater.

Another telescope is used to couple the beams into individual electro-optic modulators (EOMs). This telescope is designed such that the Rayleigh length (measured from the beam

waist) is longer than half the length of the crystal, and the beam waist at the windows is $< \frac{1}{4}$ the radius of the window. The EOM is mounted on a five-axis stage, and alignment of the laser through the EOM is done using the isogyre method described in [Thorlabs, 2012].

Prior to each EOM is a polarizer and half wave plate. The polarizer ensures that the light incident to the EOM crystal is purely linear. The half wave plate (HWP) is used to align the light polarization relative to the EOM crystal axis. The maximum and minimum voltages of the EOM drive waveform are chosen to produce $\pm\lambda/2$ retardance. The quarter wave plate at the output of the EOM converts the EOM output at $V_{max}(V_{min})$ to be $\sigma^+(\sigma^-)$. The collimated output of each EOM is coupled into the vapor cell from opposing directions.

For PM (Ch. 4), QWP A is installed and QWP B is removed. The EOMs output circularly polarized light, as described above. QWP A is used to account for any retardance introduced to the light after exiting the EOM (eg. due to the dichroic mirrors or the glass walls of the cell) such that the light incident upon the Rb is well circularly polarized.

For PDM (Ch. 5), only QWP B is installed. The EOMs are turned off, such that the light remains linearly polarized after the EOMs. The EOMs are not removed from the system, however, in order to facilitate easy transitions between different modes of operation. The linearly polarized light is then converted to circularly polarized light by QWP B.

For hybrid operation (Ch. 6), both QWP A and QWP B are installed. The EOMs are turned on and output circularly polarized light. QWP A is then used to convert the light back to linearly polarized light, with σ^+ and σ^- circular polarization being mapped onto S and P linear polarization. This is done because the dichroic mirrors after the EOM affect the polarization of linearly polarized light less than circularly polarized light. Then, after the light has reflected off of all the remaining mirrors, it is converted back to circularly polarized light by QWP B. Note that this optical setup would also work for PM, and would reduce polarization errors introduced by the mirrors. In the future, we will use this setup if we ever return to PM.

Fine tuning of each pump beam's pointing is controlled using long focal length lenses

mounted on two axis translation stages just outside the magnetic shield. The position of each steering lens is chosen to optimize the magnetometer gain.

The power and detuning of each pump laser is chosen to approximately cancel the AC Stark effect. While we could also suppress the AC Stark effect by pumping with a single laser tuned on resonance, we would not get good spatial uniformity of the Rb polarization without significantly increasing our laser power due to optical thickness effects.

To detect S_z , linearly polarized light from the output of a third distributed feedback laser diode tuned near the Rb D2 line is used as a probe beam. A telescope re-sizes the beam to a beam waist of around 1mm, such that the beam waist is $< \frac{1}{4}$ the 7mm diameter aperture through the ceramic jig heaters. Crossed polarizers are used to adjust the power to approximately 1 mW and clean up the polarization. A QWP just before the cell is used to cancel any retardance introduced by the glass cell wall. With the pumps blocked, a large AC B_y is applied and the QWP is adjusted until no AC response is seen, which ensures that the pump is linearly polarized when it reaches the Rb, and does not cause any optical pumping. The probe beam is directed through the center of the cell and parallel to \hat{z} , where any z-polarized Rb will cause a Faraday rotation. The light then passes through a HWP and Wollaston prism onto a balanced Faraday detector. The Wollaston prism is used to split the light into two beams of perpendicular linear polarization, and the HWP is used to balance the power of the two beams (we refer to this as "balancing the photodiodes"). The balanced Faraday detector then subtracts the two beams such that small variations are more easily seen. The Faraday detector signal is passed through a pre-amplifier and then sent to a DAQ for the signal to be recorded and analyzed by an FPGA.

In order to characterize the noise of our apparatus, we measure the Faraday rotation signal (without demodulation or Xe excitation) under a few different circumstances. These noise measurements are shown in Fig. 3.14. First, we block the pump beams to measure the noise inherent on our probed signal (we refer to this as the "probe noise"). We find that if we do not null the optical pumping by the probe as described above, the probe noise is

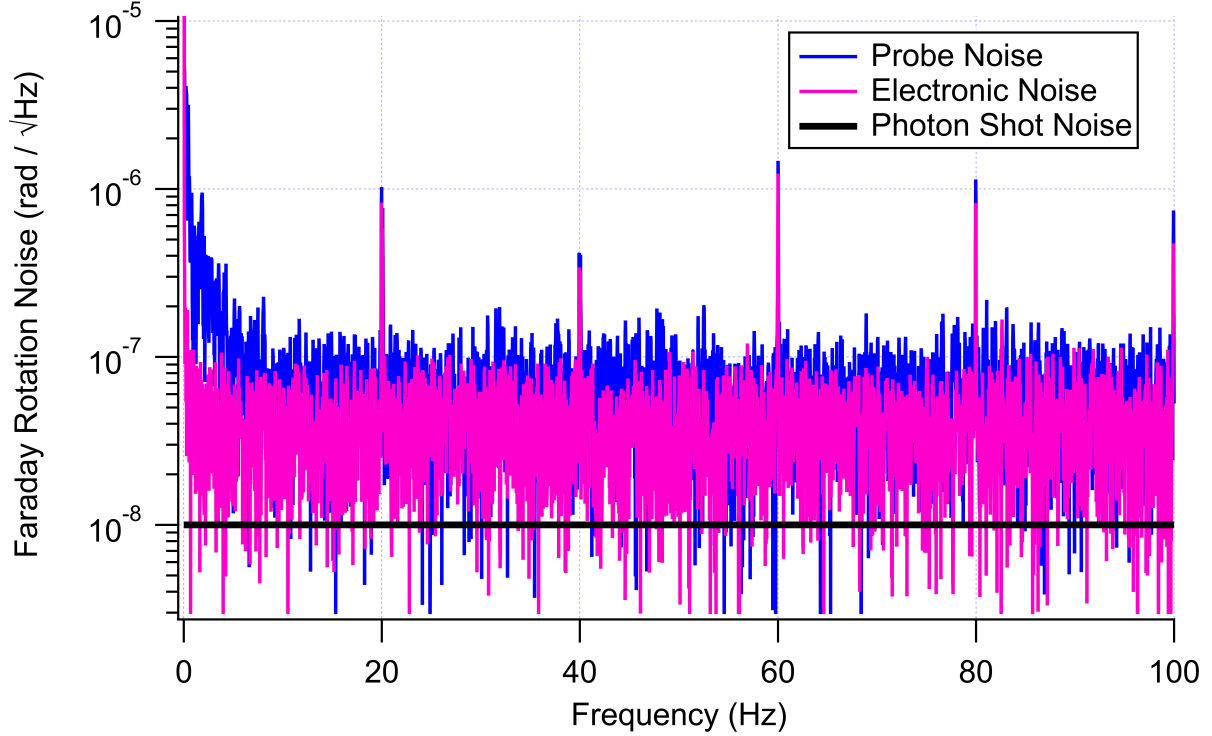


Figure 3.14: Probe noise, electronic noise, and photon shot noise on the Faraday rotation signal. These measurements were taken at a time when we were using PDM operation. The total detected photocurrent was 0.8 mA.

degraded. Next, we also block the probe so that no laser light is incident on the cell. This gives us a measure of our electronic noise. Finally, we calculate the photon shot noise, Φ , using the equation [Thrasher, 2020; Romalis, 2013] (approximating $\Delta \gg W$)

$$\Phi = nl\sigma_0 S_z S_\infty \frac{\Delta}{W(1 + 4\Delta^2/W^2)} \approx \beta_0 S_z S_\infty \frac{W}{4\Delta}, \quad (3.3)$$

where n is the Rb atomic number density, l is the probe laser path length through the cell, $\sigma_0 = \sigma_{max} W_{radiation}/W_{pressure}$ is the cross section, $\beta_0 = nl\sigma_0$ is the optical depth at line center, W is the width of the D2 resonance, Δ is the detuning from resonance, and S_∞ characterizes the degree of circular dichroism of a transition.

While it may appear that the probe noise is nearly electronic noise limited, later tests with a lower-electronic-noise detector found that the probe noise remained unchanged. We

therefore believe the probe noise is not electronic noise limited (save for electronic noise spikes at multiples of 20 Hz). The probe noise level being above the electronic noise level suggests that we are operating at sufficiently high probe power. The probe noise below 10 Hz is dominated by $1/f$ detection noise. This detection noise is why we try to keep our Xe precession frequencies (and therefore our demodulations) as high as possible. Above 10 Hz, the probe noise is mostly white, with a white noise level about an order of magnitude above the photon shot noise. It is important to note that these measurements do not mean that our system is probe noise limited. The probe noise shown here is lower than the noise we measure under operating conditions. Because of that, we did not explore reducing the probe and electronic noise further.

3.4 Digital Control

Our data acquisition, feedback, modulations, experiment control, and some analysis are performed by a combination of a field-programmable gate array (FPGA) and a host computer ("host"). All programming is done using LabVIEW. Post-processing analysis (including the production of figures) is performed in Igor Pro.

In this section, I will give an overview of the major functions of the FPGA and host computer. For more details, see [Thrasher, 2020]. In general, the FPGA is used for anything that has to happen quickly or with very precise timing, as well as for any analog-to-digital or digital-to-analog conversion. The host computer is used for anything that can be done more slowly (around 10 Hz), and is where experimental parameters can be changed in real-time by the user.

The FPGA used for this experiment is National Instruments model PXI-7852R, which has 48 BRAM units with 32 kB each. It is my recommendation that anyone shopping for an FPGA for this application find one of suitably large size, as we have sometimes run up against maximum memory usage issues. The FPGA, which is synced to an external 40 MHz

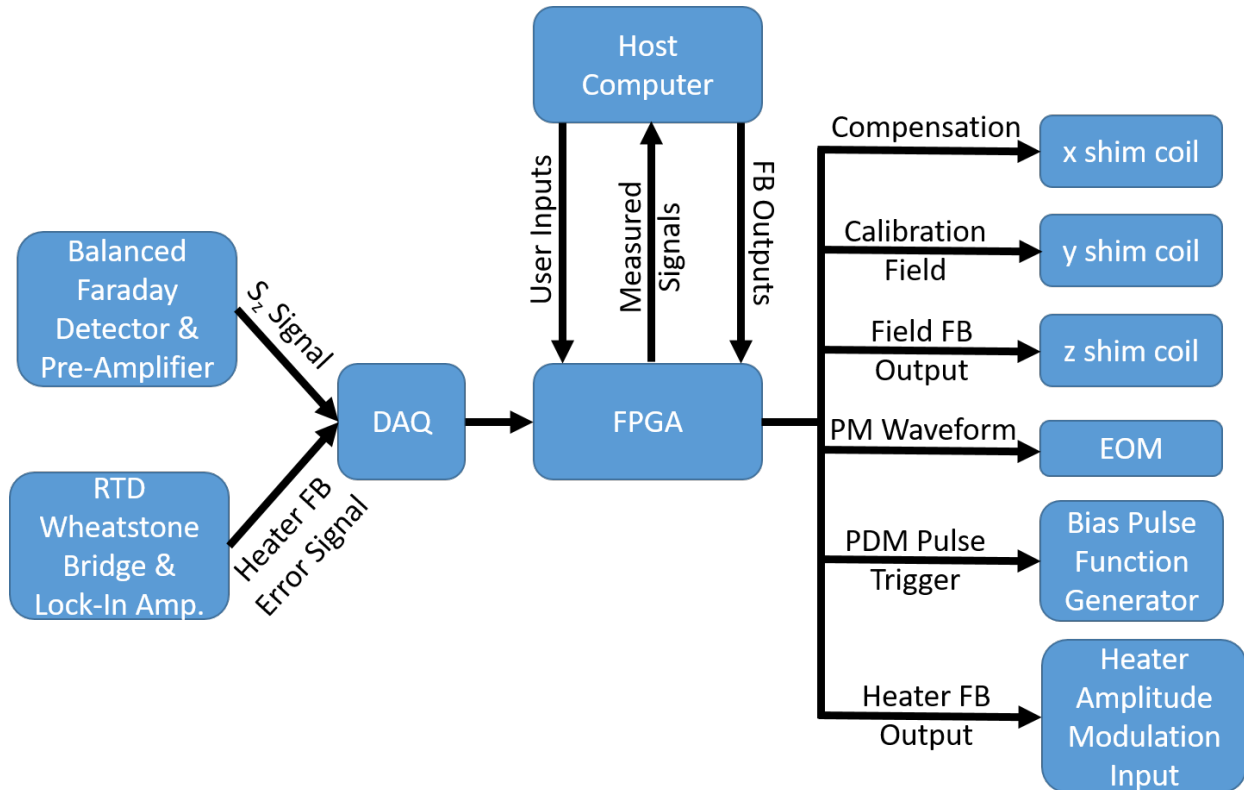


Figure 3.15: Block diagram of signal flow through digital system. FB=feedback.

commercial atomic clock, is used for fast or timing-sensitive programming.

Figure 3.15 shows a block diagram of the flow of signals through our digital system. We begin on the left with the signals we need to digitize for later use: the S_z signal from the balanced Faraday detector (after being sent through a pre-amplifier to convert from current to voltage), and the heater feedback error signal from the RTD, Wheatstone bridge circuit, and a lock-in amplifier. These signals are digitized via a 16-bit DAQ and read by the FPGA. The FPGA acquires the signals at a very high rate (of order 10 MHz) and averages the signal down to our desired sampling rate (200 Hz). The FPGA receives user inputs from the host (such as modulation frequencies and amplitudes, sampling frequencies, field sizes, etc.) which are used to produce modulation waveforms using direct digital synthesis (DDS). The measured signals and modulation waveforms are sent to the host, which performs least-squares demodulation (except in the case of PM, which used external SRS lock-in

amplifiers for demodulation). The host records and exports the demodulation output signals, and also uses them to calculate feedback outputs, which may affect the user inputs. These feedback outputs are sent back to the FPGA to be output via a 16-bit digital-to-analog converter. These outputs are used to control various aspects of the experiment.

The x-compensation field, fields used for signal amplitude calibration, and the feedback field are sent to the x, y, and z shim coil drivers (respectively). The PM waveform is sent to the EOM high-voltage driver through a follower. The FPGA determines when to produce pulses in order to produce the desired bias field via PDM, and sends trigger pulses to the function generator used to drive the pulsing circuit. The temperature stabilization feedback output is sent to the amplitude modulation input of the function generator used to heat the cell (via an amplifier and the jig heaters).

Chapter 4

Polarization Modulation

In this chapter, I describe continuous excitation of Xe in a transversely-pumped NMR gyroscope using Polarization Modulation (PM) operation. In order to excite two species of Xe, the handedness of the circularly polarized pump light is doubly-modulated (using EOMs). First, I will give the modulation waveform used, and the resulting form of the detected S_z . Then, I will demonstrate simultaneous excitation of the two Xe species, and discuss the demodulation scheme used to extract the Xe phase shifts. I will then present rotation measurements obtained for PM operation. Finally, I will discuss signal mixing present on the demodulated signals which provided motivation for transitioning to PDM operation. The work in this chapter was done in close collaboration with Daniel Thrasher, and is therefore also described in his thesis [Thrasher, 2020].

4.1 Modulation and Drive Scheme

We choose to modulate the Rb polarization as a square wave, which is accomplished by using the EOMs to switch between just two polarization states, σ^\pm . While we could theoretically use sine wave modulation, we find that it is very challenging to produce a sinusoidal PM waveform in the optically thick vapor cell. To facilitate simultaneous excitation of the two Xe

species, the PM waveform must have significant frequency content at the two Xe resonance frequencies.

In thesis, PM operation is accomplished by modulating the x-component of \mathbf{R} according to

$$R(t) = R_0 \text{sign}\left[\cos\left(\frac{\omega_d^a + \omega_d^b}{2}t + 2\cos\left(\frac{\omega_d^a - \omega_d^b}{2}t\right)\right)\right]. \quad (4.1)$$

Figure 4.1 shows the resulting light polarization, $\sigma = \pm 1$. More details on this and other modulation waveforms can be found in [Thrasher, 2020]. This waveform is advantageous in that the smallest separation between reversals is larger than the finite response times of the EOMs and of the optical pumping of the Rb atoms. We use the same modulated square waveform to apply a compensation field, the amplitude of which is set to cancel the spin-exchange field experienced by the Xe from the Rb. Doing so suppresses the production of K_z and narrows the NMR linewidths, Γ_2 [Korver et al., 2015].

For the modulation waveform given in Eq. 4.1, the modulation frequencies are $\omega_{p_1} = (\omega_d^a + \omega_d^b)/2$ and $\omega_{p_2} = (\omega_d^a - \omega_d^b)/2$. These modulation frequencies satisfy the resonance condition given in Eq. 2.2 for the drive scheme $(p_1, p_2) = (1, 1)$ for ^{129}Xe and $(1, -1)$ for ^{131}Xe .

The derivation of the Fourier coefficients for this modulation follows the concepts laid out in Ch. 2, where I assumed a generic PM waveform that could be written as Eq. 2.2. For Eq. 4.1, first consider a square wave of phase $\omega_{p_1}t + 2\cos\omega_{p_2}t$. The Fourier expansion can be written

$$\sum_{p_1} s_{p_1} e^{ip_1\omega_{p_1}t} e^{ip_1 2\cos\omega_{p_2}t} \quad (4.2)$$

where $s_{p_1} = 2/(p_1\pi)$, the well known result for a square wave. I then use the Jacobi-Anger identity (as in Eq. 2.13) on the second exponential term of Eq. 4.2, which gives a Bessel function $s_{p_2} = J_{p_2}(2p_1)$. For the chosen drive scheme, the product of the Fourier coefficients is $j = \|s_{p_1}s_{p_2}\|$. Evaluated numerically, this gives $j = 0.37$ for both Xe species. In Ch. 7, I




Figure 4.1: Measured square wave PM signals. Top: $S_+(t)$. Middle: $S_z(t)$. Bottom: Rectified $S_z(t)$. The plot on the right shows the amplitude spectral density with (red) and without (blue) rectification.

will compare this to the Fourier coefficients for PDM and hybrid operation.

4.2 Demodulation

From Equation 2.4, the detected longitudinal Rb polarization S_z is

$$S_z = -\frac{R(t)}{\Gamma^2} [b_a^S K_\perp^a \sin(\delta^a + \alpha^a - \epsilon_z) + b_b^S K_\perp^b \sin(\delta^b + \alpha^b + \epsilon_z)], \quad (4.3)$$

where $\delta = \phi - \alpha$ is the difference between the drive phase $\alpha = \int dt \omega_d$ and the Xe precession phase ϕ for each isotope, ϵ_z is the magnetometer phase shift, and $R \sim S_+$ is the optical pumping rate of the Rb.

Although lock-in detection can be accomplished on S_z as it appears in Eq. 4.3, the phase sensitivity is diminished due to the presence of $R(t)$. $R(t)$ causes signal mixing which causes Xe signal content to be spread across many different frequencies. This leads to small signal sizes at the demodulation frequencies, mixing of various noise sources onto the Xe demodulation frequencies, and cross talk (see Sec. 2.4). A more effective approach is to "rectify" the S_z signal such that $R(t)$ is removed. Removing $R(t)$ allows us to more cleanly isolate the two Fourier components of the signal due to ^{131}Xe and ^{129}Xe , and improves the Xe signal sizes at the demodulation frequencies. This is accomplished in real time by multiplying S_z by $R^{-1}(t)$ (which for our waveform is just $R(t)$). The resulting signal is then sent into two separate lock-in amplifiers, each referenced to a different isotope's drive frequency. From these demodulations, we arrive at the measured phase shifts $\delta^a - \epsilon_z$ and $\delta^b + \epsilon_z$.

While rectification does reduce signal mixing, it doesn't remove it entirely. This is because the waveform $R^{-1}(t)$ used for rectification is based on a model of $R(t)$, not a measurement. This means that the rectification signal does not perfectly match the actual $R(t)$. The actual $R(t)$ may vary due to pump laser intensity or frequency drift, changes to the bias pulse rate (see Sec. 5.2), temperature drifts, etc. Rb relaxation also means that the actual $R(t)$ experiences transients and is therefore not a perfect square wave. Even when we used a non-square rectification to account for Rb transients, we could not match the rectification signal to $R(t)$ well enough to eliminate signal mixing. Imperfect rectification leads to cross talk, as will be discussed in Sec. 4.5.

Figure 4.1 shows the time series of the measured S_z with and without rectification when we drive both isotopes near resonance simultaneously using S_+ . We see that rectification reveals the sinusoidal precession of each isotope. The outlying data on the rectified signal, which occur when the polarization is reversed, are due to the finite response time of the Rb magnetometer. Although rectification collects the many Xe signal sidebands into the two Xe carrier frequencies (see power spectrum on right), it also maps $1/f$ noise in the detection of S_z to the carrier frequencies. Even if our rectification waveform perfectly




Figure 4.2: Measured NMR lineshapes of each species. Filled circles are K_x , open circles are K_y , and lines are Lorentzian fits.

matched $R(t)$, measurement of S_z using balanced Faraday detection introduces electronic offsets and low-frequency detection noise. Rectification will then mix these low-frequency signals directly onto the Xe demodulation frequencies. The mapping of low-frequency noise on the S_z signal to the carrier frequencies can be prevented by high-pass filtering S_z with a 1 Hz corner prior to rectification. However, this high-pass filtering also contributes to cross talk, as will be discussed in Sec. 4.5.

