Nonlinear Magneto-Optical Rotation in Rubidium Vapour

Nathanial Morgan Wilson

Supervisors: Prof. André Luiten,
Dr. Christopher Perrella, Dr. Philip Light

A thesis submitted for the degree of

Doctor of Philosophy

School of Physical Sciences
The University of Adelaide

January, 2020

This thesis is dedicated to my dear friend Dr. Christopher Baker,
who instilled in me a love for physics over a decade ago.

Contents

List of Figures

List of Tables

List of Symbols

m_F Ground-state F magnetic quantum number.

F Ground-state total atomic angular momentum quantum number.

F' Excited-state total atomic angular momentum quantum number.

T Temperature.

x x-axis in the Cartesian basis.

y y-axis in the Cartesian basis.

$\hat{z}$ Unit vector along the z-axis in the Cartesian basis.

∇ Three-dimensional gradient operator.

$\boldsymbol{B}$ Magnetic-field vector.

$\boldsymbol{r}$ Position vector.

μ_0 Vacuum permeability.

$\boldsymbol{\mu}$ Magnetic dipole moment.

Φ_B Magnetic flux.

$\mathrm{d}\boldsymbol{A}$ Infinitesimal element of surface area.

V_{emf} Electromotive force.

t Time.

B Magnetic-field amplitude.

V Voltage.

A_{cr} Cross-sectional area.

μ_r Relative permeability.

δB Minimum detectable magnetic field per unit bandwidth.

ψ Quantum-mechanical wavefunction.

φ Phase.

$\Delta\varphi$ Phase difference.

e Elementary charge.

h Planck constant.

A_{vp} Magnetic vector potential.

Φ_0 Magnetic flux quantum.

I Current.

ν Linear frequency.

k_{B} Boltzmann constant.

μ_x Magnetic-dipole-moment component along the x-axis.

μ_y Magnetic-dipole-moment component along the y-axis.

μ_z Magnetic-dipole-moment component along the z-axis.

Ω_{L} Larmor frequency.

γ_{g} Gyromagnetic ratio.

T_2 Transverse spin-relaxation time.

T_1 Longitudinal spin-relaxation time.

B_z Magnetic-field amplitude along the z-axis.

B_y Magnetic-field amplitude along the y-axis.

B_x Magnetic-field amplitude along the x-axis.

z z-axis in the Cartesian basis.

ω Angular frequency.

$\Delta\omega$ Detuning.

$\hbar$ Reduced Planck constant.

g_s Electron spin g-factor.

μ_{B} Bohr magneton.

n_{v} Atomic number density.

Γ_{pump} Optical-pumping rate due to the pump beam.

Γ_{pr} Rate of depolarisation due to the probe beam.

Γ_{SD} Rate of electron spin-destruction collisions.

m_S Ground-state magnetic quantum number.

λ Wavelength.

Z Atomic number.

n Principal quantum number.

$\mathbb{Z}$ Set of all integers.

L Ground-state orbital angular momentum quantum number.

m_L Ground-state magnetic quantum number.

S Spin quantum number.

J Ground-state total electronic angular momentum quantum number.

$\mathscr{H}$ Hamiltonian operator.

m_{e} Mass of the electron.

c Speed of light in vacuum.

I Nuclear quantum number.

A_{hfs} Hyperfine structure constant.

P_{v} Vapour pressure.

N_{A} Avogadro constant.

R_{g} Ideal gas constant.

$\mathcal{R}_{\text{abs}}$ Absorption rate of photons by an atom.

σ_A	Cross-section of an atomic transition.
Φ_P	Photon flux.
ω_0	Resonant angular frequency.
r_e	Classical electron radius.
f_0	Oscillator strength.
σ_E	Uncertainty in energy.
σ_t	Uncertainty in time.
τ_{nat}	Natural lifetime of an excited state.
σ_ω	Uncertainty in angular frequency.
Γ_{nat}	Natural linewidth of an excited state.
ν_{col}	Frequency with which atoms/molecules collide.
σ_{col}	Collisional cross-section.
$\bar{v}_{th}$	Mean atomic thermal velocity.
Γ_L	Full width at half maximum of a Lorentzian profile.
$\mathcal{L}$	Lorentzian profile.
I_{int}	Optical intensity.
α_{abs}	Absorption coefficient.
$n_{|g\rangle}$	Ground-state atomic population density.
$n_{|e\rangle}$	Excited-state atomic population density.
I_{sat}	Saturation intensity.
Γ	Linewidth.
Γ_{pb}	Power-broadened linewidth.
Ω_R	Rabi frequency.
κ_2	Saturation parameter.
P_{probe}	Optical power of the probe beam.
P_{pump}	Optical power of the pump beam.
δ_s	ac Stark shift.
v_z	Velocity component along the z-axis.
m	Mass.
$\boldsymbol{v}$	Velocity vector.
$\mathscr{P}$	Probability density function.
v	Velocity magnitude.
v_x	Velocity component along the x-axis.
v_y	Velocity component along the y-axis.
Γ_D	Doppler width.
Γ_G	Width (standard deviation) of a Gaussian profile.
$\mathcal{G}$	Gaussian profile.

$\mathcal{V}$	Voigt profile.
w	Faddeeva function.
erfc	Complementary error function.
$\Gamma_{\mathcal{V}}$	Full width at half maximum of a Voigt profile.
$\mathscr{H}_{\mathrm{I}}$	Interaction Hamiltonian.
$\mathscr{H}_0$	Atomic (unperturbed) Hamiltonian.
$\Delta E^{(1)}$	First-order energy correction due to perturbation theory (linear Zeeman effect).
g_F	Landé g-factor.
$\Delta E^{(\mathrm{BR})}$	Breit-Rabi energy correction due to the nonlinear Zeeman effect.
ΔE_{hfs}	Hyperfine energy splitting.
g_I	Nuclear g-factor.
g_J	Landé g-factor.
δ^{NLZ}	Frequency splitting of the ground-state coherences due to the nonlinear Zeeman effect.
$\mathcal{S}$	Resonance slope.
$\mathcal{A}$	Resonance amplitude.
ϕ	Optical/polarisation rotation.
$\mathcal{V}_e$	Verdet constant.
Γ_{d}	Excited-state decay rate.
k_{s}	Spring constant.
$\mathcal{E}$	Electric-field vector.
$\mathscr{F}$	Fourier transform.
$\tilde{\alpha}$	Complex polarisability.
$\mathbf{P}$	Polarisation field.
n_e	Electron number density.
ϵ_0	Vacuum permittivity.
$\tilde{\chi}$	Complex susceptibility.
$\tilde{\chi}_0$	Complex-susceptibility amplitude.
$\tilde{\eta}$	Complex refractive index.
σ^+	Right-circular polarisation.
σ^-	Left-circular polarisation.
Ω_{m}	Modulation frequency.
$\mathbb{C}$	Set of all complex numbers.
$\mathbb{R}$	Set of all real numbers.
C_n	Probability amplitude.
U_n	Eigenfunction of the time-independent Schrödinger equation.
E	Energy.
ρ	Density matrix.

$\hat{\Gamma}$	Relaxation matrix.
δ	Dirac delta function.
Λ	Repopulation matrix.
$\mathcal{H}_E$	Electric-field Hamiltonian.
$\mathcal{H}_B$	Magnetic-field Hamiltonian.
$\boldsymbol{d}$	Dipole operator.
$\mathcal{E}_0$	Electric-field amplitude.
$\hat{\boldsymbol{x}}$	Unit vector along the x-axis in the Cartesian basis.
d_x	Dipole operator for the x-component in the Cartesian basis.
d_{-1}	Dipole operator for right-circularly polarised light.
d_1	Dipole operator for left-circularly polarised light.
ξ	Additional relevant quantum numbers for the ground state.
ξ'	Additional relevant quantum numbers for the excited state.
$\boldsymbol{F}$	Total atomic angular momentum vector.
F_z	Total atomic angular momentum along the z-axis.
$\widetilde{\mathcal{H}}$	Hamiltonian operator in the rotating frame.
$\tilde{\rho}$	Density matrix in the rotating frame.
U_t	Unitary transformation.
γ_t	Transit relaxation rate.
m_F'	Excited-state F magnetic quantum number.
ϵ	Electric-field ellipticity.
k	Wavevector magnitude (wavenumber).
P	Component of the polarisation field.
J'	Excited-state total electronic angular momentum quantum number.
l_0	Resonant absorption length.
$\mathcal{O}$	Quantum observable.
$\overline{\Omega}_R$	Peak Rabi frequency.
D	Duty cycle.
T_m	Modulation period.
Γ_{SE}	Rate of spin-exchange collisions.
σ_{SE}	Spin-exchange-collision cross-section.
σ_{SD}	Spin-destruction-collision cross-section.
γ_{wall}	Rate of wall collisions.
$\mathcal{D}$	Diffusion constant.
P_A	Atmospheric pressure.
T_{wall}	Time between wall collisions.
A_s	Magnetic shielding efficiency.

$\mathcal{R}$	Magnetic reluctance.
μ_{P}	Magnetic permeability.
$\hat{\boldsymbol{y}}$	Unit vector along the y-axis in the Cartesian basis.
ν_{co}	Cut-off frequency (linear) of a Butterworth filter.
B_{dc}	dc Magnetic-field amplitude.
ϕ_P	In-phase component of the demodulated optical/polarisation-rotation signal.
ϕ_Q	Out-of-phase/quadrature component of the demodulated optical/polarisation-rotation signal.
$\mathcal{L}_{\mathrm{c}}$	Complex-Lorentzian profile.
τ	Integration time.
var	Variance of a parameter.
σ	Standard deviation.
S_{PSD}	Power spectral density.
σ_B	Uncertainty in magnetic-field strength.
$\delta\phi$	Minimum detectable optical rotation per unit bandwidth.
$\delta B \sqrt{V}$	Minimum detectable magnetic field per unit spatiotemporal bandwidth.
E_{R}	Minimum detectable energy per unit bandwidth.
F_x	Total atomic angular momentum along the x-axis.
d	Wigner's (small) d-function.
Γ_B	Linewidth in magnetic units.
ΔB	ac Magnetic-field amplitude.
ν_{I}	Instantaneous frequency (linear).
$\delta\nu$	Frequency (linear) deviation.
β	Modulation index.
J	Bessel function of the first kind.
φ_{I}	Instantaneous phase.
$\mathcal{H}$	Hilbert transform.
$\Delta\nu_{\mathrm{CBR}}$	Required bandwidth for a modulated signal, according to Carson's bandwidth rule.
δI	Current noise per unit bandwidth.
δB_{J}	Johnson magnetic-field noise.
σ_{c}	Electrical conductivity.
δB_{M}	Magnetisation magnetic-field noise.
δ_{skin}	Skin depth.
$\hat{\boldsymbol{\epsilon}}_1$	Spherical-basis unit vector corresponding to left-circularly polarised light.
$\hat{\boldsymbol{\epsilon}}_0$	Spherical-basis unit vector corresponding to light polarised linearly along the z-axis.

$\hat{\epsilon}_{-1}$	Spherical-basis unit vector corresponding to right-circularly polarised light.
ϵ	Spherical-basis unit vector.
ϵ_1	Spherical-basis vector component corresponding to left-circularly polarised light.
ϵ_0	Spherical-basis vector component corresponding to light polarised linearly along the z-axis.
ϵ_{-1}	Spherical-basis vector component corresponding to right-circularly polarised light.
m_J	Ground-state magnetic quantum number.
m_J'	Excited-state magnetic quantum number.
$\widetilde{\psi}$	Quantum-mechanical wavefunction in the rotating frame.
k	Wavevector.
$\hat{\mathbf{e}}_1$	Generalised unit vector in an arbitrary basis.
$\hat{\mathbf{e}}_2$	Generalised unit vector in an arbitrary basis.
$\hat{k}$	Unit vector in the direction of the wavevector.
$\mathbf{D}$	Displacement field.
ρ_f	Free-electric-charge density.
$\mathbf{H}$	Magnetising field.
$\mathbf{J}_f$	Free-current density.
$\mathbf{M}$	Magnetisation field.
d_y	Dipole operator for the y-component in the Cartesian basis.
d_z	Dipole operator for the z-component in the Cartesian basis.
$\mathcal{T}$	Transmission function.
Y	Bessel function of the second kind.
A_{ASD}	Amplitude spectral density.
σ^2	Variance.
σ_{M}^2	M-sample variance.
σ_{A}^2	Allan variance.
σ_{A}	Allan deviation.
cov	Covariance between two parameters.
erf	Error function.
$\mathcal{I}$	Fisher information.
$\boldsymbol{\mathcal{I}}$	Fisher information matrix.
$\mathscr{D}$	Quantum-mechanical rotation operator.
J_z	Total electronic angular momentum along the z-axis.
J_y	Total electronic angular momentum along the y-axis.

Acronyms

MAD	magnetic anomaly detection.
SQUID	superconducting quantum interference device.
SNR	signal-to-noise-ratio.
FID	free-induction decay.
rf	radio-frequency.
SERF	spin-exchange relaxation-free.
NV	nitrogen-vacancy.
NMOR	nonlinear magneto-optical rotation.
rf OAM	radio-frequency optical atomic magnetometer.
HWHM	half width at half maximum.
FWHM	full width at half maximum.
cw	continuous-wave.
ECDL	external-cavity diode laser.
PZT	piezoelectric transducer.
SAS	saturated absorption spectroscopy.
AOM	acousto-optic modulator.
TE	transverse electric.
PLA	polylactic acid.
CRLB	Cramér-Rao lower bound.
rms	root mean square.
TEM	transverse electromagnetic.
AMPS	angular momentum probability surface.

Publications

The work detailed within this thesis has resulted in a number of journal publications. These publications are presented in reverse chronological order below:

1. **Nathanial Wilson**, Christopher Perrella, Russell Anderson, André Luiten, and Philip Light, *Wide-bandwidth atomic magnetometry via instantaneous-phase retrieval*, Phys. Rev. Res. **2**, 013213 (2020).

2. **Nathanial Wilson**, Philip Light, André Luiten, and Christopher Perrella, *Ultrastable Optical Magnetometry*, Phys. Rev. Appl. **11**, 044034 (2019).

3. **Nathanial Wilson**, Nicolas Bourbeau Hébert, Christopher Perrella, Philip Light, Jérôme Genest, Szymon Puselny, and André Luiten, *Simultaneous Observation of Nonlinear Magneto-Optical Rotation in the Temporal and Spectral Domains with an Electro-Optic Frequency Comb*, Phys. Rev. Appl. **10**, 034012 (2018).

Conference Proceedings

In addition to journal publications, the work detailed within this thesis has been presented at a number of domestic and international conferences. These conference proceedings are listed in reverse chronological order below, where the asterisk indicates the presenter:

1. Rujie Li, **Nathanial Wilson**, Christopher Perrella, Philip Light, André Luiten*, and Russell Anderson, *Quantum sensing of rapidly varying magnetic fields*, SPIE Photonics West, San Francisco, United States of America, 2020.

2. **Nathanial Wilson**, Rujie Li*, Christopher Perrella, Russell Anderson, Philip Light, and André Luiten, *A high-bandwidth atomic magnetometer*, The Australian and New Zealand Conferences on Optics and Photonics (ANZCOP), Melbourne, Australia, 2019.

3. Christopher Perrella*, **Nathanial Wilson**, Rujie Li, Russell Anderson, Philip Light, and André Luiten, *High-bandwidth optical magnetometry*, Workshop on Optically Pumped Magnetometers (WOPM), Mainz, Germany, 2019.

4. Christopher Perrella*, **Nathanial Wilson**, Rujie Li, Russell Anderson, Philip Light, and André Luiten, *High-bandwidth optical magnetometry*, Quantum Sensing & Magnetometry, Bad Honnef, Germany, 2019.

5. **Nathanial Wilson**, Christopher Perrella, Philip Light, Rujie Li, and André Luiten*, *Vapour Cell Magnetometry*, International Workshop on Advanced Magnetometry for Defence Applications, Adelaide, Australia, 2018.

6. **Nathanial Wilson***, Christopher Perrella, Philip Light, and André Luiten, *Highly Sensitive Optical Magnetometer with Low Drift*, Australian Institute of Physics Summer Meeting, Sydney, Australia, 2017.

Abstract

The field of magnetometry has played a pivotal role in developing our understanding of the world around us, by providing a wealth of unique information via non-invasive magnetic measurements. In many applications the characteristics of magnetic fields — their strength and temporal dependence, for example — constitute a unique fingerprint of their source, and provide a window into both their origin and any underlying physical processes that produced them.

The invention of the laser, and recent advancements in anti-relaxation coatings for vapour cells, has revolutionised the field of magnetometry and given rise to a myriad of optical magnetometers which demonstrate exquisite sensitivity. This has seen a movement away from traditional cryogenic magnetometry techniques — which utilise superconducting quantum interference devices — and has ushered in a slew of optical techniques which offer a whole host of benefits such as room-temperature operability, intrinsic accuracy, and portability. Applications which were previously inaccessible due to demanding requirements such as high sensitivity at ambient temperatures, are now permissible using these newly developed optical techniques.

This thesis details the development of an ultrastable optical magnetometer based on nonlinear magneto-optical rotation in rubidium vapour, that has been optimised for high sensitivity over long timescales. Additionally, a novel measurement technique — which relies on tracking the instantaneous phase of atomic spin precession — has been developed and demonstrated, resulting in significant improvements to both the amplitude and frequency response of the device when compared to conventional measurement techniques described in the literature. Significant attention is paid to technical and fundamental noise sources, which must be minimised in order to achieve high performance — especially over long timescales. The latter chapters discuss the development of the magnetometer into a sensitive and reliable system, capable of measuring extremely small magnetic-field fluctuations, and transient fields with arbitrarily complex temporal dependence. The capability of the device to measure minute field fluctuations, as well as rapid transient disturbances, may find use in real-world applications such as geophysical exploration, medical diagnostics and imaging, and magnetic anomaly detection. Furthermore, any fundamental physics applications which demand accurate and precise measurements of magnetic fields may find application of this device — and the novel measurement techniques developed within this thesis — extremely beneficial.

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I acknowledge that copyright of published works contained within this thesis resides with the copyright holder(s) of those works.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

_____

Nathanial Morgan Wilson
29th January, 2020

Acknowledgements

There are a multitude of people whom I wish to thank for their contributions and support throughout my candidature. Of those people, the first that come to mind are my supervisors: André Luiten, Philip Light, and Christopher Perrella. The results of this thesis would not have been possible without their ideas, expertise, and guidance. I would also like to thank my collaborators: Nicolas Bourbeau Hébert, Jérôme Genest, Szymon Pustelny, and Russell Anderson. I have enjoyed the collaborative research we have undertaken together, and I look forward to future endeavours should the opportunity arise.

I also wish to acknowledge the individuals who have fabricated bespoke experimental components. I wish to thank the enthusiastic technical staff at the Institute for Photonics and Advanced Sensing workshop — specifically Evan Johnson — for assistance with designing and 3D-printing various components for use in the experiment. I am also grateful to Michael Souza and Mikhail Balabas, for their roles in fabricating the high-quality anti-relaxation-coated vapour cells used in the experiment.

I also wish to thank my fellow postgraduate students; not only those in the Precision Measurement Group, but also those in different fields and at various universities around Australia: Ashby Hilton, Hannes Griesser, Sarah Scholten, Jonathan Hedger, Jake Clark, and Andrew Malouf. It has been a pleasure to share this wild ride with you, and I wish you all the best for the future.

I would also like to thank my friends outside of university, who have supported me through the highs and lows: Winschurill Ho, Tom Martin, Eleazar Quinto, Owen Blinn, Laura Paterson, John 'Jake' McWhae, Ben Whitfield, Joey Evans, Michael Horton, Yetchi Yeung, Kim Bui, Nikolai Kojevnikov, Jordan Summerfield, Nicholas Krippner, Alexander Kitto, and Wayde Kirby. I value my friendship with each and every one of you.

I also wish to thank my friends in the South Australian State Emergency Service for making me feel at home. Although I have only been a part of the service for the last two years, I consider you all part of my family, and serving the community with you has been both a privilege and an honour.

I would also like to extend my gratitude to the Australian Government for awarding me an Australian Postgraduate Award, which supported me financially during my candidature. Additionally, I am also grateful to the Defence Science and Technology Group for funding this project.

And last but not least, I extend my sincere gratitude to my grandparents who, over the last 17 years, have taken over the role of parents and shaped me into the man I am today. I appreciate all that you have done for me and I am lucky to have been blessed with you in my life.

Part I

Introduction and Literature Review

Applications and Motivation

The accurate and precise measurement of magnetic fields is a flourishing research area, and has a myriad of applications in both fundamental and applied physics. Determination of the amplitude, frequency, and transient response of magnetic fields can provide a wealth of unique information of the source that produced them. For example, abduction of the left index finger results in the reduction of electrophysiological activity (and hence emitted magnetic field) in the right sensorimotor cortex, which subsequently increases above baseline following cessation of the movement [1]. This chapter aims to introduce a small subset of the vast array of applications of precision magnetometry, as a motivation for the experimental work detailed within this thesis.

1.1 Applications

1.1.1 Magnetic Anomaly Detection

With research dating as far back as the late 1960's [2], magnetic anomaly detection (MAD) has been used for decades to locate and identify ferromagnetic targets [3] — even targets which are undetectable via alternate means — by exploiting the fact that most naturally occurring materials (e.g. air, water, soil, etc.) are transparent to static and low-frequency magnetic fields [4]. Magnetic anomaly detection can either be performed as an active [5–7] or a passive measurement [4,8,9], with the latter having the crucial advantage of being undetectable by the target of interest. Applications of MAD in the literature include detection, localisation and classification of: unexploded ordnance [10–17], undersea mines [18,19], vehicles [20–23], and general ferrous objects [24–30].

1.1.1.1 Working Principle

Ferromagnetic targets such as those comprised of iron, nickel, or cobalt, perturb Earth's magnetic field through generation of their own local magnetic field [31]. In the vicinity around these targets, the magnetic field has a complicated spatial dependence; however, at distances greater than three times the largest dimension of the target, it is a good approximation to treat the magnetic field as being produced by that of a point magnetic dipole[1] [2–4,8,24,32–34]:

---

[1]In the near field, the magnetic field is described by a multipole expansion. The lowest terms in the multipole expansion are the monopole, dipole and quadrupole moments; which have magnetic potentials that scale as r^{-1}, r^{-2}, and r^{-3}, respectively. Given there are no magnetic monopoles (i.e. $\nabla \cdot \boldsymbol{B} = 0$), the lowest-order term observed from magnetic sources is the dipole moment. Owing to its slow rate of decrease, the dipole term dominates at large distances and hence any magnetic source can be approximated as a dipole of the same magnetic moment in the far field.

$$B(r) = \frac{\mu_0}{4\pi}\left[\frac{3(\mu \cdot r)r}{|r|^5} - \frac{\mu}{|r|^3}\right], \tag{1.1}$$

where $B(r)$ is the magnetic field at a position r relative to the source, and μ is the target's magnetic moment.

In addition to ferromagnets, Earth's magnetic field is also perturbed by objects with high magnetic permeability due to a process called 'flux shunting' (cf. Sec. 5.3). In these objects, Earth's magnetic field induces a magnetic moment, which in turn creates its own local magnetic field [31]. When added vectorially to Earth's field, this induced magnetic field creates a non-uniform net field that can be readily measured by a magnetometer [31]. Given that Earth's magnetic field is fairly uniform — with expected gradients on the order of 10–30 nT/km [2,3,26] — it is quite easy to measure relatively small magnetic anomalies buried within Earth's large dc magnetic field of $\mathcal{O}(50\,\mu\text{T})$.

Each target possesses a magnetic moment which is relatable to its physical properties, such as its composition and dimensions. Provided that some approximations are made, the dipolar field produced by this magnetic moment can be used as a signature in order to locate and identify the target [17,31]. However, given the complicated spatial dependence of the field produced by the target, multiple field measurements from different points are required. This is further complicated by the fact that a total field sensor is sensitive to the *projection* of an anomalous field along the direction of Earth's field[2], typically necessitating the use of *arrays* of magnetometers for accurate identification and localisation.

1.1.1.2 Detection Methods

Provided that some approximations are made, the object's location can be determined via inversion techniques [31]. These techniques estimate the location by assuming that the field is produced by a magnetic dipole whose calculated magnetic field most closely resembles the set of measurements that are observed [31]. This constitutes a nonlinear optimisation problem, in which the total mismatch between the measured readings and the calculated field from an object is minimised [31].

Magnetic anomaly detection has traditionally used either superconducting quantum interference devices (SQUIDs) [2,33–38] — typically in airborne detection of underwater submarines [33] or surface ships [35] — or coil-based induction sensors [39,40] (such as the fluxgate magnetometer, etc.). However, recent demonstrations have involved using arrays of scalar atomic magnetometers (i.e. magnetometers which measure the magnitude of the magnetic field $|B|$ using atoms — typically optically pumped caesium magnetometers) to perform triangulation of airborne magnetic anomalies [8]. Scalar atomic magnetometers have also been implemented in unmanned airborne systems [17], unmanned undersea systems [17], and ground-based systems [17].

---

[2]This is a result of the relative magnitude of Earth's field to that of typical magnetic anomalies. If Earth's field dominates, the vector component of an anomaly which is *orthogonal* to Earth's field does not significantly change the magnitude of the total field, but rather only changes the direction slightly [31].

1.1.2 Medical Diagnostics and Imaging

The sensitivity and bandwidth requirements of medical imaging applications are extremely demanding, owing to the minute and often high-frequency magnetic fields produced by biological sources in the human (or animal) body. Two of the more common medical applications discussed in the literature are magnetocardiography and magnetoencephalography — that is, measurements of the magnetic field produced by the heart and brain, respectively. Of these two applications, magnetocardiography is the least demanding, requiring resolution of fields with magnitudes of $\mathcal{O}\left(100\,\mathrm{pT}\right)$ [41,42]. Magnetoencephalography on the other hand is far more difficult, requiring resolution of fields of $\mathcal{O}\left(100\,\mathrm{fT}\right)$ with a bandwidth up to 1 kHz [41].

Given their exquisite sensitivity and ability to operate without cryogenic cooling, optical magnetometers are rapidly finding application in medical diagnostics and imaging. This is in part due to an optical magnetometer's ability to be placed in close proximity to the target — yielding a three-to-five-fold improvement in sensitivity over cryogenic sensors [43,44] (cf. Sec. 2.1.2) — while also offering the ability to wear the sensor while complex tasks are being performed [1,44]. The latter is incredibly useful for magnetoencephalography measurements which require the patient to perform movements, a task which has been difficult — if not impossible — using conventional cryogenic sensors [44]. Despite the fact that for each magnetic measurement technique, an analogue electrical measurement technique often exists, the magnetic signals can often provide complementary or unique information [45]. For example, foetal magnetocardiograms are largely unaffected by the impedance of the foetal skin and *vernix caseosa*, which permits high signal-to-noise-ratio (SNR) measurements when compared with the electrical analogue [46] — foetal electrocardiograms.

1.1.2.1 Magnetoencephalography

Magnetoencephalography has been one of the main medical applications for which optical magnetometers have been developed in recent years [1,43,44,47–61]. Observations of magnetic fields originating from the human brain in response to auditory stimuli have been demonstrated [47–50,52,53,56,59–62], as well as spontaneous α-oscillations triggered by closing of the eyes [51], evoked responses from median nerve stimulation [43,53,54,56], visually-evoked fields [62], retinotopy [55], and fields arising from more complicated tasks such as head movements [1] or hippocampal-dependent tasks [63]. Magnetoencephalography has also been demonstrated on animals, via measurements of the magnetic signature of epileptiform discharges in rats [58].

1.1.2.2 Magnetocardiography

In addition to magnetoencephalography, optical magnetometers have also been fruitful in magnetocardiography, with the measurement of magnetocardiograms — the magnetic analogue of electrocardiograms — having been demonstrated [42,45,50,64–68]. Additionally, *in utero* foetal magnetocardiography has been demonstrated [69], which offers the ability to

detect birth defects such as ectopia cordis [46], and supraventricular tachycardia [70].

1.1.2.3   Miscellaneous

Although magnetoencephalography and magnetocardiography are perhaps the two most common applications of optical magnetometers in the context of medical diagnostics and imaging, there are also a number of other medical applications which have been demonstrated. These include optical magnetic induction tomography of the heart [71], and detection of animal nerve impulses [72,73].

1.1.3   Other Applications

In addition to the aforementioned applications, there are a myriad of other applications — both fundamental and applied — which benefit from accurate and precise magnetometry. These include, but are not limited to: nuclear magnetic resonance [74], the search for long-range spin-mass couplings of the proton [75], geomagnetic field measurements using mesospheric sodium [76–79], the search for local Lorentz invariance violations [80,81], magnetic resonance imaging [82], battery diagnostics [83], electromagnetic induction imaging [84–87], localisation and identification of non-metallic targets [88], remote detection of rotating machinery [89], material characterisation [88,90–92], imaging of conductive objects behind metallic and ferromagnetic barriers [93], detection of conductive targets from within saline water [94], the search for dark matter [95–100], the search for plant biomagnetism [101], real-time threat detection [17,31], magnetic surveying [102], earthquake monitoring [103–105], and magnetic communication of digitally encoded data [106,107].