4.3 NMR Excitation

Figure 4.2 shows the detected NMR signals for each isotope. These data were acquired by driving one isotope on resonance while varying the other isotope's drive frequency and recording its K_x and K_y derived using demodulation. We see that the lineshapes are nearly Lorentzian with linewidths (half-width at half-max) of 14 mHz for ^{129}Xe and 17 mHz for ^{131}Xe , and amplitudes of approximately 66 μG for ^{129}Xe and 78 μG for ^{131}Xe . The on-

resonance amplitudes are in agreement with estimates similar to those outlined in [Walker and Larsen, 2016]. The ~ 15 mHz linewidths are only possible because of two features of our experiment; (i) the use of a Rb hydride cell coating (without which T_2^b would be substantially shorter and T_2^a much longer) [Kwon et al., 1981], and (ii) the application of a magnetic compensation field B_x that cancels the Rb spin-exchange field experienced by the Xe.

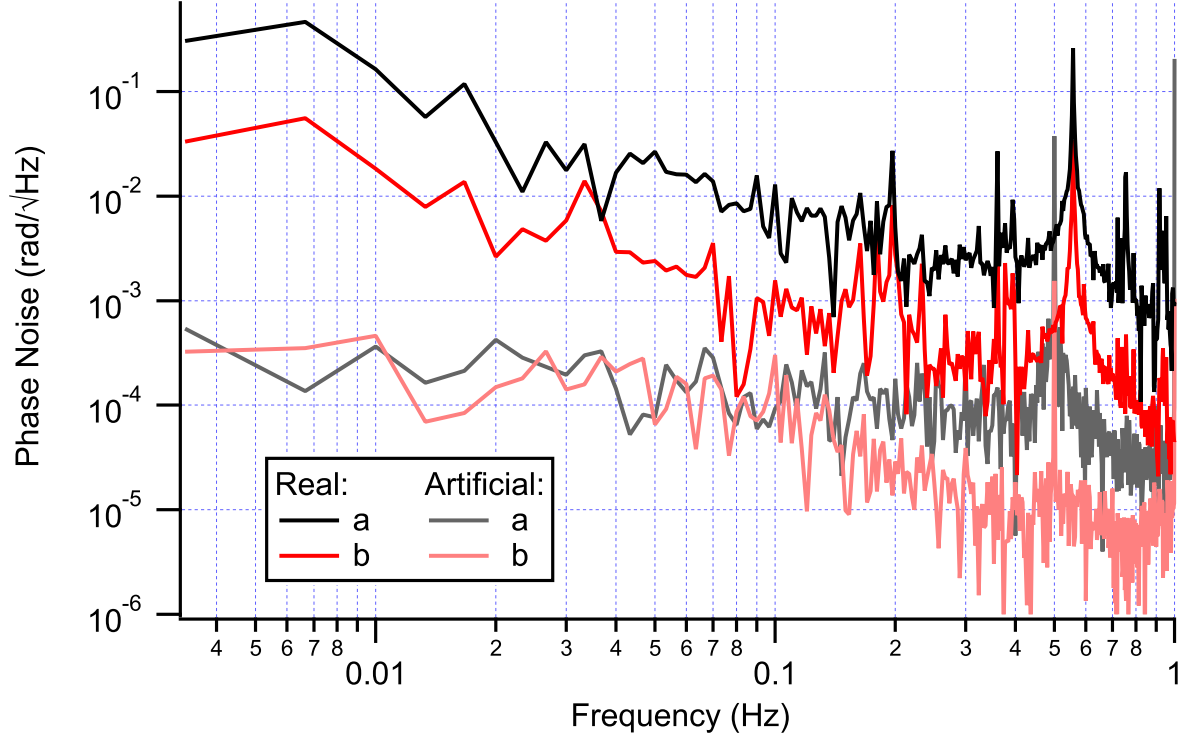


Figure 4.3: Measured phase noise of each species. Traces labeled “Real” are recorded when the Xe isotopes are excited on resonance. Traces labeled “Artificial” are recorded when the Xe isotopes are both driven off resonance (not excited) and an AC B_y is applied to the magnetometer at the off-resonance drive frequencies.

Figure 4.3 demonstrates the amplitude spectral density of the phase noise measured for each isotope under simultaneous resonant excitation conditions. We see that for frequencies less than 1 Hz the spectra are dominated by $1/f$ noise which is about $\rho = \gamma^a/\gamma^b$ greater for isotope a (black traces) than for isotope b (red traces) suggesting the dominant source of $1/f$ noise is magnetic in nature. Also shown is the phase noise measured when the Xe isotopes are driven off resonance (not excited) and a so-called “artificial” Xe signal is applied

as an AC field, B_y . The amplitudes of this signal $A^a \sin(\omega_d^a t) + A^b \sin(\omega_d^b t)$ are chosen to produce the same size magnetometer signal as real Xe. We note that this artificial Xe signal is planar, unlike the real Xe signal which rotates. The artificial signal allows us to measure the signal-to-noise ratio (SNR) of the Rb magnetometer. These signals do not show $1/f$ dependence because, unlike the real Xe phase, the SNR of the magnetometer does not depend on the bias magnetic field to first order. The detection phase noise from each artificial Xe measurement is uncorrelated and limits the possible field suppression when performing comagnetometry. The SNR for each isotope is $\sim 5000\sqrt{\text{Hz}}$. While the detection of artificial Xe is insensitive to $1/f$ bias magnetic field noise, $1/f S_z$ noise can still be mapped to the artificial Xe frequencies via rectification, and so the addition of a 1 Hz high-pass filter prior to rectification was essential for realizing such an SNR.

4.4 Comagnetometry

Transfer Function

We perform comagnetometry by subtracting magnetic field correlations between the two isotope's precession frequencies. Since our device measures phase, we need to know or measure the transfer function from phase to frequency. In Ch. 2 we derived the transfer function (see Eq. 2.18). Here, we measure the transfer function of each isotope by recording the response of the measured phase δ to sinusoidal modulation of B_z . A brief discussion of this measurement is given here, and more details can be found in [Thrasher, 2020].

Figure 4.4 shows the measured transfer function for isotope a . We use a chirp waveform to modulate the bias field $B_z^{mod}(t) = B_0 \sin(2\pi[e^{t/T_2} - 1 - t/T_2])$ (where $T_2 = 1/2\pi\Gamma_2$), the time series of which is shown in the inset of Figure 4.4. This modulation waveform allows us to measure the transfer function from 0.002 to 0.1 Hz with good SNR in a single data acquisition. The transfer function is the ratio $\gamma^K \tilde{B}_z^{mod}/\tilde{\delta}$. We fit the data according

to Equation 2.2 and find excellent agreement with the linewidth derived from the fits in Figure 4.2.




Figure 4.4: Transfer function of ^{129}Xe . Inset depicts the normalized chirp waveform used to modulate the bias field. Black crosses are measured data, while the blue line is a fit of the form $\sqrt{(\Gamma_2^a)^2 + f^2}$.

Feedback

Although conversion from phase to frequency for the measured Xe phases is possible using the measured transfer function, feedback [DeLand, 2017; Bechhoefer, 2005] is desirable because (in the high gain limit) the performance of the comagnetometry becomes insensitive to changes in the transfer function. We used the measured precession phase of isotope a to stabilize the bias field and the measured transfer function of isotope b to convert its measured phase noise to frequency noise. Under such conditions the frequency noise of

isotope b is proportional to rotation. We apply a feedback field, B_{FB} , to hold the measured phase of isotope a , $(\delta^a - \epsilon_z)$, equal to zero. The field stabilization can then be written $\tilde{B}_{FB} = \tilde{G}(\tilde{\delta}^a - \tilde{\epsilon}_z)$. In the high gain limit, this becomes

$$\lim_{G \rightarrow \infty} \tilde{B}_{tot} = \frac{1}{\gamma^a} (\tilde{\Omega}^R + i\omega \tilde{\epsilon}_z), \quad (4.4)$$

where $B_{tot} = B_{z0} + B_p + B_{FB}$, and where we made substitutions using Equation 2.2 and $\tilde{\Delta}^a = -\gamma^a \tilde{B}_z - \tilde{\Omega}^R$ (since ω_d^a is held constant). The measured phase of isotope b , $(\delta^b + \epsilon_z)$, converted to frequency is

$$\tilde{\omega}^b \equiv (\tilde{\delta}^b + \tilde{\epsilon}_z)(\Gamma_2^b + i\omega) = (\rho^{-1} + 1)(\tilde{\Omega}^R + i\omega \tilde{\epsilon}_z). \quad (4.5)$$

We see that in the high gain limit, when correcting the bias field to keep the measured phase of isotope a equal to zero, the rotation is simply $\tilde{\Omega}^R = \rho \tilde{\omega}^b / (1 + \rho)$ assuming $\tilde{\epsilon}_z$ is negligible.

The best performance we observed with bias field feedback activated is shown in Figure 4.5. The feedback consisted of two analog inverted zero gain stages. The influence of bias field feedback is dramatic from 0.1 to 200 mHz. The feedback loop suppresses $\tilde{\omega}^a$ to below 1 $\mu\text{Hz}/\sqrt{\text{Hz}}$ at low frequency, which is nearly 10^4 x less than the open loop noise. Because magnetic noise dominates each isotope's precession, using feedback to stabilize the measured phase of isotope a also greatly suppresses $\tilde{\omega}^b$.

We observe $\sim 240\times$ improvement in $\tilde{\omega}^b$ due to feedback. At this time, we were not yet applying external B_z to measure the field suppression factor (FSF) explicitly, and so ~ 240 is the lower limit for the FSF.

Stability

The modified Allan deviation [Allan and Barnes, 1981] suggests a rotation ARW sensitivity of $\sqrt{2} 15 \mu\text{Hz}/\sqrt{\text{Hz}} \frac{\rho}{1+\rho} \sim 16 \mu\text{Hz}/\sqrt{\text{Hz}}$ and a rotation bias instability of $1 \mu\text{Hz} \frac{\rho}{1+\rho} \sim 800 \text{ nHz}$. The size of ARW is within a factor of 3 of the ratio of the measured linewidths divided by the SNRs (shown in Figure 4.3). The peaks in the Allan deviation at 15 and 100 s of integration are due to low-frequency narrow-band large-amplitude noise peaks in $\tilde{\omega}^b$ which we attribute to the PM waveform. The bias instability is limited by $\tau^{1/2}$ trending noise of unknown origin.

We find that feedback causes the measured phase of isotope b to trend linearly in time. The source of this frequency bias is unknown, but I believe it may be related to the magnitude drift that will be discussed in Sec. 6.5. Although the trend is very stable over the course of a data run, the trend is not consistent between data runs. The bias instability demonstrated in Figure 4.5 was difficult to reproduce. Typically, the bias instability measured was a few μHz .

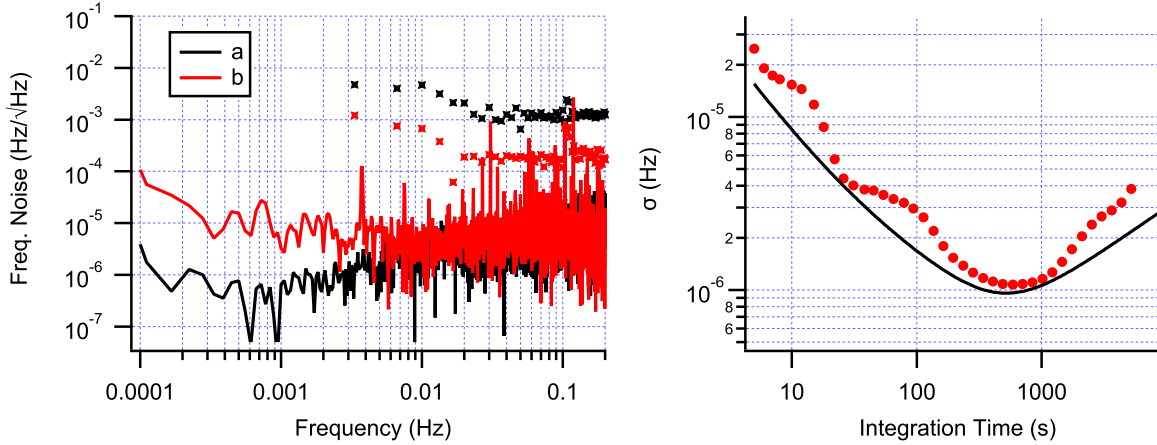


Figure 4.5: Measured comagnetometry noise and stability. Left: the amplitude spectral density of frequency noise. The cross marks indicate open loop frequency noise. The solid lines are frequency noise when the measured phase noise of isotope a is used to stabilize bias field. Right: modified Allan deviation of ω^b . Filled circles are measured data. Solid line shows the quadrature sum of $7 \times 10^{-5} \tau^{-1}$, $15 \mu\text{Hz}/\sqrt{\text{Hz}} \tau^{-1/2}$, and $30 \text{ nHz} \sqrt{\text{Hz}} \tau^{1/2}$ trends.

4.5 Cross Talk

Having measured the stability of the PM comagnetometer, we wanted to know the fidelity with which the detection separated signals from the two Xe isotopes. It is possible that the detection channel designed to measure isotope a 's phase was really measuring a linear combination of isotope a and b 's phases (see Sec. 2.4). If such ‘‘cross talk’’ were present then the scale factor (or how we convert the measured precession frequencies to rotation) would change [Thrasher, 2020].




Figure 4.6: Measurement of cross talk for detection with rectification. The Xe drive frequencies are set such that the two isotopes are not driven on resonance at the same effective bias field magnitude. The in-phase (filled symbols) and out-of-phase (un-filled symbols) components of each isotope’s detection channel (black and red correspond to isotopes a and b , respectively) are shown as the bias pulsing frequency is varied. We see that when isotope a is resonant ($f_{2\pi} \sim 15.8$ kHz) the signal on isotope b ’s detection channel is not flat despite isotope b being driven many linewidths off resonance. Similarly, when isotope b is resonant ($f_{2\pi} \sim 16$ kHz) the signal on isotope a ’s detection channel is not flat despite isotope a being driven many linewidths off resonance.

We characterize the cross talk present in the detection channels by looking for changes

in the detected signal for one isotope when the drive of the other isotope is changed. We do this by detuning one isotope's drive frequency by ~ 300 mHz and then scanning the bias pulse repetition rate. The effective resonance frequency of each Xe isotope depends on the pulsing frequency ($\omega_{2\pi}$) as $\omega^K = \omega_{2\pi}\gamma^K/\gamma^S$. The 300 mHz detuning ensures that both isotopes are not simultaneously on resonance for a given pulsing frequency. When a Xe isotope is far from resonance, we expect the measured signals for that isotope's detection channel to be relatively flat if there is no cross talk. Cross talk manifests if the signals from the off-resonant isotope are not flat when the other isotope is scanned through resonance. Figure 4.6 shows a phase sensitive measurement of cross talk. When isotope a is excited and isotope b is not, isotope b 's detection channel exhibits non-zero signal, and vice versa. We estimate $\beta = Q_{pp}^b/Q_{pp}^a = 0.17$ (when isotope a is on resonance) and $\beta' = Q_{pp}^a/Q_{pp}^b = 0.07$ (when isotope b is on resonance), where Q_{pp}^K is the peak-to-peak quadrature signal of isotope K .

We observe cross talk on the detection channels even when Xe is not excited and artificial Xe signals are applied, suggesting that the observed cross talk is not due to physical interactions between Xe isotopes. We believe the measured cross talk stems from imperfect rectification of S_z due to optical pumping transients (i.e., gain reversals) and unaccounted-for phase shifts from high-pass filtering prior to rectifying. The optical pumping transients stem from the few ms finite response time of the magnetometer. Indeed, by implementing a sample-and-hold algorithm to ignore data acquired during optical pumping transients, the cross talk is suppressed. The cross talk is further suppressed when the high-pass filter is removed. Doing so, however, reduces the detection SNR by an order of magnitude since rectification maps $1/f$ S_z noise to the Xe carrier frequencies.

4.6 Discussion

The cross talk demonstrated in Sec. 4.5 fundamentally undermines one of the desirable features of our system: a scale factor determined only by physics constants. Cross talk influences the device's effective scale factor, and therefore would result in a device which requires calibration if left un-addressed. I suspect the measured cross talk is a result of imperfect rectification of S_z . Therefore, we chose to pursue a new mode of operation which would not use rectification. As discussed in Ch. 2, Xe excitation can be accomplished using polarization or field modulations. In order to avoid cross talk, we chose to transition to PDM operation, thereby avoiding Rb gain modulations and the subsequent rectification.

Chapter 5

Pulse Density Modulation

In this chapter, I will explore the operation of the transverse comagnetometer when driven using two-frequency PDM and no PM. I will show how gated detection can be used to suppress gain modulations as a result of the varying duty cycle of the pulses. I will discuss my choice of drive and feedback schemes, and demonstrate simultaneous dual-species Xe excitation. I then present the stability of our magnetometer using this modulation scheme, and discuss a suspected limiting noise source.

5.1 Modulation and Drive Scheme

In PDM operation, the optical pumping is held fixed ($R(t) = R$). Xe excitation is then driven by modulating the bias field. Since we use a pulsed bias field, the modulation is accomplished by modulating the repetition rate of the bias pulses [Thrasher et al., 2019a]. This modulates the Xe precession frequency, while the Rb remains effectively insensitive to the modulation since each pulse results in a 2π Rb rotation. The field modulation is chosen such that the Xe precesses quickly (aka, the bias field is large) when the Xe and Rb polarizations are anti-parallel (when spin-exchange collisions would reduce \mathbf{K}), and the Xe precesses slowly (aka, the bias field is small) when the Xe and Rb polarizations are parallel

(when spin-exchange collisions would build up \mathbf{K}).

For the PDM results presented in this work, the pulsed bias field is doubly-modulated as

$$B_p(t) = B_{p0}(1 + b_1 \cos(\omega_1 t) + b_2 \cos(\omega_2 t)) \quad (5.1)$$

where B_{p0} is the average field, $\omega_{1,2}$ are the PDM modulation frequencies, and $b_{1,2}$ are the modulation depths. We also define $B_{mod}(t) = B_1 \cos(\omega_1 t) + B_2 \cos(\omega_2 t)$ as the modulated part of the field, where $B_{1,2} = B_{p0} b_{1,2}$.

Given this modulation, the steady state amplitude from Eq. 2.2 becomes

$$K_{\perp} = \frac{\Gamma_S^K S_{\perp}}{\Gamma_2} J_{q_1}\left(\frac{\gamma B_1}{\omega_1}\right) J_{q_2}\left(\frac{\gamma B_2}{\omega_2}\right) \quad (5.2)$$

We define $j_{q_1, q_2} = J_{q_1}\left(\frac{\gamma B_1}{\omega_1}\right) J_{q_2}\left(\frac{\gamma B_2}{\omega_2}\right)$ for use in comparing Xe amplitudes for various drive schemes.

Previously, we have presented results for the drive scheme ($q1^a = 3, q2^a = 1$), ($q1^b = 1, q2^b = 0$) [Thrasher et al., 2019b; Thrasher, 2020]. While this drive scheme did result in large ^{131}Xe amplitude (see Fig. 5.1), we found that when using this drive scheme that frequency feedback (see sec. 5.5) slowly drove the Xe off resonance, as shown in Fig. 5.2, and as discussed further in Sec. 6.5. The results presented in this chapter use the drive scheme ($q1^a = 0, q2^a = 1$), ($q1^b = 1, q2^b = 0$), which did not exhibit closed-loop drift.

5.2 Gated Detection

Up to this point, I have treated the Rb as being effectively unaffected by the pulsed bias field. This broad generalization was bound to rear its ugly head at some point, and here we are. When we have made this assumption in the past, we were basically ignoring anything that happened during the very short pulses and focusing only on the behavior during the




Figure 5.1: The product of the Bessel functions, j_{q_1, q_2} , vs the modulation depth $b_1 = B_0/B_1$. Orange and blue traces are ^{129}Xe and ^{131}Xe , respectively. The second modulation depth is held to $b_2 = 0.9 - b_1$, as discussed in sec. 5.2. TOP: old drive scheme with $(q_1^a = 3, q_2^a = 1)$, $(q_1^b = 1, q_2^b = 0)$. BOTTOM: new drive scheme with $(q_1^a = 0, q_2^a = 1)$, $(q_1^b = 1, q_2^b = 0)$.

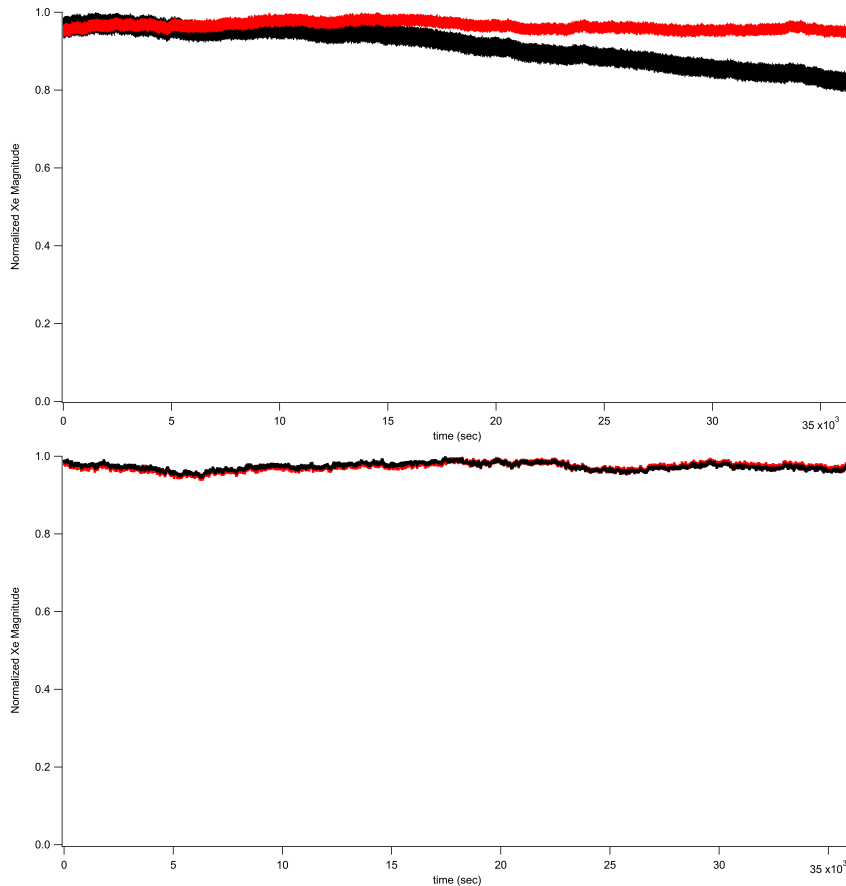


Figure 5.2: Closed loop drift of the normalized Xe magnitudes over long times. Black and red traces are ^{129}Xe and ^{131}Xe , respectively. TOP: old drive scheme with $(q1^a = 3, q2^a = 1)$, $(q1^b = 1, q2^b = 0)$. BOTTOM: new drive scheme with $(q1^a = 0, q2^a = 1)$, $(q1^b = 1, q2^b = 0)$.

much longer time between pulses. In actuality, the Rb precesses so rapidly during each pulse that for that short time the Rb is not providing useful spin-exchange polarization of the Xe nor meaningful readout of the Xe spin precession. When the pulse frequency is unmodulated, this results in a consistent slight degradation of the sensitivity of our magnetometer, which does not have a huge impact on our demodulation. However, when we modulate our pulsing frequency, the duty cycle of the pulses also changes (since the pulsing width is fixed to maintain 2π pulses of the Rb). This means the ratio of the time the Rb spends at zero field vs rapidly precessing is constantly changing with pulsing frequency, resulting in a modulation of the sensitivity of our magnetometer. This is an effect not included in our models, and one which is difficult to account for accurately in post-processing, and so we chose to implement

a new detection scheme which allows us to ignore this effect. We call this scheme "gated detection."

The basic premise behind gated detection is that, when the pulses are turned off, the Rb polarization (and hence the magnetometer gain) will settle into a steady state relatively quickly (compared to the relaxation of the Xe). By turning off the pulses periodically, and only taking measurements when the pulses have been off for some settling time, we can reduce the effect of magnetometer gain modulations from the pulsing modulation. We adjust our pulse modulation waveform to include a multiplicative gating factor, such that the pulses are applied as

$$B_p(t) = (\text{sicos}(\omega_g t) + 1) * (B_{p0} + B_{mod}(t)) \quad (5.3)$$

where ω_g is the frequency at which the pulses are turned off and on. The gating waveform is a square wave with 50 percent duty cycle, such that half of the time no bias field is applied, and half the time twice the un-gated bias field is applied. In this way, the time-averaged bias field experienced by the Xe remains unchanged, despite the gating.

Figure 5.3 shows a gated S_z signal for constant pulsing (no Xe excitation). Above the S_z signal, in blue, is a model of the bias pulses, showing when they are turned on and off due to the gating. The oscillations on S_z while the pulses are on are from interference between the pulses and AC heater drive. After the pulses are gated off, the S_z signal shows transient behavior before settling to a steady state. Data during the second half of the pulses-off period (in red) is averaged together to produce one data point per gating cycle.

Operation Frequencies

The gating frequency, ω_g , was selected to be as low as possible (to allow the Rb time to settle to a steady state) while still satisfying the Nyquist criterion (keeping in mind that ω_g sets the sampling frequency of the magnetometer). The Nyquist frequency is set by the highest frequency signal in our system. In our case, this is the ^{129}Xe precession frequency.




Figure 5.3: Example of gated S_z signal (black) with model of gated B_p (blue). Data is only recorded during the red segments. In each cycle, the red data is averaged to produce a single data point. Note that this B_p model is for visualization purposes only and does not represent the exact pulse times (as the true pulses would be too close together to be seen clearly). The pulse modulation is turned off such that Xe is not excited during this measurement. The oscillations on the black trace are due to mixing between the AC heater drive and the applied pulses.

Since PDM causes a modulation of the instantaneous Xe frequencies, the highest frequency will be $\Omega_0^a * (1 + b_1 + b_2)$, where Ω_0^a is the center frequency, and b_1, b_2 are the modulation depths. Therefore, $\omega_g > 2 * \Omega_0^a * (1 + b_1 + b_2)$.

We chose $b_1 + b_2 < 1$ so that we do not need to flip the direction of the applied field pulses. In addition, we do not want our lowest instantaneous Xe frequencies to be so low that they are sensitive to $1/f$ detection noise, so we chose $\Omega_0^b * (1 - b_1 - b_2) > 2\pi 1 \text{ Hz}$. The lowest Xe frequency also sets the bandwidth of the gyroscope, resulting in a bandwidth of about 1 Hz. The pulsing center frequency ω_{p0} determines $\Omega_0^{a,b}$. As shown in fig. 3.9, we find the pulse area produced by our pulsing circuit is relatively independent of frequency up to around 30kHz. Keeping in mind that our total modulation depth will be less than 1, we chose to operate at a bit less than half of that upper limit, at a center frequency of around 13.2kHz, such that Ω_0^b is around (but not exactly) 10 Hz and Ω_0^a is around 33.7 Hz. Therefore

$b_1 + b_2$ must be less than 0.9, and we are finally able to determine that $\omega_g > 2 * 33.7 \text{ Hz} * 1.9$, around 130 Hz.

Looking at the Rb decay after a single pulse, we found a decay constant of around 0.4 msec. Given that half of the pulses-off part of the gating cycle should be used for measurement, the Rb needs to approach a steady state in about 1/4 of a gating period. This puts an upper limit on the gating frequency of around 600 Hz.

We chose to operate at $\omega_g = 200 \text{ Hz}$. This allows us about three time-constants for settling while maintaining an appropriate sampling frequency and giving us some flexibility to increase our precession frequencies if we ever wanted to.

5.3 Demodulation

When we were demodulating PM signals, we used two single-frequency lock-in detectors and detected at one of several large discrete sideband frequencies. For PDM, however, we are constantly changing the Xe frequencies such that the signal is spread continuously over a large band of frequencies. There is no single frequency we could demodulate at that would give good signal sizes with low cross talk. Therefore, we designed a custom demodulation waveform. This demodulation scheme was developed from the idea that the Xe precesses a fixed amount during each applied pulse, no matter what the pulse repetition rate is, and we know when each pulse is applied. Therefore, we have some idea of what the expected Xe precession, α (Eq. 2.11), is at any time. From there, we initially chose to demodulate using an integral over α (see [Thrasher et al., 2019b]), but found that integrating over time was not a good enough approximation of integration over alpha. To avoid sampling unevenly in time, we decided to use a real-time least-squares fit for our demodulation.

To get our fitting function, we plug Eq. 5.1 into 2.4b. We find the Rb magnetometer

signal to be

$$\begin{aligned}
S_z = & A_{\perp}^a (\sin(\delta^a - \epsilon_z) \cos(\alpha^a) + \cos(\delta^a - \epsilon_z) \sin(\alpha^a)) \\
& + A_{\perp}^b (\sin(\delta^b + \epsilon_z) \cos(\alpha^b) + \cos(\delta^b + \epsilon_z) \sin(\alpha^b))
\end{aligned} \tag{5.4}$$

where $A_{\perp} = s\gamma^s b_K^S K_{\perp} R / \Gamma'^2$. Since we detect S_z using a balanced Faraday detector, we wrap any detection scaling factors into A_{\perp} and add an overall DC offset term, C . We then re-write our fit function in terms of a matrix product,

$$\begin{bmatrix} S_{z,1} \\ S_{z,2} \\ \vdots \\ S_{z,N} \end{bmatrix} = \begin{bmatrix} \cos(\alpha_1^a) & \sin(\alpha_1^a) & \cos(\alpha_1^b) & \sin(\alpha_1^b) & 1 \\ \cos(\alpha_2^a) & \sin(\alpha_2^a) & \cos(\alpha_2^b) & \sin(\alpha_2^b) & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \cos(\alpha_N^a) & \sin(\alpha_N^a) & \cos(\alpha_N^b) & \sin(\alpha_N^b) & 1 \end{bmatrix} \times \begin{bmatrix} A_{\perp}^a \sin(\delta^a - \epsilon_z) \\ A_{\perp}^a \cos(\delta^a - \epsilon_z) \\ A_{\perp}^b \sin(\delta^b + \epsilon_z) \\ A_{\perp}^b \cos(\delta^b + \epsilon_z) \\ C \end{bmatrix} \tag{5.5}$$

which is of the form $S_z = M.u$, where S_z is our recorded time series, M is a matrix of functions of known variables, and u is a matrix of our unknown variables. In order to improve the goodness of the fit, we measure N data points before fitting, such that each row of S_z and M is evaluated at a discrete time. S_z is then an $N \times 1$ matrix and M is $N \times 5$. For the data in this chapter, we chose $N=20$. Given our gating frequency of 200 Hz, this results in an effective sampling frequency of 10 Hz.

We can solve the matrix expression, Eq. 5.5, in real-time using matrix inversion as $u = (M^t M)^{-1} M^t S_z$. This calculation is performed using LabVIEW on the host computer, as discussed in Sec. 3.4.

As shown in Eq. 5.5, our demodulation outputs give $A_{\perp} \sin(\delta + \mathbf{s}\epsilon_z)$ and $A_{\perp} \cos(\delta + \mathbf{s}\epsilon_z)$, which can be solved to give A_{\perp} and $\delta + \mathbf{s}\epsilon_z$, but crucially NOT δ . We will show in section 5.5

how this affects our comagnetometry.

5.4 Xe Excitation

I will now present measurements obtained for simultaneous excitation of both Xe species using PDM. Figure 5.4 shows the applied (gated) pulsing frequency and resulting S_z signal. Our model for S_z (lines) from Eq. 5.4 matches our measured data (dots) well. By re-casting the S_z vs time signal to S_z vs pulse number, we clearly see the sum of two sine waves expected for dual-species precession.

Figure 5.5 shows the NMR resonance lineshapes for simultaneous excitation of the two Xe species using PDM, demodulated using our least-squares-fit. Our Xe amplitudes (in field units) are $15.8 \mu\text{G}$ for ^{131}Xe and $13.8 \mu\text{G}$ for ^{129}Xe . We measure the linewidths (Γ_2) by scanning the drive frequencies through the resonance frequency (one species at a time) while measuring the in-phase and quadrature signals, then simultaneously fitting both traces to Lorentzian lineshapes. These measurements are taken with the x-compensation field applied, which narrows the measured linewidths. Because the optical pumping is DC for PDM, the applied x-compensation field is DC, in contrast to PM and hybrid operation which both use AC x-compensation fields. We find $\Gamma_2^b = 19.1 \text{ mHz}$ and $\Gamma_2^a = 10.7 \text{ mHz}$.

Using our custom demodulation waveform, we found the cross talk on our Xe signals was reduced significantly compared to our demodulation for PM. Our measurement of cross talk for PDM uses the same general principles as the cross talk measurement for PM. We begin by exciting both Xe species on resonance and setting the demodulation phases such that the measured phase shift, $\delta + \mathbf{s}\epsilon_z$, is zero on resonance. We then detune the ^{129}Xe drive frequency, ω_2 , up by 100 mHz such that the two Xe species will be on resonance for different bias field values. We detune the ^{129}Xe drive up to avoid any accidental excitation of ^{129}Xe from the third harmonic of the ^{131}Xe modulation (even though this harmonic should be minimal given our cosine-wave field modulation). We then record the Xe amplitudes and




Figure 5.4: S_z signals obtained using PDM.

S_z signals (blue) obtained using PDM vs time (middle) or pulse number (bottom). Dots are data, lines are models. Red is model of the frequency of the applied pulses, including pulse gating.

phase shifts as we scan the pulsing frequency. We set the scan range such that each Xe species starts off resonance, passes through resonance, and ends off resonance. We convert the amplitudes and phase shifts to in-phase and quadrature signals. Rather than plotting the in-phase and quadrature signals vs the pulsing frequency, we find it useful to recast the pulsing frequency in terms of the Xe Larmor precession frequencies one would expect based on the pulsing frequency. Plotting in this way emphasizes to the reader that, although the ^{129}Xe resonance peak may appear much narrower, that is simply a result of its larger gyromagnetic ratio, such that Ω^a scans a larger range than Ω^b for the same bias field scan

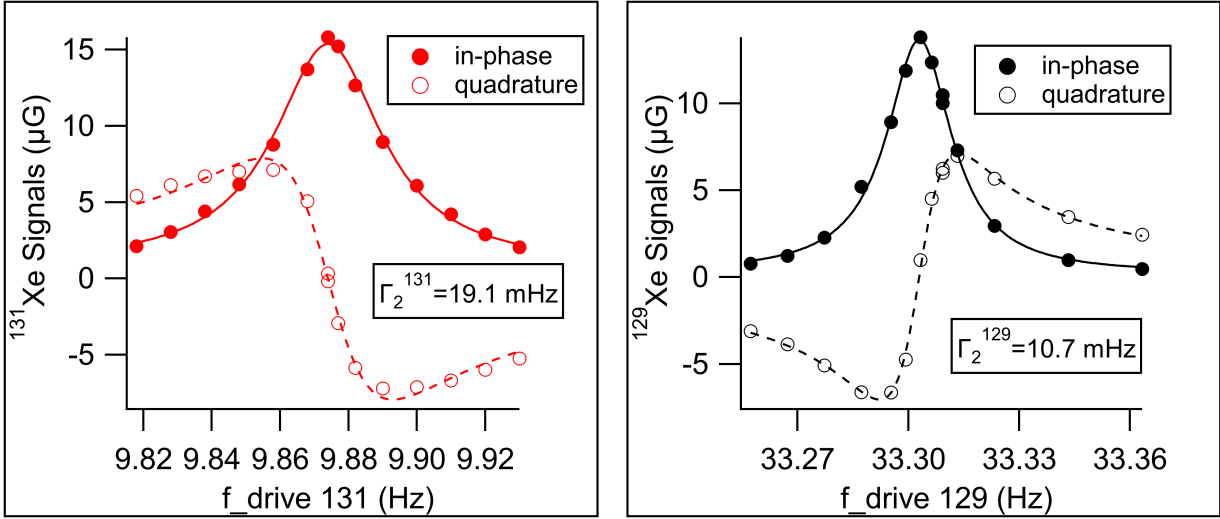


Figure 5.5: Linewidth measurements for ^{131}Xe (left) and ^{129}Xe (right). Data (dots) are fit (lines) to a Lorentzian lineshape.

range.

Figure 5.6 shows a cross talk measurement for PDM. We quantify cross talk in the same way we did for PM (see 4.5). We estimate $\beta = Q_{pp}^b/Q_{pp}^a = 0.06$ (when isotope a is on resonance) and $\beta' = Q_{pp}^a/Q_{pp}^b = 0.0004$ (when isotope b is on resonance), where Q_{pp} is the peak-to-peak quadrature signal. We find that for PDM Q_{pp} is so small for the off-resonant species that we must look at the residuals to the fit to accurately estimate Q_{pp} . These are both marked improvements over the $\beta = 0.17$ and $\beta' = 0.07$ measured for PM. We believe these improvements are the result of our least-squares fit demodulation. If we ever go back to PM operation, we suspect we could improve the PM cross talk by using a least-squares fit demodulation.

5.5 Feedback

As discussed in Sec. 4.4 and Sec. 2.3, open loop comagnetometry results in a measurement of rotation frequency which is sensitive to transfer function noise and to $i\omega\epsilon_z$. We therefore design a feedback scheme which addresses both of these issues. We utilize a two-stage

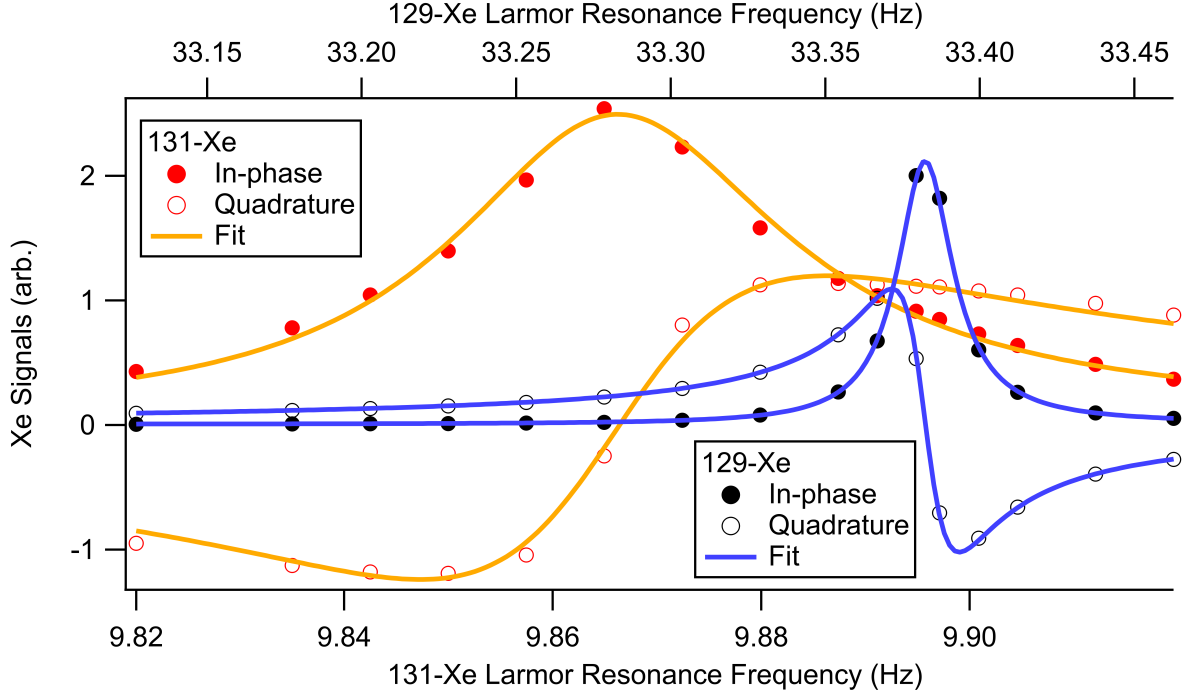


Figure 5.6: Cross talk measurement for PDM. Closed circles are in-phase data, open circles are quadrature data, and lines are fits. The quadrature and in-phase data were fit simultaneously using a global fit.

feedback scheme. First, we apply two frequency feedback loops, which adjust the two modulation frequencies (ω_1, ω_2) to hold the Xe phase shifts to zero for each species. The resulting modulation frequencies from these first-stage feedback loops are then used as the inputs of the second-stage feedback loop. The second-stage feedback adjusts the DC z-field to hold $\omega_1 + \omega_2$ constant. The first-stage feedback allows us to use the drive frequencies as the Xe precession frequencies, thereby bypassing the need to use the Xe transfer function. The second-stage feedback loop holds $i\omega\epsilon_z = 0$. I will now describe these feedback stages in greater detail.