1.2   Motivation

The work comprising this thesis was undertaken with the aim of addressing two of the main limitations of optical magnetometry: the reduction in sensitivity of dc optical atomic magnetometers over long timescales due to technical noise sources, and both the narrow amplitude and frequency response of radio-frequency optical atomic magnetometers. Both of these limitations hinder the application of optical magnetometers in MAD — of both small, fast moving and large, slow moving targets — which was the intended outcome of this research. Furthermore, medical applications such as magnetocardiography and magnetoencephalography require high sensitivity and bandwidth *simultaneously* — an exceptionally demanding criteria which has thus far proven elusive when using conventional magnetometry techniques.

Magnetometry Techniques

There are a vast array of magnetometry techniques that have been used over the last few hundred years. Prior to the invention of the laser in the 1960's, most of these techniques were non-optical; however, the technological revolution sparked by the laser has given rise to optical measurement techniques which currently dominate the field with their exquisite sensitivity. This chapter aims to discuss the more common — and highly sensitive — magnetometry techniques which fall under two distinct categories that are relevant in the context of this thesis: non-optical measurement techniques, and optical measurement techniques.

2.1 Non-Optical Measurement Techniques

Prior to the invention of the laser, non-optical measurement techniques dominated the field of magnetometry. Most of these techniques utilised induction-based or Hall-effect-based sensors, until the invention of the SQUID. Two of the most common sensitive, non-optical measurement techniques — the parallel fluxgate magnetometer, and the SQUID — are discussed below.

2.1.1 Parallel Fluxgate Magnetometers

The parallel fluxgate magnetometer, like other induction-based magnetometers, has its roots in Faraday's law of induction, and can be traced as far back as the 1930's [108]. The general idea behind induction sensors like the fluxgate magnetometer, is that the presence of an oscillating magnetic field produces a voltage in a conductor. As governed by Faraday's law of induction, for a loop of wire in a magnetic field, the magnetic flux Φ_B through a surface Σ whose boundary is the given loop, may be calculated using the surface integral

$$\Phi_B = \iint_\Sigma \boldsymbol{B} \cdot \mathrm{d}\boldsymbol{A} \,, \tag{2.1}$$

where $\boldsymbol{B}$ is the magnetic field vector, and $\mathrm{d}\boldsymbol{A}$ is the element of surface area of Σ. Faraday's law states that the electromotive force V_emf is proportional to the temporal derivative of the magnetic flux:

$$V_\mathrm{emf} = -N\frac{\mathrm{d}\Phi_B}{\mathrm{d}t} \,, \tag{2.2}$$

where N is the number of turns of the coil.

Given that the electromotive force is proportional to the magnetic flux Φ_B, which in turn is proportional to the vector inner product between the magnetic field $\boldsymbol{B}$ and element of

Figure 2.1: Diagram of a single-core (left) and dual-core (right) parallel fluxgate magnetometer [108]. An applied magnetic field B_E due to the voltage at the point E causes periodic saturation within the ferromagnetic core, subsequently modulating the core permeability and resulting in a voltage output as measured at P that is proportional to the dc component of the measured field B_M.

surface area dA, magnetometers which exploit Faraday's law of induction are intrinsically *vector* magnetometers. That is, they measure only the magnetic field *component* which is *parallel* (or *anti-parallel*) to the coil. This is in stark contrast with *scalar* magnetometers, which measure the *magnitude* (i.e. the quadrature sum of the components) of the magnetic field.

Parallel fluxgate magnetometers come in a number of configurations; however, in the simplest case depicted in the left of Fig. 2.1, the output voltage V_{fluxgate} is well approximated by [108]

$$V_{\text{fluxgate}} = -NA_{\text{cr}} \left(\frac{dB_E}{dt} + \mu_r K \frac{dB_M}{dt} + KB_M \frac{d\mu_r}{dt} \right) , \qquad (2.3)$$

where N is the number of turns of the pick-up coil, A_{cr} is the cross-sectional area of the core, B_E is the alternating excitation strength in the ferromagnetic core due to the excitation field, μ_r is the relative permeability of the core, B_M is the measured external magnetic-field strength, and K is a dimensionless coupling coefficient of the core to the field B_M.

The first term in Eq. (2.3) describes the classic output voltage one would obtain from using a standard, coreless pick-up coil (cf. Eq. 2.2), due to the applied excitation field B_E [108]. Due to the fact that this term in the output voltage arises from an *applied* field, rather than the field that is to be *measured*, the presence of this term is a disadvantage. The second term is due to an oscillating (or time-varying) component of the measured field B_M [108]. The third and final term in Eq. (2.3) is the *intrinsic benefit* of the fluxgate magnetometer. The alternating excitation field periodically causes saturation within the ferromagnetic core, subsequently modulating the core permeability and resulting in a non-zero time derivative [108]. The output voltage generated by this term is proportional to the dc component of the measured field, thereby resulting in a sensitivity of the fluxgate magnetometer to dc fields which traditional coreless pick-up coils do not possess.

In practice it can be beneficial to use *two* ferromagnetic cores, where each core faces the opposite direction to the other and each have their own excitation fields. The pick-up coil is then wrapped around *both* of the cores as shown in Fig. 2.1. Provided that the magnetic properties and physical dimensions of both cores are identical, the first term of Eq. (2.3) is

negated, subsequently removing the intrinsic disadvantage of the single-core design. If the measured magnetic field is constant, and two identical ferromagnetic cores are implemented in order to cancel the voltage output obtained from the applied excitation field, then Eq. (2.3) simplifies to [108,109]

$$V_{\text{fluxgate}} = -NA_{\text{cr}}B_{\text{M}}\frac{\mathrm{d}\mu_r}{\mathrm{d}t}\frac{1-D}{\left[1+D\left(\mu_r-1\right)\right]^2} \text{,} \qquad (2.4)$$

where D is a dimensionless 'demagnetisation factor' of the ferromagnetic core. Equation (2.4) is the 'basic fluxgate equation' commonly found in the literature [109].

Disregarding technical noise sources such as electronic noise, the sensitivity of a fluxgate magnetometer is intrinsically limited by three effects: Barkhausen noise [110,111], thermal Johnson (white) noise [110–112], and an anomalous noise source which is believed to be the result of inhomogeneous, stochastic magnetoelastic coupling of the non-zero magnetostrictive core to external stresses [108,113,114]. Barkhausen noise originates from the discrete, discontinuous changes in the magnetisation of the ferromagnetic core, which gives rise to discontinuous output voltage fluctuations and constitutes an effective noise. It has been shown that the minimum detectable magnetic-field strength $\delta B_{\text{fluxgate}}$ — given a SNR of unity — that can be detectable by a fluxgate magnetometer per unit bandwidth, can be expressed as [110]

$$\delta B_{\text{fluxgate}} = \frac{B_{\text{sat}}}{\mu_r}\sqrt{\frac{\tau_{\text{mag}}}{n_{\text{B}}}} \text{,} \qquad (2.5)$$

where B_{sat} is the saturation strength of the core, τ_{mag} is the magnetisation period (i.e. the inverse of the excitation frequency), and n_{B} is the density of Barkhausen volumes.

Typical fluxgate magnetometers exhibit a magnetic noise floor of about $10\,\text{pT}_{\text{rms}}/\sqrt{\text{Hz}}$ [108], with some demonstrating as low as $2\,\text{pT}_{\text{rms}}/\sqrt{\text{Hz}}$ at $1\,\text{Hz}$ [112]. In most cases, the white thermal Johnson noise in these devices is generally at least an order of magnitude below that of the Barkhausen noise floor [108], with some demonstrating as low as $0.39\,\text{pT}_{\text{rms}}/\sqrt{\text{Hz}}$ [112]. Bandwidths of up to $3\,\text{kHz}$ have been demonstrated in commercial devices.

2.1.2 Superconducting Quantum Interference Devices

Superconducting quantum interference devices exploit two physical phenomena in order to perform exquisitely sensitive measurements of magnetic flux: Josephson tunnelling and flux quantisation in superconductors. Flux quantisation was first predicted in the early 1950s [115], and subsequently experimentally verified by two independent groups in 1961 [116,117]. Usage of these two effects resulted in the first instance of a dc SQUID in 1963 [118], demonstration of magnetic-field-dependent critical current later that year [119], and observation of quantum interference effects the following year [120]. In the time since then — almost 6 decades — SQUIDs have set the gold standard for performing sensitive measurements of quantities such as magnetic fields and electric current.

Figure 2.2: Diagram of a conventional dc SQUID, consisting of two superconducting loops connected by two Josephson junctions [120]. Application of a bias current — which exceeds the critical current of the SQUID — results in a magnetic-field-dependent potential difference measured across the SQUID.

Contrary to conventional conductors in which an electrical current is carried by individual electrons, a supercurrent is carried by *pairs* of electrons — called Cooper pairs — which are bound together over distances much greater than the lattice spacing [121]. Each electron within the Cooper pair has the same spin but in opposite directions, yielding a net total spin of zero and therefore obeying Bose-Einstein statistics. At cryogenic temperatures, the Cooper pairs all condense into the same quantum state, and can be described by a collective wavefunction $\psi = \psi_0 \exp(i\varphi)$, where φ is a phase with both temporal and spatial dependence. Given that these Cooper pairs are described by a *single* wavefunction, the phase difference $\Delta\varphi$ along an arbitrary closed path l *must* be an integer of 2π in order to maintain continuity. For a closed path inside a thick superconductor the supercurrent density is zero, and the phase difference along the path can be expressed as [122–124]

$$\Delta\varphi = \frac{4\pi e}{h} \oint_l \boldsymbol{A}_{\mathrm{vp}} \cdot \mathrm{d}\boldsymbol{l} = \frac{4\pi e}{h} \iint_A \boldsymbol{B} \cdot \mathrm{d}\boldsymbol{A} = 2\pi \frac{\Phi_B}{\Phi_0}, \tag{2.6}$$

where $\boldsymbol{A}_{\mathrm{vp}}$ is the vector potential, $\boldsymbol{B}$ is the magnetic field, $\mathrm{d}\boldsymbol{l}$ is an infinitesimal element of the path l, $\mathrm{d}\boldsymbol{A}$ is an infinitesimal element of surface area, Φ_B is the magnetic flux threading the superconducting loop, $\Phi_0 = h/2e$ is the flux quantum, e is the charge of an electron, and h is the Planck constant. Given that continuity demands $\Delta\varphi = 2\pi n$ for $n \in \mathbb{Z}^*$, Eq. (2.6) implies that

$$\Phi_B = n\Phi_0. \tag{2.7}$$

That is, the magnetic flux Φ_B threading a superconducting loop is quantised in integer multiples of the flux quantum Φ_0 [125].

Josephson junctions are comprised of two superconductors that are coupled by a weak, insulating layer [126]. Provided that the insulating barrier is sufficiently thin, the supercon-
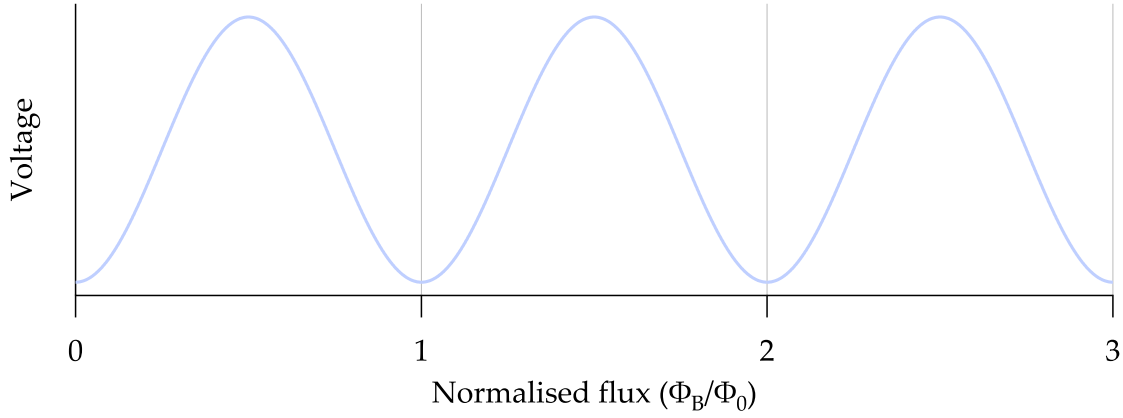
Figure 2.3: Flux-to-voltage transfer function of a dc SQUID with a constant bias current [127].

ducting wavefunctions in both electrodes overlap and allow Cooper pairs to tunnel through the junction despite a lack of applied voltage [126]. This subsequently gives rise to a supercurrent I_{CP} flowing through the junction that can be expressed as

$$I_{CP} = I_C \sin \Delta \varphi \, , \tag{2.8}$$

where I_C is the critical current (i.e. the maximum supercurrent), and $\Delta \varphi = \varphi_1 - \varphi_2$ is the phase difference between the junction [126–128]. If the critical current is exceeded, a dc voltage V_J appears across the Josephson junction, and causes the phase difference to change in time according to [122–124,127]

$$\frac{\partial \Delta \varphi}{\partial t} = \frac{2\pi}{\Phi_0} V_J \, . \tag{2.9}$$

Conventional dc SQUIDs — the design of which was first introduced in 1964 [120] — consist of a superconducting loop with inductance L_S separated by two Josephson junctions (cf. Fig. 2.2). Assuming the physical and magnetic properties of the two loops are *identical*, the critical current of the SQUID is the sum of the critical currents of the two Josephson junctions — I_{C1} and I_{C2} — given by [122]

$$I = I_C \left(\sin \Delta \varphi_1 + \sin \Delta \varphi_2 \right) = 2 I_C \cos \left(\frac{\Delta \varphi_1 - \Delta \varphi_2}{2} \right) \sin \left(\frac{\Delta \varphi_1 + \Delta \varphi_2}{2} \right) , \tag{2.10}$$

where $\Delta \varphi_1$ and $\Delta \varphi_2$ are the phase differences across the two junctions. The difference of these phase differences, i.e. $\Delta \varphi_1 - \Delta \varphi_2$, is dependent upon the magnetic flux in the superconducting loop via [122]

$$\Delta \varphi_1 - \Delta \varphi_2 = 2\pi \frac{\Phi_B}{\Phi_0} = 2\pi \left(\frac{\Phi_{\text{ext}} + L_S \left[I_{C1} - I_{C2} \right]}{\Phi_0} \right) , \tag{2.11}$$

where the total flux Φ_B has been substituted with the sum of the external flux Φ_{ext} due to the

external magnetic field, and the flux induced by a current which circulates within the loop in order to maintain flux quantisation (cf. Eq. 2.7). When the external flux is *already quantised* (i.e. $\Phi_B = n\Phi_0$), no circulating current flows within the superconductor, and the critical current of the SQUID is simply $I = I_{C1} + I_{C2} = 2I_C$. However, in the case when $\Phi_B \neq n\Phi_0$, the critical current of the SQUID is suppressed. Two critical current extrema are therefore observed — when $\Phi_{\text{ext}} = n\Phi_0$ and $\Phi_{\text{ext}} = n\Phi_0/2$ — with a periodic dependence upon Φ_{ext}. In practice however, a dc SQUID is typically biased with a dc current of $I_B \approx 2I_C$ [122], which gives rise to a periodic voltage across the SQUID which is nonlinearly dependent upon the magnetic field[1] [127] (cf. Fig. 2.3). This enables highly sensitive measurements of small magnetic fields to be performed. Note that since SQUIDs measure magnetic *flux*, they are intrinsically *vector* magnetometers.

The magnetic noise of dc SQUIDs is typically dominated by the Johnson noise of the resistors in the surrounding electronics. Under optimal conditions, the amplitude spectral density (cf. App. H.1) of the magnetic noise is given by [132,133]

$$\sqrt{S_B(\nu)} = \frac{4L_S^{3/4}C^{1/4}\sqrt{2k_BT}}{A_{\text{eff}}}, \tag{2.12}$$

where L_S is the inductance of the SQUID, C is the total capacitance of one Josephson junction, k_B is the Boltzmann constant, T is the absolute temperature, and A_{eff} is the effective area [133]. Current state-of-the-art dc SQUIDs have demonstrated white noise floors of as low as $\sqrt{S_B(\nu)} \approx 0.33\,\text{fT}_{\text{rms}}/\sqrt{\text{Hz}}$ above $1\,\text{kHz}$ [133].

Although SQUID magnetometers are *exquisitely* sensitive — being able to measure down to a small fraction of a single flux quantum Φ_0 — and have large frequency response (up to and beyond $1\,\text{GHz}$ in an open-loop configuration), they have three major caveats: they *must* be cryogenically cooled in order to achieve superconductivity, they have a small amplitude response in an open-loop configuration, and they can only perform *relative* magnetic-field measurements. The temperature to which the superconducting material must be reduced is dependent upon the material, but is commonly between $4\,\text{K}$ and $80\,\text{K}$ (or $-269.2\,°\text{C}$ to $-193.2\,°\text{C}$, respectively). This necessitates the use of cryogenics such as liquid helium or liquid nitrogen, which hampers the applicability of these devices — especially in use cases which demand either portability or room-temperature operability.

Due to the periodic and nonlinear flux-to-voltage transfer function of an open-loop dc SQUID (cf. Fig. 2.3), two limitations immediately arise from this: the linear dynamic range is limited to less than $\Phi_0/2$ [131] (on the order of $1\,\text{nT}$ in typical high-performance devices [133]) and, since each flux quantum state is indistinguishable from the next, magnetic-field measurements are *relative* rather than *absolute*. Fortunately in the former case, the dynamic range of a dc SQUID can be dramatically increased — to several thousand flux quanta [131] — by linearising the flux-to-voltage transfer function through the use of a flux-locked loop [131]; however, due to the electronics required to operate in this manner, the frequency-

---

[1]In practical devices, the transfer function of the SQUID is linearised using a flux-locked-loop feedback circuit. However, this is outside the scope of this thesis and will not be discussed in detail here — see Refs. [129], [130] and [131] for further discussion.

response bandwidth is reduced to $\mathcal{O}\left(1\,\mathrm{MHz}\right)$ [131].

The aforementioned caveats associated with dc SQUIDs — the need for cryogenics to achieve superconductivity, requiring complicated readout electronics to linearise the periodic flux-to-voltage transfer function, and the lack of an absolute measurement — herald the use of optical techniques discussed in the following section. These techniques can achieve similar — or in some cases, *superior* — performance to that of dc SQUIDs, without the need for cryogenics, and all while yielding *absolute*, all-optical magnetic-field measurements.

2.2 Optical Measurement Techniques

As mentioned in Sec. 2.1.2, SQUIDs reigned, for a number of decades, as the most sensitive devices for the measurement of magnetic fields. However, the invention of the laser in 1960 [134] saw an influx of optical measurement techniques which — although initially had poor sensitivity relative to their SQUID counterparts — have incrementally caught up with, rivaled, and even *surpassed*, the absolute magnetic sensitivity of modern SQUIDs. This is by and large due to advancements in anti-relaxation coatings for vapour cells, as well as the development of new optical measurement techniques.

There are a vast number of optical magnetometry techniques, each with their own intrinsic advantages and disadvantages. This section aims to discuss a number of the more commonly used optical techniques, their advantages and disadvantages, and the performance of state-of-the-art devices found in the literature.

2.2.1 Free-Induction Decay Magnetometers

Free-induction decay (FID) magnetometers are one of the oldest style of optical technique, with instances of these devices dating as far back as 1969 [135]. The first instance of FID was demonstrated in Ref. [135], in which ^{87}Rb vapour was optically pumped by circularly polarised light to produce atomic orientation, and subsequently used to measure the magnetic field produced by the nucleus of ^{3}He atoms undergoing FID. In doing so, a nuclear relaxation time of 2 hours and 20 minutes was observed [135]. Modern FID magnetometers on the other hand typically use single-species atomic vapour, with pump and/or probe light tuned to electronic resonances, though recent measurements of nuclear-polarised ^{3}He have been performed with excellent precision [136]. Although the transverse spin-relaxation time of alkali electronic spins is significantly shorter than that of the ^{3}He nuclear spin, the ground-state gyromagnetic ratio — upon which the magnetic sensitivity also depends — is typically many orders of magnitude larger than that of the ^{3}He nucleus.

The principle of an FID magnetometer is quite simple: generate ground-state atomic polarisation via optical pumping (using either linearly polarised or circularly polarised light) until the system has reached steady state, cease optical pumping by switching off or blocking the pump beam, and subsequently measure the temporal evolution of the ground-state polarisation — via its effect on a weak probe beam — as it undergoes FID (cf. Fig. 2.4). The frequency of the oscillating probe signal is proportional to the Larmor frequency (or a

Figure 2.4: Diagram depicting a free-induction decay measurement of a magnetic field. Sinusoidal regression to the FID yields an estimate of the time-averaged Larmor frequency, which can be used to determine a dc magnetic field using a single pulse (left), or an ac magnetic field using a train of pulses (right).

harmonic thereof), and hence sinusoidal regression to the FID signal enables very precise measurement of the *time-averaged* dc magnetic field over the decay time. Vector FID magnetometers using optically pumped Cs atoms have demonstrated an Allan deviation of $40-80\,\mathrm{fT}$ in a $1\,\mu\mathrm{T}$ field, integrated over a few seconds [137,138]. Furthermore, FID magnetometers using Cs atoms have demonstrated $200\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mu\mathrm{T}$ [139], and $16\,\mathrm{pT_{rms}}/\sqrt{\mathrm{Hz}}$ in a $1.5\,\mathrm{mm}$-thick microfabricated vapour cell [140]. In cold atomic vapours, all three components of the magnetic field, as well as one gradient component, have been measured in a single decay with a spatial resolution of about $20\,\mu\mathrm{m}$ [141].

Recent work has seen FID magnetometry extended to the measurement of ac magnetic fields by performing a train of pulses, each of which yields a time-averaged dc magnetic field, and subsequently mapping out the time-dependent magnetic field (cf. Fig. 2.4). This technique works sufficiently well provided that the magnetic-field frequency is significantly less than the pulse repetition rate; however, magnetic fields with frequencies above this are aliased. In a spinor Bose-Einstein condensate, pulsed FID has enabled time-varying magnetic fields to be measured with $100\,\mathrm{pT}$ accuracy every $5\,\mathrm{ms}$, corresponding to a Nyquist-limited bandwidth of $100\,\mathrm{Hz}$ [142]. In warm atomic vapours, an FID pulse train with a repetition rate of $1\,\mathrm{kHz}$ was used to track magnetic fields oscillating orthogonal to the dc field at frequencies up to $100\,\mathrm{Hz}$ [143], and parallel to the dc field with a bandwidth up to $1\,\mathrm{kHz}$ [144]. Free-induction decay has also been applied to a spin-exchange relaxation-free magnetic gradiometer, in which a pulse train yielded a quantum-noise-limited performance of $14\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $50\,\mu\mathrm{T}$ [145].

2.2.2 M_x and M_z **Magnetometers**

M_x and M_z magnetometers date back as far as FID magnetometers, with initial measurements of magnetic resonances using their effect on the absorption of a transverse light beam dating back to 1957 [146–148].

As their name may suggest, M_x and M_z magnetometers measure the x or z components of the magnetic moment vector $\boldsymbol{\mu} = (\mu_x, \mu_y, \mu_z)$, respectively. For alkali atoms situated within an ambient magnetic field $\boldsymbol{B}$, the atoms will undergo Larmor precession due to their magnetic moment, with an angular frequency given by

$$\Omega_{\mathrm{L}} = \gamma_{\mathrm{g}} |\boldsymbol{B}| , \tag{2.13}$$

where γ_{g} is the ground-state gyromagnetic ratio. Therefore, given a measurement of Ω_{L}, one can infer the value of the *magnitude* of the magnetic field.

Optical detection of the Larmor frequency is achieved by using a laser beam to probe the magnetic moment of an atomic ensemble. The evolution of the magnetic moment, with a relaxation time $\boldsymbol{T}_{\mathrm{rel}} = (T_2, T_2, T_1)$, is described by the Bloch equation [149]

$$\frac{\mathrm{d}\boldsymbol{\mu}}{\mathrm{d}t} = \gamma_{\mathrm{g}} \boldsymbol{\mu} \times \boldsymbol{B} - \frac{\boldsymbol{\mu}}{\boldsymbol{T}_{\mathrm{rel}}} , \tag{2.14}$$

which, for a Cartesian coordinate system, yields the following coupled differential equations for each component [147,149]:

$$\frac{\mathrm{d}\mu_x}{\mathrm{d}t} = \gamma_{\mathrm{g}} \left(\mu_y B_z - \mu_z B_y \right) - \frac{\mu_x}{T_2} , \tag{2.15}$$

$$\frac{\mathrm{d}\mu_y}{\mathrm{d}t} = \gamma_{\mathrm{g}} \left(\mu_z B_x - \mu_x B_z \right) - \frac{\mu_y}{T_2} , \tag{2.16}$$

$$\frac{\mathrm{d}\mu_z}{\mathrm{d}t} = \gamma_{\mathrm{g}} \left(\mu_x B_y - \mu_y B_x \right) - \frac{\mu_z - \mu_{\mathrm{in}}}{T_1} , \tag{2.17}$$

$$\tag{2.18}$$

where μ_{in} is the *initial* magnetic moment at time $t = 0$ along the z-axis, generated by optical pumping with circularly polarised light. In the case of a static field B_0 along the z-axis, and a radio-frequency (rf) field with an amplitude B_1 applied in the transverse plane, the total magnetic-field vector is given by $\boldsymbol{B} = (B_1 \cos(\omega t), -B_1 \sin(\omega t), B_0)$ and hence the Bloch equations for each magnetisation component become:

$$\frac{\mathrm{d}\mu_x}{\mathrm{d}t} = \gamma_{\mathrm{g}} \left[\mu_y B_0 + \mu_z B_1 \sin(\omega t) \right] - \frac{\mu_x}{T_2} , \tag{2.19}$$

$$\frac{\mathrm{d}\mu_y}{\mathrm{d}t} = \gamma_{\mathrm{g}} \left[\mu_z B_1 \cos(\omega t) - \mu_x B_0 \right] - \frac{\mu_y}{T_2} , \tag{2.20}$$

$$\frac{\mathrm{d}\mu_z}{\mathrm{d}t} = -\gamma_{\mathrm{g}} \left[\mu_x B_1 \sin(\omega t) + \mu_y B_1 \cos(\omega t) \right] - \frac{\mu_z - \mu_{\mathrm{in}}}{T_1} . \tag{2.21}$$

It can be seen in Eqs. (2.19) – (2.21) that the magnetisation vector precesses at the frequency ω in the laboratory frame. It is therefore convenient to work in a rotating frame, in which the magnetisation vector at equilibrium appears stationary. In order to transform from the laboratory frame to the rotating frame, the transverse components must be transformed from $\mu_x \rightarrow \mu_x{}'$ and $\mu_y \rightarrow \mu_y{}'$ via

$$\mu_x{}' = \mu_x \cos(\omega t) - \mu_y \sin(\omega t) \, , \tag{2.22}$$

$$\mu_y{}' = \mu_x \sin(\omega t) + \mu_y \cos(\omega t) \, . \tag{2.23}$$

By differentiating Eqs. (2.22) and (2.23), and substituting Eqs. (2.19) – (2.21) into the result, one obtains the rotating-frame differential equations for the magnetisation components:

$$\frac{\mathrm{d}\mu_x{}'}{\mathrm{d}t} = \Delta\omega\mu_y{}' - \frac{\mu_x{}'}{T_2} \, , \tag{2.24}$$

$$\frac{\mathrm{d}\mu_y{}'}{\mathrm{d}t} = \omega_1\mu_z - \Delta\omega\mu_x{}' - \frac{\mu_y{}'}{T_2} \, , \tag{2.25}$$

$$\frac{\mathrm{d}\mu_z}{\mathrm{d}t} = -\omega_1\mu_y{}' - \frac{(\mu_z - \mu_{\mathrm{in}})}{T_1} \, , \tag{2.26}$$

where $\Delta\omega = \omega_0 - \omega = \gamma_{\mathrm{g}}B_0 - \omega$ is the detuning and $\omega_1 = \gamma_{\mathrm{g}}B_1$. Solving the rotating-frame differential equations for the steady-state behaviour yields the following solutions [149,150]:

$$\mu_x{}' = \mu_{\mathrm{in}} \frac{\Delta\omega T_2{}^2\omega_1}{1 + (\Delta\omega T_2)^2 + \omega_1^2 T_1 T_2} \, , \tag{2.27}$$

$$\mu_y{}' = \mu_{\mathrm{in}} \frac{T_2\omega_1}{1 + (\Delta\omega T_2)^2 + \omega_1^2 T_1 T_2} \, , \tag{2.28}$$

$$\mu_z = \mu_{\mathrm{in}} \frac{1 + (\Delta\omega T_2)^2}{1 + (\Delta\omega T_2)^2 + \omega_1^2 T_1 T_2} \, . \tag{2.29}$$

M_z magnetometers generally consist of a single, circularly polarised beam. By applying a resonant rf field in the transverse plane, the population difference induced by the optical pumping is reduced. Modulation of the rf-field frequency therefore results in changes in the absorption of the pumping light. By tuning the mean frequency of the rf field to the frequency of the resonance line centre, one is able to infer the Larmor frequency and hence the local magnetic-field strength.

M_x magnetometers work slightly differently in that, rather than measuring the longitudinal magnetisation, the transverse magnetisation is observed instead. After optically pumping with circularly polarised light and generating atomic orientation along the dc magnetic field, a transverse magnetic field oscillating with a frequency $\omega \approx \gamma_{\mathrm{g}}B_0 = \Omega_{\mathrm{L}}$ is applied to the atoms. This results in a reduction of longitudinal magnetisation, as well as the appearance of a transverse component of magnetisation which precesses with angular

frequency ω. Subsequent probing by a circularly polarised beam directed orthogonally to the pump beam, results in modulated absorption due to periodic changes in population of the ground-state Zeeman sublevels (given a quantisation axis transverse to the dc field). As one might expect, the frequency of the periodic change in absorption is simply the Larmor frequency.

The main advantage of using an M_z magnetometer is that, unlike many other optical magnetometers, the centre of the resonance does not depend on the phase of the signal under observation [150]. Any errors in demodulation phase simply result in a homogeneous reduction in resonance amplitude, rather than a distortion of the line shape [150]. This results in a highly *accurate* measurement of the Larmor frequency and hence the magnetic-field strength — useful in applications which demand an accurate estimate of the field, rather than *purely* excellent precision.

In the literature, M_x magnetometers have demonstrated measured field fluctuations of 240 fT for integration times of 0.7 s, with a projected intrinsic noise of 29 fT in a 1-Hz bandwidth due to optical power instability, and 10 fT in a 1-Hz bandwidth due to photon shot noise [151]. In a single-beam configuration — where the beam is at $45°$ to the dc magnetic field — M_x magnetometers have demonstated a projected shot-noise-limited floor of $600\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ [152] at low fields. Calculated spin-projection-limited sensitivities of $1.8\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ – $12\,\mathrm{pT_{rms}}/\sqrt{\mathrm{Hz}}$ have been demonstrated [153–159], corresponding to volume-normalised sensitivities (cf. Sec. 6.3.3) of $3.9\,\mathrm{fT\,cm^{3/2}}/\sqrt{\mathrm{Hz}}$ – $537\,\mathrm{fT\,cm^{3/2}}/\sqrt{\mathrm{Hz}}$ [153–158].

M_z magnetometers on the other hand have recently demonstrated a projected photon shot-noise-limited floor of about $10\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at Earth's magnetic field [160]. This was achieved by using two 4-mm diameter vapour cells, each heated to $130°$ to increase vapour density.

2.2.3 Spin-Exchange Relaxation-Free

Spin-exchange collisions are one of the fundamental performance limitations of an optical magnetometer. As discussed in further detail in Sec. 5.2.2.1, spin-exchange collisions reduce the transverse spin-relaxation time T_2 by destroying the coherence between ground-state Zeeman sublevels. This leads to a broader magnetic resonance, and therefore reduced sensitivity. As the density of the atomic vapour is increased, the rate of spin-exchange collisions increases linearly. Since the number density of an atomic vapour is related its temperature, spin-exchange collisions often become a dominant performance limitation above room temperature.

However, it was discovered that if the number density was increased sufficiently high, and the ambient magnetic field reduced to near zero, magnetic resonance lines would — rather counterintuitively — *decrease* in width [161]. This effect was later analysed using the density matrix formalism, indicating that for low atomic polarisation, and in the limit of high spin-exchange collision rates in zero field, there would be *zero* broadening [162]. In subsequent work, it was shown that spin-exchange collisions could be reduced at high po-

larisation, too [163]. Significantly reduced resonance width at high vapour densities paved the way for ultrasensitive atomic magnetometers — so-called spin-exchange relaxation-free (SERF) magnetometers — which were first demonstrated in Ref. [164].

Spin-exchange relaxation-free magnetometers rely on measuring the Faraday effect — that is, the rotation of the plane of polarisation of linearly polarised light as it traverses the atomic vapour. By pumping the vapour with circularly polarised light, the atomic spins are oriented along the propagation direction of the light. An applied magnetic field, orthogonal to the pump, causes the spins to precess in the plane of the pump beam. Transmission of a weak, linearly polarised probe beam — which is orthogonal to both the pump beam and the magnetic field direction — measures the projection of the spins along its propagation trajectory [150]. The resulting polarisation rotation is sinusoidal, with a frequency given by the Larmor frequency. Subsequent measurement of this frequency enables the local magnetic-field strength to be inferred with both great accuracy and precision.

Spin-exchange relaxation-free magnetometers are typically operated at atomic densities on the order of $10^{20} \, \text{m}^{-3}$, which requires the atomic vapour to be heated in the range of 100 – 180°C [165,166]. Typical anti-relaxation coatings comprised of paraffin — with a melting point on the order of 60°C — are irreversibly damaged at these temperatures, requiring buffer gas in order to suppress spin-destruction collisions with the cell walls (cf. Sec. 5.2.2.4).

In the absence of technical noise, the sensitivity of a SERF magnetometer is fundamentally limited by the two sources common to all optical magnetometers: spin-projection noise, and photon shot noise. Mathematically, the magnetic noise per unit bandwidth for a SERF magnetometer that is limited by spin-projection noise and photon shot noise, is given by [166,167]

$$\delta B_{\text{SERF}} = \frac{\hbar}{g_s \mu_B P_z \sqrt{n_v V_m}} \sqrt{2 \left(\Gamma_{\text{pump}} + \Gamma_{\text{pr}} + \Gamma_{\text{SD}} \right) + \frac{4 \left(\Gamma_{\text{pump}} + \Gamma_{\text{pr}} + \Gamma_{\text{SD}} \right)^2}{\Gamma_{\text{pr}} \left(\text{OD} \right)_0}}, \qquad (2.30)$$

where $\hbar$ is the reduced Planck constant, g_s is the electron spin g-factor, μ_B is the Bohr magneton, P_z is the electron spin polarisation along the z axis, n_v is the atomic number density, V_m is the measurement volume, Γ_{pump} is the optical pumping rate due to the pump beam, Γ_{pr} is the rate of depolarisation due to the probe beam, Γ_{SD} is the electron spin-destruction rate, and $(\text{OD})_0$ is the on-resonance optical depth.

Spin-exchange relaxation-free magnetometers frequently demonstrate magnetic sensitivities on the order of $10 \, \text{fT}_{\text{rms}} / \sqrt{\text{Hz}}$ or less [45,48,49,60,100,167–176], with sensitivities below $100 \, \text{fT}_{\text{rms}} / \sqrt{\text{Hz}}$ demonstrated in millimetre-scale microfabricated vapour cells [176,177]. To date, the highest demonstrated sensitivity — limited by magnetic-field noise within the shielded volume — is $160 \, \text{aT}_{\text{rms}} / \sqrt{\text{Hz}}$ in a measurement volume of $0.45 \, \text{cm}^3$, corresponding to a volume-normalised sensitivity of $107 \, \text{aT} \, \text{cm}^{3/2} / \sqrt{\text{Hz}}$ and energy resolution per unit time of $44\hbar$ [168] (cf. Secs. 6.3.3 and 6.3.4, respectively). In the absence of probe beam technical noise, the projected photon shot-noise limit was $50 \, \text{aT}_{\text{rms}} / \sqrt{\text{Hz}}$ [168]. Spin-exchange relaxation-free magnetometers have also been demonstrated to work in unshielded envi-

ronments — albeit with significantly reduced sensitivity — provided that appropriate field-cancellation coils are implemented to ensure sufficiently low ambient magnetic-field strength within the measurement volume [178].

2.2.4 Nitrogen-Vacancy Centres in Diamond

Diamonds with nitrogen-vacancy (NV) colour centres enable magnetometry to be performed with exceptionally high spatial resolution. The NV centre is a point defect within the crystal lattice of diamond, consisting of a substitutional nitrogen atom and a neighbouring vacancy (i.e. a missing carbon atom) [179]. The NV centre can either have negative (NV^-), positive (NV^+), or neutral (NV^0) charge states — corresponding to the electron distribution within the NV centre — with negative charge states typically being used for magnetometry and other applications [179–181]. The negative charge state corresponds specifically to five of the electrons being contributed from the three neighbouring carbon atoms and the nitrogen atom, with an extra electron captured from an electron donor [179].

The NV centre has spatial symmetry since its structure is symmetric with respect to rotations of 0, $2\pi/3$ and $4\pi/3$ around the NV axis, and to reflections in the mirror planes defined by the NV axis and one of the three neighbouring carbon atoms [179]. This gives rise to energy levels of NV^- which can be constructed using linear combinations of atomic orbitals [179,182,183]. Each electronic configuration can contain spin-triplet states ($^3A_2$ and ^3E), as well as spin-singlet states ($^1A_1$ and ^1E) [179,182], with the spin-triplet states each possessing three sublevels with magnetic quantum numbers $m_S = 0, \pm 1$ [179] and where the quantisation axis is set by the NV axis — the line connecting the nitrogen atom and the vacancy [179]. An energy-level diagram of the spin-triplet and spin-singlet states is shown in Fig. 2.5, along with a diagram of the NV centre defect in the diamond crystal lattice.

The optical transition between $^3A_2$ and ^3E has a wavelength of $637\,\text{nm}$, while the transition between $^1A_1$ and ^1E has a wavelength of $\lambda = 1042\,\text{nm}$. Both of these transitions have phonon sidebands due to vibrations in the diamond lattice, which broadens the NV absorption and fluorescence spectra by hundreds of nanometres [179]. The $^3A_2 \rightarrow {}^3E$ transition can therefore be excited with light within a wavelength of $\lambda \in [450, 637]\,\text{nm}$ [179], with a fluorescence from the $^1E \rightarrow {}^1A_1$ decay containing wavelengths $\lambda \in [637, 800]\,\text{nm}$ [179].

By optically pumping the NV centre on the $^3A_2 \rightarrow {}^3E$ transition, the ground-state population preferentially accumulates in the $m_S = 0$ sublevel after several excitation cycles, due to the probability of population originating in $m_S = \pm 1$ states undergoing decay via non-radiative transitions [179]. Once the NV centre is optically pumped into the $m_S = 0$ sublevel, and a high level of fluorescence is detected, microwave radiation is applied to the NV centre. When the microwave frequency is tuned such that it is on resonance with one of the $m_S = 0 \leftrightarrow \pm 1$ transitions, the amount of fluorescence *decreases* since population is pushed from $m_S = 0$ to $m_S = \pm 1$. By scanning the microwave frequency across the $m_S = 0 \leftrightarrow \pm 1$ transitions, two dips in the fluorescence spectrum are obtained. In the limit where the component of the magnetic field transverse to the NV symmetry axis is negligible, the frequencies of the $m_S = 0 \rightarrow \pm 1$ microwave transitions are given by [184,185]

Figure 2.5: Left: Diagram of the nitrogen-vacancy centre and diamond lattice [186]. The vacancy V (grey) is joined to three nearest-neighbour carbon atoms C (green), and a substitutional nitrogen atom N (red). Next-to-nearest carbon atoms (black) are also shown. Right: Energy-level diagram of the NV centre [179], showing the transition wavelengths for the $^3A_2 \rightarrow \, ^3E$ (spin-triplet) and $^1E \rightarrow \, ^1A_1$ (spin-singlet) transitions. Dashed black arrows indicate intersystem crossing between the spin-triplet and spin-singlet states [179].

$$\nu_{0\pm} \approx \Delta_0 + \beta_T \delta_T \pm \left(\frac{\gamma_g}{2\pi} \right) B_{NV}, \tag{2.31}$$

where $\Delta_0 \approx 2.87\,\text{GHz}$ is the zero-field splitting at room temperature [186], $\beta_T \approx -74\,\text{kHz/K}$ near room temperature [187], δ_T is the temperature offset from $300\,\text{K}$ during the measurement, $\gamma_g/2\pi \approx 28\,\text{GHz/T}$ is the NV gyromagnetic ratio [185], and B_{NV} is the projection of the applied magnetic field along the NV symmetry axis. The frequency separation of the fluorescence dips correspond to the separation of the $m_S = \pm 1$ ground states, which using Eq. (2.31), is simply given by $\Delta\nu_{0\pm} = \nu_{0+} - \nu_{0-} = 2\gamma_g B_{NV}/(2\pi)$. By accurately and precisely measuring the frequency separation between the fluorescence dips, one can determine the local magnetic-field strength B_{NV}.

The ability to precisely determine the magnetic-field strength will depend on two quantities: contrast and linewidth. Contrast $\mathcal{C}$ is the fractional difference between the fluorescence obtained on- and off-resonance, with $\mathcal{C} = 0.20$ typical of single NV centres [179]. The linewidth is related to the transverse spin-relaxation time T_2, with typical values of $T_2 \approx 100\,\text{ns}$ for nitrogen-rich NV centres, and a few μs for NV centres with low nitrogen concentration [179].

The above discussion was couched in terms of a single NV centre, with the magnetic field aligned along the quantisation axis defined by the NV axis. If instead one has an ensemble of NV centres and an arbitrarily oriented magnetic field, there may be in total eight magnetic resonances due to the four possible alignments of the NV centre — though some of these may be degenerate [179]. It is possible to select *one* NV alignment by applying a bias field along the NV axis so that changes in the magnetic field projection along this axis affect the

resonance frequencies approximately linearly [179]. Alternatively, by interrogating all four NV alignments — and subsequently yielding eight resonances — one can obtain full vector information about the magnetic field [179,188].

As with other optical magnetometry techniques described in previous sections, the magnetic field sensitivity of NV centres is fundamentally limited by spin-projection noise and photon shot noise. However, given the experimental difficulty associated with capturing fluorescence originating within a material with high refractive index ($n_\mathrm{d} \approx 2.419$ [179]), photon shot noise often limits the magnetic-field sensitivity. The highest sensitivity demonstrated with an NV ensemble is $15\,\mathrm{pT_{rms}}/\sqrt{\mathrm{Hz}}$ [73], with a projected standard quantum limit of about $10\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ [73]. Given a measurement volume of $(13 \times 200 \times 2000)\,\mathrm{\mu m^3}$ [73], the shot-noise limit and spin-projection-noise limit correspond to spatiotemporal resolutions of $34.2\,\mathrm{fT\,cm^{3/2}}/\sqrt{\mathrm{Hz}}$ and $22\,\mathrm{aT\,cm^{3/2}}/\sqrt{\mathrm{Hz}}$, respectively [73].

Given their extraordinary spatiotemporal resolution, NV centres have found application in: eddy-current imaging [189], biomedical imaging [73,190,191], high-resolution thermometry [192,193], nanoscale magnetic sensing [194], and general magnetic imaging [195–198].

2.2.5 Nonlinear Magneto-Optical Rotation

Nonlinear magneto-optical rotation (NMOR) is the nonlinear contribution to the overall magneto-optical rotation (Faraday) signal — that is, a nonlinear rotation of the plane of polarisation of linearly polarised light as it traverses a medium situated within a magnetic field (cf. Fig. 2.6). It yields signals which have a nonlinear dependence upon both the light intensity and magnetic-field strength, with the latter dependence enabling precision magnetometry to be performed [150,199–203]. The physical processes which give rise to NMOR are described in great detail in Secs. 3.5–3.7; however, it should be noted here that the main advantages of NMOR magnetometry are: technical simplicity, high accuracy, and wide dynamic range [199,204,205].

As a strictly optical technique, NMOR magnetometry is fundamentally limited by both photon shot noise and spin-projection noise. To date, the highest shot-noise-limited sensitivity demonstrated with an NMOR magnetometer is $15\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ [206], with a spin-projection-noise limit of $3.1\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ [206]. There are also numerous examples of NMOR magnetometers achieving sensitivities on the order of $500\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ or less [202,204,207].

Nonlinear magneto-optical rotation magnetometers are intrinsically *scalar* magnetometers, as they rely on measuring the Larmor frequency. However, NMOR magnetometers *can* operate as vector magnetometers by modulating the magnetic field along orthogonal axes [208].

2.2.6 Radio-Frequency Optical Atomic Magnetometers

Until now, all of the aforementioned magnetometers — even the non-optical techniques — have been dc magnetometers. That is, they primarily measure the value of the dc magnetic field, with a finite frequency response to time-varying fields that is governed by a relatively
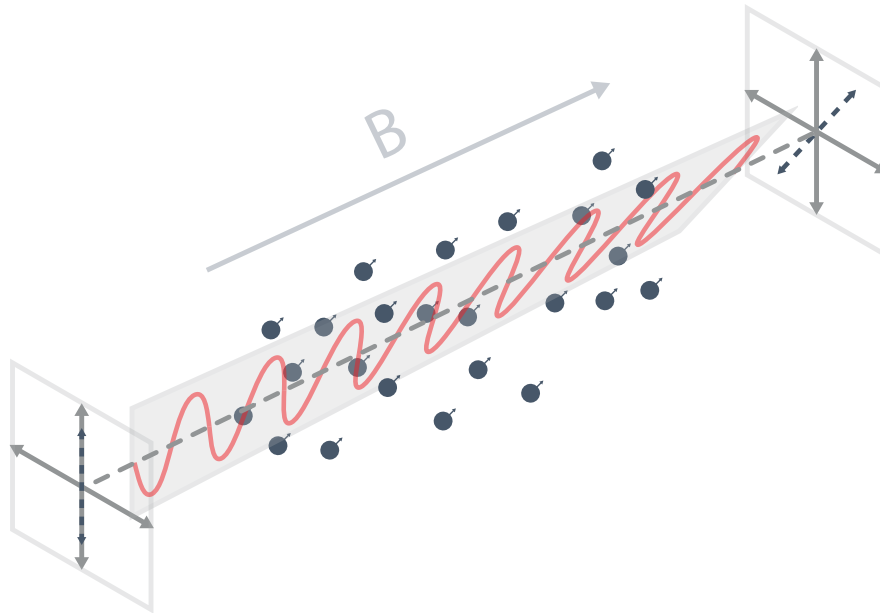
Figure 2.6: Diagram of linearly polarised light undergoing polarisation rotation due to circular birefringence in an atomic vapour — nonlinear magneto-optical rotation.

small bandwidth. This bandwidth — at least in the context of the optical-resonance-based techniques — is typically *inversely related* to the transverse spin-relaxation time T_2, meaning that high sensitivity devices yield narrow bandwidths.

One pathway to ac magnetometry is via so-called radio-frequency optical atomic magnetometers (rf OAMs) [85,88–94,209–217]. These devices operate by producing atomic orientation (i.e. longitudinal magnetisation) along a static magnetic field using circularly polarised light. Once the longitudinal polarisation has been established, a weak oscillating rf field is applied transverse to the dc field. Provided that the rf frequency is approximately equal to the Larmor frequency of the static field (i.e. provided that the rf field drives Zeeman transitions), the application of this rf field tilts the magnetisation away from the quantisation axis defined by the static field, causing the magnetisation to precess around the static field in accordance with Eq. (2.14). The transverse spin component — which precesses about the dc magnetic field at the Larmor frequency — is then monitored via the optical rotation of a weak, linearly polarised probe beam that travels both orthogonal to the dc magnetic field, and the rf field.

These devices are suitable when the signal is relatively narrowband, about a frequency which is known *a priori*, and provided that the dc magnetic-field strength can be tuned to produce a Larmor frequency that is sufficiently close to the rf-field frequency of interest. These are quite restrictive requirements for a measurement device, which impedes their use in many applications. Furthermore, in contrast to their dc counterparts, rf OAMs are *not* calibration-free, and must be amplitude calibrated by artificially broadening the magnetic resonance with increasing rf power [212].

The sensitivity of these devices is typically on the same order as their dc counterparts, with many exhibiting noise floors of less than about $100\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ [85,209,213,216], with

some even reaching below $1\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ [211,214,215,217]. Cold-atom rf OAMs have also been demonstrated, yielding increased spatial resolution at the expense of sensitivity [210].

2.2.7 Summary

As mentioned previously, there are a vast number of magnetometry techniques, each of which possess their own characteristic advantages and disadvantages. There is hence no *best* magnetometer *per se*, but rather, there will exist a magnetometer *which best suits the application*. Choosing the appropriate magnetometer is therefore a matter of balancing pros and cons in consultation with the application criteria. A summary of the techniques discussed in this chapter — including their sensitivity, advantages and disadvantages — is presented in Table. 2.1.

The magnetometry technique chosen as the research topic for this thesis was nonlinear magneto-optical rotation, as it offered the best trade-off between pros and cons for the applications which motivated this research. Of particular interest was magnetic anomaly detection, which demands high sensitivity at ambient temperature and at relatively high ambient field strength — typically Earth's field or above. Given that spatial resolution — the primary drawback of NMOR magnetometers — is generally not of immediate concern in the context of magnetic anomaly detection, this magnetometry technique was found to be the best candidate upon which to base this research.

Technique	Sensitivity	Advantages	Disadvantages
Fluxgate	2,000	Large dynamic range, room-temperature operability	Low bandwidth, low sensitivity, requires calibration
SQUID	0.33	Large dynamic range, highest bandwidth	Requires cryogenics, expensive, requires calibration, can only perform relative measurements
FID	200	Large dynamic range, high sensitivity, high accuracy	Low bandwidth, large measurement dead time
M_z	10	Large dynamic range, high sensitivity, high accuracy, single-beam setup	Low bandwidth, requires applied rf field, generally requires heating
M_x	245	Large dynamic range, high sensitivity	Generally requires two orthogonal beams, requires applied rf field, generally requires heating
SERF	0.16	Highest sensitivity, high spatial resolution	Low dynamic range, low bandwidth, requires heating
NV Centres	15,000	Highest spatial resolution	Low sensitivity
rf OAMs	1	High sensitivity, can measure rapidly varying ac fields	Requires calibration, low bandwidth, requires *a priori* knowledge of field frequency, requires control over background dc field
NMOR	15	Large dynamic range, high sensitivity, room-temperature operability, can operate in a single-beam setup, technical simplicity	Low spatial resolution

Table 2.1: Magnetometry techniques, their sensitivity, and their advantages and disavantages. Sensitivity values correspond to state-of-the-art *absolute* sensitivity (cf. Sec. 6.3.1) found in the literature, and are listed in units of $fT_{rms}/\sqrt{Hz}$.

Part II

Background and Theory

Introduction

3.1 Rubidium

At the heart of many optical magnetometers lies an atomic vapour — typically an alkali metal such as potassium, caesium, or rubidium[1] — which acts as a mediator between the interaction of the light and the magnetic field. The primary objective of the work detailed within this thesis was to investigate nonlinear magneto-optical rotation in *rubidium*, and utilise its extreme sensitivity to external magnetic fields in order to perform sensitive magnetometry. This section briefly summarises the important characteristics of rubidium — both physical and quantum mechanical — which form the basic foundation required to grasp the concepts introduced throughout this body of work.

3.1.1 Energy-Level Structure and Spectroscopic Notation

Rubidium has 32 isotopes, only two of which occur naturally: ^{85}Rb (72.2% abundance) and the slightly radioactive ^{87}Rb (27.8% abundance)[2]. Due to the high atomic number of Rb ($Z = 37$), its energy-level structure is relatively complex; and, in order to completely describe the electronic configuration of an atom, each of the electronic states must be defined.

3.1.1.1 Gross Structure

The gross structure of the atomic energy levels is characterised by the *principal quantum number*, $n \in \mathbb{Z}^*$. This quantum number loosely describes both the *size* and *energy* of an electronic shell[3]. Within each electron shell, there are n sub-shells which are represented by an additional quantum number, $L \in \mathbb{Z}^*$. The quantum number L is called the *orbital quantum number*[4], and takes integer values between $0 \leq L \leq n - 1$. For historical reasons, however, the orbital quantum numbers are assigned letters for each value. Letters for the first ten states are presented in Table 3.1.

---

[1]By definition, an optical magnetometer is a magnetometer in which light is used as part of its measurement. This is typically performed using alkali atoms; however, there are also optical magnetometers which use helium vapour [218,219], as well as those which utilise nitrogen-vacancy centres in diamond [179–181,185].

[2]Rubidium-87 decays via β^- decay to strontium-87 with a half-life of 4.976×10^{10} years [220] via:

$$^{87}\text{Rb} \rightarrow \,^{87}\text{Sr} + \text{e}^- + \bar{\nu}_\text{e} \,, \tag{3.1}$$

where e^- denotes an electron, and $\bar{\nu}_\text{e}$ denotes an electron antineutrino.

[3]Two relevant examples of this are the wavefunctions obtained when solving the Schrödinger equation for the infinite potential well Hamiltonian, as well as that of the hydrogen atom.

[4]In the same way that the principal quantum number describes the energy and size of the electronic shell, the orbital quantum number describes the shape of the sub-shell's wavefunction.

L	0	1	2	3	4	5	6	7	8	9	...
	s	p	d	f	g	h	i	k	l	m	...

Table 3.1: Letters assigned to the orbital quantum numbers.

3.1.1.2 Fine Structure

Within each sub-shell there exist a number of orbitals which correspond to various *orientations* of the electron's wavefunction around the nucleus. These orientations are known as 'projection states', and are characterised by the *magnetic quantum number*, m_L. For a given orbital quantum number, there are $2L + 1$ projection states which take integer values between $-L \leq m_L \leq L$. In each of these sub-shell orbitals, there can exist two electrons in accordance with the Pauli exclusion principle[5].

Using all of the above quantum numbers, the full electronic configuration for an atom in the ground state can be written in the form nL^e, where e is the number of electrons within the corresponding sub-shell. For a rubidium atom in the ground state, the complete energy-level configuration is given by $1s^2 2s^2 2p^6 3s^2 3p^6 3d^{10} 4s^2 4p^6 5s^1$. This notation is cumbersome; however, there is only one electron in the valence shell and hence the strongest electromagnetic interaction occurs with this unpaired electron. This enables the energy-level structure to be greatly simplified by considering *only* the unpaired valence electron.

The spectroscopic notation for the valence electron is given by $n^{2S+1}L_J$, where S is the spin quantum number, and $J = L + S$ is the *total electronic angular momentum quantum number*. The ground state of the valence electron in rubidium has quantum numbers $n = 5$, $S = 1/2$, $L = 0$ and $J = 1/2$. In spectroscopic notation, this is written as $5^2 S_{1/2}$. The origin of the total electronic angular momentum quantum number J is the interaction between the spin of the electron, S, and its orbital angular momentum, L. This interaction is known as the spin-orbit interaction, and is described by the following Hamiltonian [221–228]:

$$\mathcal{H}_{\text{S-O}} = \frac{\mu_{\text{B}} (g_s - 1)}{\hbar m_e e c^2 r} \frac{\partial V(r)}{\partial r} \mathbf{L} \cdot \mathbf{S}, \tag{3.2}$$

where μ_{B} is the Bohr magneton, $\hbar$ is the reduced Planck constant, m_e is the electron mass, e is the elementary charge, c is the speed of light, g_s is the electron spin g-factor, and $V(r)$ is the potential energy in the radial direction. The result of the spin-orbit interaction is to give rise to a new set of sublevels characterised by J, taking unity-spaced values in the range $|L - S| \leq J \leq L + S$.

3.1.1.3 Hyperfine Structure

Further energy-level complexity arises due to the hyperfine interaction, which is a result of the interplay between the total electronic angular momentum J, and the *total nuclear angular momentum I*. The hyperfine interaction is described by the Hamiltonian [221,224,227,228]

---

[5]The Pauli exclusion principle states that two or more identical fermions (particles with half-integer spin, e.g. electrons), cannot simultaneously occupy the same quantum state.

$$\mathscr{H}_{\text{hfs}} = \frac{A_{\text{hfs}}}{\hbar^2} \boldsymbol{I} \cdot \boldsymbol{J}, \tag{3.3}$$

where A_{hfs} is the hyperfine structure constant and $\hbar$ is the reduced Planck constant. The result of this interaction is to give rise to *additional* structure, characterised by the *total angular momentum quantum number F*, taking unity-spaced values between $|J - I| \leq F \leq J + I$.

3.1.1.4 Zeeman Structure

Finally, the projection states m_L discussed previously, are converted to projection states $|m_F\rangle$ as a result of the hyperfine interaction (i.e. $m_L \to m_F$)[6]. In the absence of an external magnetic field, these projection states $|m_F\rangle$ are energy degenerate. However, in a finite external magnetic field, the Zeeman effect splits the energy degeneracy of these states (cf. Sec. 3.3).

A table summarising the physical origins of the atomic energy-level structure, including the associated quantum numbers and the values they assume, is presented in Table. 3.2.