Using our drive scheme $\omega^b = \omega_1$, $\omega^a = \omega_2$, we can construct very simple frequency feedback loops. Since our demodulation gives $\delta + s\epsilon_z$ rather than just δ , our frequency




Figure 5.7: The gain function used for the drive frequency feedback. Inverted zero with lead compensation gain.

feedback loops are of the form

$$\begin{aligned}\omega_1 &= G^b(\delta^b + \epsilon_z) \\ \omega_2 &= G^a(\delta^a - \epsilon_z)\end{aligned}\quad (5.6)$$

where $G^{a,b}$ are frequency-dependent gain functions. In order to have high gain at low frequency while obeying the stability conditions for feedback (see [Bechhoefer, 2005]), we utilize inverted zero with lead compensation gain, as shown in figure 5.7. The gain function for frequency feedback is

$$G^{a,b} = P\left(1 + \frac{\omega_0}{i\omega}\right)\left(\frac{1 + A\frac{i\omega}{\omega_1}}{1 + \frac{i\omega}{\omega_1}}\right)\quad (5.7)$$

where P and A are amplitude coefficients and ω_0 and ω_1 are corner frequencies. This feedback is applied digitally in LabVIEW.

If we only applied the frequency feedback loops given in Eq. 5.6, our comagnetometry calculation in the high gain limit ($G \rightarrow \infty$) would not give the rotation frequency directly,

but rather

$$\frac{\rho\omega^b - \omega^a}{1 + \rho} = \Omega^R + i\omega\epsilon_z \quad (5.8)$$

such that our measurement of rotation is sensitive to the derivative of the magnetometer phase shift, ϵ_z . This would negatively impact our FSF, and so we apply another feedback loop to eliminate the effect of ϵ_z .

In order to design a feedback loop to hold ϵ_z to zero, we don't actually measure ϵ_z . Instead, we remember that ϵ_z is proportional to the z-field experienced by the Rb. Therefore, stabilizing the non-pulsed z-field will hold ϵ_z to a constant ($\epsilon_z \ll 1$) such that $i\omega\epsilon_z$ is zero. In order to apply such a field stabilization, we wish to find an error signal which is dependent on B_z , but NOT on ϵ_z . This is easily constructed from a sum of the Xe drive frequencies determined by the frequency feedback, since the sign of ϵ_z is species-dependent. We are then left with a field feedback of the form

$$B_{FB} = G^{Bz}(\omega_1 + \omega_2 - \omega_{sum}^0) \quad (5.9)$$

where B_{FB} is the applied field correction, G^{Bz} is the gain of the field feedback, and ω_{sum}^0 is the set point of the sum of the modulation frequencies. The total field experienced by the Rb is then $B_{tot} = B_{z0} + B_{FB}$ where B_{z0} is the stray z-field. Figure 5.8 shows G^{Bz} , which is inverted zero gain of the form

$$G^{Bz} = P\left(1 + \frac{\omega_0}{i\omega}\right) \quad (5.10)$$

where we did not need the lead compensation term to attain stable, high gain feedback to B_z .

Using all three of the feedback loops (Eqs. 5.6 and 5.9) together in the high gain limit, the field feedback holds $B_{FB} \rightarrow -B_{z0}$, such that B_{tot} is held constant. As a result, $i\omega\epsilon_z$ goes to zero in our comagnetometry calculation Eq. 5.8, and we are left with just Ω^R .

Having implemented feedback loops to suppress field noise on our measure of rotation,




Figure 5.8: The gain function used for field feedback. Inverted zero gain.

we must now quantify how well we did by measuring our field suppression factor (FSF). As described in section 2.4, we apply an AC z-field, B_{app} , at 5 mHz and calculate the ratio of the applied frequency noise $\gamma^b B_{\text{app}}$ to the measured rotation frequency noise (at the frequency of the applied field). We find an FSF of approximately 4700 for the measurement shown in Fig. 5.9.

5.6 SNR

In order to understand one of the breakthroughs that lead to the PDM SNR and stability improvements presented in this chapter, I must explain the motivation for a feature of our demodulation scheme that ended up temporarily limiting our results.

Two suspected sources of instability in our system were the magnetometer phase shift, ϵ_z , and the transverse pumping rate. Fluctuations in the transverse pumping rate could be caused by power or frequency fluctuations in our pump lasers. Unless otherwise stated, the pump and probe lasers in our system are run open loop (apart from the constant temperature and constant current controls built into the laser drivers). Fluctuations in the pumping rate




Figure 5.9: Field Suppression Factor measurement for PDM for an applied field at 5 mHz. Blue: the expected frequency noise on ^{131}Xe due to the applied field modulation. Red: the measured ^{131}Xe frequency noise. This trace is lower than the blue trace due to the field feedback. Green: the measured rotation frequency noise. This is lower than the red trace as a result of comagnetometry.

would affect our spin-exchange optical pumping and the magnetometer gain.

In order to monitor both our magnetometer gain and ϵ_z simultaneously, we applied a known rotating field, B_{perp} , and demodulated this applied field. We applied B_{perp} at a frequency close to the average precession frequencies of the two Xe species in order to monitor the quality of the magnetometer near the frequencies of interest. The amplitude of the demodulated signal, A_{perp} , would be proportional to the magnetometer gain, while the demodulated phase, δ_{perp} , would give a direct measurement of ϵ_z . We could, theoretically, use A_{perp} to stabilize the pump lasers and δ_{perp} to stabilize B_{z0} . In practice, these additional feedback loops did not improve our performance. The SNR on our measurement of δ_{perp} was poor enough that other feedback methods did a better job of stabilizing B_{z0} . The A_{perp}

measurement did tell us something about the magnetometer gain, but this turned out to be an insufficient way to build a feedback loop as we had only one observable and four control points (the currents and temperatures of our two pump lasers). In the end, we stopped applying B_{perp} , but left in the B_{perp} demodulation in case we decided to use it again in the future. This turned out to be a mistake.

We wrongly assumed that the unused demodulation would have no affect on the quality of our demodulation of the Xe signal. However, the B_{perp} demodulation frequency was left within the modulation band of the Xe, meaning that as the bias pulse frequency changed, $\Omega^{a,b}$ swept through the frequency of this extra unused demodulation. This degraded the quality of our Xe demodulation because, when $\Omega^{a,b} = \omega_{perp}$, some of the Xe signal power ended up in the B_{perp} demodulation channels. When I moved the B_{perp} demodulation frequency out of the Xe modulation frequency range, our measured SNR improved immediately!

Figure 5.10 shows how moving the B_{perp} demodulation affected the open-loop phase noise (and hence SNR). The SNR is obtained from the average phase noise above 0.2 Hz, as described in Sec. 2.4. When the B_{perp} demodulation frequency was within the Xe modulation band ($\omega_{perp} = 15.1$ Hz), the SNR was $1460 \sqrt{Hz}$ for ^{131}Xe and $2627 \sqrt{Hz}$ for ^{129}Xe . After moving the B_{perp} demodulation frequency out of the Xe modulation band ($\omega_{perp} = 65.1$ Hz), we found $\text{SNR}^b = 2650 \sqrt{Hz}$ and $\text{SNR}^a = 2680 \sqrt{Hz}$. While this is only a marginal improvement for ^{129}Xe , SNR^b improved by almost a factor of two.

It is important to note that while the B_{perp} demodulation was also limiting our bias instability (to around 1 μHz), the improvement seen after moving the B_{perp} demodulation frequency was the result of other ongoing stability improvement efforts. These efforts, and the resulting bias instability improvement, are presented in the following section.

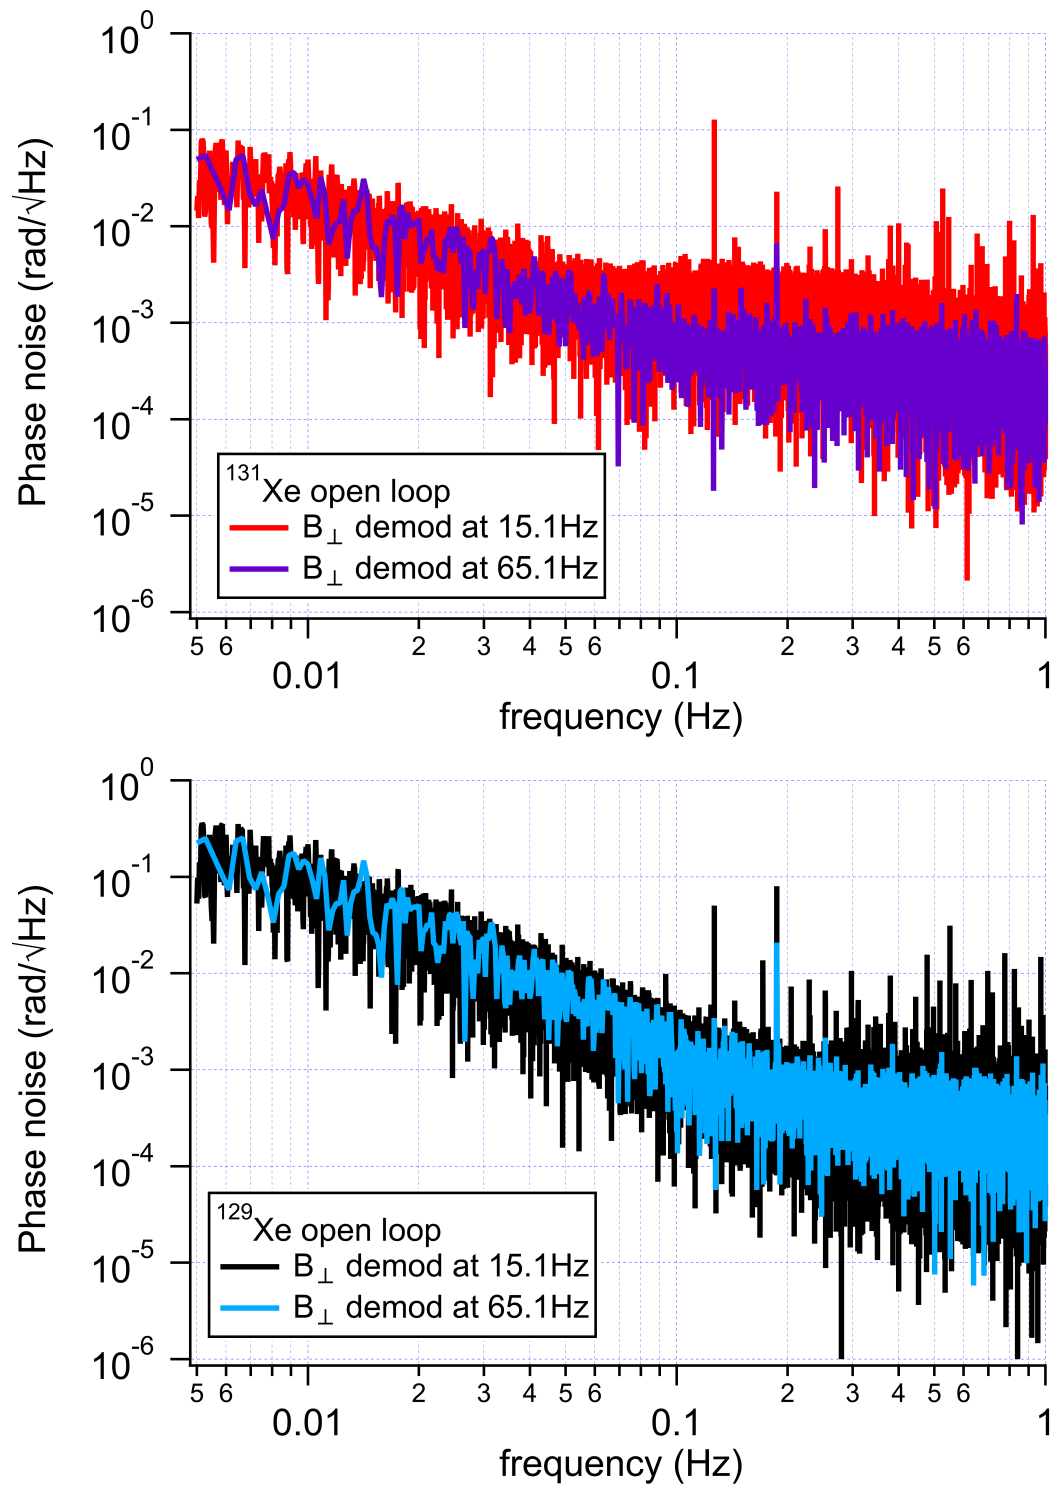


Figure 5.10: Open loop phase noise measurements used to calculate SNR. Top: phase noise for ^{131}Xe . Bottom: phase noise for ^{129}Xe . Each graph shows the phase noise at two different B_{\perp} frequencies.

5.7 Transverse Field Sensitivity and Stability

Based on our Bloch Equation derivation, we expect our calculation of rotation to be sensitive to $i\omega\epsilon_z$, but not to transverse fields. However, we have observed experimentally that our measurement of rotation has transverse field sensitivity that goes as the product $\Omega_x\Omega_y$. It is likely that one or more of the assumptions we made in the Bloch Equation derivation was too broad, and caused us to discard one or more term that is not negligible.

Indeed, in my original derivation of the transverse Rb polarization, I dropped the $1/\Gamma^3$ term in Eq. 2.2 when writing S_+ in Eq. 2.4a. If I include this term, S_+ becomes

$$S_+ = \frac{R(t)}{\Gamma'} + i\frac{R(t)\Omega_z}{\Gamma'^2} - \frac{R(t)}{\Gamma'}(\Omega_y^2 + \Omega_z^2) + i\frac{R(t)}{\Gamma'}\Omega_x\Omega_y \quad (5.11)$$

which includes a term $\Omega_x\Omega_y$. Continuing the derivation mostly as described in Ch. 2 but solving for K_x and K_y rather than K_\perp , I find

$$\delta = \frac{K_y}{K_x} = \frac{(\omega - i\Gamma_1)(S_y(\omega - i\Gamma_2) + iS_x\Omega_z)}{S_x(\omega - i\Gamma_1)(\omega - i\Gamma_2) - S_x(j\Omega_x)^2 - S_y(j\Omega_x\Omega_y + (i\omega + \Gamma_1)\Omega_z)} \quad (5.12)$$

which includes both $\Omega_x\Omega_y$ and Ω_x^2 dependence.

In order to verify this form of δ , I make the assumption $S_y \ll 0$ (which is equivalent to assuming small ϵ_z). This does get rid of the $\Omega_x\Omega_y$ term, but it also simplifies the expression significantly to a more manageable

$$\delta = \frac{-\Omega_z}{(i\omega + \Gamma_2) + (j\Omega_x)^2/(i\omega + \Gamma_1)}. \quad (5.13)$$

We can easily test Eq. 5.13 by measuring the transfer function for a mis-set x-compensation field (such that Ω_x is non-zero). We measure the transfer function by applying a B_z chirp such that Ω_z is known. We simultaneously record the response to the chirp by measuring δ as usual (assuming negligible ϵ_z). The transfer function is then $-\Omega_z/\delta$. Without including

the higher-order terms in S_+ , we previously found the transfer function in Eq. 2.18 to be $i\omega + \Gamma_2$. Including the higher-order terms in S_+ , we now find the transfer function from Eq. 5.13 to be $(i\omega + \Gamma_2) + \Omega_x^2/(i\omega + \Gamma_1)$. In Fig. 5.11, I fit the measured transfer function to these two forms of the transfer function. The updated transfer function fits the measurement much better (especially for ^{129}Xe) than the transfer function from Eq. 2.18 for a non-zero Ω_x . For $\Omega_x = 0$, the two two transfer functions are the same.

I have developed the following procedure to minimize sensitivity to transverse field drifts on the measurement of rotation frequency. The basic idea is to use the rotation sensitivity to transverse fields as the observable when nulling the transverse field setpoints. This is in contrast to our usual nulling procedure which uses S_z as the observable, as described in Ch. 3.

We take a series of measurements to scan for the B_x and B_y values which minimize sensitivity to transverse fields. We first set up our rotation measurement, including f_{sum} feedback to B_z . We then apply a low frequency oscillating B_y . If the DC component of B_x is non-zero, we see the oscillation from B_y on our measurement of rotation (since we observe the rotation to have a $B_x * B_y$ dependence). We plot rotation vs B_y to produce Lissajous figures. For the Lissajous figures shown in Fig. 5.12, a bandpass filter has been applied to the rotation to show only the component of the signal at the frequency of the applied oscillating transverse field.

We then vary the DC B_x to find the point where the B_y oscillation is no longer present on the rotation (aka, when the slope of the Lissajous figure is zero). At this point, we have set the DC B_x to zero such that we are minimally sensitive to noise on B_y .

We then repeat this procedure for the other transverse direction, this time applying an AC B_x and varying the DC B_y between measurements to minimize the rotation sensitivity to B_x noise. We then plot the slope $(\frac{d\Omega^R}{dB_{x,y}})$ vs DC field and fit a line to find the DC B_x and B_y that minimize sensitivity to transverse field noise.

After we go through this procedure, we turn off the slowly oscillating field and take a




Figure 5.11: Comparison of the transfer function with and without higher-order transverse field effects. Solid lines are measurements of the transfer function for a purposely mis-set x-compensation field, dashed are fits to one of the two proposed transfer functions. Top: magnitude of the transfer function. Bottom: phase of the transfer function. (Note that this experiment was actually performed later using hybrid operation, but the transfer function has the same form for PDM.)




Figure 5.12: (a): Rotation vs AC B_y for different values of DC B_x . (b): Rotation vs AC B_x for different values of DC B_y . A bandpass filter is applied the rotation in both figures. A slope of zero indicates minimal sensitivity of rotation to transverse field noise. DC field values have an arbitrary offset.

long measurement of rotation. Figure 5.13 shows MDEVs of rotation taken before and after the transverse field nulling procedure. We see that, when the DC transverse fields are mis-set, we are limited by transverse field noise.

This represents benchmark stability performance for PDM drive, as shown in figure 5.14 comparing our new MDEV to our previously presented result [Thrasher et al., 2019b]. We reach a bias instability of 420 nHz after 1700 seconds. This is more than a factor of two improvement! Based on the slope of the $\tau^{-1/2}$ trend (taking into account a factor of $\sqrt{2}$ as a result of using an MDEV, rather than an ADEV, see Sec. 2.4), our ARW is $17 \mu\text{Hz}/\sqrt{\text{Hz}}$. This ARW is actually more than a factor of two worse than our previously reported ARW, meaning that we must measure longer to get to the same bias instability. This decrease is likely due to the reduced ^{131}Xe signal size when we changed our drive scheme to avoid long-term drifts. What we have lost in SNR we have gained in long-term performance.




Figure 5.13: Rotation MDEVs before (purple) and after (green) nulling transverse fields.

5.8 Chapter Summary

In this chapter, I have discussed PDM operation of the transverse NMR gyroscope. I have presented gated detection, which is necessitated by the modulation of the pulse frequency. I have presented the drive scheme used for this work, and described our new least-squares fit matrix inversion demodulation. I demonstrate simultaneous excitation of both Xe species using PDM, and show that our new detection reduces cross talk significantly. I apply feedback to the drive frequencies and the z-field. I present best-to-date bias instability for PDM operation.

Based on our explorations of transverse field noise sensitivity, I suspect the bias instability is currently limited by transverse field noise. That being the case, we look to the largest transverse field in our system: the Rb spin-exchange field. Though we apply a compensation




Figure 5.14: Comparison of previously reported PDM MDEV and new best data from this thesis.

field to cancel the Rb spin-exchange field, we fix the compensation field size before starting a measurement. That means that any changes to the Rb polarization will change how well compensated we are. Since the power and frequency of our pump lasers are uncontrolled, it would not be surprising to find noise on the Rb polarization. Therefore, in deciding how to move forward, we wanted to consider a mode of operation which would move the Rb polarization to AC. It is possible that our demodulated signals will be less sensitive to drifts on an AC amplitude than to DC drifts of the Rb polarization. From there, we decided to move forward with a method that involves both PM and PDM, which we call "hybrid" operation.

Chapter 6

Hybrid Modulation

In this chapter, I describe a new approach to continuous operation of a transverse spin-exchange optically pumped NMR gyroscope which utilizes modulation of both the applied bias field and the optical pumping [Sorensen and Walker, 2023]. I demonstrate simultaneous, continuous excitation of ^{131}Xe and ^{129}Xe using this hybrid modulation approach, and real-time demodulation of the Xe precession using a custom least-squares fitting algorithm. I present rotation rate measurements with this device, with a common field suppression factor of ~ 1400 , angle random walk of $21 \mu\text{Hz}/\sqrt{\text{Hz}}$, and bias instability of $\sim 480 \text{ nHz}$ after ~ 1000 seconds. I present results of comagnetometry under three conditions: open loop, magnetic field feedback, and combined field and drive frequency feedback.

6.1 Modulation and Drive Scheme

As discussed in Ch. 2, the spin dynamics for the Xe nuclear polarization (\mathbf{K}) and the Rb electron polarization (\mathbf{S}) are described by the Bloch equations. The equation for the transverse Xe polarization for species a (with analogous equation for species b) can be written

$$\frac{dK_+^a}{dt} = -(s^a i\Omega_z^a + \Gamma_2^a)K_+^a + \Gamma_S^a S_+ + s^a i\Omega_+^a K_z^a \quad (6.1)$$

where Ω is the Xe resonance frequency as given in Eq. 1.5, Γ_2 is the transverse relaxation rate, and Γ_3^g is the spin-exchange rate constant.