Structure	Origin	Quantum number(s)	Range of values
Gross	Bohr levels	n, L	$n \in \mathbb{Z}^*, 0 \leq L \leq n - 1$
Fine	Spin-orbit interaction	J	$|L - S| \leq J \leq L + S$
Hyperfine	Hyperfine interaction	F	$|J - I| \leq F \leq J + I$
Zeeman	Zeeman effect	m_F	$-F \leq m_F \leq F$

Table 3.2: Atomic energy-level structure and its origin.

3.1.2 The D$_1$ and D$_2$ Transitions

The first set of excited states of rubidium are the $5^2P_{1/2}$ and $5^2P_{3/2}$ states, both of which have direct optical transitions from the ground state. These two transitions are $5^2S_{1/2} \to 5^2P_{1/2}$ and $5^2S_{1/2} \to 5^2P_{3/2}$, which for historical reasons, are known as the D$_1$ and D$_2$ transitions, respectively. A few important properties of the D$_1$ and D$_2$ transitions for ^{85}Rb and ^{87}Rb are presented in Table 3.3, with a corresponding energy-level diagram shown in Fig. 3.10.

Atom	Excited state	Transition wavelength (nm)	Lifetime (ns)	Decay rate (MHz)
^{85}Rb	$5^2P_{1/2}$	794.979014933(96)	27.679(27)	5.7500(56)
	$5^2P_{3/2}$	780.241368271(27)	26.2348(77)	6.0666(18)
^{87}Rb	$5^2P_{1/2}$	794.978851156(23)	27.679(27)	5.7500(56)
	$5^2P_{3/2}$	780.241209686(13)	26.2348(77)	6.0666(18)

Table 3.3: Key values corresponding to the $5^2P_{1/2}$ and $5^2P_{3/2}$ excited states of ^{85}Rb and ^{87}Rb, presented in Ref. [229].

---

[6]Analogous to the case with the $|m_L\rangle$ projection states, there are $2F + 1$ projection states corresponding to the values of m_F. These states are typically referred to as 'Zeeman sublevels', or 'Zeeman states'.

3.1.3 Vapour Pressure and Number Density

Rubidium has a melting point of 312.45 K (39.3°C) and a boiling point of 961 K (688°C) [230]; both of which are significantly greater than room temperature. One would thus expect that Rb atoms are *purely* in the solid phase in room-temperature experiments. However; the atomic energies will be distributed around a median value[7], resulting in a small percentage of atoms possessing sufficient energy to transition into the liquid phase. Out of these liquid-phase atoms, some will possess sufficient energy to overcome electrostatic forces [231] and evaporate into the surrounding space[8]. These atoms therefore undergo a phase transition into the gaseous phase, with a vapour pressure that is determined by thermodynamic equilibrium — which in turn, depends upon the ambient temperature. The vapour pressure P_v of rubidium in units of Torr, as a function of temperature T, is described by the following equation [229]:

$$\log_{10} P_v = \begin{cases} 2.881 + 4.857 - \frac{4215}{T} & \text{for } T < 312.45\,\text{K}, \\ 2.881 + 4.312 - \frac{4040}{T} & \text{for } T \geq 312.45\,\text{K}. \end{cases} \tag{3.4}$$

The number density n_v of the vapour can then be calculated via [232]

$$n_v = \left(\frac{101325}{760}\right) \frac{P_v N_A}{R_g T}, \tag{3.5}$$

where $N_A = 6.022140857 \times 10^{23}\,\text{mol}^{-1}$ is the Avogadro constant, $R_g = 8.3144598\,\text{J}\,\text{mol}^{-1}\,\text{K}^{-1}$ is the gas constant, and the factor of $101325/760$ converts P_v from units of Torr to Pa.

At room temperature ($T \approx 300\,\text{K}$), the number density of rubidium — calculated using Eqs. (3.4) and (3.5) — is $n_v \approx 1.57 \times 10^{16}\,\text{m}^{-3}$, which is sufficiently high to permit observation of nonlinear magneto-optical rotation at room temperature in millimetre-scale vapour cells. The number density of Rb over the range $T \in [200, 400]\,\text{K}$ is presented in Fig. 3.1.

3.2 Optical Absorption and Lineshape

Optical magnetometers require resonant or near-resonant light to both polarise (via optical pumping) and subsequently probe an atomic vapour. The rate $\mathcal{R}_{abs}(\omega)$ at which an atom absorbs photons of angular frequency ω is calculated by summing the incident photon flux over all atomic transitions:

$$\mathcal{R}_{abs}(\omega) = \sum_{\text{trans}} \sigma_A(\omega) \Phi_P(\omega), \tag{3.6}$$

where $\sigma_A(\omega)$ is the optical cross-section of the atomic transition and $\Phi_P(\omega)$ is the flux of

---

[7]Provided that the liquid can be approximated as an ideal gas, this distribution of atomic energies will be a Maxwell-Boltzmann distribution. However, any long-range interaction forces (such as those in water, for example) will result in a deviation from the Maxwell-Boltzmann distribution.

[8]In addition to the condition that the atoms require sufficient energy to overcome electrostatic forces, they must also be at the surface of the liquid, and their velocity must have a non-zero projection in the direction orthogonal to the surface plane.

Figure 3.1: Number density of Rb as a function of temperature, calculated using Eqs. (3.4) and (3.5). The melting point occurs at 312.45 K (39.3°C) and is indicated by the grey dashed line.

photons of angular frequency ω incident on the atom. The photon absorption cross-section is determined by the frequency response about the atomic transition frequency, ω_0, and is generally dependent upon three effects: the finite lifetime of the excited state, pressure broadening due to collisions, and Doppler broadening due to thermal motion of the atomic vapour. However; regardless of the form of the frequency response, the integral of the optical cross-section associated with a given atomic transition is a constant [222]:

$$\int_0^\infty \sigma_A(\omega)\,\mathrm{d}\omega = 2\pi^2 r_e c f_0\,, \tag{3.7}$$

where r_e is the classical electron radius, c is the speed of light, and f_0 is the oscillator strength[9].

3.2.1   The Natural Linewidth and Pressure Broadening

A fundamental result of quantum mechanics is the fact that atomic energy levels are discrete. One would thus expect that the optical spectrum of an atom is comprised of Dirac delta functions, where each function corresponds to the absorption of a photon with *exactly* the correct energy required to excite the transition. However; due to the Heisenberg uncertainty principle, each transition has a non-zero spread of energies for which the transition may be excited. The time-energy uncertainty principle states that the uncertainties in energy, σ_E,

---

[9]The oscillator strength is the fraction of the total classical integrated cross-section associated with the given atomic transition.

and time, σ_t, are related by the inequality[10]

$$\sigma_E \sigma_t \geq \frac{\hbar}{2},\tag{3.9}$$

where $\hbar$ is the reduced Planck constant. In the context of optical transitions, the uncertainty in time is equal to the natural lifetime of the excited state, $\sigma_t = \tau_{\text{nat}}$, while the uncertainty in energy, σ_E, is directly proportional to the uncertainty in optical frequency, σ_ω, through the Planck-Einstein relation $\sigma_E = \hbar \sigma_\omega$. The uncertainty in optical frequency is simply the half width at half maximum (HWHM) of the natural linewidth, i.e. $\Gamma_{\text{nat}} = 2\sigma_\omega$. This enables the natural linewidth to be calculated via Eq. (3.9):

$$\Gamma_{\text{nat}} = \frac{1}{\tau_{\text{nat}}}.\tag{3.10}$$

Collisions between rubidium atoms and other atoms/molecules result in both a shift and broadening of the optical resonance due to electromagnetic interactions. The magnitude of these effects is directly proportional to the number density of the atoms or molecules, and hence they are referred to as pressure shifts and pressure broadening, respectively. The frequency ν_{col} at which these collisions occur is given by [221]

$$\nu_{\text{col}} = n_{\text{v}} \sigma_{\text{col}} \bar{v}_{\text{th}},\tag{3.11}$$

where n_{v} is the number density of the colliding atoms or molecules, σ_{col} is the collisional cross-section, and $\bar{v}_{\text{th}}$ is the mean relative thermal velocity of the atoms or molecules.

The linewidths due to the natural lifetime and pressure broadening sum together to give a single linewidth, $\Gamma_{\text{L}} = \Gamma_{\text{nat}} + 2\nu_{\text{col}}$ [221]. The atomic frequency response around the resonant frequency, ω_0, takes the form of a Lorentzian lineshape, where Γ_{L} is the full width at half maximum (FWHM):

$$\mathcal{L}(\omega - \omega_0) = \frac{1}{\pi} \frac{\left(\frac{\Gamma_{\text{L}}}{2}\right)}{(\omega - \omega_0)^2 + \left(\frac{\Gamma_{\text{L}}}{2}\right)^2},\tag{3.12}$$

where the lineshape defined by Eq. (3.12) obeys the normalisation condition

$$\int_{-\infty}^{\infty} \mathcal{L}(\omega - \omega_0)\, d\omega = 1.\tag{3.13}$$

---

[10]The time-energy uncertainty principle is derived from the more general Schrödinger uncertainty principle between two observables A and B [233]:

$$\sigma_A \sigma_B \geq \sqrt{\left|\frac{1}{2}\langle\{\hat{A},\hat{B}\}\rangle - \langle\hat{A}\rangle\langle\hat{B}\rangle\right|^2 + \left|\frac{1}{2i}\langle[\hat{A},\hat{B}]\rangle\right|^2},\tag{3.8}$$

where $\hat{A}$ and $\hat{B}$ are the operators corresponding to the observables A and B respectively, $[\hat{A},\hat{B}]$ denotes the commutator of A and B, and $\{\hat{A},\hat{B}\}$ denotes the anti-commutator.

3.2.2 Power Broadening

Power broadening is the broadening of atomic resonances due to the presence of high-intensity light. The physical origin of power broadening is a population redistribution — either to the excited state, or to other ground states — due to optical pumping. In order for a photon to be absorbed, an electron must be excited from the ground state to the excited state. However, in doing so, the availability of atoms in the ground state is subsequently reduced[11]. Reduced ground-state availability results in saturation behaviour; i.e. as the intensity is *increased*, the amount of absorption *decreases*.

The simplest context in which power broadening arises is a two-level atom[12]. In this case, mathematical derivation of power broadening is readily achieved using the Beer-Lambert law, which states that the absorption of light by the atoms is described by the following differential equation [221]:

$$\frac{\mathrm{d}I_{\mathrm{int}}}{\mathrm{d}z} = -\alpha_{\mathrm{abs}}\left(\omega\right) I_{\mathrm{int}} = -\left(n_{|\mathrm{g}\rangle} - n_{|\mathrm{e}\rangle}\right) \sigma_{\mathrm{A}}\left(\omega\right) I_{\mathrm{int}}, \tag{3.14}$$

where $\alpha_{\mathrm{abs}}\left(\omega\right)$ is an absorption coefficient, $\sigma_{\mathrm{A}}\left(\omega\right)$ is the optical absorption cross-section, and $n_{|\mathrm{g}\rangle}$ and $n_{|\mathrm{e}\rangle}$ are the population densities of ground-state and excited-state atoms, respectively. In steady state, the net rate of energy absorbed per unit volume *must* equal the rate of energy spontaneously emitting from the excited state back down to the ground state. Mathematically, this can be written as

$$\left(n_{|\mathrm{g}\rangle} - n_{|\mathrm{e}\rangle}\right) \sigma_{\mathrm{A}}\left(\omega\right) I_{\mathrm{int}} = n_{|\mathrm{e}\rangle} A_{21} \hbar\omega, \tag{3.15}$$

where A_{21} is the Einstein A coefficient. Equation (3.15) can be written as $\left(n_{|\mathrm{g}\rangle} - n_{|\mathrm{e}\rangle}\right) r = n_{|\mathrm{e}\rangle}$, where $r = \sigma_{\mathrm{A}}\left(\omega\right) I_{\mathrm{int}}/\hbar\omega A_{21}$ is a dimensionless ratio. Combined with the fact that $n_{\mathrm{tot}} = n_{|\mathrm{g}\rangle} + n_{|\mathrm{e}\rangle}$, the difference in populations can be written as

$$n_{|\mathrm{g}\rangle} - n_{|\mathrm{e}\rangle} = \frac{n_{\mathrm{tot}}}{1 + 2r}. \tag{3.16}$$

At saturation, the difference between the ground-state and excited-state populations must be half of the total population. This enables Eq. (3.16) to be written as

$$n_{|\mathrm{g}\rangle} - n_{|\mathrm{e}\rangle} = \frac{n_{\mathrm{tot}}}{1 + \frac{2\sigma_{\mathrm{A}}(\omega)I_{\mathrm{int}}}{\hbar\omega A_{21}}} = \frac{n_{\mathrm{tot}}}{1 + \frac{I_{\mathrm{int}}}{I_{\mathrm{sat}}(\omega)}}, \tag{3.17}$$

where $I_{\mathrm{sat}}\left(\omega\right) \equiv 2\sigma_{\mathrm{A}}\left(\omega\right)/\hbar\omega A_{21}$ is the saturation intensity at which $n_{|\mathrm{g}\rangle} - n_{|\mathrm{e}\rangle} = n_{\mathrm{tot}}/2$[13].

---

[11]Due to stimulated emission, the maximum amount of population that can be in the excited state of a two-level atom is 50%.

[12]It should be noted that the general problem of optical pumping in the presence of multiple ground states, each with varying degrees of coupling to the light, is significantly more complicated. However, the general result remains unchanged; in the presence of high-intensity light, atomic resonances are broadened.

[13]When $\Delta\omega = \omega_0 - \omega = 0$, the on-resonance saturation intensity is given by [221]

$$I_{\mathrm{sat}}\left(\Delta\omega = 0\right) = \left(\frac{\pi hc}{3\lambda^3}\right)\Gamma_{\mathrm{nat}}, \tag{3.18}$$

where h is the Planck constant, c is the speed of light, and λ is the transition wavelength. For a given transition,

By inspection of Eq. (3.14), it is clear that $\alpha_{\text{abs}}(\omega) = \left(n_{|g\rangle} - n_{|e\rangle}\right) \sigma_{\text{A}}(\omega)$. In conjunction with Eq. (3.17), this enables us to write

$$\alpha_{\text{abs}}\left(I_{\text{int}}, \Delta\omega\right) = n_{\text{tot}}\sigma_{\text{A}}\left(0\right) \frac{\left(\frac{\Gamma}{2}\right)^2}{\Delta\omega^2 + \left(\frac{\Gamma}{2}\right)^2} \frac{1}{1 + \frac{I_{\text{int}}}{I_{\text{sat}}(\Delta\omega)}} , \tag{3.20}$$

where Γ is the linewidth of the optical absorption cross-section, and $\Delta\omega = \omega_0 - \omega$ is the detuning from the resonant angular frequency ω_0. Using the definition for the frequency-dependent saturation intensity, it is possible to write

$$\frac{I_{\text{sat}}\left(\Delta\omega\right)}{I_{\text{sat}}\left(0\right)} = \frac{\sigma_{\text{A}}\left(0\right)}{\sigma_{\text{A}}\left(\Delta\omega\right)} . \tag{3.21}$$

Upon inserting Eq. (3.21) into Eq. (3.20) and simplifying, we obtain the following expression for the normalised absorption coefficient [221]:

$$\frac{\alpha_{\text{abs}}\left(I_{\text{int}}, \Delta\omega\right)}{\alpha_{\text{abs}}\left(I_{\text{int}}, 0\right)} = \frac{\left(\frac{\Gamma}{2}\right)^2}{\Delta\omega^2 + \left(\frac{\Gamma}{2}\right)^2 \left(1 + \frac{I_{\text{int}}}{I_{\text{sat}}}\right)} . \tag{3.22}$$

Equation (3.22) is a Lorentzian profile with a FWHM of:

$$\Gamma_{\text{pb}} = \Gamma_{\text{nat}}\sqrt{1 + \frac{I_{\text{int}}}{I_{\text{sat}}}} . \tag{3.23}$$

The result of power broadening is therefore an intensity-dependent broadening of the natural linewidth[14].

Equation (3.23) is the mathematical description of power broadening, which can be understood as follows; the absorption coefficient is maximised on resonance (i.e. $\Delta\omega = 0$), which means that for a fixed intensity I_{int}, more ground-state population is excited when the light is resonant with the transition. This results in saturation, i.e. the rate of absorption slows down with higher intensity. If the light is detuned (i.e. $\Delta\omega \neq 0$), the absorption coefficient is reduced and hence less population is transferred to the excited state. The result of this is that the absorption in the wings increases faster than the absorption near the centre of the resonance, which yields a broadening of the resonance. The broadening predicted by Eq. (3.23), as well as the dependence of the normalised absorption coefficient on light intensity, Eq. (3.22), is presented in Fig. 3.2.

3.2.2.1 Power Broadening of Magneto-Optical Resonances

Analogous to the power broadening of the natural linewidth due to high-intensity light, magneto-optical resonances are also affected by power broadening, albeit due to a slightly

---

the on-resonance saturation intensity $I_{\text{sat}}\left(\Delta\omega = 0\right)$ is related to the Rabi frequency Ω_{R} and decay rate Γ_{nat} via [221]

$$\frac{I_{\text{int}}}{I_{\text{sat}}\left(\Delta\omega = 0\right)} = 2\left(\frac{\Omega_{\text{R}}}{\Gamma_{\text{nat}}}\right)^2 . \tag{3.19}$$

[14]Note that the natural linewidth is retrieved in the limit as $I_{\text{int}} \to 0$.
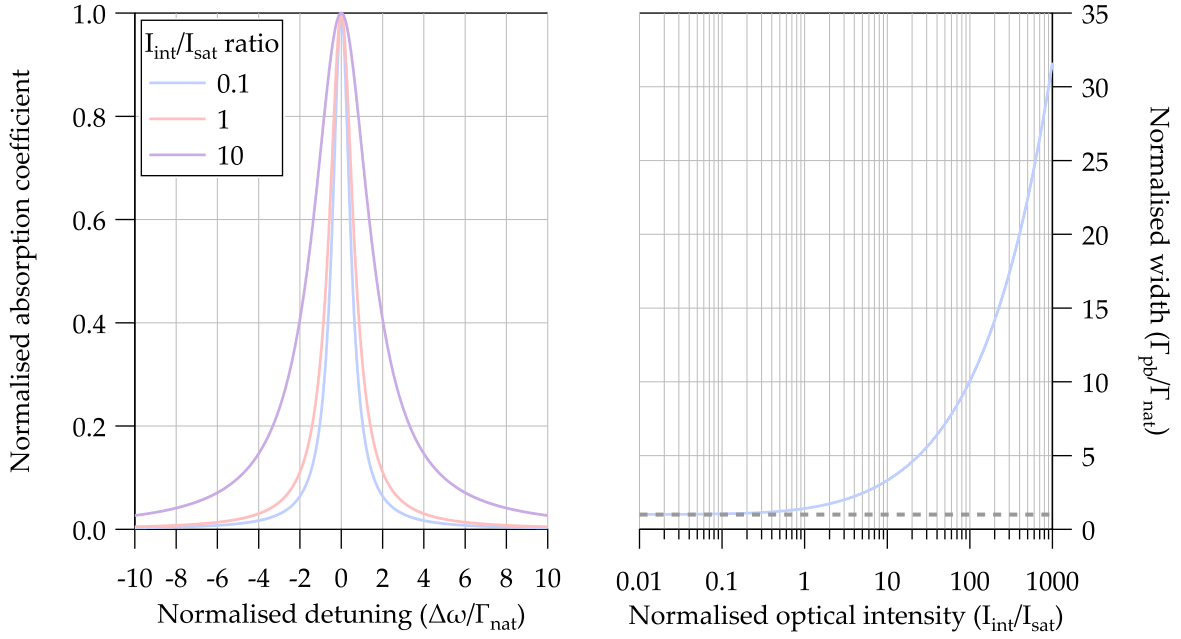
Figure 3.2: Left: The normalised absorption coefficient, calculated using Eq. (3.22) and scaled to an amplitude of unity, as a function of normalised detuning. Right: Normalised power-broadened width as a function of normalised optical intensity, calculated using Eq. (3.23). The grey dashed horizontal line indicates zero power broadening (i.e. $\Gamma_{pb} = \Gamma_{nat}$), which is achieved in the limit as $I_{int} \to 0$.

different physical mechanism. As discussed in Sec. 3.6, the observed magneto-optical resonance is due to the formation of coherences between ground-state Zeeman sublevels as a result of optical pumping. The FWHM of the magneto-optical resonance is related to the tranverse spin-relaxation time, T_2, via $\Gamma = 2/T_2$. Increases in optical powers, especially the probe beam, result in destruction of ground-state coherence via an effective relaxation due to out-of-phase optical pumping. This leads to a subsequent broadening of the resonance width due to reduced transverse spin-relaxation time.

When the optical pumping rate $\Gamma_{pump} = \Omega_R{}^2/\Gamma_{nat}$ is sufficiently small, the resonance is described by a zero-power resonance width Γ_0[15], and any additional relaxation due to the light can be safely ignored. However, at high light powers where the optical pumping rate is comparable to or greater than the standard ground-state relaxation rate (i.e. $\Gamma_{pump} \gtrsim \Gamma_0$), atoms which have been polarised in prior pumping cycles may have their polarisation destroyed by an additional pumping cycle before they would otherwise relax due to non-power-induced relaxation mechanisms [242][16]. This effectively reduces the transverse spin-

---

[15]Depending on the experimental circumstances, this can be due to such mechanisms as magnetic-field gradients [234], transit relaxation [150,165,200,202,235–238], wall collisions [150,236,239] (dominant in the case of no anti-relaxation coatings or buffer gases), spin-exchange collisions [150,240], spin-destruction collisions [241], etc.

[16]Relaxation due to optical pumping with high optical power only occurs if the polarisation of the atom differs from that which would be induced from the incident light. If these two polarisations are the same, then there will be no apparent effect [242]. However, given the spatial dependence of the Larmor frequency (due to magnetic-field gradients, for example), there will in general exist a subset of atomic population possessing a polarisation which differs from that which the light would induce. This ensures that power broadening always occurs, to some extent, under typical experimental conditions.

relaxation time of the vapour which, for aforementioned reasons, increases the resonance width. This is known as *power broadening* in the context of magneto-optical resonances.

Power broadening of magneto-optical resonances is typically characterised using a dimensionless saturation parameter κ_2 — defined as the ratio of the excitation rate to the relaxation rate[17] — which is related to the optical pumping rate Γ_{pump} and zero-power resonance width Γ_0 via [150,165,242,243]

$$\kappa_2 = \frac{\Gamma_{\text{pump}}}{\Gamma_0} = \frac{\Omega_{\text{R}}{}^2}{\Gamma_{\text{nat}}\Gamma_0} . \tag{3.24}$$

Typically, if relaxation processes are isotropic (i.e. they affect all atoms *equally*, regardless of their quantum state), the total effective relaxation rate is simply the *sum* of the individual relaxation rates. In this case, the total relaxation rate would be $\Gamma_0 + \Gamma_{\text{pump}} = \Gamma_0 (1 + \kappa_2)$. However, optical-pumping-induced relaxation is *not* isotropic, as linearly polarised light along one axis *only* results in relaxation of atoms polarised in the orthogonal axis. This results in a weaker dependence on κ_2, which has been shown to be $\Gamma_0 \sqrt{1 + \kappa_2}$ [165,242]. The total resonance FWHM in the case of power broadening is therefore given by[18] [165,242]

$$\Gamma = \Gamma_0 \sqrt{1 + \kappa_2} . \tag{3.25}$$

Under the optimal experimental conditions detailed within this thesis, the power broadening of the magneto-optical resonance width due to both pump and probe powers has been demonstrated in Ref. [206] to take the form $\Gamma \propto \sqrt{\left(P_{\text{probe}}\right)^2 + 3\left(P_{\text{pump}}\right)^2}$. The prefactor on the pump power term is the result of the power-broadening being more sensitive to increases in probe power, relative to pump power. This is attributed to the fact that the probe beam is continuous-wave (cw), and hence any optical pumping by the probe beam is generally out-of-phase and subsequently leads to destruction of ground-state coherence.

3.2.2.2 Broadening of Magneto-Optical Resonances due to the ac Stark Effect

In addition to the power broadening described above — which results from the destruction of atomic polarisation — the presence of optical power also results in resonance broadening due to the ac Stark effect.

A comprehensive treatment of the ac Stark shift requires a density matrix calculation to account for complicated processes [204] such as optical pumping, magnetic-field-induced evolution of the Zeeman sublevels, Doppler broadening, modulation of the light frequency, and the transit of atoms in and out of the light. However, by ignoring these complexities, the ac Stark shift of the $|m_F\rangle$ Zeeman state for a two-level system is well approximated by [204]

$$\delta_{\text{s}}{}^{(m_F)} = \frac{\Omega_{\text{R}}{}^2}{4} \frac{\Delta\omega + m_F\Omega_{\text{L}}}{(\Delta\omega + m_F\Omega_{\text{L}})^2 + \left(\frac{\Gamma_{\text{pb}}}{2}\right)^2} , \tag{3.26}$$

---

[17]Based on the definition of κ_2, power broadening becomes an issue when $\kappa_2 \gtrsim 1$.

[18]Note that Eq. (3.25) has the same form as the power broadening of the natural linewidth given by Eq. (3.23).

where Ω_R is the Rabi frequency, Ω_L is the Larmor frequency, $\Delta\omega$ is the optical detuning, and Γ_{pb} is the power-broadened natural linewidth. The effect of the ac Stark shift for an $F = 2$ hyperfine ground state is to modify the resonance frequencies $\omega_0{}^{(m_F, m_F)}$ for the three different $|\Delta m_F| = 2$ coherences [204]:

$$\omega_0{}^{(2,0)} = 2\Omega_L + \left(\delta_s{}^{(2)} - \delta_s{}^{(0)} \right) , \tag{3.27}$$

$$\omega_0{}^{(1,-1)} = 2\Omega_L + \left(\delta_s{}^{(1)} - \delta_s{}^{(-1)} \right) , \tag{3.28}$$

$$\omega_0{}^{(0,-2)} = 2\Omega_L - \left(\delta_s{}^{(-2)} - \delta_s{}^{(0)} \right) . \tag{3.29}$$

Given that the three $|\Delta m_F| = 2$ coherences constructively interfere at low magnetic-field strength — and partially so in the geophysical range — the result of the ac Stark shift is to broaden the resonance by splitting the coherences away from each other.

3.2.3 Doppler Broadening

Further to the natural broadening arising from a finite excited-state lifetime, and power broadening resulting from optical pumping and the ac Stark effect, an additional important broadening mechanism for a room-temperature atomic vapour is Doppler broadening. When light is incident on an atomic ensemble, the thermal motion of absorbing atoms leads to small variations in the observed frequency, ω, due to the Doppler effect. If an atom has a velocity component v_z along the propagation direction $\hat{z}$ of the incident light, then the frequency of the light as experienced by the atom is given by

$$\omega = \omega_0 \left(1 + \frac{v_z}{c} \right) , \tag{3.30}$$

where ω_0 is the 'true' frequency of the light — typically tuned to the atomic resonance — and c is the speed of light. The Doppler effect results in broadening of atomic resonance lines, as light of frequency ω *appears* to be at the resonant frequency ω_0 to atoms with a velocity component v_z along the light's propagation direction which satisfies the condition

$$v_z = c \left(\frac{\omega - \omega_0}{\omega_0} \right) . \tag{3.31}$$

In an ideal gas, atoms with a mass m and temperature T have a non-zero velocity v which is described by the Maxwell-Boltzmann distribution [244]. The probability that an atom has a velocity in the interval v to $v + dv$ is given by [244]

$$\mathscr{P}(v)\, d^3 v = \left(\frac{m}{2\pi k_B T} \right)^{3/2} \exp\left(\frac{-mv^2}{2k_B T} \right) d^3 v , \tag{3.32}$$

where k_B is the Boltzmann constant. If we take $d^3 v = dv_x dv_y dv_z$, then Eq. (3.32) factorises into three separate identical distributions for each Cartesian component. The velocity dis-

tribution in the z-component, for example, is therefore given by

$$\mathscr{P}(v_z)\,\mathrm{d}v_z = \sqrt{\frac{m}{2\pi k_B T}} \exp\left(\frac{-mv_z^2}{2k_B T}\right) \mathrm{d}v_z. \tag{3.33}$$

From an experimental viewpoint, it is far more convenient to express the probability density function as a function of the optical frequency rather than the atomic velocity. It is possible to write Eq. (3.33) in terms of optical frequency by employing a change of variables:

$$\mathscr{P}(\omega)\,\mathrm{d}\omega = \mathscr{P}(v_z)\frac{\mathrm{d}v_z}{\mathrm{d}\omega}\mathrm{d}\omega. \tag{3.34}$$

From Eq. (3.31), we find that $\mathrm{d}v_z = (c/\omega_0)\,\mathrm{d}\omega$. If we now substitute both our expression for $\mathrm{d}v_z$ and Eq. (3.31) into Eq. (3.33), we obtain the following frequency-dependent probability density function:

$$\mathscr{P}(\omega)\,\mathrm{d}\omega = \frac{c}{\omega_0}\sqrt{\frac{m}{2\pi k_B T}} \exp\left(-\frac{(\omega-\omega_0)^2}{2\Gamma_D^2}\right) \mathrm{d}\omega, \tag{3.35}$$

where Γ_D is the Doppler width defined as

$$\Gamma_D \equiv \frac{\omega_0}{c}\sqrt{\frac{k_B T}{m}}. \tag{3.36}$$

The resulting frequency response around the resonant frequency ω_0, due to Doppler broadening, takes the form of a Gaussian distribution with a $1/e$ width of $\Gamma_G = \Gamma_D$:

$$\mathcal{G}(\omega-\omega_0) = \frac{1}{\Gamma_G\sqrt{2\pi}} \exp\left(-\frac{(\omega-\omega_0)^2}{2\Gamma_G^2}\right), \tag{3.37}$$

where the lineshape defined by Eq. (3.37) obeys the normalisation condition

$$\int_{-\infty}^{\infty} \mathcal{G}(\omega-\omega_0)\,\mathrm{d}\omega = 1. \tag{3.38}$$

The Doppler width (FWHM) of the D_1 transition of ^{87}Rb, as a function of temperature, is shown in Fig. 3.3. By inspection of Fig. 3.3 it is evident that the Doppler width is *significantly* greater than the natural linewidth — even at cryogenic temperatures. It is therefore a good approximation to consider Doppler broadening as the dominant broadening mechanism in most room-temperature experiments.