From Eq. 6.1, we see that Xe can be continuously excited by modulating Ω_+ [Volk et al., 1980], S_+ [Sorensen et al., 2020], and/or Ω_z [Thrasher et al., 2019a]. (Since the transverse geometry of the device results in negligible K_z , we do not modulate Ω_+ in this work.) In previous chapters, I have focused on Xe excitation by modulating only one of S_+ (Ch. 4 on PM operation) or Ω_z (Ch. 5 on PDM operation). Hybrid operation builds upon these works by modulating both S_+ and Ω_z simultaneously, thereby combining PM and PDM. The modulations for hybrid operation are applied as follows.

The pumping rate, $R(t)$ (and hence S_+) is modulated along \hat{x} as a square wave (the sign of a cosine), such that

$$R(t) = R \text{sign}(\cos \omega_{PM} t) \quad (6.2)$$

where ω_{PM} is the frequency of the applied modulation. This PM waveform is shown in Fig. 6.1. Note that hybrid operation uses only a single PM frequency, whereas pure PM operation doubly-modulated $R(t)$ (see Eq. 4.1).

We also modulate the repetition frequency of the applied bias pulses (and hence Ω_z). Because of the slow response of the Xe (compared to the Rb), the short pulses applied at a rate $\omega_p(t)$ result in an effective field $B_p(t) = \omega_p(t)/\gamma^S$ experienced by the Xe. The Rb is effectively insensitive to this field since each pulse results in a 2π Rb rotation [Korver et al., 2013]. We modulate ω_p sinusoidally, resulting in the Xe nuclei experiencing a time-averaged pulsed z-field

$$B_p(t) = B_{p0}(1 + b_1 \cos(\omega_{PDM} t)) \quad (6.3)$$

where B_{p0} is the average field, b_1 is the modulation index, and ω_{PDM} is the frequency of the applied modulation. This PDM waveform is shown in Fig. 6.1. Hybrid operation uses only a single PDM frequency, whereas pure PDM operation doubly-modulated B_p (see Eq. 5.1).

Plugging these modulations into the Bloch equation (see Ch. 2), we obtain a resonance




Figure 6.1: Applied modulation waveforms. Top: The pump light polarization is modulated as a square wave. Bottom: The bias field is modulated sinusoidally with a DC offset.

condition (see Eq. 2.2) for exciting ^{129}Xe in hybrid operation (with analogous equation for ^{131}Xe)

$$\omega_d^a = p^a \omega_{PM} + q^a \omega_{PDM} \quad (6.4)$$

where ω_d is the drive frequency of the Xe, p and q are integers, and p is odd (as a consequence of square wave modulation). In principle, the NMR can be driven with any combinations of p and q that satisfy this equation for both isotopes. Our choice of p and q is informed by the desire to maximize our simultaneous Xe polarizations. From the steady state solution to the Xe Bloch equation, we know that the magnitude of the transverse Xe polarization K_{\perp} is proportional to the product of the Fourier coefficients from the two modulations, such that for species a (and similarly for species b)

$$K_{\perp}^a = \frac{R\Gamma_S^a}{\Gamma'\Gamma_2^a} \frac{2}{p^a\pi} J_{q^a} \left(\frac{\omega_d^a b_1}{\omega_{PDM}} \right) \quad (6.5)$$

where Γ' is the magnetic width of the Rb magnetometer, the factor $2/p\pi$ comes from the Fourier decomposition of the polarization modulation, and J_q (the q -th Bessel function of the first kind) arises from the sinusoidal modulation of the bias field (through the Jacobi-Anger identity, see 2.13). We define an amplitude coefficient $j_{pq}^a = \frac{2}{p^a\pi} J_{q^a} \left(\frac{\gamma^a B_{p0} b_1}{\omega_{PDM}} \right)$ (with analogous result for j^b). In general, j^a and j^b will not maximize at the same modulation index, but we look for a scheme that gives large j for both species at the same b_1 . Furthermore, we require $0 < b_1 < 1$ so that we do not need to switch the direction of the applied z-pulses. Figure 6.2 shows j versus b_1 for our chosen drive scheme, $(p^a, q^a) = (1, 1)$ and $(p^b, q^b) = (1, -1)$. In this chapter, we use $b_1 = 0.78$, which approximately maximizes the sum of the amplitude coefficients. We also chose this drive scheme to enable the dual species, dual output feedback scheme to be discussed in section 6.5.




Figure 6.2: Amplitude coefficient (j) vs PDM modulation index (b_1) for ^{129}Xe (orange) and ^{131}Xe (blue) given a drive scheme $\omega_d^a = \omega_{PM} + \omega_{PDM}$, $\omega_d^b = \omega_{PM} - \omega_{PDM}$.

6.2 Demodulation

In order to monitor the Xe precession, we use a z-probe laser and a balanced Faraday detector to measure a signal proportional to S_z . From the Bloch equation derivation (Eq. 2.4b),

$$S_z = -\frac{R(t)}{\Gamma'^2}(\Omega_y^S - \frac{\Omega_z^S}{\Gamma'}\Omega_x^S) = \frac{-R(t)}{\Gamma'^2}\text{Im}[\gamma^S b_K^S K_+ e^{-i\epsilon_z}]. \quad (6.6)$$

where we define

$$\epsilon_z = \tan^{-1}(\gamma^S B_{z0}/\Gamma') \quad (6.7)$$

as the magnetometer phase shift that results from low frequency z-field drift (B_{z0}), as discussed in 2.2. We make the substitution $K_+ = K_\perp e^{-si\phi}$ where K_\perp and ϕ are the amplitude and phase of the transverse Xe precession, respectively. Then we can see that the

imaginary term in Eq. 6.6 will include the term $-\sin(\mathbf{s}^a \phi^a + \epsilon_z)$, which can be re-written as $-\mathbf{s} \sin(\phi^a + \mathbf{s}^a \epsilon_z)$ (plus analogous terms for species b). We then make the substitutions $\phi^{a,b} = \delta^{a,b} + \alpha^{a,b}$. The variable

$$\alpha = \int (\omega_d + \gamma b_1 B_{p0} \cos(\omega_{PDM} t)) dt \quad (6.8)$$

is the expected Xe precession phase if the only field present was an ideal pulsed bias field and we were driving perfectly on resonance. Note that α has the same form as in the pure PDM case (see Eq. 2.11 and Eq. 5.1), but with one fewer field modulation. The variable δ is the Xe precession phase shift (aka, how far the measured precession deviates from the expected phase α). We then find that our measured S_z is of the form

$$\begin{aligned} S_z = & [A^S \text{sign}(\cos(\theta_{PM})) + C^S] \times \\ & [A^a \sin(\alpha^a) \cos(\delta^a - \epsilon_z) + A^a \cos(\alpha^a) \sin(\delta^a - \epsilon_z) \\ & + A^b \sin(\alpha^b) \cos(\delta^b + \epsilon_z) + A^b \cos(\alpha^b) \sin(\delta^b + \epsilon_z) \\ & + C^y] + C^{PD} \end{aligned} \quad (6.9)$$

where θ_{PM} is $\omega_{PM} t$ plus an arbitrary phase offset, $A^{S,a,b}$ are amplitude constants, and $C^{S,y,PD}$ are DC offsets. If our experimental apparatus was perfect, all of the C constants would be zero, but alas we must allow for imperfections. Any asymmetry in the polarization modulation is represented as C^S , C^y is from any DC y-fields, and C^{PD} is a detection offset (eg. from differential photodiode mis-balance).

Figure 6.3 shows our recorded Faraday signal (dots) and our model waveform from Eq. 6.9 (lines). In order to see how this signal is a combination of the precession of two Xe species, we can plot this signal in ways that account for our modulations. First, we multiply by the PM modulation square wave to obtain a "rectified" signal. Then, we plot the rectified signal vs pulse number (rather than time) to account for the pulse density modulation. We are left




Figure 6.3: Measured (dots) and theoretical (lines) S_z measurements, converted to field units. Top: raw signal vs time. Middle: rectified signal vs time. Bottom: rectified signal vs pulse number.

with a plot that is clearly a sum of two sine waves. This treatment of the S_z signal is used for visualization purposes only, not for demodulation.

When we were demodulating purely PM signals [Sorensen et al., 2020], we used two single-frequency lock-in detectors and detected at one of several large discrete sideband frequencies. Now that we have added PDM, however, we are constantly changing the Xe frequencies such that the signal is spread continuously over a large band of frequencies. There is no single frequency we could demodulate at that would give large signal content. Further, any single frequency that showed reasonable signal content from one Xe species would also show signal from the other Xe species, due to the overlap in their frequency ranges, meaning that we could not isolate the signal from a single Xe species. Therefore, we designed a custom least-squares fit [Press et al., 1992] waveform based off of the matrix-inversion demodulation method developed for PDM in Sec. 5.3. This demodulation scheme was developed from the idea that the Xe precesses a fixed amount during each applied pulse no matter what the pulse repetition rate is. We know when each pulse is applied, and so we know what the expected Xe precession (α) is at any given time. Those conclusions remain unchanged when we add PM to PDM for hybrid operation, though the fit function does become more complicated.

We use the known form of our measured signal to design our custom fitting function. If we expand Eq. 6.9, each term in the resulting expression will include one known factor and one unknown factor. Known variables include α^a , α^b , and θ_{PM} . Unknowns include A^S , A^a , A^b , C^S , C^y , C^{PD} , and sine and cosine of $\delta^a - \epsilon_z$ and $\delta^b + \epsilon_z$. On the other side of the expression, S_z is our measured Faraday rotation signal. We can re-cast the expanded expression as a matrix equation, $S_z = M.u$, where M is a matrix where each row contains all of our various known functions and variables evaluated at a specific time, and u is a column matrix of our unknowns. We can solve this expression for u through the matrix-inversion $u = (M^t M)^{-1} M^t S_z$.

Given that there are ten fit parameters, we measure multiple points before taking a

least-squares fit [Press et al., 1992] to improve the goodness of fit. Through matrix inversion, we can solve for our unknowns at a series of discrete points, and then make a least-squares fit to minimize the root-mean-square difference between the data and our fitting function. We measure 50 points before calculating u , such that S_z is 50×1 and M is 50×10 . Given our gated (see Sec. 5.2) sampling frequency of 200 Hz, this results in an effective data acquisition rate of 4 Hz. The full matrix expression is given in Eq. 6.10, where $\Theta = \text{sign}(\cos(\theta_{PM}))$, $s(\alpha) = \sin(\alpha)$, $c(\alpha) = \cos(\alpha)$, and subscripts 1, 2, ...50 represent points taken at different times.

$$\begin{bmatrix} S_{z,1} \\ S_{z,2} \\ \vdots \\ S_{z,50} \end{bmatrix} = \begin{bmatrix} s(\alpha_1^a) & c(\alpha_1^a) & s(\alpha_1^b) & c(\alpha_1^b) & 1 & s(\alpha_1^a)\Theta_1 & c(\alpha_1^a)\Theta_1 & s(\alpha_1^b)\Theta_1 & c(\alpha_1^b)\Theta_1 & \Theta_1 \\ s(\alpha_2^a) & c(\alpha_2^a) & s(\alpha_2^b) & c(\alpha_2^b) & 1 & s(\alpha_2^a)\Theta_2 & c(\alpha_2^a)\Theta_2 & s(\alpha_2^b)\Theta_2 & c(\alpha_2^b)\Theta_2 & \Theta_2 \\ \vdots & & & & & & & & & \\ s(\alpha_{50}^a) & c(\alpha_{50}^a) & s(\alpha_{50}^b) & c(\alpha_{50}^b) & 1 & s(\alpha_{50}^a)\Theta_{50} & c(\alpha_{50}^a)\Theta_{50} & s(\alpha_{50}^b)\Theta_{50} & c(\alpha_{50}^b)\Theta_{50} & \Theta_{50} \end{bmatrix} \cdot \begin{bmatrix} C^S A^a c(\delta^a - \epsilon_z) \\ C^S A^a s(\delta^a - \epsilon_z) \\ C^S A^b c(\delta^b + \epsilon_z) \\ C^S A^b s(\delta^b + \epsilon_z) \\ C^S C^y + C^{PD} \\ A^S A^a c(\delta^a - \epsilon_z) \\ A^S A^a s(\delta^a - \epsilon_z) \\ A^S A^b c(\delta^b + \epsilon_z) \\ A^S A^b s(\delta^b + \epsilon_z) \\ A^S C^y \end{bmatrix} \quad (6.10)$$

6.3 Xe Excitation

Figure 6.4 demonstrates simultaneous ^{129}Xe - ^{131}Xe excitation using hybrid operation and no feedback. We measure linewidths of 14 mHz and 16 mHz for ^{129}Xe and ^{131}Xe , respectively, with peak field sizes of 26 μG for ^{129}Xe and 11 μG for ^{131}Xe . These linewidths were measured simultaneously by scanning the center pulsing frequency (thus changing the resonance




Figure 6.4: Simultaneous measurements of the Xe linewidths using hybrid operation. Top: ^{129}Xe linewidth = 14 mHz. Bottom: ^{131}Xe linewidth = 16 mHz. In-phase and quadrature signals are acquired from the measured magnitude and phase.

frequencies) without changing the drive frequencies. The pulsing frequency was converted to effective Xe Larmor frequency and plotted on the bottom axis. From open-loop phase noise measurements on resonance, we find the Signal-to-Noise Ratio (SNR) is $2900 \sqrt{\text{Hz}}$ for ^{129}Xe and $1100 \sqrt{\text{Hz}}$ for ^{131}Xe .

6.4 Cross talk

In order to quantify our cross talk, we perform an experiment similar to the resonance curve scan shown in figure 6.4. That is, we scan the bias pulse center frequency, $f_{2\pi}^0$, and record K_x and K_y as the Larmor frequency of the Xe passes through resonance with the drive frequency. The cross talk is very small, which makes it difficult to identify a small cross talk contribution on top of the large resonance curve for each species. We therefore change ω_{PDM} such that ^{131}Xe and ^{129}Xe are driven on resonance for different values of $f_{2\pi}^0$. Introducing an offset between the two resonance peaks makes it easier to identify cross talk because we can look for small deviations in the relatively flat off-resonance species' demodulation channel (which is easier than looking for small deviations in a large on-resonant signal) while the other species is on resonance. The cross talk measurement is shown in Fig. 6.5.

We quantify cross talk as before (Sec. 4.5, Sec. 5.4) with the ratio $\beta = Q_{pp}^{off}/Q_{pp}^{on}$ where Q_{pp}^{off} and Q_{pp}^{on} are the peak-to-peak signals for the off- and on-resonance species, respectively. Since the cross talk is not easily identifiable on the data, we use the residuals to the fit to identify Q_{pp} for the off-resonant species. We find that when ^{131}Xe is on resonance, $\beta' = 0.01$. When ^{129}Xe is on resonance, $\beta = 0.002$. Cross talk is dramatically better for both PDM and hybrid operation than it was for PM operation (see Ch. 7). This suggests that the improvement in cross talk we observed when we switched from PM to PDM operation was not due to the mode of operation, but rather due to the matrix inversion demodulation (which was used for PDM and hybrid operation, but not for PM).

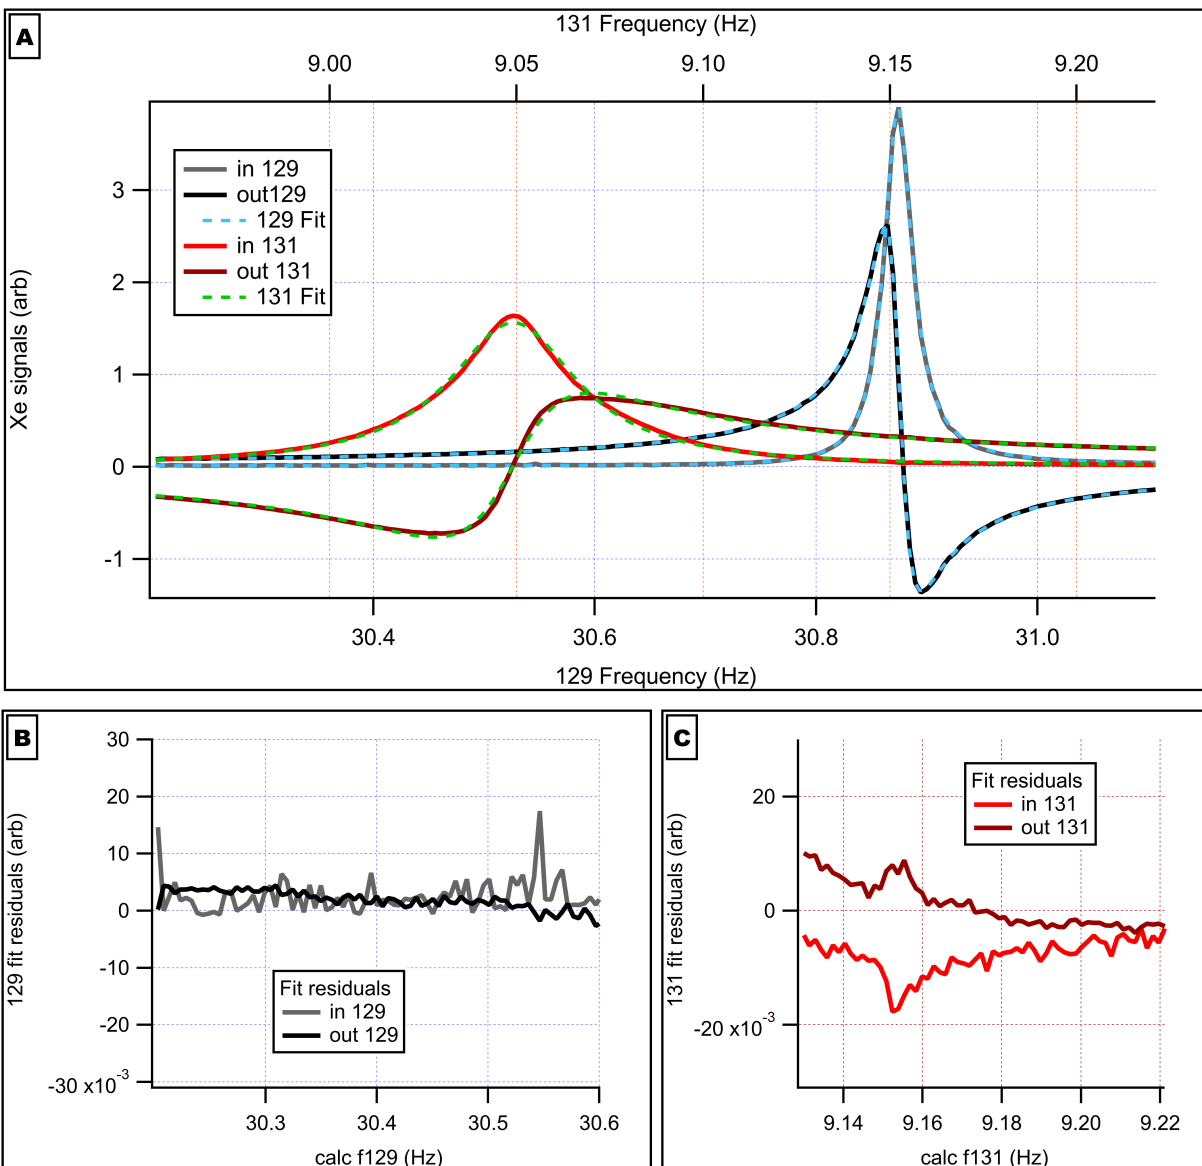


Figure 6.5: A: Measurement of cross talk for hybrid operation. B: Residuals to the fit for ^{129}Xe when ^{131}Xe is on resonance. C: Residuals to the fit for ^{131}Xe when ^{129}Xe is on resonance.

6.5 Comagnetometry

We will now perform comagnetometry using our measured signals with the goal of obtaining Ω^R independent of magnetic fields. There are a few different approaches to this, depending on how we run our gyroscope. In this section, we will show how Eq. 1.5 is implemented in practice for open loop, single feedback, and dual feedback operation.

Open loop operation is simply a lack of feedback. We use the term open loop to describe operation without field or drive frequency feedback, though we do still have feedback to hold the temperature of the cell constant. When it comes to designing the single and dual feedback schemes, there are a number of considerations involved. The available feedback outputs that could be used to hold the Xe species on resonance are ω_{PM} , ω_{PDM} , $f_{2\pi}^0$, and DC B_z . We opt not to adjust ω_{PM} because we want to maintain a fixed phase relationship between ω_{PM} and our pulse gating frequency. In order to hold both Xe species on resonance in the presence of correlated (eg. field) and anti-correlated (eg. rotation) noise on the precession frequencies, we choose to use one magnetic field feedback output and ω_{PDM} (see Eq. 6.3). For the magnetic field feedback, we must choose whether to apply corrections to the pulsed or DC z-field. This choice boils down to whether or not you want the Rb phase to be affected by the field feedback. In our case, we choose to feedback to the DC field using a sum of the Xe phases. This will allow us to stabilize the magnetometer phase shift, ϵ_z , as in Sec. 5.5 for PDM.