In the context of the experiments detailed within this thesis, Doppler broadening determines the resonance width for *linear* magneto-optical effects [245,246]. Fortunately, however, the width of *nonlinear* magneto-optical effects described here are largely unaffected.

3.2.4 The Voigt Profile

The general frequency response of an atomic resonance is dependent upon *all* of the aforementioned effects: the natural lifetime, pressure broadening, power broadening, and Doppler broadening. The Lorentzian lineshape that results from the natural lifetime, pres-

Figure 3.3: Doppler width (FWHM) of the D_1 transition of ^{87}Rb (blue) as a function of temperature, calculated by $2\sqrt{2\ln 2}\Gamma_D/2\pi$ using Eq. (3.36). The grey dashed line indicates room temperature (22.5°C).

sure broadening and power broadening is further broadened by the Gaussian lineshape that the Doppler effect produces. The general frequency response is therefore a convolution of the Lorentzian and Gaussian lineshapes, and is known as the Voigt profile:

$$\mathcal{V}\left(\omega - \omega_0\right) = \int_{-\infty}^{\infty} \mathcal{L}\left(\omega' - \omega_0\right) \mathcal{G}\left(\omega - \omega'\right) \mathrm{d}\omega'. \tag{3.39}$$

This integral cannot be expressed in terms of elementary functions and must therefore be calculated numerically. Unfortunately, the convolution operation is *incredibly* computationally expensive [247–249], rendering this definition of the Voigt profile cumbersome — especially when fitting experimental data with a high number of samples. However, provided that $\Gamma_L > 0$, the Voigt profile can be calculated via [250,251]

$$\mathcal{V}\left(\omega - \omega_0\right) = \frac{\mathrm{Re}\left\{w\left(z\right)\right\}}{\sqrt{2\pi}\Gamma_G}, \tag{3.40}$$

where $w\left(z\right)$ is the Faddeeva function, and is defined as [250,251]

$$w\left(z\right) = e^{-z^2}\mathrm{erfc}\left(-iz\right) \tag{3.41}$$

$$= e^{-z^2}\left(1 + \frac{2i}{\sqrt{\pi}}\int_0^z e^{t^2}\mathrm{d}t\right), \tag{3.42}$$

and where z is evaluated at

Figure 3.4: Comparison between Voigt, Gaussian and Lorentzian lineshapes. The FWHM
of the Gaussian and Lorentzian lineshapes are $\Gamma_L/2\pi = 2\sqrt{2\ln 2}\Gamma_G/2\pi =$
500 MHz. Upon convolving the Gaussian and Lorentzian lineshapes, the result-
ing Voigt profile has a FWHM of $\Gamma_\mathcal{V}/2\pi \approx 819$ MHz according to Eq. (3.45).

$$z = \frac{(\omega - \omega_0) + i\left(\frac{\Gamma_L}{2}\right)}{\sqrt{2}\Gamma_G}. \tag{3.43}$$

The Faddeeva function may be calculated using routines for the complementary error func-
tion, erfc (z). Although the error function *also* requires numerical computation, it is orders
of magnitude *quicker* than the convolution operation.

Since the Lorentzian and Gaussian profiles, as defined via Eqs. (3.12) and (3.37) respec-
tively, are normalised to *unit area* in accordance with Eqs. (3.13) and (3.38), so too is the Voigt
profile. This is true regardless of whether the Voigt profile is calculated using Eq. (3.39) or
Eq. (3.40). Explicitly, the normalisation condition for the Voigt profile is given by

$$\int_{-\infty}^{\infty} \mathcal{V}(\omega - \omega_0)\,\mathrm{d}\omega = 1. \tag{3.44}$$

A comparison between the Gaussian, Lorentzian and Voigt profiles is presented in Fig.
3.4. The FWHM of the resulting Voigt profile is a complicated function of the FWHM of
both the Lorentzian and Gaussian components. However, it can be approximated to within
0.02%, for all Γ_L and Γ_G, via [252]

$$\Gamma_\mathcal{V} \approx 0.5346\Gamma_L + \sqrt{\left(\sqrt{0.2166}\Gamma_L\right)^2 + \left(2\sqrt{2\ln 2}\Gamma_G\right)^2}. \tag{3.45}$$

In a room-temperature vapour cell which contains no buffer gas (i.e anti-relaxation coated

with paraffin), there is very little collisional broadening and hence $\Gamma_L \ll \Gamma_G$. In the case of ^{87}Rb, the Doppler width (FWHM) at room temperature is $2\sqrt{2\ln 2}\Gamma_D/2\pi \approx 500\,\text{MHz}$, while the natural linewidth is $\Gamma_{nat}/2\pi = 5.75\,\text{MHz}$ for the D_1 transition [229]. The resulting Voigt profile thus has a FWHM of $\Gamma_V/2\pi \approx 500.07\,\text{MHz}$ according to Eq. (3.45). In this scenario, it is therefore a *very good* approximation to treat the atomic resonances as Gaussian line-shapes[19]. For this reason, all atomic resonances will be treated merely as Gaussian profiles throughout this thesis unless explicitly stated otherwise.

3.3 The Zeeman Effect

The Zeeman effect arises due to the interaction of an atom with an external magnetic field. Given a magnetic dipole moment $\boldsymbol{\mu}$ in a magnetic field $\boldsymbol{B}$, the magnetic potential energy is given by $\mathscr{H}_I = -\boldsymbol{\mu} \cdot \boldsymbol{B}$ [221]. This gives rise to a total Hamiltonian given by

$$\mathscr{H} = \mathscr{H}_0 - \boldsymbol{\mu} \cdot \boldsymbol{B}, \tag{3.47}$$

where $\mathscr{H}_0$ is the unperturbed atomic Hamiltonian (cf. Sec. 4.1). The result of this interaction is to remove the energy degeneracy of the atomic hyperfine levels, by splitting the energies of the magnetic projection states $|m_F\rangle$. This splitting is dependent upon the strength of the magnetic field, and can be separated into linear and nonlinear regimes.

3.3.1 Linear Zeeman Effect

If the external magnetic-field strength, $B = |\boldsymbol{B}| = \sqrt{B_x{}^2 + B_y{}^2 + B_z{}^2}$, is sufficiently small, first-order perturbation theory can be used to calculate the correction to the energy levels. The first-order correction is simply the expectation value of the perturbation Hamiltonian (i.e. $\mathscr{H}_I$), while the system is in the unperturbed state. Explicitly, this is given by

$$\begin{aligned}
\Delta E^{(1)} &= \langle F\,m_F | \mathscr{H}_I | F\,m_F \rangle \\
&= \frac{m_F g_F \mu_B}{\hbar} B \\
&= m_F \Omega_L,
\end{aligned} \tag{3.48}$$

where m_F is the magnetic quantum number, $\Omega_L = g_F \mu_B B/\hbar$ is the Larmor frequency, g_F is the Landé g-factor, μ_B is the Bohr magneton, and $\hbar$ is the reduced Planck constant. The result is that, for small magnetic-field strength, the energies of the $|m_F\rangle$ Zeeman states (with the exception of $|m_F = 0\rangle$) are shifted linearly in the magnetic-field strength, and proportional to the magnetic quantum number m_F. Of course, if the magnetic-field strength is too large, first-order perturbation theory is no longer sufficient and higher-order terms must be considered.

---

[19]By inspection of Eq. (3.45), it is evident that $\Gamma_V \to 2\sqrt{2\ln 2}\Gamma_G$ as $\Gamma_L \to 0$. Therefore, for small values of Γ_L, the Voigt profile tends towards a Gaussian:

$$\lim_{\Gamma_L \to 0} \mathcal{V}(\omega - \omega_0) = \mathcal{G}(\omega - \omega_0). \tag{3.46}$$

3.3.2   Nonlinear Zeeman Effect

If the magnetic-field strength is too large to be covered by the linear Zeeman effect, yet too small to be accounted for by the Paschen-Back effect [221,224,229] (i.e. the strong-field regime), the field fits into the intermediate regime[20]. Calculating the energy-level shifts in this regime is difficult in general; however, provided that $L = 0$ and $J = 1/2$, the shift in hyperfine energy levels, $\Delta E^{(\mathrm{BR})}$, is given by the Breit-Rabi formula [204,224,229,253–257]:

$$\Delta E^{(\mathrm{BR})} = -\frac{\Delta E_{\mathrm{hfs}}}{2(2I+1)} + g_I \mu_{\mathrm{B}} m_F B \pm \frac{\Delta E_{\mathrm{hfs}}}{2} \sqrt{1 + \frac{4 m_F x}{2I+1} + x^2}, \qquad (3.49)$$

where $\Delta E_{\mathrm{hfs}} = A_{\mathrm{hfs}} \left(I + 1/2 \right)$ is the hyperfine splitting, A_{hfs} is the magnetic dipole constant, g_I is the nuclear g-factor, μ_{B} is the Bohr magneton, B is the magnetic-field strength, I is the nuclear quantum number, and x is a dimensionless 'field-strength parameter' defined by

$$x \equiv \frac{(g_J - g_I)\, \mu_{\mathrm{B}} B}{\Delta E_{\mathrm{hfs}}}, \qquad (3.50)$$

where g_J is the ground-state Landé g-factor. It should be noted that the sign of the third term is *positive* for the higher-number F level, and *negative* for the lower-number level. Additionally, for $m_F = -\left(I + 1/2 \right)$, the square root is an exact square, and should be interpreted as $+\left(1 - x \right)$. The $5^2\mathrm{S}_{1/2}$ ground state hyperfine levels of rubidium have $L = 0$ and $J = 1/2$, which enables the shifts in the Zeeman states to be calculated via the Breit-Rabi formula, and is shown in Fig. 3.5.

Due to the quadratic nature of the nonlinear Zeeman effect, the three $|\Delta m_F| = 2$ Zeeman resonances which comprise the NMOR signal in the $F = 2$ ground state will begin to diverge from each other as the magnetic-field strength is increased. This becomes problematic at geophysical fields where the nonlinear Zeeman splitting is sufficient to partially resolve each of the three individual resonances, as it results in reduced resonance slope and hence magnetic sensitivity[21] [204,255,256]. The frequency splitting of the three resonances, δ^{NLZ}, about the central resonance, is given by [256] (cf. App. J)

$$\delta^{\mathrm{NLZ}} \approx \frac{(\mu_{\mathrm{B}} B)^2}{\hbar \Delta E_{\mathrm{hfs}}}. \qquad (3.51)$$

The resonance frequencies of the three $|\Delta m_F| = 2$ coherences, due to both the nonlinear Zeeman effect and the ac Stark effect (cf. Sec. 3.2.2.2), are given by [204]

$$\omega_0^{(2,0)} = 2\Omega_{\mathrm{L}} - \delta^{\mathrm{NLZ}} + \left(\delta_{\mathrm{s}}^{(2)} - \delta_{\mathrm{s}}^{(0)} \right), \qquad (3.52)$$

$$\omega_0^{(1,-1)} = 2\Omega_{\mathrm{L}} + \left(\delta_{\mathrm{s}}^{(1)} - \delta_{\mathrm{s}}^{(-1)} \right), \qquad (3.53)$$

---

[20]The intermediate regime, along with the Paschen-Back effect (strong-field regime), are typically collectively referred to as the 'nonlinear Zeeman effect'.

[21]The nonlinear Zeeman effect can also give rise to a phenomenon known as 'heading error' [255,256,258–260], which is a *spurious dependence* of the magnetic-field strength measurement on the relative orientation of the sensor in the magnetic field, despite the sensor being a scalar magnetometer.
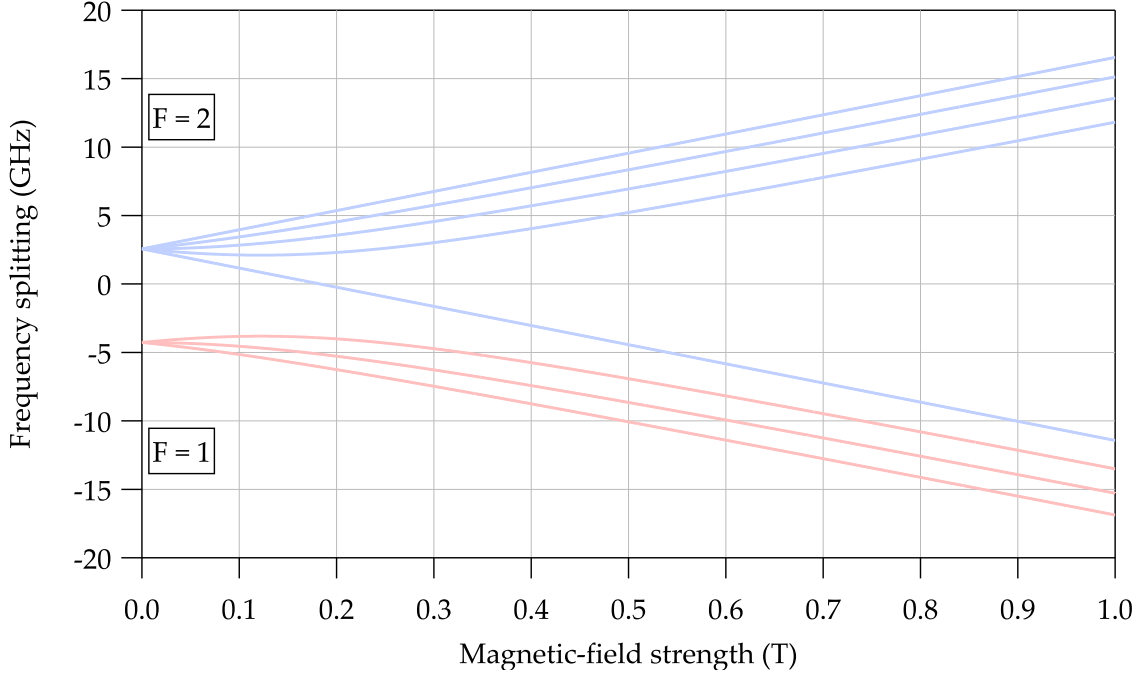
Figure 3.5: Frequency splitting of the $5^2S_{1/2}$ hyperfine ground states of ^{87}Rb due to the Zeeman effect — showing both linear and nonlinear contributions — calculated using Eq. (3.49).

$$\omega_0^{(0,-2)} = 2\Omega_\text{L} + \delta^\text{NLZ} - \left(\delta_\text{s}^{(-2)} - \delta_\text{s}^{(0)} \right) . \tag{3.54}$$

The nonlinear Zeeman splitting of the ground-state Zeeman resonances is shown in Fig. 3.6.

For the $F = 2$ ground state of ^{87}Rb, there are three distinct resonances associated with coherences between Zeeman states with $|\Delta m_F| = 2$. The NMOR signal is therefore comprised of three separate resonances which, when measured in the frequency domain using conventional phase-sensitive detection techniques, take the form of a complex Lorentzian (cf. Sec. 6.1.2). At low magnetic-field strength, $\delta^\text{NLZ} \approx 0$ and the resonances constructively interfere to yield one effective resonance. As the magnetic-field strength is increased however, the resonances split *quadratically* in accordance with Eq. (3.51), resulting in decreased resonance slope and therefore reduced magnetic sensitivity. For the $F = 2$ ground state, the zero-crossing resonance slope, as a function of nonlinear Zeeman shift, is given by

$$
\begin{aligned}
\mathcal{S}_3 = {} & \frac{2\mathcal{A}_{1,-1}}{\Gamma} + \frac{\mathcal{A}_{2,0}\Gamma}{2\left[\left(\frac{\Gamma}{2}\right)^2 + (\delta^\text{NLZ})^2\right]} + \frac{\mathcal{A}_{0,2}\Gamma}{2\left[\left(\frac{\Gamma}{2}\right)^2 + (\delta^\text{NLZ})^2\right]} \\
& - \frac{\mathcal{A}_{2,0}\Gamma\,(\delta^\text{NLZ})^2}{\left[\left(\frac{\Gamma}{2}\right)^2 + (\delta^\text{NLZ})^2\right]^2} - \frac{\mathcal{A}_{0,2}\Gamma\,(\delta^\text{NLZ})^2}{\left[\left(\frac{\Gamma}{2}\right)^2 + (\delta^\text{NLZ})^2\right]^2} ,
\end{aligned}
\tag{3.55}
$$

where $\mathcal{A}_{i,j}$ are the amplitudes of the resonance between $|m_F = i\rangle$ and $|m_F = j\rangle$, and Γ is the resonance width. Equation (3.55) is plotted in Fig. 3.6 for $2\mathcal{A}_{2,0} = 2\mathcal{A}_{0,2} = \mathcal{A}_{1,-1}$ and $\Gamma = 10$ — that is, for a magnetic field aligned directly along the pump/probe axis. There are two
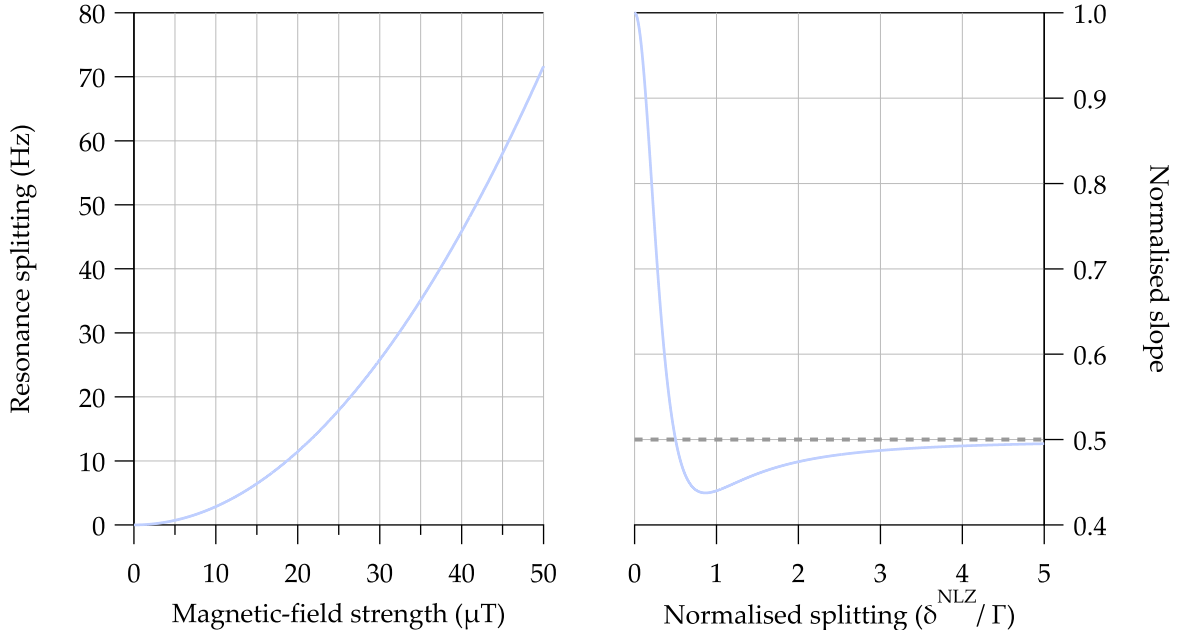
Figure 3.6: Left: Frequency splitting of the $|\Delta m_F| = 2$ ground-state resonances due to the nonlinear Zeeman effect, calculated using Eq. (3.51) for the $F = 2$ ground state of ^{87}Rb. Right: The effect of the nonlinear Zeeman shift on the $F = 2$ ground-state resonance slope, calculated using Eq. (3.55) when the pump/probe axis lies directly along the magnetic field. The dashed grey line corresponds to the limit as $\delta^{\mathrm{NLZ}} \to \infty$, in which case the resonance slope is half of its maximum value.

interesting limits to observe here: the limit as $\delta^{\mathrm{NLZ}} \to \infty$, and the global minimum in the resonace slope. In the former case, the resonance slope reduces to half of its maximum value for large values of δ^{NLZ}, and is indicated by the grey dashed line in Fig. 3.6. In the latter case, the slope is minimised when

$$\delta^{\mathrm{NLZ}} = \frac{\sqrt{3}}{2}\Gamma\,, \tag{3.56}$$

which, for the parameters presented in Fig. 3.6, yields just over 40% of its maximum value. Note that although Eq. (3.55) is true in general, Eq. (3.56) and the graph presented in the right of Fig. 3.6 are both valid *only* for a magnetic field that is aligned along the pump/probe axis. The relative amplitudes of the ground-state resonances are, in general, dependent upon the orientation of the magnetometer within the field [255,256,258–260]. If the magnetic field is *not* aligned along the pump/probe axis, there will be an asymmetry in the resonance [255], and the dependence of the resonance slope on the nonlinear Zeeman shift will differ from that presented here.

Although the resonance splitting which arises from the nonlinear Zeeman effect results in reduced resonance slope and hence sensitivity, it is not necessarily a significant detrimental effect *per se*. As shown in Fig. 3.6, the maximum reduction in slope — for a magnetic field aligned along the pump/probe axis — is about 60%, with a limiting reduction of a factor of 2 in the case of large δ^{NLZ}. In this case, the significant reduction in sensitivity at high magnetic-field strength is *not* due to the increased δ^{NLZ}, but rather it is often due to

gradient broadening. As will be discussed in Sec. 5.4, the transverse spin-relaxation time T_2 — which is related to the magnetic sensitivity via. Eq. (6.18) — scales quadratically in the first-order field gradient. For a fixed fractional field uniformity, as the magnetic-field strength is increased, the first order gradient also increases and therefore reduces sensitivity.

3.4   The Faraday Effect

The Faraday effect — first discovered by Michael Faraday circa 1846 [261] — is a magneto-optical effect in which the plane of polarisation of linearly polarised light is rotated as it traverses a medium situated within a magnetic field. This is a linear effect, in that the optical rotation ϕ experienced by the light is linear in the magnetic-field strength:

$$\phi = \mathcal{V}_e l B , \tag{3.57}$$

where $\mathcal{V}_e$ is the material-dependent Verdet constant, l is the length of active medium, and B is the magnetic-field strength. The Verdet constant is effectively a proportionality constant, which describes the amount of rotation of the incident light per unit path length and per unit magnetic-field strength.

In most materials — such as solids and liquids — the Verdet constant is quite small, resulting in little optical rotation. However, in low-density alkali vapour (such as rubidium, potassium and caesium) the Verdet constant is *many* orders of magnitude greater than most other materials. A comparison of Verdet constants for some common materials, as well as rubidium vapour, is presented in Table 3.4.

3.5   The Macaluso-Corbino Effect

The Macaluso-Corbino effect — named after its discoverers Damiano Macaluso and Orso Corbino — is the Faraday effect in the vicinity of resonance absorption lines. This effect is readily derived through the Lorentz electron oscillator model, in which an electron is harmonically bound to the nucleus. The atomic response is determined by solving the following second-order ordinary differential equation:

$$m_{\mathrm{e}} \frac{\mathrm{d}^2 \boldsymbol{r}(t)}{\mathrm{d}t^2} + m_{\mathrm{e}} \Gamma_{\mathrm{d}} \frac{\mathrm{d}\boldsymbol{r}(t)}{\mathrm{d}t} + k_{\mathrm{s}} \boldsymbol{r}(t) = e \boldsymbol{\mathcal{E}}(t) , \tag{3.58}$$

where m_{e} and e are the mass and charge of the electron respectively, Γ_{d} is the relaxation rate of the excited state, $k_{\mathrm{s}} = m_{\mathrm{e}} \omega_0{}^2$ is the spring constant of the resonance, $\boldsymbol{r}(t)$ is the time-dependent position of the electron, and $\boldsymbol{\mathcal{E}}(t)$ is the applied electric field. In order to solve this differential equation for the frequency response, Fourier transforms (cf. App. G) are used. Computing the Fourier transform of both sides of Eq. (3.58), yields the following:

$$m_{\mathrm{e}} \mathscr{F}\left\{ \frac{\mathrm{d}^2 \boldsymbol{r}(t)}{\mathrm{d}t^2} \right\} + m_{\mathrm{e}} \Gamma_{\mathrm{d}} \mathscr{F}\left\{ \frac{\mathrm{d}\boldsymbol{r}(t)}{\mathrm{d}t} \right\} + k_{\mathrm{s}} \mathscr{F}\left\{ \boldsymbol{r}(t) \right\} = e \mathscr{F}\left\{ \boldsymbol{\mathcal{E}}(t) \right\} . \tag{3.59}$$

Material	$\mathcal{V}_e$ (rad T^{-1} m^{-1})	Comments	Reference(s)
Lithium fluoride	3	$\lambda \approx 500\,\mathrm{nm}$	[262]
Olive oil	3.35	$\lambda = 659.2\,\mathrm{nm}$	[263]
Water	$5-7$	$\lambda \approx 500\,\mathrm{nm}$	[262,264]
Terbium-doped phosphate fibre	6.2	$\lambda = 1053\,\mathrm{nm}$	[265]
Ethanol	7	$\lambda \approx 500\,\mathrm{nm}$	[264]
Methanol	7	$\lambda \approx 500\,\mathrm{nm}$	[264]
BK7 glass	9	$\lambda \approx 500\,\mathrm{nm}$	[264]
Fused quartz (SiO_2)	9.62	$\lambda = 404.7\,\mathrm{nm}$	[262]
Dense flint glass (SF-57)	21	$\lambda = 633\,\mathrm{nm}$	[266]
Terbium gallium garnet ($Tb_3Ga_5O_{12}$)	290	$\lambda = 457.9\,\mathrm{nm}$	[267]
Rubidium (Linear Faraday)	1.4×10^3	Warm atoms	[268]
Rubidium (Linear Faraday)	3.5×10^6	Cold atoms, optically thick, 150° of rotation	[269]
Rubidium (Nonlinear Faraday)	$10^5 - 10^{10}$	Near-vacuum density	[203,205,206]

Table 3.4: Verdet constants of various materials. Note that these values are only approximate, as in general they are dependent upon variables such as temperature, density and wavelength.

We know that the Fourier transform of our function $r(t)$ is simply $\mathscr{F}\{r(t)\} = r(\omega)$, but what are the Fourier transforms of these time derivatives? In general, the Fourier transform of a temporal derivative is given by

$$\mathscr{F}\left\{\frac{\mathrm{d}^n f(t)}{\mathrm{d}t^n}\right\} = (i\omega)^n \mathscr{F}\{f(t)\}. \tag{3.60}$$

By using the above Fourier transform identity, Eq. (3.59) simplifies to

$$\begin{aligned} r(\omega) &= \frac{(e/m_\mathrm{e})}{(\omega_0{}^2 - \omega^2) + i\omega\Gamma_\mathrm{d}} \mathcal{E}(\omega) \\ &= \frac{\tilde{\alpha}(\omega)}{e} \mathcal{E}(\omega), \end{aligned} \tag{3.61}$$

where $\tilde{\alpha}(\omega)$ is the complex polarisability. This now enables an expression for the polarisation response to be derived:

$$\begin{aligned} \mathbf{P}(\omega) &= n_e e r(\omega) \\ &= e\tilde{\alpha}(\omega)\mathcal{E}(\omega), \end{aligned}$$

where n_e is the number density of electrons. Recall for a linear material, the polarisation can

be written as $\mathbf{P}(\omega) = \epsilon_0 \tilde{\chi}(\omega) \boldsymbol{\mathcal{E}}(\omega)$, where $\tilde{\chi}(\omega)$ is the complex susceptibility and ϵ_0 is the vacuum permittivity. This allows us to write an expression for the complex susceptibility, in terms of the complex polarisability:

$$
\begin{aligned}
\tilde{\chi}(\omega) &= \frac{n_e \tilde{\alpha}(\omega)}{\epsilon_0} \\
&= \left(\frac{n_e e^2}{m_e \epsilon_0} \right) \frac{1}{(\omega_0{}^2 - \omega^2) + i\omega\Gamma_d} \\
&\approx \left(\frac{n_e e^2}{m_e \epsilon_0} \right) \frac{1}{2\omega_0(\omega_0 - \omega) + i\omega\Gamma_d} \ ,
\end{aligned}
$$

where the approximation $\omega_0{}^2 - \omega^2 = (\omega_0 + \omega)(\omega_0 - \omega) \approx 2\omega_0(\omega_0 - \omega)$ has been made. Further simplifications can be made by noticing that, near resonance, $\omega \approx \omega_0$. This enables ω_0 to be factored out of the expression for the complex susceptibility. By defining $\tilde{\chi}_0 \equiv n_e e^2 / 2m_e \omega_0 \epsilon_0$ to be the complex-susceptibility *amplitude*, and by noting that $(\omega_0 - \omega) = \Delta\omega$ is the detuning:

$$
\begin{aligned}
\tilde{\chi}(\omega) &\approx \left(\frac{n_e e^2}{2m_e \omega_0 \epsilon_0} \right) \frac{1}{(\omega_0 - \omega) + i\frac{\Gamma_d}{2}} \\
&= \frac{\tilde{\chi}_0}{\Delta\omega + i\frac{\Gamma_d}{2}} \ .
\end{aligned}
$$

Now that an expression for the complex susceptibility has been derived, it enables the complex refractive index to be obtained through the relation

$$
\tilde{\eta}(\omega) = \sqrt{1 + \tilde{\chi}(\omega)} \approx 1 + \frac{\tilde{\chi}(\omega)}{2} \ , \tag{3.62}
$$

where an approximation has been made by truncating $\sqrt{1 + \tilde{\chi}(\omega)}$ to its first-order Taylor expansion[22].

In the presence of a non-zero magnetic field, the resonances are shifted by $\pm\Omega_L$, the Larmor frequency:

---

[22]A Taylor series, or Taylor expansion, is a representation of a function $f(x)$ as an infinite sum of terms which are calculated from the values of the function's derivatives at a single point a [270]:

$$
f(x) = \sum_{n=0}^{\infty} \frac{f^{(n)}(a)}{n!}(x - a)^n \ , \tag{3.63}
$$

where $f^{(n)}(a)$ denotes the n^{th} derivative of f, with respect to x, evaluated at the point a. Using the above definition, the Taylor series of $\sqrt{1 + \tilde{\chi}(\omega)}$ is given by

$$
\sqrt{1 + \tilde{\chi}} = 1 + \frac{\tilde{\chi}}{2} - \frac{\tilde{\chi}^2}{8} + \frac{\tilde{\chi}^3}{16} - \frac{5\tilde{\chi}^4}{128} + \mathcal{O}\left(\tilde{\chi}^5\right) \ , \tag{3.64}
$$

which, to first order in $\tilde{\chi}$, is simply:

$$
\sqrt{1 + \tilde{\chi}} \approx 1 + \frac{\tilde{\chi}}{2} \ . \tag{3.65}
$$

$$\Omega_{\mathrm{L}} = \frac{g_F \mu_{\mathrm{B}}}{\hbar} B \,, \tag{3.66}$$

where g_F is the Landé g-factor, μ_{B} is the Bohr magneton and $B = |\boldsymbol{B}|$ is the magnetic-field strength. The resonance frequency ω_0 then shifts to

$$\omega_0 \rightarrow \omega_0 \pm \Omega_{\mathrm{L}} \,. \tag{3.67}$$

This shift in resonance frequencies removes the degeneracy of the complex refractive indices, subsequently yielding two distinct refractive indices for each circular polarisation:

$$\tilde{\eta}_{\pm}(\omega) = 1 + \frac{\tilde{\chi}_0}{2} \frac{1}{(\Delta\omega \mp \Omega_{\mathrm{L}}) + i\frac{\Gamma_{\mathrm{d}}}{2}} \,. \tag{3.68}$$

Linearly polarised light can be thought of, and mathematically represented as, a superposition of right- and left-circularly polarised components. These components are denoted σ^+ and σ^-, respectively; and carry spin angular momentum of $-\hbar$ and $+h$, respectively. Given that the refractive indices, $\tilde{\eta}_{\pm}$, are now different, the circularly polarised components which constitute linearly polarised light will experience different refractive indices. The result of this is that one component of the light will propagate faster than the other, resulting in a relative phase shift between the two components at the output of the medium. This phase shift causes a rotation of the plane of polarisation of the linearly polarised light, which is given mathematically by

$$\phi = \frac{\omega l}{2c} \left(\mathrm{Re}\left\{ \tilde{\eta}_+ \right\} - \mathrm{Re}\left\{ \tilde{\eta} \right\} \right) \,, \tag{3.69}$$

where l is the length of active material, and c is the speed of light. By expressing Eq. (3.68) in the form $\tilde{\eta}_{\pm} = \alpha_{\pm} + i\beta_{\pm}$ and taking the real part, the difference in refractive indices is given by

$$\mathrm{Re}\left\{ \tilde{\eta}_+ \right\} - \mathrm{Re}\left\{ \tilde{\eta} \right\} = \frac{\tilde{\chi}_0 \Omega_{\mathrm{L}} \left(\Delta\omega^2 - \frac{\Gamma_{\mathrm{d}}^2}{4} - \Omega_{\mathrm{L}}^2 \right)}{\left(\Delta\omega^2 - \frac{\Gamma_{\mathrm{d}}^2}{4} - \Omega_{\mathrm{L}}^2 \right)^2 + (\Delta\omega \Gamma_{\mathrm{d}})^2} \,. \tag{3.70}$$

This can then substituted into the equation for the polarisation rotation, which yields

$$\phi = \frac{\tilde{\chi}_0 \omega l}{2c} \Omega_{\mathrm{L}} \frac{\left(\Delta\omega^2 - \frac{\Gamma_{\mathrm{d}}^2}{4} - \Omega_{\mathrm{L}}^2 \right)}{\left(\Delta\omega^2 - \frac{\Gamma_{\mathrm{d}}^2}{4} - \Omega_{\mathrm{L}}^2 \right)^2 + (\Delta\omega \Gamma_{\mathrm{d}})^2} \,. \tag{3.71}$$

Equation (3.71) essentially has two free variables: the detuning, $\Delta\omega$, and the Larmor frequency, Ω_{L}, which is directly proportional to the magnetic-field strength. The dependence of the polarisation rotation on these two variables is shown in Fig. 3.7. As seen in Fig. 3.7, the dependence of the polarisation rotation on the magnetic-field strength exhibits a resonance behaviour, in which there is a relatively large change in polarisation-rotation amplitude in response to a relatively small change in magnetic-field strength. This hints towards the idea that this effect may prove useful for sensitive magnetometry applications.
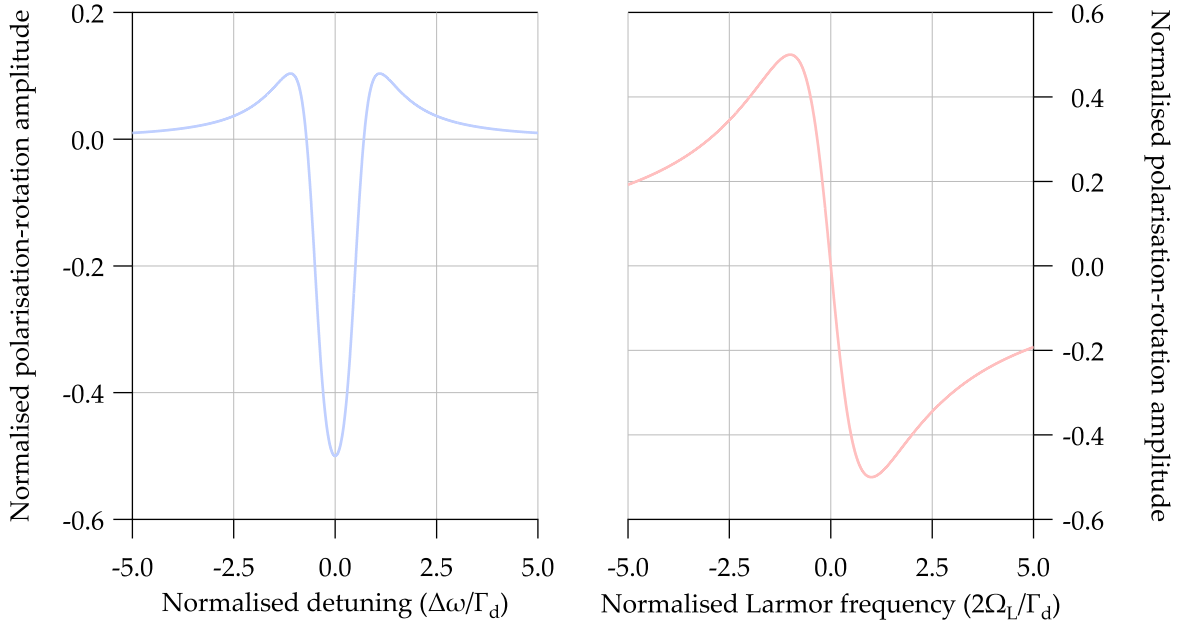
Figure 3.7: Dependence of the polarisation-rotation amplitude — due to the Macaluso-Corbino effect — on detuning (blue trace) and Larmor frequency (red trace), calculated using Eq. (3.71). The trace on the left was calculated for $2\Omega_L = \Gamma_d = 1$, while the trace on the right was calculated for $\Delta\omega = 0$.

Although the derivation of Eq. (3.71) does not take into account Doppler broadening (cf. Sec. 3.2.3), the predicted features are *qualitatively* the same as those seen in experiments in which the atoms are undergoing thermal motion. The features observed in the Macaluso-Corbino effect are a stark contrast to those observed in the traditional Faraday effect in solids. In the traditional Faraday effect in solids, the amount of polarisation rotation is *weakly* dependent upon the optical frequency. However, as shown in Fig. 3.7, the Macaluso-Corbino effect exhibits a *strong* dependence on the optical detuning (i.e. the optical frequency *near* the resonant frequency). Another major difference between the two effects is that, in the linear Faraday effect, the amount of optical rotation is *linearly* dependent upon the magnetic-field strength. In the case of the Macaluso-Corbino effect however, there is a strong resonant behaviour — described by the imaginary part of a complex Lorentzian profile — in which only a small range around zero magnetic field yields a linear response[23].

3.6 Nonlinear Magneto-Optical Rotation

The physical outcome of nonlinear magneto-optical rotation is *essentially* the same as both that of the Macaluso-Corbino effect, and the traditional Faraday effect. That is, the plane of polarisation of linearly polarised light is rotated due to circular birefringence. However, there is one major difference. Nonlinear magneto-optical rotation, as the name suggests[24],

---

[23]In experiments, this region of linear response is typically determined by the Doppler width [199,245,246], which is $\mathcal{O}\,(500\,\text{MHz})$ at room temperature (cf. Fig. 3.3).

[24]By definition, a nonlinear optical process is one in which the optical properties of the medium are dependent upon the intensity of the traversing light [203].

Figure 3.8: Illustration of the creation of ground-state coherence between Zeeman states with $|\Delta m_F| = 2$ using linearly polarised light propagating along the quantisation axis in the $F = 1 \rightarrow F' = 0$ hyperfine transition.

is a nonlinear process and is therefore strongly dependent upon *both* the intensity of the traversing light and the magnetic-field strength.

There are a number of underlying physical processes which can give rise to nonlinear magneto-optical rotation. Under certain circumstances, one of the major contributions to the NMOR signal is due to the velocity-selective optical pumping that occurs when a narrowband laser is tuned to a Doppler-broadened transition. This optical pumping gives rise to 'peaks' and 'holes' in the velocity distribution due to the redistribution of atomic population [199,203,237,271]. This population redistribution (known as the Bennett effect [271]) changes the optical properties of the medium, thereby giving rise to polarisation rotation [199,203,237]. Polarisation rotation induced in this way is known as Bennett-structure NMOR [199,203,237]. Given that this effect relies on velocity-selective optical pumping, it is more dominant in uncoated cells, or those which contain a buffer gas (due to the restricted mean free path). However, in anti-relaxation-coated cells — such as the ones used in this thesis — atoms can undergo *many* velocity-changing collisions without losing polarisation, which tends to diminish NMOR signals arising from the Bennett effect [237].

Although there are multiple contributions to NMOR signals, the narrowest — and hence most useful in the context of magnetometry — resonances are related to ground-state coherence[25] [272]. Optical pumping with linearly polarised light generates quantum coherence between ground-state Zeeman sublevels which satisfy the condition $|\Delta m_F| = 2$ (cf. Fig. 3.8). When exposed to a weak longitudinal magnetic field, the energy levels of the Zeeman states are shifted by the Larmor frequency. This energy splitting induces temporal evolution (i.e. a quantum beat[26]) of the $|\Delta m_F| = 2$ coherences at a frequency of $2\Omega_L$[27], in accordance with the time-dependent Schrödinger equation. Given that the optical-rotation observable is pro-

---

[25]Nonlinear magneto-optical resonances associated with ground-state coherence have demonstrated widths (FWHM) as low as $\Gamma/2\pi = 1\,\text{Hz}$ [236], whereas the smallest achievable linewidth for Bennett-structure NMOR is limited by the natural linewidth of the optical transition [203,237] — typically $\Gamma/2\pi = 1$–$10\,\text{MHz}$ [203,237].

[26]A quantum beat is the temporal evolution of a coherent superposition of nondegenerate energy eigenstates, at a frequency determined by the energy splitting [254].

[27]This temporal evolution is also known as Larmor precession.
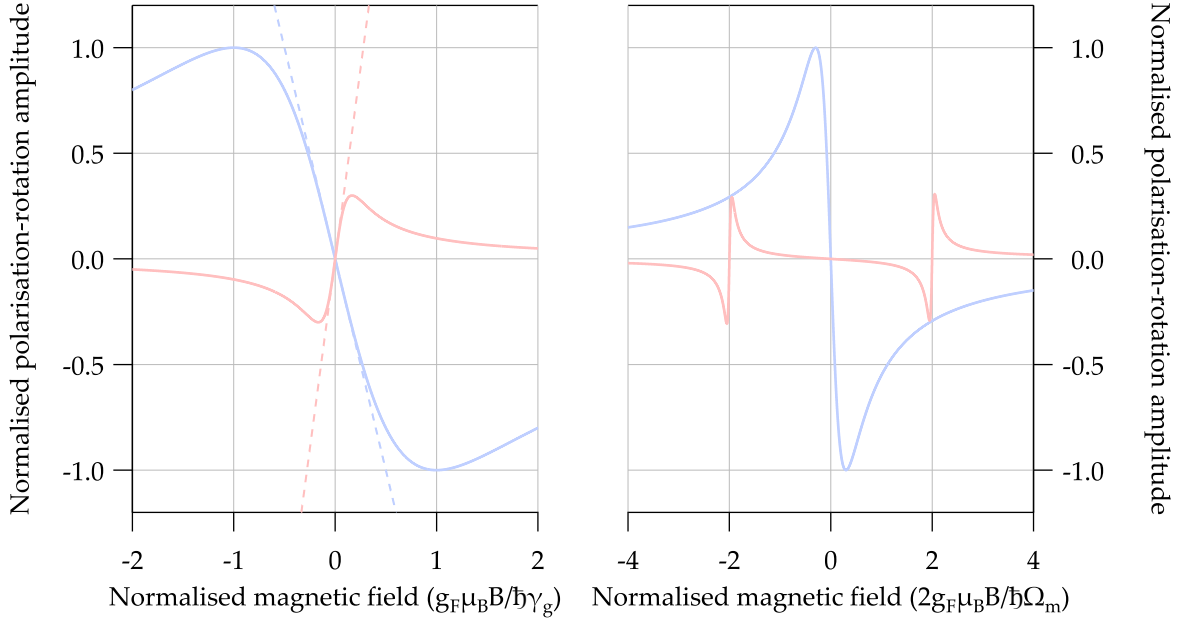
Figure 3.9: Left: Qualitative comparison between the complex-Lorentzian resonances obtained via the Macaluso-Corbino effect (blue) and cw NMOR (red). In the case of the Macaluso-Corbino effect, the peak polarisation rotation is observed when $B = \hbar\gamma_g/g_F\mu_B$. The slope of the cw NMOR resonance (red, dashed) is, in general, greater than that of the Macaluso-Corbino effect (blue, dashed). Right: Qualitative comparison between the complex-Lorentzian resonances obtained via the Macaluso-Corbino effect (blue) and modulated NMOR (red). The Macaluso-Corbino effect occurs only for a small range of fields around $B \approx 0$, whereas modulated NMOR resonances are observed when $B = \pm\hbar\Omega_m/2g_F\mu_B$.

portional to the sum of the ground-state coherences (cf. Sec. 4.8), the optical-rotation signal oscillates sinusoidally at $2\Omega_L$.

In the regime where the quantum-beat frequency is lower than the relaxation rate (i.e. $2\Omega_L/2\pi \lesssim T_2^{-1}$), each polarised particle will be unable to undergo an entire quantum-beat cycle before relaxing. The quantum beats from various atoms will therefore be unable to completely destructively interfere, leaving a finite residual steady-state polarisation. It is this steady-state polarisation that yields anisotropy (circular birefringence) in the medium, which subsequently gives rise to the rotation of polarisation of a weak, linearly polarised probe beam — i.e. nonlinear magneto-optical rotation.

However; if the magnetic field is too strong (i.e. $2\Omega_L/2\pi \gtrsim T_2^{-1}$), each particle undergoes an entire quantum-beat cycle (or more) before relaxing, and hence continuous optical pumping tends to 'wash out' the steady-state polarisation via destructive interference in this regime. For this reason, cw NMOR can only be observed in a small range of magnetic-field strength such that $B \lesssim \hbar\gamma_g/g_F\mu_B$ [150,199,236,245]. In order to extend the dynamic range of NMOR, modulation techniques must be implemented.

3.7   Nonlinear Magneto-Optical Rotation with Modulated Light

As eluded to in the previous section, if the magnetic-field strength is too large (i.e. $2\Omega_{\mathrm{L}}/2\pi \gtrsim T_2^{-1}$), the amplitude of the cw NMOR signal rapidly vanishes due to destructive interference. However; by modulating the pumping process — through either frequency [150,204,235], amplitude [150,202,206,207] or polarisation [150] modulation — it is possible to overcome this limitation. If the atoms are optically pumped approximately once every oscillation cycle, i.e. $\Omega_{\mathrm{m}} \approx 2\Omega_{\mathrm{L}}$, atoms which are pumped in one cycle will be in phase with atoms which were pumped in previous cycles, and those that will be pumped in future cycles. This yields a resonance condition, in which the polarisation is maximised when $\Omega_{\mathrm{m}} = 2\Omega_{\mathrm{L}}$[28]. This technique is known as 'synchronous optical pumping' [150] or 'optically driven spin precession' [150,273], and dramatically increases the dynamic range of this measurement technique by removing the constraint that $B \lesssim \hbar\gamma_{\mathrm{g}}/g_F\mu_{\mathrm{B}}$. A qualitative comparison between the complex-Lorentzian resonances obtained via the Macaluso-Corbino effect, cw NMOR, and modulated NMOR, is presented in Fig. 3.9.

In principle the dynamic range when using modulated pumping should have no limit; however, complications at high fields (such as the nonlinear Zeeman effect and alignment-to-orientation conversion [274]) may provide practical upper limits.

In order to adequately model and describe NMOR, the density matrix formalism must be utilised. The following chapter will introduce the density matrix formalism, and use it to model NMOR signals.

---

[28]In fact, it is possible to observe resonances at any subharmonic of the quantum-beat frequency, i.e. $\Omega_{\mathrm{m}} = 2\Omega_{\mathrm{L}}/n$, where n is the harmonic order (cf. Sec. 6.1.1.3). In the context of magnetometry however, the magnetic-field sensitivity is maximised when $\Omega_{\mathrm{m}} = 2\Omega_{\mathrm{L}}$, as opposed to any other harmonic order.

Figure 3.10: Energy levels of the D_1 and D_2 transitions for the ^{87}Rb (left) and ^{85}Rb (right) atoms, using data presented in Ref. [229]. For each F hyperfine state, the number of corresponding Zeeman states is labelled in red. The splitting in each hyperfine manifold is drawn to scale.

The Theory of Light-Atom Interactions

This chapter introduces the mathematical framework used to model light-atom interactions; beginning with a brief discussion of 'Schrödinger quantum mechanics', and then delving into the density matrix formalism. The $F = 1 \rightarrow F' = 0$ transition will be considered for simplicity[1]; however, more complicated transitions can be readily modelled by adapting the content discussed in this chapter. Derivations given here closely follow those presented in Ref. [242], though are repeated here in standard international units. Long derivations have been excluded from this chapter in order to maintain readability, but can be found in full detail in App. A.

4.1   The Schrödinger Formalism

Consider a quantum particle (or system of particles) whose quantum state at a position $\boldsymbol{r}$, and at a time t, can be described by a wavefunction $\psi\left(\boldsymbol{r}, t\right) \in \mathbb{C}$. The wavefunction is a complex-valued probabiltiy *amplitude*, the squared modulus of which, $\left|\psi\left(\boldsymbol{r}, t\right)\right|^{2} = \psi^{*}\left(\boldsymbol{r}, t\right)\psi\left(\boldsymbol{r}, t\right) \in \mathbb{R}$, describes the probability *density* of a particle being measured at a given position and at a given time. The temporal evolution of the quantum state is governed by the time-dependent Schrödinger equation, and can be calculated from its wavefunction via [221,222,224–226,242,275–282]

$$i\hbar \frac{\partial}{\partial t} \left|\psi\left(\boldsymbol{r}, t\right)\right\rangle = \mathscr{H} \left|\psi\left(\boldsymbol{r}, t\right)\right\rangle , \tag{4.1}$$

where $\mathscr{H}$ is the Hamiltonian of the system, and $\hbar$ is the reduced Planck constant. The Hamiltonian is an operator which corresponds to the total energy of the system, and will typically be separated into two (or more) independent components in most cases [242]:

$$\mathscr{H} = \mathscr{H}_0 + \mathscr{H}_{\mathrm{I}}, \tag{4.2}$$

where $\mathscr{H}_0$ is the unperturbed Hamiltonian, and $\mathscr{H}_{\mathrm{I}}$ is the interaction Hamiltonian. The unperturbed Hamiltonian, as its name suggests, describes the system *without* the presence of external influences (e.g. the energy levels of an atom). The interaction Hamiltonian therefore describes any *external* forces applied to the system (e.g. an electric or magnetic field).

In order to solve the time-dependent Schrödinger equation, one method is to assume that solutions are 'seperable', i.e. that the temporal and spatial dependence of the wavefunction

---

[1]The $F = 1 \rightarrow F' = 0$ hyperfine transition is the simplest transition that will support $\left|\Delta m_F\right| = 2$ ground-state coherences, which give rise to nonlinear magneto-optical rotation.

can be separated into two independent factors, in which case the wavefunction takes the form [279]

$$\psi\left(\boldsymbol{r},t\right) = \sum_n C_n U_n\left(\boldsymbol{r}\right) e^{-i\omega_n t}, \tag{4.3}$$

where C_n is the probability amplitude (normalised such that $\sum_n |C_n|^2 = 1$ [279]) for the quantum state $U_n\left(\boldsymbol{r}\right)$, which is a stationary solution (eigenfunction) of the time-independent Schrödinger equation with an energy E [279]:

$$\mathscr{H}\left|U_n\left(\boldsymbol{r}\right)\right\rangle = E\left|U_n\left(\boldsymbol{r}\right)\right\rangle. \tag{4.4}$$

In most circumstances, each quantum particle must be described by its own individual wavefunction. In principle it should be possible to keep track of and perform calculations on *each* wavefunction in an atomic vapour; however, in practice this is *immensely* computationally expensive and is therefore unfeasible in the majority of cases.

4.2   The Density Matrix Formalism

As briefly discussed in the previous section, keeping track of and performing calculations on each individual atomic wavefunction in a warm vapour is *immensely* computationally expensive. It is therefore incredibly useful to describe the system not by wavefunctions, but rather by a 'density matrix' which describes the *statistical state* of the quantum system.

If a fraction P_j of N atoms is in state ψ_j at time t, then the mixed-state density matrix for the entire ensemble is given by [225,275,276]

$$\rho = \sum_{j=1}^{N} P_j\left|\psi_j\right\rangle\left\langle\psi_j\right|. \tag{4.5}$$

Consider the $F = 1 \rightarrow F' = 0$ transition. This specific transition consists of three Zeeman sublevels in the $F = 1$ ground state, and one Zeeman sublevel in the $F' = 0$ excited state. The density matrix describing this atomic transition is therefore given explicitly by

$$\rho = \begin{pmatrix} \rho_{1,1} & \rho_{1,0} & \rho_{1,-1} & \rho_{1,0'} \\ \rho_{0,1} & \rho_{0,0} & \rho_{0,-1} & \rho_{0,0'} \\ \rho_{-1,1} & \rho_{-1,0} & \rho_{-1,-1} & \rho_{-1,0'} \\ \rho_{0',1} & \rho_{0',0} & \rho_{0',-1} & \rho_{0',0'} \end{pmatrix}. \tag{4.6}$$

The elements on the diagonal of the density matrix correspond to the populations of their respective quantum states, while the off-diagonal elements correspond to 'coherences'. Note that the coherences are related to each other symmetrically via $\rho_{i,j} = \rho_{j,i}^*$, while the populations are constrained by the normalisation condition $\mathrm{tr}\left(\rho\right) = 1$ [225].

4.3 The Liouville Equation

The temporal evolution of the density matrix is calculated using the equation of motion of the system, which is given elegantly by the commutation of the Hamiltonian and the density matrix (cf. App. A.1) [225]:

$$\frac{d\rho}{dt} = -\frac{i}{\hbar} \left[\mathscr{H}, \rho \right] .$$

(4.7)

Equation (4.7) is known as the Liouville equation, and is the density-matrix analogue of the time-dependent Schrödinger equation which governs the temporal evolution of wavefunctions [279]. As it stands, the Hamiltonian does not incorporate *any* relaxation or repopulation processes and hence neither does the Liouville equation. If relaxation processes are to be modelled they must be incorporated *phenomenologically* via additional terms appended to the Liouville equation. These phenomenological terms can be separated into two distinct categories: relaxation terms which account for the decay of atomic populations and coherences, and repopulation terms which describe the transfer of populations and coherences into other states.

Relaxation processes are incorporated into the Liouville equation through the appendage of a diagonal relaxation matrix, $\hat{\Gamma}$, whose elements are given by $\langle \rho_n | \hat{\Gamma} | \rho_m \rangle = \Gamma_d \delta_{nm}$ [242,276,283]. This matrix implicitly assumes that the relaxation processes are exponential in time[2], with a relaxation rate given by Γ_d. Typical examples of relaxation processes are spontaneous decay from an excited state to a ground state, and transit relaxation due to atoms leaving the experimental volume (cf. App. L.5). As the populations and coherences decay, an additional term must be appended to the Liouville equation to ensure proper normalisation of the density matrix. This term, a diagonal matrix denoted by Λ, describes the repopulation of states. Similar to the case of the relaxation processes, typical repopulation processes are transit repopulation (from atoms entering the experimental volume), and repopulation of the ground states due to spontaneous decay from the excited states. After appending these two phenomenological terms to Eq. (4.7), the modified Liouville equation is obtained [242,283]:

$$\frac{d\rho}{dt} = -\frac{i}{\hbar} \left[\mathscr{H}, \rho \right] - \frac{1}{2} \left\{ \hat{\Gamma}, \rho \right\} + \Lambda ,$$

(4.8)

where $\left\{ \hat{\Gamma}, \rho \right\} = \hat{\Gamma}\rho + \rho\hat{\Gamma}$ denotes the anti-commutation of the relaxation matrix with the density matrix.

4.4 The Hamiltonian

The total Hamiltonian, $\mathscr{H}$, as mentioned in Sec. 4.1, consists of two separate terms: an unperturbed atomic Hamiltonian, $\mathscr{H}_0$, and an interaction Hamiltonian, $\mathscr{H}_I$. The atomic Hamil-

---

[2]In the case of transit relaxation, this turns out to be a good approximation. However, there is some discrepancy between the assumed exponential decay, and that calculated using a Monte Carlo simulation. Refer to App. L.5 for further details.

tonian is used to describe the atomic energy levels, whereas the interaction Hamiltonian is used to describe the atomic interaction with fields and can be further separated into two (or more) components:

$$\mathscr{H}_{\mathrm{I}} = \mathscr{H}_{\mathrm{E}} + \mathscr{H}_{\mathrm{B}} \,, \tag{4.9}$$

where $\mathscr{H}_{\mathrm{E}}$ is the electric-field Hamiltonian (usually an optical field), and $\mathscr{H}_{\mathrm{B}}$ is the magnetic-field Hamiltonian.