Open Loop Comagnetometry

From 6.9, we see that our demodulation will not give us the Xe precession phase shift $\delta^{a,b}$ directly, but rather the term $\delta^{a,b} + \mathbf{s}^{a,b}\epsilon_z$. The Bloch equation gives an expression for $\dot{\delta}$, which we can convert into the Fourier domain (such that $\tilde{f} = f(\omega)$) and solve to obtain

$$\tilde{\delta}^a + \mathbf{s}^a \tilde{\epsilon}_z = \frac{-\tilde{\omega}_d^a + \gamma^a(\tilde{B}_{z0} + \tilde{B}_{p0}) + \mathbf{s}^a \tilde{\Omega}^R + \mathbf{s}^a i\omega \tilde{\epsilon}_z}{i\omega + \Gamma_2^a} \quad (6.11)$$




Figure 6.6: MDEV of comagnetometry calculation using open loop Xe measurements.

for species a , with analogous expression for species b .

We use our measurements of $\delta^a - \epsilon_z$ and $\delta^b + \epsilon_z$ (as given by Eq. 6.11) to obtain a comagnetometer signal which eliminates γB . For fixed drive ($\tilde{\omega}_d \rightarrow 0$),

$$\frac{\rho(\tilde{\delta}^b + \tilde{\epsilon}_z)(i\omega + \Gamma_2^b) - (\tilde{\delta}^a - \tilde{\epsilon}_z)(i\omega + \Gamma_2^a)}{1 + \rho} = \tilde{\Omega}^R + i\omega\tilde{\epsilon}_z. \quad (6.12)$$

One drawback of this method of comagnetometry is that it requires knowledge of Γ_2 , which we only measure separately from our rotation measurements and which can drift over time. Figure 6.6 shows the modified Allan deviation (MDEV) [Allan and Barnes, 1981; Vanier and Audoin, 1989] for this comagnetometry calculation using open loop Xe measurements.

We have seen in [Thrasher et al., 2019b] that $i\omega\tilde{\epsilon}_z$ is non-negligible, so in order to get a measure of only $\tilde{\Omega}^R$ we must either measure $i\omega\tilde{\epsilon}_z$ and include it in our analysis, or hold it to

zero. In the following section, we choose to hold $i\omega\tilde{\epsilon}_z$ to zero using feedback [Bechhoefer, 2005]. I will show that holding $i\omega\tilde{\epsilon}_z = 0$ leads to dramatic improvement of the bias instability.

Field Feedback

The magnetometer phase shift ϵ_z depends on z-field drifts, as shown in Eq. 6.7. Therefore, if we can stabilize the z-fields, we can hold ϵ_z constant (thereby holding $i\omega\tilde{\epsilon}_z$ to zero). We apply a z-field correction, B_{FB} , such that

$$\tilde{B}_{FB} = \frac{1}{\gamma^a + \gamma^b} \times [\tilde{G}^a(\tilde{\delta}^a - \tilde{\epsilon}_z) + \tilde{G}^b(\tilde{\delta}^b + \tilde{\epsilon}_z)]. \quad (6.13)$$

The gains, $G^{a,b}$, are calculated as inverted zero gain with lead compensation (see Eq. 5.7). This shape is chosen to give high gain at low frequency and avoid feedback instability from the phase shift at Γ_2 . By using a weighted sum of the two Xe signals, we obtain an error signal which is independent of both rotation and ϵ_z . Since we use two signals to apply a single correction, we refer to this as dual-species single output (DSSO) feedback. The field experienced by the Rb (which does not include the pulses) is then $B'_{z0} = B_{z0} + B_{FB}$. In the high gain limit ($G \rightarrow \infty$), $\tilde{B}_{FB} \rightarrow -\tilde{B}_{z0}$ such that B'_{z0} goes to zero. When B'_{z0} is held to zero, $\tilde{\epsilon}_z$ will be held constant.

Having held $\tilde{\epsilon}_z$ constant and $i\omega\tilde{\epsilon}_z$ to zero using field feedback, the comagnetometry calculation from Eq. 6.12 becomes

$$\frac{\rho\tilde{\delta}^b(i\omega + \Gamma_2^b) - \tilde{\delta}^a(i\omega + \Gamma_2^a)}{1 + \rho} = \tilde{\Omega}^R. \quad (6.14)$$

Thereby, DSSO field feedback allows us to deduce rotation independent of z-fields and magnetometer phase shifts.

This feedback accomplishes the same thing as the field feedback used for PDM (see Sec. 5.5) in a slightly different way. In both cases, the form of the feedback is a sum of

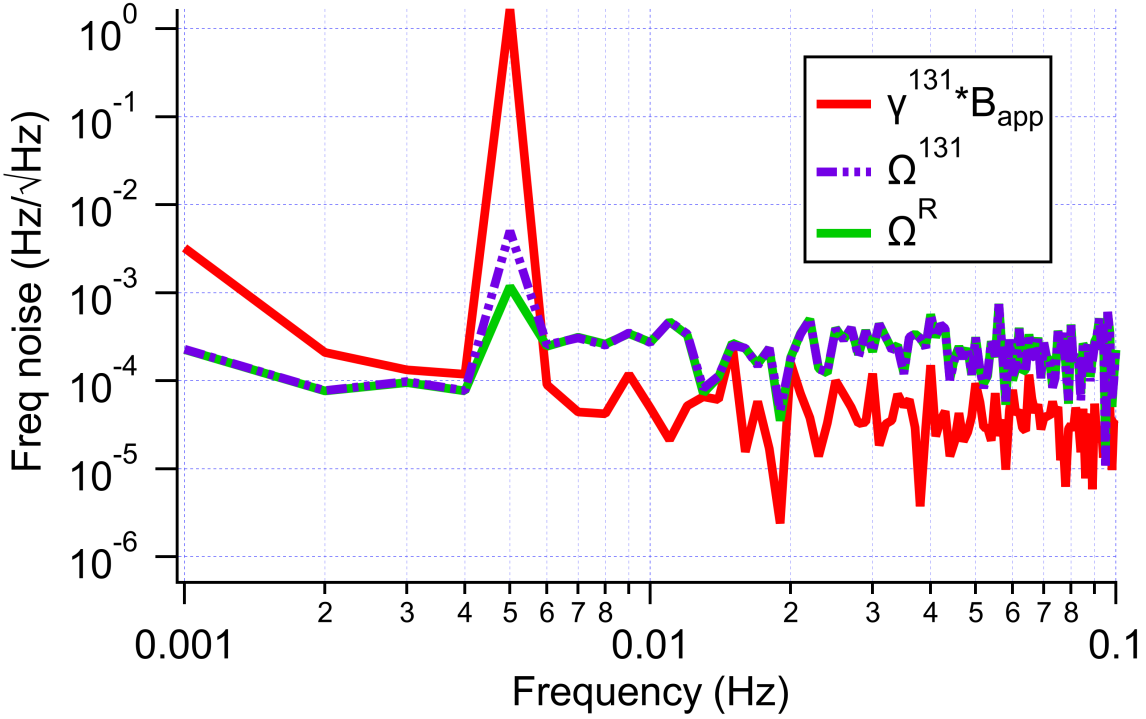


Figure 6.7: Suppression of the effect of an applied 5 mHz z-field (red) due to field feedback (purple) and comagnetometry (green).

signals from each species. This results in the suppression of the magnetometer phase shift. For PDM, the inputs were the Xe resonant drive frequencies, and for hybrid operation, the inputs are the measured phase signals for each species.

In Fig. 6.7, we demonstrate that the DSSO field feedback results in substantial suppression of common magnetic field noise. By applying an ancillary AC B_{app} along \hat{z} , we can see how much that field is suppressed in our final measure of rotation. For this test, we applied B_{app} at 5 mHz with an amplitude of just over $200\mu\text{G}$. Since our final measure of rotation is in units of Hz, we cast our applied field in terms of the frequency noise it would cause on a single species of Xe. For this calculation, we use the gyromagnetic ratio of the less magnetically sensitive species, ^{131}Xe . Our applied field would therefore be expected to produce a frequency oscillation $\gamma^b * B_{app}$. We apply DSSO feedback and measure the resulting $\delta + \epsilon_z$. If we use the transfer function to convert the ^{131}Xe phase shift to frequency, we see that the feedback

alone suppresses the applied signal by a factor of 330. We then use our measured ^{129}Xe signal to perform comagnetometry and deduce rotation. We see that the oscillation on the rotation is a factor of 1400 smaller than the applied oscillation.

Applying our DSSO FB and taking rotation measurements, we found that our bias instability was very dependent on our DC transverse fields. Our rotation measurements show an apparent dependence on transverse fields that goes as the product $B_x B_y$. If B_x or B_y are nonzero, we will have increased sensitivity to field noise along the other transverse direction. In order to find the optimal settings, we take three sets of measurements, one with varying DC B_y , one with varying DC B_x , and one with varying AC x-compensation field amplitude [Korver et al., 2015], as shown in Fig. 6.8. We compare the MDEVs to find the transverse field settings which give the lowest bias instability.

We repeat these scans every time we want to take a stability measurement, as low-frequency field drifts cause the optimal settings to drift over time. Indeed, we find that long measurements (even those taken after scanning the transverse fields) show variation over time. For example, we took a 10-hour measurement which did not even reach a bias instability of $1 \mu\text{Hz}$. We cut this measurement into 5 ksec long segments and analyzed the segments individually. As shown in Fig. 6.9, the stability of the magnetometer varied widely over the course of the long measurement. This means that our measurements are not perfectly repeatable, but rather depend on how far the transverse fields drift over the course of a given measurement. In the future, we will need to find a way to stabilize - or become less sensitive to - transverse fields.

Figure 6.10 shows our best-to-date results for both hybrid operation and pure PDM (no PM). The PDM data was obtained with the same apparatus as the hybrid data. The apparatus is able to be easily switched between PDM and hybrid operation by inserting/removing two quarter wave plates and changing the modulations applied to the pulsed field and optical pumping. Both measurements also follow the same matrix inversion least-squares fit principle for demodulation. For hybrid operation, we find an Angle Random Walk (ARW)

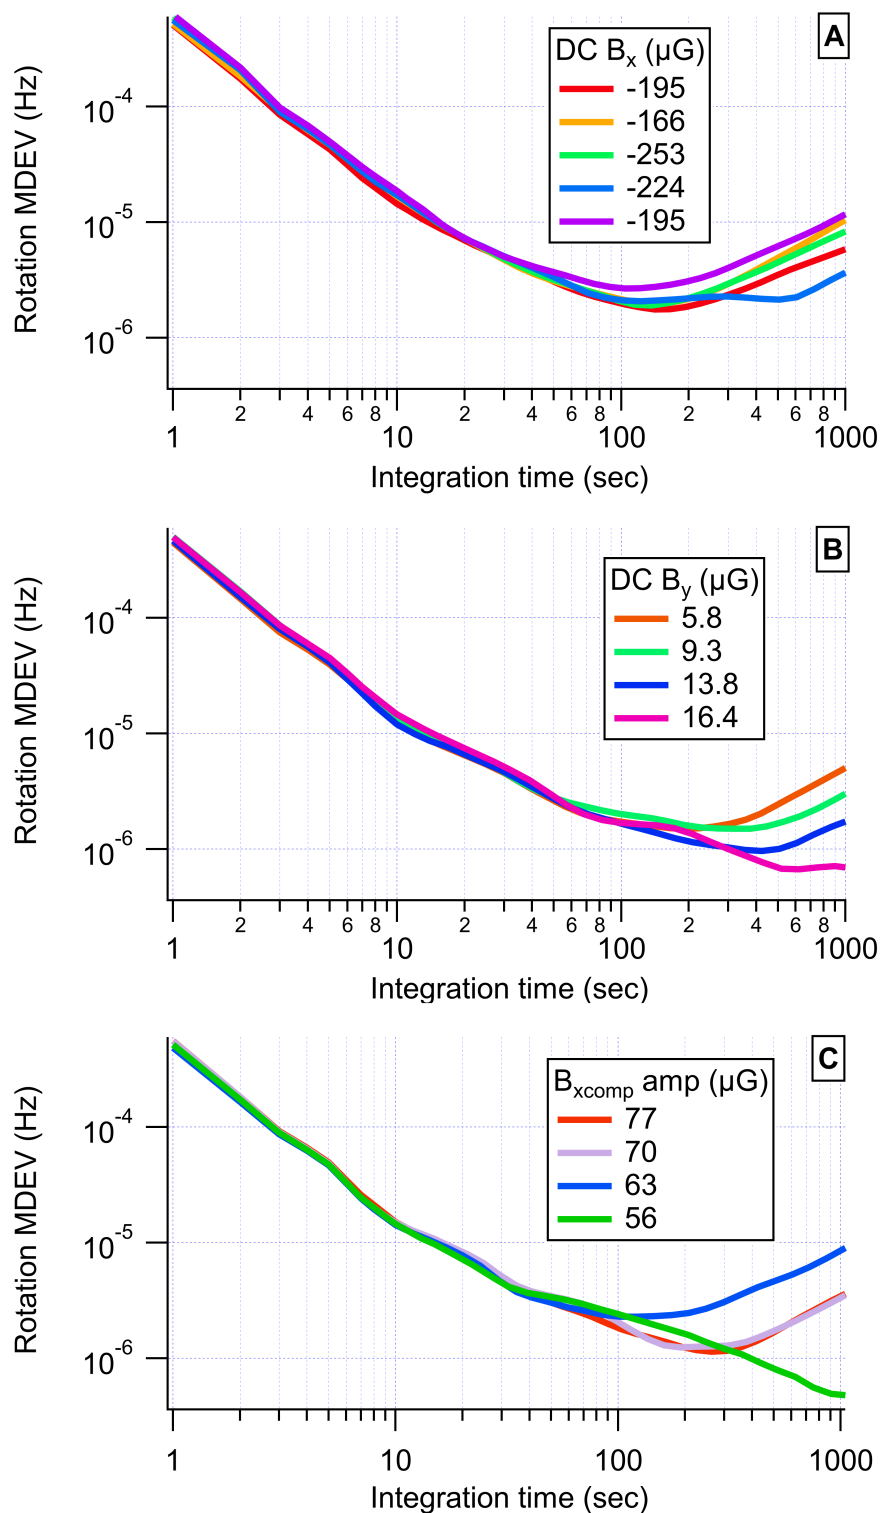


Figure 6.8: Rotation MDEVs for transverse field scans. We compare MDEVs to find the lowest bias instability for scans of DC B_x (A), DC B_y (B), and the AC x-compensation amplitude (C). We observe some drift over time, as evident for DC $B_x = -195\mu\text{G}$.

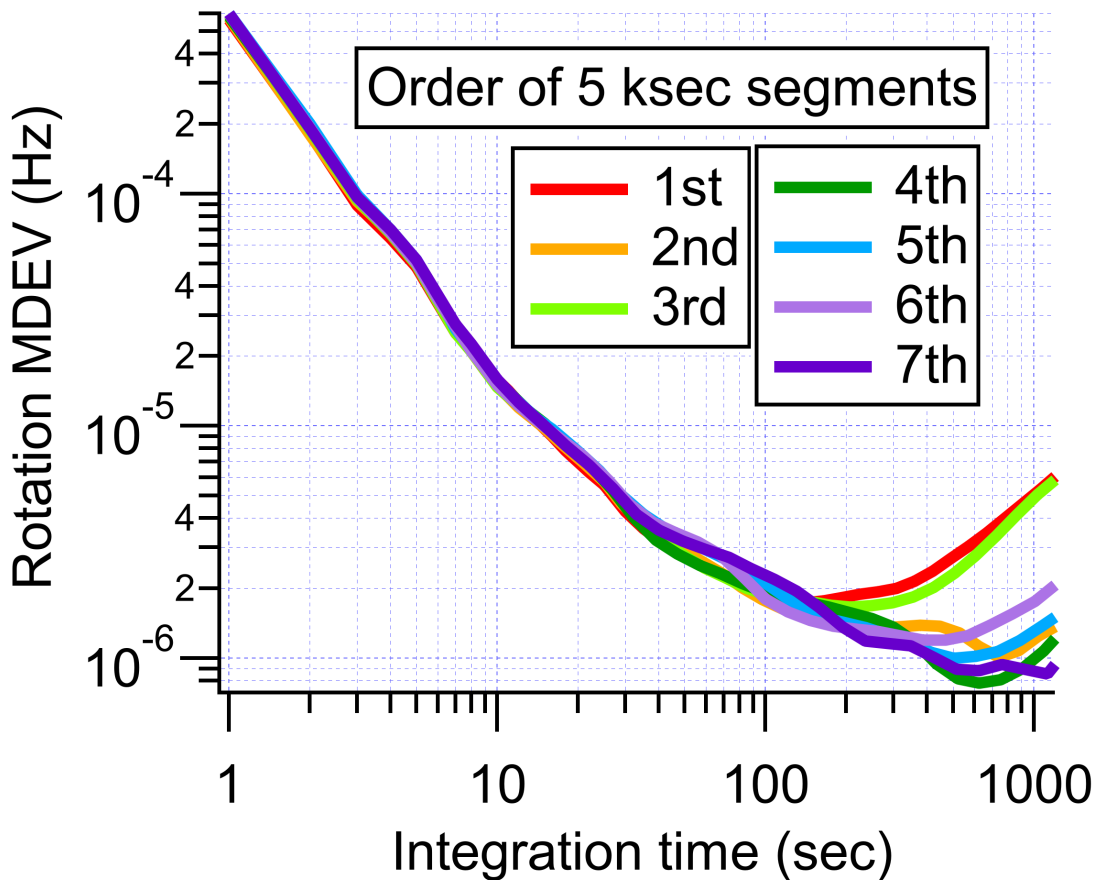


Figure 6.9: MDEVs of single 10 hour long rotation measurement cut into 5 ksec segments.

of $21 \mu\text{Hz}/\sqrt{\text{Hz}}$, and a bias instability of 480 nHz after about 1000 seconds. Comparing the two MDEVs, we see similar stability performance for PDM and hybrid. In both cases, we present bias instability improvements over our previously published results [Thrasher et al., 2019a]. At this time there is not an obvious benefit to operating with hybrid operation rather than pure PDM. It seems likely that both hybrid and PDM drive share the same limiting noise source. We suspect that transverse field noise is limiting the stability of our rotation measurements in both cases. In the future, when we can suppress transverse field noise (or reduce our sensitivity to it), we should revisit both hybrid and pure PDM drive to evaluate the best drive scheme moving forward.

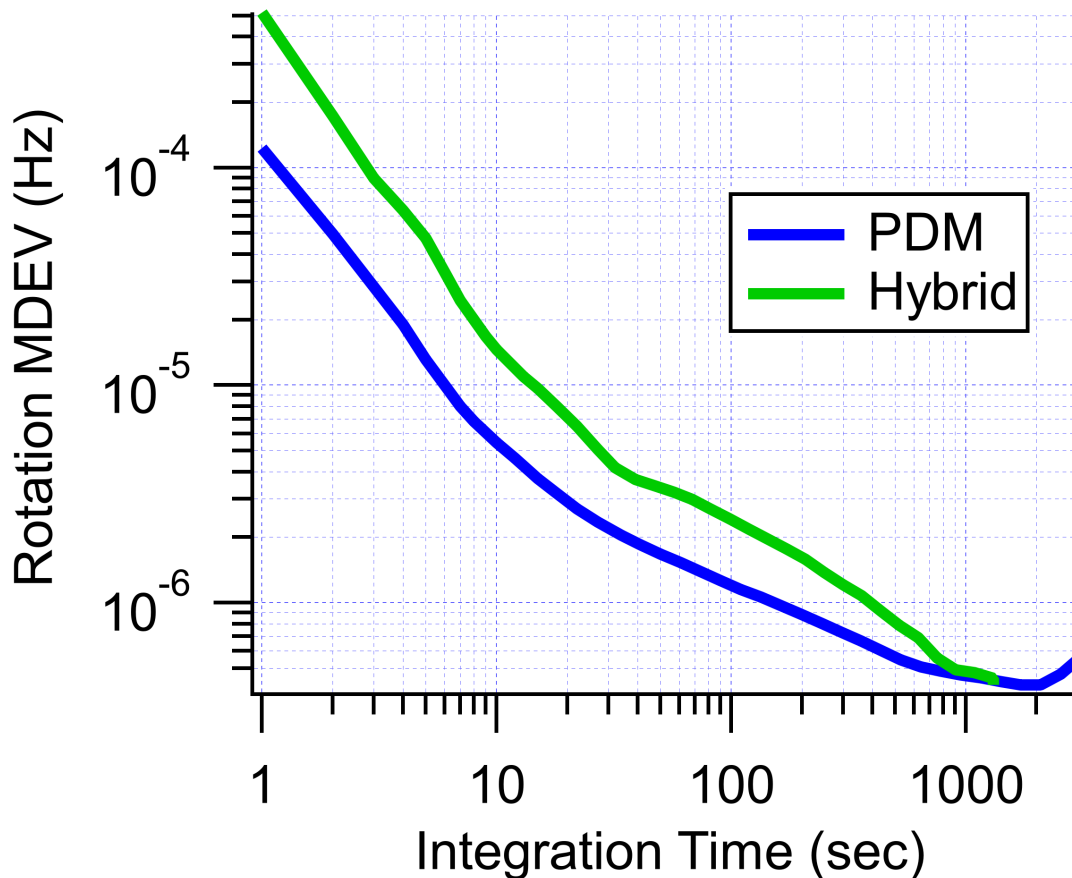


Figure 6.10: MDEVs for best-to-date stability results for PDM (blue) and hybrid (green) operation.