4.4.1   The Unperturbed Hamiltonian

The role of the unperturbed Hamiltonian is to specify the relative energies of the Zeeman sublevels for both the ground and excited states. By defining the energy of the lower state to be zero, and the transition frequency of the excited state to be ω_0, the unperturbed atomic Hamiltonian for the $F = 1 \to F' = 0$ transition is given by

$$\mathscr{H}_0 = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \hbar\omega_0 \end{pmatrix} \,, \tag{4.10}$$

where $\hbar$ is the reduced Planck constant.

4.4.2   The Electric-Field Hamiltonian

The electric-field Hamiltonian is simply the potential energy of an electric dipole in an electric field, which is given explicitly by $\mathscr{H}_{\mathrm{E}} = -\boldsymbol{d} \cdot \boldsymbol{\mathcal{E}}$, where $\boldsymbol{d}$ is the dipole operator and $\boldsymbol{\mathcal{E}}$ is the electric-field vector. Even in the case of no applied external electric field, $\mathscr{H}_{\mathrm{E}}$ is still included in the interaction Hamiltonian as it is also used to describe the interaction with an *optical* field.

In the case of a linearly polarised optical field along the x-axis, the electric field vector $\boldsymbol{\mathcal{E}}$ can be written as

$$\boldsymbol{\mathcal{E}} = \mathcal{E}_0 \cos\left(\omega t\right) \hat{\boldsymbol{x}} \,, \tag{4.11}$$

where $\mathcal{E}_0$ is the electric-field amplitude, ω is the optical (angular) frequency, and $\hat{\boldsymbol{x}}$ is the unit vector along the x-axis. Evaluating the electric-field Hamiltonian for the electric-field vector defined in Eq. (4.11) yields

$$
\begin{aligned}
\mathscr{H}_{\mathrm{E}} &= -\boldsymbol{d} \cdot \boldsymbol{\mathcal{E}} && \text{(4.12)} \\
&= -\mathcal{E}_0 \cos\left(\omega t\right) d_x && \text{(4.13)} \\
&= -\frac{1}{\sqrt{2}} \mathcal{E}_0 \cos\left(\omega t\right) \left(d_{-1} - d_1\right) \,, && \text{(4.14)}
\end{aligned}
$$

where d_{-1} and d_1 are the two spherical-basis components of the dipole operator (cf. App. A.2). The matrix elements for the two spherical-basis components of the dipole operator, d_{-1} and d_1, can be written using the Wigner-Eckart theorem (cf. App. A.3) — specifically Eqs. (A.10) and (A.11). Upon invoking the Wigner-Eckart theorem, the two spherical-basis components of the dipole operator for the $F = 1 \rightarrow F' = 0$ transition are given explicitly in matrix form by

$$d_{-1} \;=\; \langle \xi 1 \| d \| \xi' 0' \rangle \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ -1 & 0 & 0 & 0 \end{pmatrix}, \tag{4.15}$$

$$d_{1} \;=\; \langle \xi 1 \| d \| \xi' 0' \rangle \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \end{pmatrix}, \tag{4.16}$$

where $\langle \xi 1 \| d \| \xi' 0' \rangle$ is the reduced dipole matrix element for the $F = 1 \rightarrow F' = 0$ transition, which is an experimentally measured quantity. Substituting Eqs. (4.15) and (4.16) into Eq. (4.14) yields the matrix representation of the electric-field Hamiltonian:

$$\mathscr{H}_{\mathrm{E}} = \frac{\hbar \Omega_{\mathrm{R}} \cos\left(\omega t \right)}{\sqrt{2}} \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 1 & 0 & -1 & 0 \end{pmatrix}, \tag{4.17}$$

where $\Omega_{\mathrm{R}} = \langle \xi 1 \| d \| \xi' 0' \rangle \mathcal{E}_0 / \hbar$ is the Rabi frequency for the $F = 1 \rightarrow F' = 0$ transition.

4.4.3 The Magnetic-Field Hamiltonian

The magnetic-field Hamiltonian, in analogy with the electric-field Hamiltonian described in Sec. 4.4.2, is simply the potential energy of a magnetic dipole in a magnetic field. The magnetic potential energy is given explicitly by $\mathscr{H}_{\mathrm{B}} = \boldsymbol{\mu} \cdot \boldsymbol{B}$, where $\boldsymbol{\mu}$ is the magnetic dipole moment, and $\boldsymbol{B}$ is the magnetic field. Assuming that the magnetic field is weak (such that there is negligible contribution from the nonlinear Zeeman effect — cf. Sec. 3.3), and is directed along $\hat{z}$ (the designated quantisation axis), the magnetic-field Hamiltonian is given by

$$\mathscr{H}_{\mathrm{B}} \;=\; \boldsymbol{\mu} \cdot \boldsymbol{B} \tag{4.18}$$

$$=\; m_F g_F \mu_{\mathrm{B}} \boldsymbol{F} \cdot \boldsymbol{B} \tag{4.19}$$

$$=\; m_F \hbar \Omega_{\mathrm{L}} F_z \,, \tag{4.20}$$

where $\Omega_{\mathrm{L}} = g_F \mu_{\mathrm{B}} B / \hbar$ is the Larmor frequency, m_F is the magnetic quantum number, g_F is the Landé g-factor, and μ_{B} is the Bohr magneton. This can then be written in matrix form for the $F = 1 \rightarrow F' = 0$ transition via

$$
\mathscr{H}_{\mathrm{B}} = \hbar\Omega_{\mathrm{L}} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} .
\tag{4.21}
$$

As the magnetic-field Hamiltonian is a diagonal matrix, its effect is to shift the energy levels by the Larmor frequency. At low magnetic-field strength, this shift is linear (the linear Zeeman effect); however, at high magnetic-field strength, the shift becomes quadratic (the nonlinear Zeeman effect). The Zeeman effect is discussed in further detail in Sec. 3.3.

4.4.4 The Total Hamiltonian

Now that all three Hamiltonian contributions have been calculated — the unperturbed Hamiltonian, the electric-field Hamiltonian, and the magnetic-field Hamiltonian — it is possible to calculate the total Hamiltonian. By summing the components in their respective matrix forms, the total Hamiltonian is given by

$$
\mathscr{H} = \begin{pmatrix} \hbar\Omega_{\mathrm{L}} & 0 & 0 & \frac{\hbar\Omega_{\mathrm{R}}}{\sqrt{2}}\cos(\omega t) \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -\hbar\Omega_{\mathrm{L}} & -\frac{\hbar\Omega_{\mathrm{R}}}{\sqrt{2}}\cos(\omega t) \\ \frac{\hbar\Omega_{\mathrm{R}}}{\sqrt{2}}\cos(\omega t) & 0 & -\frac{\hbar\Omega_{\mathrm{R}}}{\sqrt{2}}\cos(\omega t) & \hbar\omega_0 \end{pmatrix} .
\tag{4.22}
$$

4.5 The Rotating-Wave Approximation

By inspecting the Hamiltonian given by Eq. (4.22), it is clear that there is temporal dependence at the optical frequency in the off-diagonal elements. This oscillation is vital in order to observe the resonant behaviour of the system; however, solving the Liouville equation for this Hamiltonian would require *incredibly* high temporal resolution, which would come at the cost of great computational expense. In order to conserve the resonant behaviour of the system, while removing the rapid oscillatory terms from the calculations, a unitary transform $\mathscr{H} \rightarrow \widetilde{\mathscr{H}}$ is performed which transforms the Hamiltonian to the reference frame of the light. The rotating-wave approximation, discussed in detail in App. A.4, yields the following static rotating-wave Hamiltonian:

$$
\widetilde{\mathscr{H}} = \begin{pmatrix} \hbar\Omega_{\mathrm{L}} & 0 & 0 & \frac{\hbar\Omega_{\mathrm{R}}}{2\sqrt{2}} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -\hbar\Omega_{\mathrm{L}} & -\frac{\hbar\Omega_{\mathrm{R}}}{2\sqrt{2}} \\ \frac{\hbar\Omega_{\mathrm{R}}}{2\sqrt{2}} & 0 & -\frac{\hbar\Omega_{\mathrm{R}}}{2\sqrt{2}} & -\hbar\Delta\omega \end{pmatrix} ,
\tag{4.23}
$$

where $\Delta\omega = \omega_0 - \omega$ is the optical detuning. The resultant Hamiltonian under the rotating-wave approximation is both simpler and easier (in terms of computational complexity) to solve. Once solved, the density-matrix elements in the laboratory frame, $\rho_{i,j}$, can be calculated from the rotating-frame density-matrix elements, $\widetilde{\rho}_{i,j}$, via $\rho = U_t \widetilde{\rho} U_t^{\dagger}$, where U_t is the unitary transform matrix given by Eq. (A.12). Explicitly, this is given by

$$
\rho = \begin{pmatrix}
\widetilde{\rho}_{1,1} & \widetilde{\rho}_{1,0} & \widetilde{\rho}_{1,-1} & e^{i\omega t}\widetilde{\rho}_{1,0'} \\
\widetilde{\rho}_{0,1} & \widetilde{\rho}_{0,0} & \widetilde{\rho}_{0,-1} & e^{i\omega t}\widetilde{\rho}_{0,0'} \\
\widetilde{\rho}_{-1,1} & \widetilde{\rho}_{-1,0} & \widetilde{\rho}_{-1,-1} & e^{i\omega t}\widetilde{\rho}_{-1,0'} \\
e^{-i\omega t}\widetilde{\rho}_{0',1} & e^{-i\omega t}\widetilde{\rho}_{0',0} & e^{-i\omega t}\widetilde{\rho}_{0',-1} & \widetilde{\rho}_{0',0'}
\end{pmatrix} .
\tag{4.24}
$$

4.6 Relaxation and Repopulation

As discussed in Sect. 4.4, relaxation and repopulation processes must be added phenomenologically to the density matrix via the modified Liouville equation. Relaxation processes are incorporated via the relaxation matrix, $\hat{\Gamma}$, while repopulation processes incorporated via the repopulation matrix, Λ.

4.6.1 The Relaxation Matrix

Under typical experimental conditions, there are two main ways in which the atomic population can relax: via transit relaxation due to atoms leaving the beam, and spontaneous decay of population from the excited state back down to the ground state. These two processes are characterised by the decay rates γ_t and Γ_d, respectively. The relaxation matrix is therefore given by

$$
\hat{\Gamma} = \begin{pmatrix}
\gamma_t & 0 & 0 & 0 \\
0 & \gamma_t & 0 & 0 \\
0 & 0 & \gamma_t & 0 \\
0 & 0 & 0 & \gamma_t + \Gamma_d
\end{pmatrix} .
\tag{4.25}
$$

4.6.2 The Repopulation Matrix

Just as the atoms in each state are *exiting* the beam at a rate of γ_t (i.e. transit *relaxation*), there are also new atoms *entering* the beam via the same mechanism (i.e. transit *repopulation*). These 'new' atoms are assumed to be unpolarised (i.e. no coherence between states), and equally distributed amongst the Zeeman sublevels of the hyperfine ground state. This implies that the matrix describing transit repopulation is diagonal, with each ground state increasing in population at a rate of $\gamma_t/3$:

$$
\frac{d}{dt}\begin{pmatrix}
\rho_{1,1} & \rho_{1,0} & \rho_{1,-1} & \rho_{1,0'} \\
\rho_{0,1} & \rho_{0,0} & \rho_{0,-1} & \rho_{0,0'} \\
\rho_{-1,1} & \rho_{-1,0} & \rho_{-1,-1} & \rho_{-1,0'} \\
\rho_{0',1} & \rho_{0',0} & \rho_{0',-1} & \rho_{0',0'}
\end{pmatrix} = \frac{1}{3}\begin{pmatrix}
\gamma_t & 0 & 0 & 0 \\
0 & \gamma_t & 0 & 0 \\
0 & 0 & \gamma_t & 0 \\
0 & 0 & 0 & 0
\end{pmatrix} .
\tag{4.26}
$$

In addition to transit repopulation, there is also repopulation of ground states and coherences due to spontaneous emission from the excited states. The rate of change of the ground-state density-matrix elements due to spontaneous emission is a random process, which is caused by fluctuations of the vacuum field [275]. Using results derived from quantum electrodynamics, spontaneous emission can be represented by a spontaneous-emission operator F_{mn}^{sr} [283]. A detailed derivation, and subsequent calculation of, the spontaneous-emission operator is discussed in App. A.5, yielding the following ground-state density-matrix derivatives for the $F = 1 \rightarrow F' = 0$ hyperfine transition:

$$\frac{\mathrm{d}}{\mathrm{d}t} \begin{pmatrix} \rho_{1,1} & \rho_{1,0} & \rho_{1,-1} \\ \rho_{0,1} & \rho_{0,0} & \rho_{0,-1} \\ \rho_{-1,1} & \rho_{-1,0} & \rho_{-1,-1} \end{pmatrix} = \Gamma_{\mathrm{d}} \rho_{0',0'} \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}. \tag{4.27}$$

Combining Eqs. (4.26) and (4.27) now yields a total repopulation matrix Λ for the $F = 1 \rightarrow F' = 0$ transition:

$$\Lambda = (\gamma_{\mathrm{t}} + \Gamma_{\mathrm{d}} \rho_{0',0'}) \begin{pmatrix} \frac{1}{3} & 0 & 0 & 0 \\ 0 & \frac{1}{3} & 0 & 0 \\ 0 & 0 & \frac{1}{3} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{4.28}$$

This is now all that is needed in order to solve the Liouville equation for a given set of parameters.

4.7   Optical Pumping

Perhaps the simplest density-matrix calculation that can be performed, is the redistribution of atomic population amongst the Zeeman sublevels (both the ground and excited states) due to optical pumping by the incident light. The population of the Zeeman states can be calculated by solving the modified Liouville equation — given by Eq. (4.8) — for either the steady-state or time-dependent solution.

Presented in Fig. 4.1 are the steady-state and time-dependent ground-state atomic populations, upon being optically pumped with light that is linearly polarised along $\hat{x}$ and propagating parallel to the magnetic field along $\hat{z}$. The parameters used in this calculation are $\Gamma_{\mathrm{d}} = 2\pi \times 1$, $\gamma_{\mathrm{t}} = 2\pi \times 0.05$, $\Omega_{\mathrm{L}} = 2\pi \times 0.05$, and $\Omega_{\mathrm{R}} = 2\pi \times 0.5$; with the initial condition that $\rho_{i,i} = 1/3$ for all ground-state Zeeman sublevels, and $\rho_{i,j} = 0$ for all i, j (i.e. the atomic vapour is unpolarised and in thermal equilibrium). In the case of the steady-state solution, the detuning is varied; however, for the time-dependent solution, the detuning is fixed at $\Delta\omega = 0$.

It can be seen by inspection of the steady-state population distribution presented in Fig. 4.1 that, when pumping on resonance, *most* of the population ends up in the $|m_F = 0\rangle$ ground state. This is a result of the incident linearly polarised light coupling the $|m_F = \pm 1\rangle$ ground states to the $|m_F' = 0\rangle$ excited state (cf. Fig. 3.8), and effectively pumping the popu-

Figure 4.1: Populations of the $F = 1$ ground-state Zeeman sublevels in the absence of Doppler broadening, calculated using the modified Liouville equation given by Eq. (4.8), for the rotating-wave Hamiltonian given by Eq. (4.23). The $|m_F = 0\rangle$ state is shown in blue, while the $|m_F = -1\rangle$ and $|m_F = 1\rangle$ states are shown in red and purple, respectively. Left: Zeeman-state populations as a function of normalised time for $\Delta\omega = 0$. Right: Steady-state population distribution as a function of normalised optical detuning. The parameters used in both calculations are discussed in Sec. 4.7.

lation *out* of these states and *into* the $|m_F = 0\rangle$ state which is not coupled by light.

Optically pumping an initially unpolarised atomic ensemble with linearly polarised light leads to the generation of atomic alignment (a rank-2 polarisation moment) along the axis of polarisation [242,283] (cf. App. M.3.2). This is readily seen by visualising the density matrix using angular momentum probability surfaces [242,283,284] (cf. App. M.3). Figure 4.2 shows the temporal evolution of the angular momentum distribution corresponding to the $F = 1$ ground-state density matrix presented in Fig. 4.1. Evident in Fig. 4.2 is a transition from an initially unpolarised atomic ensemble at $t/\gamma_t^{-1} = 0$ — characterised by a spherically symmetric angular momentum distribution — to a state which is aligned along a preferred axis. This atomic alignment is the result of $|\Delta m_F| = 2$ coherences in the density matrix [242,245,257,283,285,286] (cf. Sec. M.3.3) and, as will be shown in future sections, gives rise to modulation of light-field parameters.

4.8 Observables

As the incident light traverses the atomic vapour, not *only* is there a redistribution of population between the ground-state Zeeman sublevels (i.e. optical pumping) and generation of ground-state coherence(s), but the light-field parameters in turn are affected. This can be considered as a three-stage process [245]. The first stage involves the incident light mod-

Figure 4.2: Angular momentum probability surfaces of an initially unpolarised $F = 1$ ground state (at $t/\gamma_t^{-1} = 0$), that is optically pumped into an aligned state via light which is linearly polarised along $\hat{x}$ and propagating along $\hat{z}$. These surfaces correspond to the data presented in Fig. 4.1, and were calculated via Eq. (M.7) using the solution to Eq. (4.8) for the parameters discussed in Sec. 4.7. Note that the atomic alignment axis is slightly rotated from $\hat{x}$ in the xy-plane due to the presence of a non-zero magnetic field (i.e. $\Omega_L \neq 0$) directed along $\hat{z}$.

ifying the optical properties of the atomic ensemble. During the second stage, the optical properties of the medium undergo temporal evolution as a result of interaction with external electric and magnetic fields. Finally, the altered (and evolving) optical properties of the medium give rise to changes of the incident light itself. In reality all three of these stages will be occurring simultaneously, giving rise to interesting nonlinear phenomena [245].

The light-field parameters that can be affected by the atomic vapour include the electric-field amplitude $\mathcal{E}_0$, optical phase φ, optical rotation ϕ, and optical ellipticity ϵ. The changes in these light-field parameters are derived using Maxwell's equations, yielding the general result (cf. App. A.6):

$$\frac{1}{\mathcal{E}_0}\frac{d\mathcal{E}_0}{dl} = \frac{k}{2\mathcal{E}_0\epsilon_0}\left[\sin\phi\left(\cos\epsilon P_4 - \sin\epsilon P_1\right) + \cos\phi\left(\cos\epsilon P_2 + \sin\epsilon P_3\right)\right], \qquad (4.29)$$

$$\frac{d\phi}{dl} = \frac{k}{2\mathcal{E}_0\epsilon_0}\sec 2\epsilon\left[\cos\phi\left(\sin\epsilon P_1 + \cos\epsilon P_4\right) - \sin\phi\left(\cos\epsilon P_2 - \sin\epsilon P_3\right)\right], \qquad (4.30)$$

$$\frac{d\epsilon}{dl} = -\frac{k}{2\mathcal{E}_0\epsilon_0}\left[\sin\phi\left(\cos\epsilon P_1 + \sin\epsilon P_4\right) + \cos\phi\left(\sin\epsilon P_2 - \cos\epsilon P_3\right)\right], \qquad (4.31)$$

$$\frac{d\varphi}{dl} = \frac{k}{2\mathcal{E}_0\epsilon_0}\sec 2\epsilon\left[\cos\phi\left(\cos\epsilon P_1 + \sin\epsilon P_4\right) + \sin\phi\left(\cos\epsilon P_3 - \sin\epsilon P_2\right)\right], \qquad (4.32)$$

where l is the path length of the light through the atomic vapour, k is the wavevector magnitude, ϵ_0 is the vacuum permittivity, P_1 and P_3 are the in-phase components of the polarisation of the medium, and P_2 and P_4 are the out-of-phase (quadrature) components.

Equations (4.29) to (4.32) enable the light-field parameters to be evaluated in terms of polarisation components of the medium. These polarisation components are calculated from the density-matrix elements by evaluating the expectation value of the dipole operator (cf. App. A.6):

$$\mathbf{P} = n_v\langle\mathbf{d}\rangle, \qquad (4.33)$$

where n_v is the number density of the atomic vapour. Upon calculating the polarisation

components, and relating them back to the form of the incident electric field, it is possible to calculate the light-field parameters in terms of density-matrix elements. For incident light which is linearly polarised along $\hat{x}$ and propagating along $\hat{z}$, the differential equations for the light-field parameters for the $F = 1 \rightarrow F' = 0$ hyperfine transition simplify to (cf. App. A.6)

$$\frac{1}{\mathcal{E}_0} \frac{\mathrm{d}\mathcal{E}_0}{\mathrm{d}z} = \left(\frac{3 n_\mathrm{v} \Gamma_\mathrm{d} \lambda^2}{4\sqrt{2}\pi \Omega_\mathrm{R}} \frac{2J'+1}{2J+1} \right) \mathrm{Im} \left(\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'} \right), \tag{4.34}$$

$$\frac{\mathrm{d}\phi}{\mathrm{d}z} = - \left(\frac{3 n_\mathrm{v} \Gamma_\mathrm{d} \lambda^2}{4\sqrt{2}\pi \Omega_\mathrm{R}} \frac{2J'+1}{2J+1} \right) \mathrm{Re} \left(\widetilde{\rho}_{-1,0'} + \widetilde{\rho}_{1,0'} \right), \tag{4.35}$$

$$\frac{\mathrm{d}\epsilon}{\mathrm{d}z} = \left(\frac{3 n_\mathrm{v} \Gamma_\mathrm{d} \lambda^2}{4\sqrt{2}\pi \Omega_\mathrm{R}} \frac{2J'+1}{2J+1} \right) \mathrm{Im} \left(\widetilde{\rho}_{-1,0'} + \widetilde{\rho}_{1,0'} \right), \tag{4.36}$$

$$\frac{\mathrm{d}\varphi}{\mathrm{d}z} = \left(\frac{3 n_\mathrm{v} \Gamma_\mathrm{d} \lambda^2}{4\sqrt{2}\pi \Omega_\mathrm{R}} \frac{2J'+1}{2J+1} \right) \mathrm{Re} \left(\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'} \right), \tag{4.37}$$

where λ is the wavelength of the incident light.

4.9   Calculated Observables

Although Eqs. (4.34) to (4.37) can in principle be used to *quantitatively* model the effect of an optically polarised atomic vapour on the light-field parameters, this is quite often an arduous undertaking. For this reason, it can be beneficial to simply consider *qualitative* results. This can be achieved by normalising Eqs. (4.34) to (4.37) by the resonant absorption length $l_0 = (n_\mathrm{v} \sigma_{\mathrm{A}0})^{-1}$, where $\sigma_{\mathrm{A}0}$ is the photon-absorption cross-section on resonance. Assuming that the resonance is *homogeneously* broadened (i.e. $\Gamma = \Gamma_\mathrm{nat}$), the on-resonance photon-absorption cross-section is given by [165,221,224,243]

$$\sigma_{\mathrm{A}0} = \frac{\lambda^2}{2\pi} \frac{2J'+1}{2J+1}. \tag{4.38}$$

Upon multiplying Eqs. (4.34) to (4.37) by l_0, the following normalised observable equations are obtained:

$$\frac{l_0}{\mathcal{E}_0} \frac{\mathrm{d}\mathcal{E}_0}{\mathrm{d}z} = \left(\frac{3}{2\sqrt{2}} \frac{\Gamma_\mathrm{d}}{\Omega_\mathrm{R}} \right) \mathrm{Im} \left(\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'} \right), \tag{4.39}$$

$$l_0 \frac{\mathrm{d}\phi}{\mathrm{d}z} = - \left(\frac{3}{2\sqrt{2}} \frac{\Gamma_\mathrm{d}}{\Omega_\mathrm{R}} \right) \mathrm{Re} \left(\widetilde{\rho}_{-1,0'} + \widetilde{\rho}_{1,0'} \right), \tag{4.40}$$

$$l_0 \frac{\mathrm{d}\epsilon}{\mathrm{d}z} = \left(\frac{3}{2\sqrt{2}} \frac{\Gamma_\mathrm{d}}{\Omega_\mathrm{R}} \right) \mathrm{Im} \left(\widetilde{\rho}_{-1,0'} + \widetilde{\rho}_{1,0'} \right), \tag{4.41}$$

$$l_0 \frac{\mathrm{d}\varphi}{\mathrm{d}z} = \left(\frac{3}{2\sqrt{2}} \frac{\Gamma_\mathrm{d}}{\Omega_\mathrm{R}} \right) \mathrm{Re} \left(\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'} \right). \tag{4.42}$$
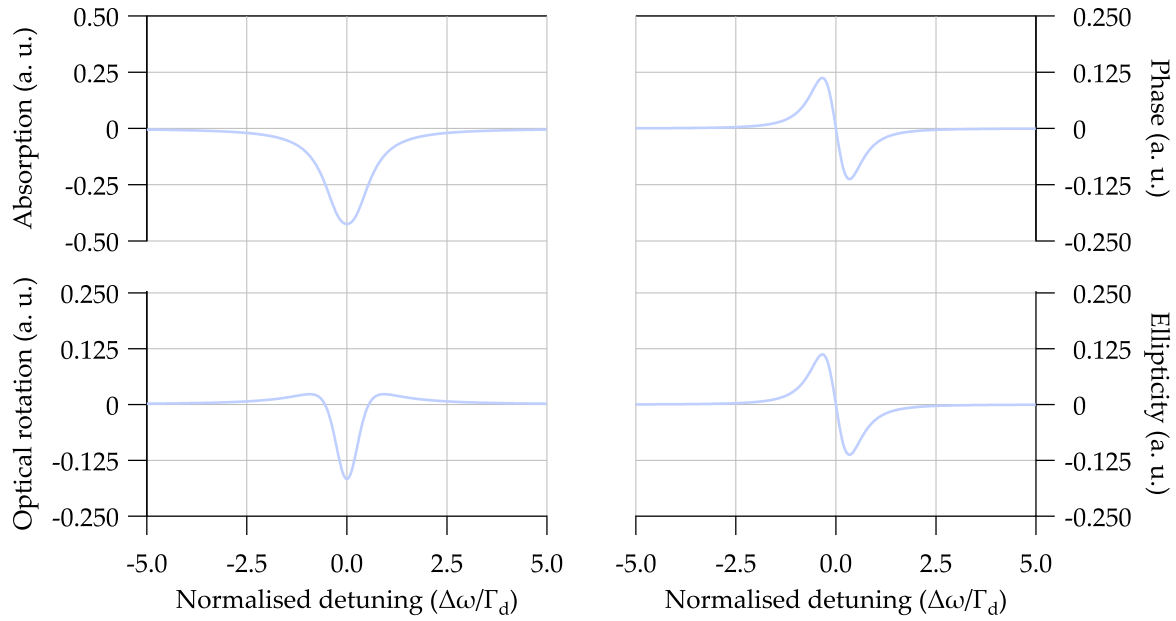
Figure 4.3: Detuning dependence of the observables of the modified Liouville equation (Eq. 4.8), calculated using Eqs. (4.39) to (4.42), in response to the Hamiltonian given by Eq. (4.23). Parameters used in these calculations are discussed in Sec. 4.9.1.

4.9.1   Detuning

For a fixed spontaneous-decay rate Γ_d and Rabi frequency Ω_R, the observables are dependent upon two variables: the detuning $\Delta\omega$ and the Larmor frequency Ω_L. In order to investigate the detuning dependence of the observables, the modified Liouville equation (Eq. 4.8) was solved for the steady-state solution, using the Hamiltonian given by Eq. (4.23). The normalised observable equations, given by Eqs. (4.39) to (4.42), are plotted in Fig. 4.3 using the following parameters: $\Gamma_d = 2\pi \times 1$, $\gamma_t = 2\pi \times 0.01$, and $\Omega_L = 2\pi \times 0.2$. The observables presented in Fig. 4.3 are in the linear regime — i.e. when the saturation parameter $\kappa_2 \to 0$ — calculated by taking the limit as $\Omega_R \to 0$.

The detuning dependence of the optical rotation presented in Fig. 4.3 shows excellent qualitative agreement with that expected from the Macaluso-Corbino effect (cf. Fig. 3.7).