Dual Output Feedback

Performing comagnetometry as given in Eq. 6.14 still has the drawback of requiring us to know Γ_2 . Since we do not measure Γ_2 simultaneously with our measures of $\delta + s\epsilon_z$, any noise on Γ_2 is mapped onto the rotation. We can theoretically circumvent this drawback using a second feedback loop to adjust the drive frequencies and hold the measured phase shifts equal to zero.

The second feedback loop uses a weighted difference of the two Xe signals to obtain an error signal (as opposed to the weighted sum used for the field feedback). Because of our choice of drive scheme, we can apply this correction to our drive frequencies by adjusting

ω_{PDM} while keeping ω_{PM} fixed. The new feedback expressions for dual-species dual output (DSDO) feedback are then

$$\begin{aligned}\gamma^a B_{FB} - \omega_{PDM} &= G^a(\delta^a - \epsilon_z) \\ \gamma^b B_{FB} + \omega_{PDM} &= G^b(\delta^b + \epsilon_z)\end{aligned}\tag{6.15}$$

When these two feedback loops are used together, the total field noise is still held to zero in the high gain limit, and we also get that $\tilde{\omega}_{PDM} \rightarrow -\tilde{\Omega}^R$. Thus we are able to directly obtain a measure of rotation noise without using Γ_2 in calculations.

When we apply the DSDO feedback, we see long timescale (1000s of seconds) drifts. For reasons we don't understand, rather than holding the drive frequencies to the Xe resonance, the feedback holds the detuning to a point that starts on resonance and then slowly drifts off resonance. If the feedback was holding our drive frequencies on resonance, the measured Xe amplitudes would remain constant. Instead, we observe that the measured Xe amplitude decreases over time, though the measured $\delta + \mathbf{s}\epsilon_z$ is still held to zero. The effect is most dramatic on ^{129}Xe . The speed of the drift is inconsistent from run-to-run, where we have sometimes observed that the speed with which the magnitude drifts depends upon the initial Xe detuning of the feedback setpoint. Figure 6.11 shows the amplitude drift on ^{129}Xe during DSDO operation. We have seen similar drifting problems using various drive and feedback schemes, even when using pure PDM (rather than hybrid) drive. Comparing many instances of drift over the course of several years, the drift seems to be present when both of the following two conditions are met: 1) we are driving either Xe species using a combination of two modulation frequencies, and 2) we are applying feedback to a modulation frequency (with an exception observed for PM operation that will be discussed below, see Sec. 4.4).

Condition 1) means that, in a generic drive scheme $\omega_d = \sum_n p_n \omega_{PM,n} + \sum_m q_m \omega_{PDM,m}$, more than one of the drive coefficients p_n and q_n are nonzero. For example, a drive scheme $\omega_d^a = \omega_{PDM,1}$, $\omega_d^b = \omega_{PDM,2}$ would not exhibit drift (even if condition 2) were met), but

a drive scheme $\omega_d^a = \omega_{PDM,1} + 3\omega_{PDM,2}$, $\omega_d^b = \omega_{PDM,2}$ would exhibit drift (if condition 2) were also met). Condition 2) means that we are applying a feedback correction to one or more ω_{PM} , ω_{PDM} (for PDM or hybrid operation). In contrast, drift is not present if the only feedback applied is a correction to B_z while using PDM or hybrid operation. During PM operation, with a drive scheme that meets condition 1), we observed a Xe phase drift over time when applying only B_z feedback (see Sec. 4.4). The feedback applied during PM operation was performed using an analog feedback circuit, as opposed to the feedback for PDM and hybrid operation which was performed digitally. It is possible that the PM phase drift is unrelated to the magnitude drift observed for PDM and hybrid operation.

Since the drift issue seems to have something to do with drive and/or modulation frequencies, and since all of our frequencies and feedback (with the exception of feedback for PM) are calculated via LabVIEW on the FPGA and host computer, we suspect there may be some kind of accumulated rounding error in our code. While our investigation is ongoing, we hope to in the future identify and eliminate the source of drift, allowing us to use DSDO feedback to further improve our comagnetometry calculation of rotation.

6.6 Chapter Summary

In this chapter, I have demonstrated a new approach to continuous dual-species Xe excitation in the transverse NMR gyroscope. Hybrid operation builds upon the foundations of PM and PDM operation described in Ch. 4 and Ch. 5. By utilizing matrix inversion demodulation (based on the demodulation developed for PDM) we have seen that cross talk is not inherent to drive schemes that utilize PM, and can be avoided using least-squares fitting. I have shown that comagnetometry is drastically improved by using \hat{z} field feedback which is carefully designed to suppress ϵ_z . By carefully scanning the transverse fields, the gyroscope has achieved similar bias instability performance using hybrid operation as was attained for PDM operation. This suggests that both PDM and hybrid operation are limited by transverse




Figure 6.11: An example of magnitude drift when applying DSDO feedback. The amplitude (arb. units) drifts dramatically over the course of several thousand seconds (especially for ^{129}Xe - black), while the measured phase shifts (inset) are held to zero. This implies the feedback is driving Xe further and further off resonance. A low pass filter at 0.1 Hz is applied to the shown phase shifts.

field drifts. I believe that future work to stabilize the transverse field drifts would lead to further improvement in the bias instability.

Chapter 7

Outlook

7.1 Comparing PM, PDM, and Hybrid Operation

In this thesis, I have presented work toward the development of a transverse NMR gyroscope. Progress has been quantified in terms of several performance metrics, which were reported separately for each mode of operation. Here, I will directly compare these metrics for PM, PDM, and hybrid operation. Figure 7.1 shows a summary of the major results presented in this thesis for PM, PDM, and hybrid operation.

Linewidths and Xe Field Sizes

The linewidths (Γ_2) obtained for both species are around 15 mHz for PM, PDM, and hybrid. It is not too surprising that the linewidths do not change much between modes of operation because the minimum linewidths are determined by wall relaxation (which is independent of mode of operation). Furthermore, we scanned the x-compensation to minimize the linewidths in each case.

To help understand the field sizes obtained, let us consider our theoretical field sizes. We start with the steady state Bloch equation solutions for K_\perp from Eq. 2.2, and then find the predicted Xe field sizes measured by the embedded Rb magnetometer according to $b_K^S K_\perp$.

		unit	PM	PDM	Hybrid
Linewidth	Γ_2^b	mHz	17	19	16
	Γ_2^a	mHz	14	11	14
Spin-exchange field amplitude	$b_b^S K^b$	μG	78	16	11
	$b_a^S K^a$	μG	66	14	26
Fourier coefficient product	j^b		0.37	0.2	0.19
	j^a		0.37	0.14	0.35
Cross talk coefficient	β'		0.07	0.0004	0.01
	β		0.17	0.06	0.002
Signal-to-noise ratio	SNR^b	$\sqrt{\text{Hz}}$	5000	2650	1100
	SNR^a	$\sqrt{\text{Hz}}$	5000	2680	2900
Field suppression factor	FSF		>240	4700	1400
Angle-random walk	ARW	$\mu\text{Hz}/\sqrt{\text{Hz}}$	16	17	21
Bias instability	Bias inst.	nHz	800	420	480

Figure 7.1: Table summarizing the major results of this thesis. Species "a" is ^{129}Xe and species "b" is ^{131}Xe .

The predicted field size is

$$B_K^S = \frac{b_K^S \Gamma_S^K R j}{\Gamma \Gamma_2}. \quad (7.1)$$

The values of j and Γ_2 are given in Fig. 7.1.

First, let us consider only the expected field amplitude changes due to changes in the Fourier coefficient product, j , and the measured linewidths, Γ_2 . Therefore, we assume that $b_K^S \Gamma_S^K R / \Gamma'$ is the same for each mode of operation. Let B_{PDM} be the measured PDM spin-exchange amplitudes, which I will use as a reference point. Then the predicted amplitudes for PM and hybrid operation would be $B_{PDM} * (j_{PM} / j_{PDM}) * (\Gamma_{2,PDM} / \Gamma_{2,PM})$ and $B_{PDM} * (j_{hybrid} / j_{PDM}) * (\Gamma_{2,PDM} / \Gamma_{2,hybrid})$, respectively. We can then compare the predicted vs measured spin-exchange field amplitudes. I find that the measured hybrid field is in good agreement (within about 5%) with the prediction for ^{129}Xe , but about 40%

smaller than predicted for ^{131}Xe . The measured PM spin-exchange fields are about a factor of $2.3\times$ larger than the predicted PM fields (based on PDM). The discrepancy between measured and predicted values must then be attributed to the terms I previously assumed were constant: the spin-exchange field factor b_K^S , the spin-exchange rate Γ_S^K , and the DC Rb polarization R/Γ' .

A detailed discussion of how to estimate b_K^S , Γ_S^K , and Γ' in our system can be found in Appendix A of [Thrasher, 2020]. Each of these terms depend on multiple factors. To explain the predicted and measured fields, we are looking for something that changed between PM and PDM operation (due to the large discrepancy between predicted and measured fields) and did not change much between PDM and hybrid operation (based on the measured ^{129}Xe field amplitude being consistent with prediction). Based on the discussion in [Thrasher, 2020], b_K^S should only depend on properties of the vapor cell which remained constant for PM, PDM, and hybrid operation. For our cell with 40 Torr Xe pressure, we estimate $b_a^S = 0.002$ G and $b_b^S = 0.019$ G. From the measurement shown in Fig. 3.5, we find $\Gamma' = 49000$ rad/sec. This measurement was taken using PM operation, but we have consistently measured similar Γ' throughout the work of this thesis. We tune the pumping to achieve a DC Rb polarization of $R/\Gamma' \approx 0.5$ for all modes of operation.

We then look to Γ_S^K as our main suspect. Γ_S^K has a linear dependence on the Rb density [Walker and Larsen, 2016]. For example, at a density of 10^{13} cm^{-3} , we would find $\Gamma_S^a = 6$ mrad/sec, and $\Gamma_S^b = 0.5$ mrad/sec. As mentioned in Ch. 3, we operated at various temperatures (and hence various densities) over the course of the work of this thesis. For the PM data presented in 4, the Rb density was 3.36×10^{13} cm^{-3} . We operated at a lower density of 1.7×10^{13} cm^{-3} for both PDM and hybrid operation in this thesis. (The choice to lower the density was based primarily in fear after another cell was heated too much and the Rb was lost due to chemical reactions with the cell wall. We therefore chose to operate at a lower density in order to have a bigger buffer against accidentally overheating the cell.) Density is therefore a prime suspect to account for the discrepancy between the measured

		unit	PM	PDM	Hybrid
Calculated field amplitude	\mathbf{B}_b^S	μG	85	14	15
	\mathbf{B}_a^S	μG	55	21	41
Measured field amplitude	$b_b^S K^b$	μG	78	16	11
	$b_a^S K^a$	μG	66	14	26
Measured/Calc.	$b_b^S K^b / \mathbf{B}_b^S$		1.41	1.18	0.72
	$b_a^S K^a / \mathbf{B}_a^S$		0.77	0.68	0.64

Figure 7.2: A comparison of the measured and calculated spin-exchange field of the Xe as measured by the Rb for each mode of operation.

and predicted spin-exchange field amplitudes.

We can account for the change in density by re-defining the predicted fields to be $B_{PDM} * (j_{PM}/j_{PDM}) * (\Gamma_{2,PDM}/\Gamma_{2,PM}) * ([Rb]_{PM}/[Rb]_{PDM})$, with analogous expression for hybrid. Since the density was the same for PDM and hybrid operation, our predicted hybrid field remains unchanged. Accounting for the change in density for PM operation, we now find that the predicted field is only 20% larger than the measured field.

Given these updates to the predicted field, we see that the discrepancies in the measured field sizes between the three modes of operation are mostly accounted for just by considering the differences in Rb density and in the Fourier coefficients.

Given the estimated values discussed above, I can now calculate the Xe spin-exchange field given by Eq. 7.1. Figure 7.2 shows a comparison of the calculated and measured field sizes. As you can see, there is likely still some error in our estimated values, but the discrepancy between measured and calculated fields is less than 50% in all cases.

Cross Talk

We quantify cross talk with the ratio $\beta = Q_{pp}^{off}/Q_{pp}^{on}$ where Q_{pp}^{off} and Q_{pp}^{on} are the peak-to-peak signals for the off- and on-resonance species, respectively. We denote the cross talk coefficient as β when ^{129}Xe is on resonance, and β' when ^{131}Xe is on resonance. The PM cross talk

is significantly worse than the cross talk for PDM and hybrid operation (which is why we originally moved away from PM operation). In fact, the cross talk for PDM and hybrid operation were so small that cross talk was only visible when looking at the residuals to Lorentzian fits.

The PM data were taken using traditional lock-in detection, whereas the PDM and hybrid data were taken using matrix-inversion demodulation. I suspect that switching to matrix-inversion demodulation would improve the cross talk for PM, as will be discussed later in this chapter.

In this thesis, I have treated cross talk as being strictly a demodulation error, where signal content from one Xe species is present on the detection of the other species. It is important to note that the method of measuring cross talk presented in this thesis may also include another effect. Through signal mixing, it may be possible that the supposedly off-resonant species could be slightly excited unintentionally. Such excitation would be included in our quantification of cross talk. Because of the drastic improvements in cross talk when we switched to matrix-inversion demodulation, I do not believe unexpected excitation to be the dominant source of the measured cross talk in our system.

If someone wanted explicitly test whether measured cross talk was due to unexpected excitation, I would propose the following experiment. First, operate under conditions where the T_2 relaxation times are measurably different for the two noble gas species. Then, perform a cross talk measurement as described in this thesis, but stop when there is cross talk present on the off-resonant species. Then, abruptly turn off the modulations which drive the Xe while continuing to measure the signals as they decay, and fit those signals for the decay rates. If the decay of the off-resonant signal matches the T_2 of the on-resonant species, the cross talk is an artefact of the demodulation. If the decay of the off-resonant signal matches the T_2 of the off-resonant species, there is unexpected excitation.

SNR

The measured SNRs vary quite a bit between the three modes of operation, with SNR^{PM} almost a factor of two larger than SNR^{PDM} . In turn SNR^{PDM} is more than a factor of two larger than SNR^{hybrid} for ^{131}Xe , though they are similar for ^{129}Xe .

Of course, a reduction in SNR could be due to a reduction in the signal, an increase in the noise, or a combination of the two. As discussed above, the Xe spin-exchange field amplitudes ($b_K^S K$) for each species did vary between the different modes of operation. I check whether the changes in SNR can be attributed to the changes in the field amplitudes by calculating $\delta B = b_K^S K / \text{SNR}$ for each mode of operation, which estimates the white noise level (δB) implied by the field sizes and SNR measurements. For PM, $\delta B^a = 13 \text{ nG}/\text{sqr}tHz$ and $\delta B^b = 16 \text{ nG}/\text{sqr}tHz$. For PDM, $\delta B^a = 5 \text{ nG}/\text{sqr}tHz$ and $\delta B^b = 6 \text{ nG}/\text{sqr}tHz$. For hybrid operation, $\delta B^a = 9 \text{ nG}/\text{sqr}tHz$ and $\delta B^b = 10 \text{ nG}/\text{sqr}tHz$. The implied noise levels for hybrid operation were similar for both species of Xe, which suggests that the discrepancy between SNR^a and SNR^b is primarily due to the difference in the spin-exchange field amplitudes for the two species.

Though the SNR was the highest for PM operation, the implied noise was actually better for PDM and hybrid operation. This is not too surprising, given the system developments made after PM operation. For example, we made changes to the demodulation that we would expect to result in less mixing of background noise onto the detected signals. We also changed the field feedback to suppress noise from magnetometer phase shifts. If we were able to increase the Xe spin-exchange field amplitudes while maintaining the improved noise performance of PDM, we would expect to measure SNRs higher than those obtained for PM in Fig. 7.1. I will discuss my proposal to accomplish this in Sec. 7.2.

FSF

When comparing the FSFs obtained for PM, PDM, and hybrid operation, we must keep a few things in mind. First off, we were not yet measuring FSF explicitly at the time when we were operating with PM. As such, the FSF listed for PM is only a lower-end estimate found from the closed-loop and open-loop data available. The true FSF for PM may be much better. When comparing PDM and hybrid operation, note that while the PDM FSF was higher, the bias instability was not limited by FSF for PDM or hybrid. This means that the higher FSF for PDM is not necessarily an argument for using PDM over hybrid operation. If we were to get to a point where residual field noise limits the stability of our measure of rotation frequency (either by reducing other noise sources or by operating in environments with higher field noise), FSF would hold more weight in deciding which mode of operation to use.

ARW

The lower the ARW, the better. Not only does it allow us to reach our limiting noise source at a shorter integration time, but if the limiting noise source trends positively with integration time, a lower ARW would correspond to a lower bias instability. PM and PDM had similar ARW performance, but the ARW for hybrid operation is slightly worse. It is possible this worse ARW relates to the worse SNR^b for hybrid operation. If our ARW were SNR-limited, we would expect to measure $ARW^{SNR} = \Gamma_2/SNR$. Since we use two species to measure rotation, we can bet an estimate by calculating ARW^{SNR} for each species, and then we consider the larger of the two results. We find that $ARW_{PM}^{SNR} = 3\mu Hz/\sqrt{Hz}$, $ARW_{PDM}^{SNR} = 7\mu Hz/\sqrt{Hz}$, and $ARW_{hybrid}^{SNR} = 15\mu Hz/\sqrt{Hz}$. Comparing these to our measured results, the hybrid ARW is within a factor of two of the predicted SNR-limited ARW (whereas the difference is greater than a factor of two for PM and PDM). The measured ARW for PM and PDM are very similar, whereas the predicted SNR-limited ARW for PM and PDM are about a factor of

two different. I suspect that for PM and PDM, the ARW is not SNR limited, but for hybrid operation it is.

Bias Instability

The bias instabilities reported in this thesis for PDM and hybrid operation show the best-to-date stability performance measured on this device. In deg/Hr, our best bias instability is about 0.5 deg/Hr. While this performance would be acceptable for some navigation applications, such as automotive sensing, it is still insufficient for commercial airlines (which require stability of order 0.01 deg/hr) [Donley and Kitching, 2013] and other very sensitive applications. That said, I believe that we have identified the limiting noise source in the device, which should pave the way for future stability improvements.

The PM bias instability was worse than what we got for PDM and hybrid operation, though I predict we could attain a similar bias instability for PM with a few changes (which will be discussed in the next section). The bias instabilities for PDM and hybrid are very similar, around 450 nHz. In this thesis, I have presented evidence to suggest that the bias instability of our device is limited by transverse field noise for PDM and hybrid operation.

7.2 Revisiting PM

While PM presents the worst bias instability presented in this thesis, we have learned a lot from PDM and hybrid operation since then. I believe that if we were to return to PM operation and apply the lessons we have learned, we have the potential to improve PM beyond what we have been able to obtain using PDM and hybrid operation.

The main reason we shifted away from PM operation was because of cross talk on our detection. We were unable to isolate the signals from the two Xe species very well, and so the detected amplitude and phase of each species contained some information from the other species. This cross talk was a result of our demodulation scheme. As described in Ch. 4,

we demodulate by multiplying the high-pass filtered S_z by our polarization modulation square wave, then send that signal into two lock-in amplifiers, each referenced to a different Xe isotope's drive frequency. Cross talk enters into the system via phase shifts from the high-pass filter and imperfect rectification (meaning that we assume the optical pumping exactly matches the modulation waveform). Even at the time, we saw evidence that our demodulation scheme was to blame for the high observed cross talk, but it wasn't until we moved to PDM operation that we found a solution.

For PDM operation, it didn't make sense to use single-frequency lock-in amplifiers for demodulation, since the Xe precession was swept through a range of frequencies by the modulation. We therefore developed a new demodulation scheme as described in Ch. 5 and Ch. 6 that involves using matrix inversion to take a least-squares fit, using the form of S_z to design the fit function. This new demodulation scheme did not show significant signs of cross talk for either PDM or hybrid operation. The results for hybrid operation are especially promising for the potential application of this demodulation method for PM since hybrid operation also involves modulation of the Rb polarization. I believe that with least-squares fitting, we could see a dramatic reduction of cross talk in PM operation. The application of this demodulation scheme to PM would be fairly straightforward, following along the lines of the demodulation matrices for hybrid operation, but without the bias modulation.

The cross talk improvements seen when we switched to PDM cannot be exclusively attributed to the demodulation change, because it was accompanied by a change to the detection, and thus the effects of the two changes cannot be completely isolated. The detection change in question was the move to a gated detection scheme, as discussed in Ch. 5. This gating of the pulses was originally developed to reduce the influence of magnetometer gain changes from the changing bias pulse frequency. I believe that this method could be equally effective in reducing the influence of gain changes due to Rb polarization transients following steps in the modulation waveform. In order to use gated detection effectively for PM operation, the gating would not need to turn off the bias pulses (though it could

if deemed necessary), but rather stop the Rb polarization from changing direction during a measurement period. As with gating for PDM, this would ensure the embedded Rb magnetometer gain was approaching a steady state before taking a measurement. While I do believe there would be benefits to utilizing gated detection for PM, it may not be worth the reduction in Xe signal size that would accompany that change.

My final proposal for PM operation relates to the transverse field sensitivity we have observed for both PDM and hybrid operation. For both PDM and hybrid, our best bias instability results were obtained immediately after taking many scans to optimize the transverse field settings. In fact, we often saw that the bias instabilities obtained before taking transverse field scans bottomed out around $1 \mu\text{Hz}$, which is the bias instability we measured for PM. This leads me to suspect that we may see similar improvements (around a factor of 2) to the bias instability for PM by scanning the transverse fields before taking long measurements.

Taking transverse field scans has helped improve our measured bias instabilities, but it is still only a band-aid solution. Without addressing our uncontrolled transverse field drifts, our long-term stability will still be limited by transverse fields, as is evidenced by Fig. 6.9. I believe that we need to find a way to actively stabilize our transverse fields using feedback, so that we can consistently remain at an operating point that minimizes our sensitivity to transverse fields. In order to do that, we will likely need to apply parametric modulations to our system to monitor transverse fields. In addition to feeding back to the transverse field magnetic shim coils, we will also need to do something to stabilize the cancellation of the Rb spin-exchange field by the x-compensation field. This is because the Rb spin-exchange field is a major potential source of transverse field drift. While we do apply the x-compensation field to cancel the Rb spin-exchange field, if the optical pumping changes, the optimal x-compensation field also changes. While we could stabilize this via feedback to the x-compensation amplitude, I believe should consider instead using stabilization of our pump laser frequencies and/or powers. This would address the source of the Rb spin-exchange

field drift, rather than addressing its effect on the optimal x-compensation. Furthermore, stabilizing the Rb optical pumping would reduce drifts of the magnetometer gain.

I also have a recommendation regarding the feedback scheme used for PM operation. Previously, we used only ^{129}Xe to stabilize B_z drifts. We then used the Xe transfer function to convert $\delta^b + \epsilon_z$ to a frequency, and scaled that frequency to obtain a rotation signal. Since then, we have come to more fully understand the effects of the magnetometer phase shift ϵ_z in our system, and how to avoid it. I propose that we use a weighted sum of the two Xe phases to stabilize B_z , as we have for PDM and hybrid operation. This should eliminate ϵ_z from our measures of rotation. This feedback scheme would follow along the lines of the DSSO scheme used for hybrid operation. From there, we could attempt to perform feedback to one or more of the polarization modulation frequencies, though we will need to keep an eye out for Xe magnitude drift like what was seen for the DSDO scheme used for hybrid operation.

7.3 Second Apparatus

The apparatus used in this thesis has been developed bit-by-bit over more than a decade. Bits and pieces of the apparatus have been swapped out or adjusted in various ways, but nothing so major as to halt our investigations for an extended period of time. As such, there were some choices made in the early days of the project that I believe could be improved upon in the apparatus. I would not wish to take apart our current apparatus, however, as I believe moving back to PM operation to be the most crucial next step. Because of this, we have started building a second apparatus with some major changes informed by things we have learned over the lifetime of the project. The apparatus presented in Ch. 3 will be referred to as the "current" apparatus.

First, the second apparatus is being built on a Lazy Susan (turntable), such that the whole apparatus can be rotated about the y axis. Once loaded up with optical and magnetic

elements, the turntable we are using does not move very smoothly, and so the apparatus should not be rotated during measurements, but can be rotated between measurements. The purpose for this rotation is to enable measurements of Earth's rotation. We generally orient our apparatus such that the z-axis is oriented along East-West. This means that there is no significant component of Earth's rotation about our sensitive rotation axis (z). If we were to rotate our apparatus such that the z-axis was oriented along North-South, there would then be a component of Earth's rotation about our sensitive rotation axis, and we could make measurements of Earth's rotation. Obviously, the rotation frequency of the Earth is well known, and the purpose of this study is not simply to prove what is already known. Rather, by comparing our measurements of Earth's rotation to known values, we can study the accuracy of rotation measurements made with a transverse NMR gyroscope. All of the measurements done with the apparatus presented in Ch. 3 have been studies of the precision of the device. Since there was no expected rotation to measure, we have been essentially studying how well can we measure zero. With the second apparatus, we will measure non-zero rotation and be able to study the accuracy of the device.

Another major change on the second apparatus relates to our pump lasers. In the current apparatus, we use two overlapped pump beams, each detuned to either side of the Rb resonance. Once combined, the beams are re-split to enable dual-sided pumping. This pumping setup was largely informed by the fact that the laser diodes used in this apparatus do not produce enough power to fully penetrate our optically thick vapor when pumping on resonance. At the time, two-sided pumping with two lasers detuned to cancel the AC Stark shift was an elegant solution. However, over time we have learned about our sensitivity to transverse field drifts. The largest transverse field present in our system is the Rb spin-exchange field. Though we strive to cancel the average Rb spin-exchange field using x-compensation, we do not compensate for drifts in the Rb polarization. In our current pumping setup, there are a lot of variables that can contribute to drifts. Each pump laser may be subject to drifts in frequency, power, and pointing. Having two lasers doubles the

sources of possible drift. These many degrees of freedom, along with limited modulation capabilities in our current laser drivers, have made it complicated to design feedback to stabilize our optical pumping. Thus, the frequencies and powers of the lasers in our current system are all open loop.

I believe we would be better served to operate with a single pump laser on resonance. Since a single pump laser still does not produce enough power to penetrate through the optically thick vapor, we install a tapered amplifier to increase the power of the pump laser. This will allow us to pump on resonance, thereby avoiding AC Stark shifts. The tapered amplifier should provide enough power to pump from a single side without bad Rb polarization gradients, but if we find it necessary we can use a mirror after the cell to back-reflect the pump light for a second pass through the cell. We are also switching to laser drivers which enable more robust laser control. With the addition of a spectrum analyzer unit, we can actively stabilize the frequency of the laser. Laser stabilization is important when using a tapered amplifier, and has the added bonus of reducing Rb spin-exchange field drifts. Since transverse field sensitivity has been shown to be a factor limiting the performance of the transverse NMR gyroscope, I believe these changes will lead to improved performance.

Appendix A

Acronyms, Symbols, and Abbreviations

A.1 Acronyms and Abbreviations

The following is a list of acronyms and abbreviations that appear throughout this paper, in alphabetical order.

ADEV	Allan deviation
ARW	angle-random walk
DM	dichroic mirror
DSDO	dual-species, dual output
DSSO	dual-species, single output
EOM	electro-optic modulator
FB	feedback
FID	free-induction decay
FPGA	field-programmable gate array
FSF	field suppression factor

HWP	half-wave plate
MDEV	modified Allan deviation
NMR	Nuclear Magnetic Resonance
PBS	polarizing beam splitter
PD	photodiode
PDM	pulse density modulation
PM	polarization modulation
Pol	linear polarizer
QWP	quarter-wave plate
RTD	resistance temperature detector
SE	spin-exchange
SEOP	spin-exchange optical pumping
SNR	signal-to-noise ratio
TS	translation stage
WP	Wollaston prism

A.2 Symbols

The following is a list of symbols that appear throughout this paper, in approximate order of appearance.

γ	gyromagnetic ratio
s	sign of noble gas gyromagnetic ratio
B	magnetic field
X	spin-dependent non-magnetic frequency shift
Ω	precession frequency

a	denotes ^{129}Xe
b	denotes ^{131}Xe
Ω^R	rotation frequency
$\rho = \frac{\gamma^a}{\gamma^b}$	ratio of gyromagnetic ratios
S	alkali-metal electron spin polarization
K	noble gas nuclear spin polarization
$b_K^S K$	Xe spin-exchange field as experienced by Rb
$b_S^K S$	Rb spin-exchange field as experienced by Xe
κ	Fermi-contact enhancement factor
B_p	the pulsed bias magnetic field
ω_p	the repetition frequency of the bias pulses
R	optical pumping rate
Γ'	Rb relaxation rate
$B_w = \frac{\Gamma'}{\gamma^S}$	magnetic width of Rb magnetometer
$S_+ = S_x + iS_y$	transverse Rb spin polarization
$K_+ = K_x + iK_y$	transverse Xe spin polarization
S_\perp	magnitude of transverse Rb spin polarization
K_\perp	magnitude of transverse Xe spin polarization
$\epsilon_z = \tan^{-1}\left(\frac{S_y}{S_x}\right)$	magnetometer phase shift
B_{ext}	external magnetic fields, not intentionally applied
Γ_S^K	spin-exchange pumping rate from the Rb electrons to the Xe nuclei
$\Gamma_1 = 1/T_1$	longitudinal Xe relaxation rate
$\Gamma_2 = 1/T_2$	transverse Xe relaxation rate
ϕ	phase of transverse Xe spin precession
$s_{p,n}$	Fourier coefficients of the n-th polarization modulation

$\omega_{PM,n}$	modulation frequency of the n-th polarization modulation
$\omega_{PDM,m}$	modulation frequency of the m-th pulse density modulation
ω_d	Xe drive frequency
p_n	integer describing the contribution of the n-th polarization modulation to the drive frequency
q_m	integer describing the contribution of the m-th pulse density modulation to the drive frequency
B_{p0}	the un-modulated part of the applied pulsed bias field
$B_{mod}(t)$	the modulated part of the applied pulsed bias field
B_m	amplitude of the m-th pulse density modulation
$b_m = B_m/B_{p0}$	modulation depth of m-th pulse density modulation
$f_{2\pi}^0$	the center pulsing frequency
δ	Xe precession phase shift
$\alpha = \int (\omega_d + \gamma B_{mod}(t)) dt$	ideal expected Xe precession phase
J_q	q-th Bessel function of the first kind
j	product of the Fourier coefficients
$\Delta = \omega_d - \Omega$	detuning of the drive from the Xe resonance frequency
$\Delta_0 = \omega_d - (\Omega - \gamma B_{mod})$	detuning, not including modulated fields
\tilde{f}	a variable f represented in the Fourier domain, in other words the frequency noise of f
β, β'	cross talk coefficients
G	feedback gain
σ^\pm	right or left hand circularly polarized laser light
σ^π	linearly polarized laser light

B_{FB}	z-field correction applied by a feedback loop
B_{z0}	low-frequency external z-field
$scos$	sign of a cosine
ω_g	pulse gating frequency for gated detection
A	detected amplitude on S_z signal
C	a constant offset

references

Allan, D.W., and J.A. Barnes. 1981. A modified "allan variance" with increased oscillator characterization ability. *Proceedings of the 35th Annual Frequency Control Symposium, Philadelphia, 1981* 470–475.

Allmendinger, F., I. Engin, W. Heil, S. Karpuk, H.-J. Krause, B. Niederländer, A. Offenhäuser, M. Repetto, U. Schmidt, and S. Zimmer. 2019. Measurement of the permanent electric dipole moment of the ^{129}Xe atom. *Phys. Rev. A* 100:022505.

Allmendinger, F., W. Heil, S. Karpuk, W. Kilian, A. Scharth, U. Schmidt, A. Schnabel, Yu. Sobolev, and K. Tullney. 2014. New limit on lorentz-invariance- and *CPT*-violating neutron spin interactions using a free-spin-precession ^3He - ^{129}Xe comagnetometer. *Phys. Rev. Lett.* 112:110801.

Bechhoefer, J. 2005. Feedback for physicists: A tutorial essay on control. *Rev. Mod. Phys.* 77:783–836.

Bell, William E., and Arnold L. Bloom. 1957. Optical detection of magnetic resonance in alkali metal vapor. *Phys. Rev.* 107(6):1559–1565.

Bouchiat, M. A., T. R. Carver, and C. M. Varnum. 1960. first ^3He pumping. *PRL* 5:373.

Brown, J. M., S. J. Smullin, T. W. Kornack, and M. V. Romalis. 2010. New limit on lorentz- and *CPT*-violating neutron spin interactions. *Phys. Rev. Lett.* 105:151604.

Bulatowicz, M., R. Griffith, M. Larsen, J. Mirijanian, C. B. Fu, E. Smith, W. M. Snow, H. Yan, and T. G. Walker. 2013. Laboratory search for a long-range T -odd, P -odd interaction from axionlike particles using dual-species nuclear magnetic resonance with polarized ^{129}Xe and ^{131}Xe gas. *Phys. Rev. Lett.* 111:102001.

DeLand, Zachary J. 2017. Advances in fetal magnetocardiography using serf atomic magnetometers. Ph.D. thesis, University of Wisconsin-Madison. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2023-03-04.

Donley, E. A., and J. Kitching. 2013. Nuclear magnetic resonance gyroscopes. In *Optical magnetometry*, chap. 19, 369–386. Cambridge University Press.

Donley, E.A. 2010. Nuclear magnetic resonance gyroscopes. In *Sensors, 2010 ieee*, 17 –22.

Feng, Y.-K., D.-H. Ning, S.-B. Zhang, Z.-T. Lu, and D. Sheng. 2022. Search for monopole-dipole interactions at the submillimeter range with a ^{129}Xe – ^{131}Xe –Rb comagnetometer. *Phys. Rev. Lett.* 128:231803.

Glenday, A. G., C. E. Cramer, D. F. Phillips, and R. L. Walsworth. 2008. Limits on anomalous spin-spin couplings between neutrons. *Phys. Rev. Lett.* 101:261801.

Grover, Bruce C., Edward Kanegsberg, John G. Mark, and Roger L. Meyer. 1979. Nuclear magnetic resonance gyro. U.S. Patent 4,157,495.

Horowitz, P., and W. Hill. 2015. *The art of electronics*. Cambridge University Press.

Jiang, Liwei, Wei Quan, Rujie Li, Wenfeng Fan, Feng Liu, Jie Qin, Shuangai Wan, and Jiancheng Fang. 2018. A parametrically modulated dual-axis atomic spin gyroscope. *Applied Physics Letters* 112(5):054103. <https://doi.org/10.1063/1.5018015>.

Karwacki, F. A. 1980. Nuclear magnetic resonance gyro development. *Navigation* 27:72.

- Kornack, T. W., R. K. Ghosh, and M. V. Romalis. 2005. Nuclear spin gyroscope based on an atomic comagnetometer. *Physical review letters* 95(23):230801.
- Kornack, T. W., and M. V. Romalis. 2002. Dynamics of two overlapping spin ensembles interacting by spin exchange. *Phys. Rev. Lett.* 89:253002.
- Korver, A. 2015. Towards an nmr oscillator. Ph.D. thesis, University of Wisconsin-Madison.
- Korver, A., D. Thrasher, M. Bulatowicz, and T. G. Walker. 2015. Synchronous spin-exchange optical pumping. *Phys. Rev. Lett.* 115:253001.
- Korver, A., R. Wyllie, B. Lancor, and T. G. Walker. 2013. Suppression of spin-exchange relaxation using pulsed parametric resonance. *Phys. Rev. Lett.* 111:043002.
- Kwon, T. M., J. G. Mark, and C. H. Volk. 1981. Quadrupole nuclear spin relaxation of ^{131}Xe in the presence of rubidium vapor. *Phys. Rev. A* 24:1894–1903.
- Kwon, Tae M., and Charles H. Volk. 1984. Magnetic resonance cell. U.S. Patent 4,446,428.
- Lancor, Brian R. 2011. Studies of the efficiency of spin exchange optical pumping. Ph.D. thesis, University of Wisconsin-Madison.
- Lee, J., A. Almasi, and M. V. Romalis. 2018. Improved limits on spin-mass interactions. *Phys. Rev. Lett.* 120:161801.
- Limes, M. E., D. Sheng, and M. V. Romalis. 2018. ^3He – ^{129}Xe comagnetometry using ^{87}Rb detection and decoupling. *Phys. Rev. Lett.* 120:033401.
- Nahlawi, A. I., Z. L. Ma, M. S. Conradi, and B. Saam. 2019. High-precision determination of the frequency-shift enhancement factor in $\text{Rb} - ^{129}\text{Xe}$. *Phys. Rev. A* 100:053415.
- Press, William H., Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. 1992. *Numerical recipes in c (2nd ed.): the art of scientific computing*. New York, NY, USA:Cambridge University Press.

Romalis, M. V., D. Sheng, B. Saam, and T. G. Walker. 2014. Comment on “new limit on lorentz-invariance- and *CPT*-violating neutron spin interactions using a free-spin-precession ^3He - ^{129}Xe comagnetometer”. *Phys. Rev. Lett.* 113:188901.

Romalis, M.V. 2013. Quantum noise in atomic magnetometers. In *Optical magnetometry*, chap. 2, 25–39. Cambridge University Press.

Rosenberry, M. A., and T. E. Chupp. 2001. Atomic electric dipole moment measurement using spin exchange pumped masers of ^{129}Xe and ^3He . *Phys. Rev. Lett.* 86:22–25.

Sachdeva, N., I. Fan, E. Babcock, M. Burghoff, T. E. Chupp, S. Degenkolb, P. Fierlinger, S. Haude, E. Kraegeloh, W. Kilian, S. Knappe-Grüneberg, F. Kuchler, T. Liu, M. Marino, J. Meinel, K. Rolfs, Z. Salhi, A. Schnabel, J. T. Singh, S. Stüiber, W. A. Terrano, L. Trahms, and J. Voigt. 2019. New limit on the permanent electric dipole moment of ^{129}Xe using ^3He comagnetometry and squid detection. *Phys. Rev. Lett.* 123:143003.

Smiciklas, M., J. M. Brown, L. W. Cheuk, S. J. Smullin, and M. V. Romalis. 2011. New test of local lorentz invariance using a ^{21}Ne -Rb-K comagnetometer. *Phys. Rev. Lett.* 107:171604.

Sorensen, Susan S., Daniel A. Trasher, and Thad G. Walker. 2020. A synchronous spin-exchange optically pumped nmr-gyroscope. *Appl. Sci.* 10:7099.

Sorensen, Susan S., and Thad G. Walker. 2023. Combined polarization/magnetic modulation of a transverse nmr gyroscope. *Sensors* 23(10):4649.

Terrano, W A, and M V Romalis. 2021. Comagnetometer probes of dark matter and new physics. *Quantum Science and Technology* 7(1):014001.

Thorlabs. 2012. Eo-am-nr-cx eo-am-r-20-cx electro-optic amplitude modulator user guide.

Thrasher, D. A., S. S. Sorensen, J. Weber, M. Bulatowicz, A. Korver, M. Larsen, and T. G. Walker. 2019a. Continuous comagnetometry using transversely polarized xe isotopes. *Phys. Rev. A* 100:061403(R).

Thrasher, Daniel A. 2020. Continuous comagnetometry using transversely polarized xe isotopes. Ph.D. thesis, University of Wisconsin-Madison.

Thrasher, Daniel A., Susan S. Sorensen, and Thad G. Walker. 2019b. Dual-species synchronous spin-exchange optical pumping. 1912.04991.

Vanier, J., and C. Audoin. 1989. In *The Quantum Physics of Atomic Frequency Standards*, vol. 1, chap. Appendix 2F, 216–256. IOP Publishing Ltd, Bristol.

Volk, C. H., T.M. Kwon, and J. G. Mark. 1980. Measurement of the $^{87}\text{Rb} - ^{129}\text{Xe}$ spin-exchange cross section. *Phys. Rev. A* 21(5):1549–1555.

Walker, T.G., and M.S. Larsen. 2016. Spin-exchange-pumped nmr gyros 65:373 – 401.

Wyllie, R. 2012. The development of a multichannel atomic magnetometer array for fetal magnetocardiography. Ph.D. thesis, University of Wisconsin-Madison.

Zapasskii, V. S. 1982. Highly sensitive polarimetric techniques (review). *Journal of Applied Spectroscopy* 37(2):857–869.

Zhang, Da-Wei, Zheng-Yi Xu, Min Zhou, and Xin-Ye Xu. 2017. Parameter analysis for a nuclear magnetic resonance gyroscope based on $^{133}\text{Cs} - ^{129}\text{Xe} / ^{131}\text{Xe}^*$. *Chinese Physics B* 26(2):023201.

Zhang, Ke, Nan Zhao, and Yan-Hua Wang. 2020. Closed-loop nuclear magnetic resonance gyroscope based on rb-xe. *Scientific Reports* 10(2258).