4.9.2   Larmor Frequency

In a similar fashion to the investigation of the detuning dependence in Sec. 4.9.1, the Larmor-frequency dependence of the observables can also be explored. Plotted in Fig. 4.4 are the normalised observable equations, Eqs. (4.39) to (4.42), calculated in the same way as per Sec. 4.9.1. There are two different regimes of interest here: the linear regime ($\kappa_2 \to 0$), and the nonlinear regime ($\kappa_2 \gtrsim 1$). The Larmor-frequency depdendent observables for these two regimes are presented in Fig. 4.4, with the dashed lines corresponding to the linear regime, and the solid lines corresponding to the nonlinear regime with $\kappa_2 = 5$. The explicit parameters used in these density-matrix calculations are: $\Gamma_d = 2\pi \times 1$, $\gamma_t = 2\pi \times 0.01$, and $\Delta\omega = 2\pi \times 0.2$ for the linear regime; and $\Gamma_d = 2\pi \times 1$, $\gamma_t = 2\pi \times 0.05$, and $\Delta\omega = 2\pi \times 0.5$ for
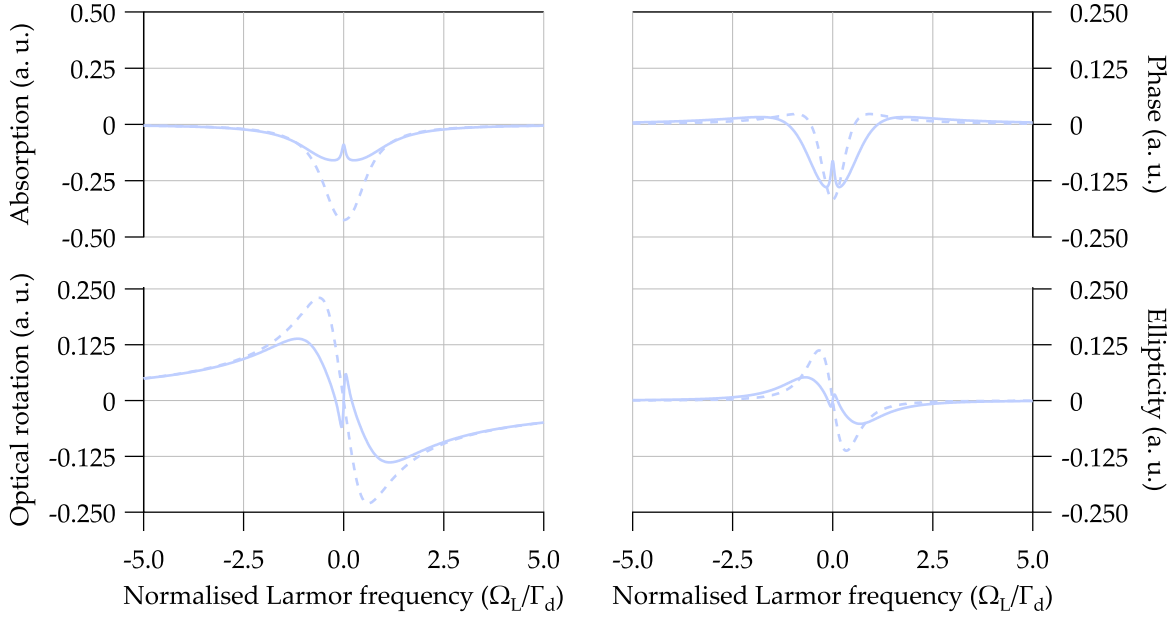
Figure 4.4: Larmor-frequency dependence of the observables of the modified Liouville equation (Eq. 4.8) in the linear ($\kappa_2 \to 0$, dashed) and nonlinear ($\kappa_2 = 5$, solid) regimes, calculated using Eqs. (4.39) to (4.42), in response to the Hamiltonian given by Eq. (4.23). Parameters used in these calculations are discussed in Sec. 4.9.2.

the nonlinear regime. As with the detuning dependence of the optical rotation discussed in Sec. 4.9.1, the Larmor-frequency dependence in the linear regime ($\kappa_2 \to 0$) presented in Fig. 4.4 shows excellent qualitative agreement with that expected from the Macaluso-Corbino effect (cf. Fig. 3.7).

4.9.3 Doppler-Broadened Observables

Most magnetometers exploiting NMOR will use a room-temperature vapour — though cold-atom magnetometers, operating at cryogenic temperatures, do exist [205,269,287]. As described in Sec. 3.2.3, the optical transitions of a room-temperature vapour are *significantly* Doppler broadened due to thermal motion of the atoms. In order to reproduce the detuning-dependent observables $\mathscr{O}(\Delta\omega)$ in a room-temperature experiment, they must therefore be *convolved* with a Doppler profile (cf. Sec. 3.2.3):

$$\mathscr{O}_{\text{DB}}(\Delta\omega) = \int_{-\infty}^{\infty} \mathscr{O}(\Delta\omega - \Delta\omega_v)\,\mathcal{G}(\Delta\omega_v)\,\mathrm{d}\Delta\omega_v, \tag{4.43}$$

where $\mathcal{G}(\Delta\omega_v)$ is the normalised Doppler profile given by Eq. (3.37). Upon convolving the observables calculated in Sec. 4.9.1 — specifically those presented in Fig. 4.3 — with a Doppler profile of width Γ_{D}, and subsequently normalising into sensible (albeit arbitrary) units, the Doppler-broadened observables presented in Fig. 4.5 are obtained. It can be seen by comparison of Fig. 4.5 with Fig. 4.3 that the observables are *qualitatively* the same, though they now have a width which is dominated by Γ_{D} rather than Γ_{d} or γ_{t}.

Figure 4.5: Detuning dependence of the normalised Doppler-broadened observables of the modified Liouville equation (Eq. 4.8), calculated using Eqs. (4.39) to (4.42) and subsequently convolved using Eq. (4.43), in response to the Hamiltonian given by Eq. (4.23). Parameters used in these calculations are discussed in Sec. 4.9.3.

4.10 Amplitude Modulation

In order to improve the dynamic range of magnetometry measurements using NMOR, it is conventional to employ modulation techniques: either amplitude, frequency, or polarisation modulation. In doing so, a resonance is created when $\Omega_m = 2\Omega_L$, as described in Sec. 3.7. The most common modulation technique is amplitude modulation, in which the pump beam is switched on and off — typically with a square wave — with a 100% modulation depth. Mathematically, amplitude modulation can be modelled by incorporating a time dependence into the Rabi frequency: $\Omega_R \rightarrow \Omega_R(t)$. The simplest form of periodic amplitude modulation is sinusoidal[3], which is described by

$$\Omega_R(t) = \left[\frac{1}{2} + \frac{\cos(\Omega_m t)}{2}\right]\overline{\Omega}_R, \qquad (4.45)$$

where $\overline{\Omega}_R$ denotes the peak Rabi frequency (i.e. the unmodulated value). In order to im-

---

[3]In most amplitude-modulated NMOR experiments, the pump beam is modulated with a low-duty-cycle square wave, rather than a sine wave. In order to model an arbitrary square-wave amplitude modulation, one could use a Fourier-series expansion for the time-dependent Rabi frequency:

$$\Omega_R(t) = \overline{\Omega}_R\left\{(-D) + 2D\sum_{n=1}^{\infty}\text{sinc}(n\pi D)\cos\left(2\pi n\left[\frac{\Omega_m t}{2\pi} - \frac{D}{2}\right]\right)\right\}, \qquad (4.44)$$

where the linear modulation frequency is $\Omega_m/2\pi$ and the duty cycle is $D \in (0, 1)$. However; the summation in Eq. (4.44) must be evaluated for a *large* number of terms in order to converge to a square wave, which would become *exceptionally* computationally expensive when solving the Liouville equation. Furthermore, numerical instabilities may become an issue due to discontinuities introduced by the square-wave modulation. For these reasons, one might consider exploiting the fact that the system parameters are piecewise constant, and solve for periodic conditions instead (see Ref. [283] for further discussion).

plement the time-dependent Rabi frequency, the Hamiltonian after performing the rotating-wave approximation (i.e. Eq. 4.23) must be updated by replacing $\Omega_R \to \Omega_R(t)$. After making this change for a sinusoidal amplitude modulation using Eq. (4.45), the new rotating-wave Hamiltonian is given by

$$
\widetilde{\mathscr{H}} = \begin{pmatrix}
\hbar\Omega_L & 0 & 0 & \frac{\hbar\overline{\Omega}_R}{2\sqrt{2}}\left[\frac{1}{2} + \frac{\cos(\Omega_m t)}{2}\right] \\
0 & 0 & 0 & 0 \\
0 & 0 & -\hbar\Omega_L & -\frac{\hbar\overline{\Omega}_R}{2\sqrt{2}}\left[\frac{1}{2} + \frac{\cos(\Omega_m t)}{2}\right] \\
\frac{\hbar\overline{\Omega}_R}{2\sqrt{2}}\left[\frac{1}{2} + \frac{\cos(\Omega_m t)}{2}\right] & 0 & -\frac{\hbar\overline{\Omega}_R}{2\sqrt{2}}\left[\frac{1}{2} + \frac{\cos(\Omega_m t)}{2}\right] & -\hbar\Delta\omega
\end{pmatrix} .
\tag{4.46}
$$

Through modulation of the incident light, the observables in turn, become modulated. This is a direct result of the time-dependent Rabi frequency, giving rise to oscillatory time dependence of the observables. This gives rise to a situation in which the observables (typically the optical rotation) can be readily measured using frequency-demodulation techniques (cf. Sec. 6.1.2 and App. K). The expected observables after demodulation with a lock-in amplifier are calculated and discussed in Sec. 4.11.

4.11   Signal Demodulation

As will be discussed further in Sec. 6.1.2, most sensitive magnetometers perform measurements in the frequency domain by demodulating the optical-rotation signal at the modulation frequency Ω_m using a lock-in amplifier. In order to reproduce the two output channels of a lock-in amplifier, X and Y, one must multiply the observable $\mathscr{O}(t)$ — in this case the optical rotation $\phi(t)$ — by either $\cos(\Omega_m t + \varphi)$ or $\sin(\Omega_m t + \varphi)$ respectively, and integrate over a modulation period T_m:

$$
X = \frac{1}{T_m}\int_t^{t+T_m} \phi(t)\cos(\Omega_m t + \varphi)\,dt ,
\tag{4.47}
$$

$$
Y = \frac{1}{T_m}\int_t^{t+T_m} \phi(t)\sin(\Omega_m t + \varphi)\,dt .
\tag{4.48}
$$

In order to calculate the expected demodulated lock-in amplifier observable — the optical rotation — it was necessary to solve the modified Liouville equation for the Hamiltonian given by Eq. (4.46). However, given the sinusoidal time dependence of the Rabi frequency, it is not possible to solve for the steady-state solution directly. Therefore, the time-dependent behaviour must be numerically calculated instead. Upon demodulation of the time-dependent optical rotation using Eqs. (4.47) and (4.48), a steady-state value is obtained for both X and Y. By plotting the steady-state X and Y values as a function of Ω_m, it is possible to reproduce — at least qualitatively — the lock-in amplifier resonances observed in experiment.

Figure 4.6: Example of the expected *X* and *Y* output channels of a lock-in amplifier when measuring the optical-rotation signal, obtained by applying Eqs. (4.47) and (4.48) to the normalised optical-rotation observable (Eq. 4.40) of the modified Liouville equation (Eq. 4.8) in response to the amplitude-modulated Hamiltonian given by Eq. (4.46). Numerical data is shown in grey markers, while complex-Lorentzian fits — given by Eq. (6.2) — are shown in blue (in-phase), red (quadrature), and purple (magnitude). Parameters used in these calculations are discussed in Sec. 4.11. The residuals of the fit to the resonance magnitude (purple markers) shows residual structure, indicating that the complex-Lorentzian function is only an approximation to the true theoretical lineshape.

Calculated lock-in amplifier resonances are presented in Fig. 4.6, using the following parameters: $\Gamma_d = 500$, $\gamma_t = 0.01$, $\Delta\omega = 500$, $\Omega_L = 1$, and $\Omega_R = 2$. This combination of parameters corresponds to a saturation parameter of $\kappa_2 = 0.8$, which is *just* within the linear regime. The observed resonances are well approximated by a complex-Lorentzian profile — defined by Eq. (6.2) — as evidenced by the fit residuals in Fig. 4.6.

As will be shown in the following chapters, the theoretically predicted resonance signals derived using the density-matrix formalism (Fig. 4.6) show *superb* agreement with experimental measurements (Fig. 6.4). The density-matrix formalism therefore provides an excellent mathematical framework for modelling light-atom interactions and, more specifically, nonlinear magneto-optical effects.

Part III

Experimental Design and Methods

Experimental Setup

The objective of an optical magnetometer, as its name suggests, is to perform magnetometry using optical techniques. One based on nonlinear magneto-optical rotation — the technique utilised throughout this thesis — accomplishes this by measuring the Larmor frequency of an ensemble of alkali atoms via the nonlinear Faraday effect. Knowledge of this frequency then enables magnetometry to be performed via a direct relation between the local magnetic-field strength and the Larmor frequency through the Zeeman effect (cf. Sec. 3.3).

The experimental setup of an NMOR magnetometer is conceptually rather simple, and is essentially comprised of five major components: a laser source, vapour cell, magnetic shields, magnetic coils, and a balanced polarimeter. The role of the laser source is to provide the light which interrogates the atoms within the vapour cell. These atoms act as the sensor, and enable the local magnetic-field strength to be determined based on their interaction *with* the field, and subsequently their effect on the properties of the light. Given the extreme sensitivity of the atoms to the field, magnetic shields are often used to provide a quiet magnetic environment in which to perform the measurements, while magnetic coils situated within the shields produce a *homogeneous* and *stable* sensing field. Finally, the light properties — specifically, the polarisation state — are measured at the output of the vapour using a balanced polarimeter.

This chapter serves as an in-depth overview of the experimental apparatus, explains why particular components/parameters were chosen, and discusses the present performance limitations. Where possible, improvements have been suggested that would yield superior magnetic sensitivity.

5.1   Laser Source

The laser source used in all variations of the experiment is a cateye external-cavity diode laser (ECDL), which has a tunable wavelength centred around $\lambda \approx 795\,\text{nm}$. As the name would suggest, the ECDL utilises a cateye reflector rather than a typical diffraction grating (as is the case in the traditional Littrow or Littman-Metcalf configurations). In the cateye configuration, the external cavity is formed between the rear reflecting surface of the semiconductor diode and the cateye reflector in front of the diode. An ultranarrow bandpass filter inside the cavity enables a single cavity mode to be selected, as the filter transmission wavelength is a function of tilt angle. Due to the inherent self-aligning nature of the cateye configuration, this laser has the intrinsic benefit of a high degree of mechanical/acoustic immunity. In order to further maximise suppression of environmental fluctuations, the laser chassis is machined from a solid aluminium block, and the cavity is hermetically sealed.

Coarse wavelength adjustment is performed by rotating the angle of the bandpass filter, whereas fine adjustments are performed by adjusting the cavity length via a piezoelectric transducer (PZT). The cavity length variation is on the order of 10 nm/V, which produces a frequency shift of 70 MHz/V, with a range of 10 GHz for a 150 V drive voltage. Frequency modulation of the laser can be achieved by driving the PZT with a periodic (or otherwise) voltage, with a mechanical-resonance-limited bandwidth of around 25 kHz.

The laser frequency is also dependent upon the injection current and diode temperature, with sensitivities of around 3 MHz/μA and 30 GHz/K respectively. This necessitates temperature control of the laser cavity, as well as the injection current of the semiconductor diode in order to achieve sub-MHz linewidths. A stable cavity temperature is achieved through the use of a Peltier thermoelectric device, whereas the injection current is controlled through the use of suitably low-noise electronics.

5.1.1 Frequency Stabilisation

Despite the efforts made to maximise passive frequency stability, there remain significant long-term temporal fluctuations of the laser frequency which are too large for the sensitive measurements detailed within this thesis (cf. App. D). This necessitates active stabilisation using negative feedback, in order to ensure that the laser frequency does not drift away from atomic resonance over the course of a measurement.

There exist multiple ways to lock a laser to an atomic transition; however, the technique employed in the experiment is known as 'ac locking', or 'FM demodulation'. Inside the laser driver is an internal oscillator which provides a stable 250 kHz reference frequency. This reference frequency is then used to sinusoidally modulate the diode injection current at 250 kHz. As discussed previously, changing the injection current results in a frequency shift of the laser. Therefore, dithering the injection current results in sinusoidal modulation of the laser frequency.

In order to stabilise the laser frequency to an atomic transition, a reference cell in a saturated absorption spectroscopy (SAS) configuration (cf. App. B) is utilised to mitigate the effects of Doppler broadening and provide a narrow spectral feature (around 100 times narrower than the Doppler-broadened resonance) with which to lock to. At the output of the cell, the laser transmission is measured using a photodetector, and subsequently demodulated using a lock-in amplifier referenced to the injection-current dither frequency. The demodulated signal becomes the error signal — which is approximately the derivative of the transmission signal — from which a control signal is generated and used as negative feedback to adjust the laser frequency via both the injection current (fast) and PZT (slow). This enables the laser frequency to be locked to an atomic transition, at the expense of residual amplitude modulation at 250 kHz due to the dither of the injection current.

In each experimental measurement performed in this thesis, the laser is frequency locked to the $F = 2 \rightarrow F' = 1$ hyperfine transition of the ^{87}Rb D$_1$ line. An example of an SAS transmission spectrum, and the corresponding error signals, is shown in Fig. 5.1.
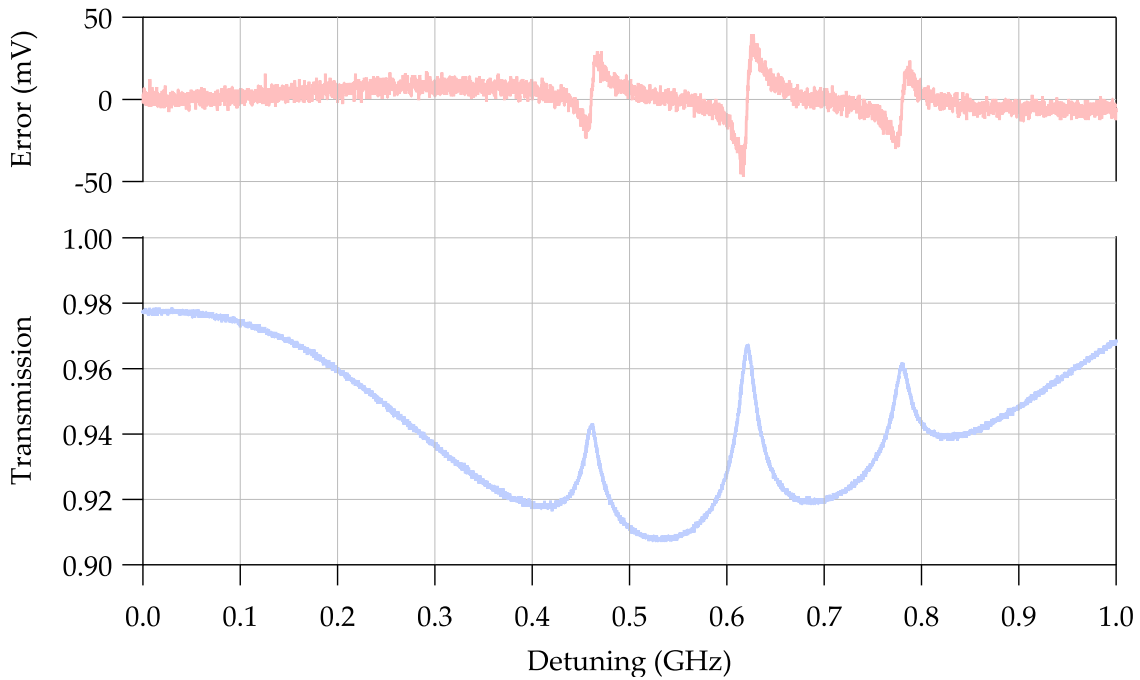
Figure 5.1: Saturated absorption spectroscopy transmission spectrum (blue) with corresponding error signals (red) used for laser-frequency stabilisation. The two outer resonances of the blue trace correspond to the $F = 3 \rightarrow F' = 2$ (left) and $F = 3 \rightarrow F' = 3$ (right) hyperfine transitions of the ^{85}Rb D$_1$ line, while the centre resonance is a crossover transition (cf. App. B.2).

5.1.2 Optical-Power Stabilisation

In addition to active stabilisation of the laser frequency, the optical power — of both the pump and probe beams — is also stabilised. However, rather than directly feeding back to the laser, optical-power stabilisation is achieved through the use of acousto-optic modulators (AOMs) — cf. App. C.

Given that the sensitivity of the magnetometer exhibits a strong dependence upon the optical powers of both the pump and probe beams (cf. Sec. 6.4.3), this necessitates the stabilisation of these powers to their optimal values, as close to the input of the vapour cell as possible. This is achieved by using an optical wedge, which siphons a small fraction of the beam via a Fresnel reflection from the surface. The optical power in the reflected beam is measured using a photodetector, and subsequently delivered to a high-speed servo controller. The servo controller sets a 'target voltage' for the photodetector output, and uses a proportional-integral controller to calculate a control signal based on the error signal (the difference between the target voltage and the measured value). This control signal is used to stabilise the optical power in the beam by adjusting the rf driving power — and therefore first-order diffraction efficiency — of the AOM (cf. App. C). This is realised by using a frequency mixer to heterodyne the 80 MHz local oscillator with the control signal produced by the servo controller. An experimental diagram showing the key components of the optical-power servo is presented in Fig. 5.2.
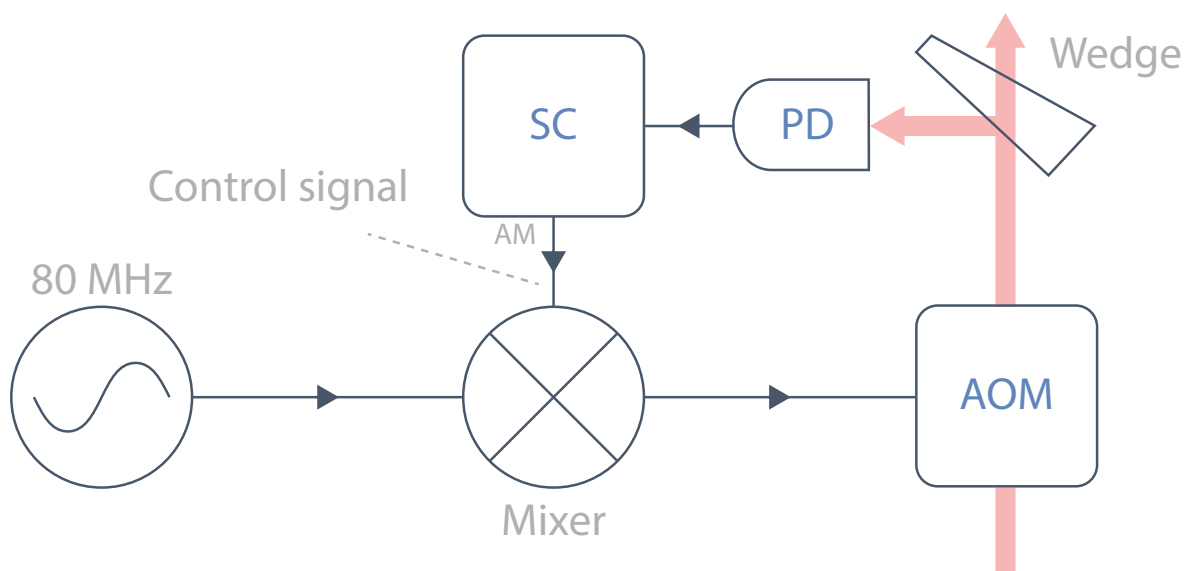
Figure 5.2: Key components of an optical-power servo: AOM — acousto-optic modulator, PD — photodetector, SC — servo controller. The amplitude of the AOM driving signal (and hence the optical power) is controlled by using a frequency mixer to heterodyne the 80 MHz local oscillator with the control signal produced by the servo controller. Preferably, a voltage-controlled attenuator should take the place of the mixer, as mixers can exhibit nonlinear response when driven at either high or low power.

When implementing optical-power control systems, great care must be taken to avoid optical etalons in the beam path *after* the lock point. An optical etalon is typically comprised of a transparent plate with two reflective surfaces (similar to a Fabry-Pérot interferometer). When inadvertently inserted into the beam path of a laser, an etalon acts as an optical resonator (cavity), and therefore results in periodically varying transmission as a function of optical frequency. Fluctuations in the optical frequency of the laser will therefore result in optical-power fluctuations at the output of the etalon, which is problematic in experiments which are nonlinearly dependent upon optical powers — as is the case here. Even in the case that the optical frequency is stabilised, thermal fluctuations in the effective cavity length will *also* give rise to optical power fluctuations due to variations in the transmission through the etalon.

Optical etalons prior to the cell can be inadvertently created through the use of optical attenuation elements such as variable attenuators or neutral density filters. To mitigate this, anti-reflection coatings can be layered onto these optical elements in order to reduce the amount of light that is reflected and hence minimise the effect of the etalon. Examples of fractional optical-power fluctuations with and without an etalon in the beam path are shown in Fig. 5.3. Calculating the Allan deviation (cf. Sec. H.2) of this data reveals that, without the presence of an etalon, the probe beam has a fractional stability of less than 30 ppm over a period of 30 minutes, while the pump beam has a fractional stability of less than 10 ppm over the same time frame. The *effective* magnetic noise that this optical-power instability introduces is discussed in detail in Sec. 8.2.1.

(a) Probe beam with etalon.

(b) Pump beam with etalon.

(c) Probe beam without etalon.

(d) Pump beam without etalon.
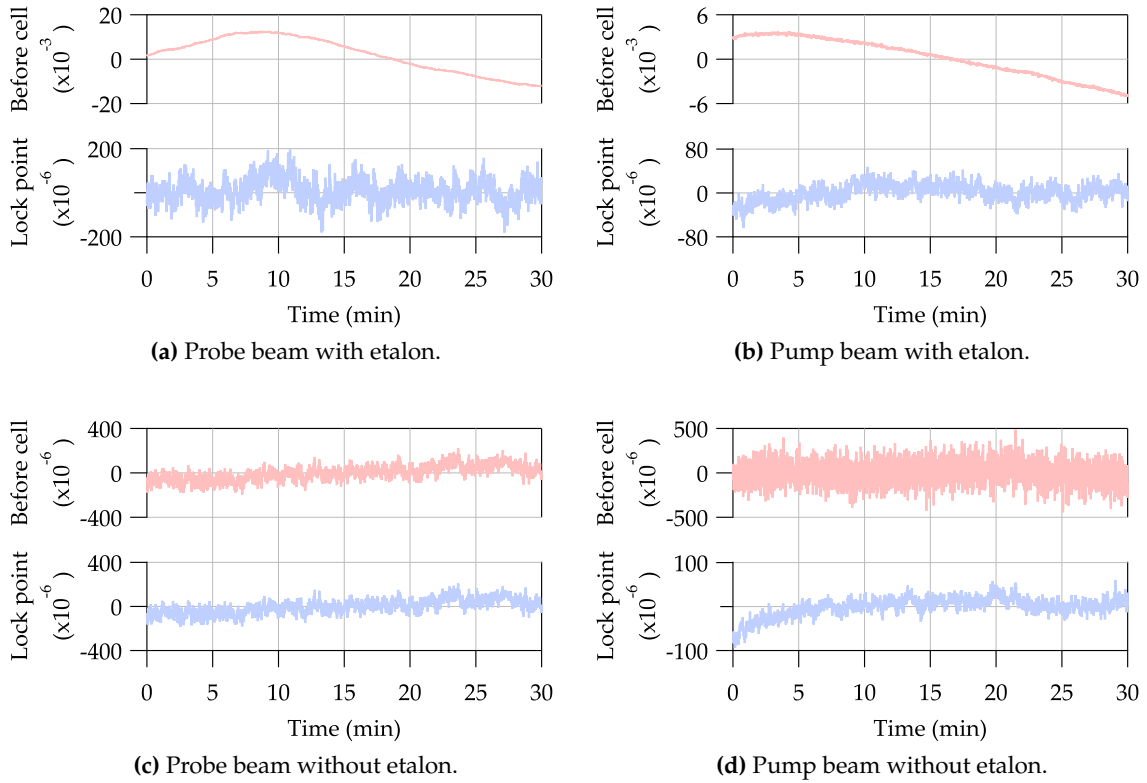
Figure 5.3: Fractional optical-power fluctuations measured at both the lock point and at the input to the vapour cell over a 30-min period. Figures (a) and (b) show the optical-power fluctuations due to the presence of an optical etalon, for the probe and pump beams respectively. Figures (c) and (d) show the optical-power fluctuations with the etalons removed. Note that the measurement before the cell in (d) was limited by electronic noise of the photodetector.

5.1.3 Polarisation

The output beam of the laser is elliptical with a transverse electric (TE) polarisation — that is, the light is linearly polarised along the minor axis of the ellipse. Given that the NMOR technique relies on determining the local magnetic-field strength by measuring small changes in polarisation angle (cf. Secs. 3.6 and 3.7), the polarisation *must* be both pure *and* stable in order to achieve high sensitivity — especially over long timescales. Despite the fact that the NMOR signal is demodulated at high frequencies (typically $\Omega_\mathrm{m}/2\pi \gtrsim 30\,\mathrm{kHz}$), slow fluctuations in polarisation yield asymmetry in the quadrature component of the demodulated resonance, which is *indistinguishable* from real magnetic noise (cf. Sec. 8.2.2). To ensure that both pump and probe beams are linearly polarised with high purity, Glan-Thompson prisms with a high extinction ratio (10^5:1) are used to set the polarisation of both beams — typically immediately prior to the vapour cell. The amount of polarisation drift, and its effect on the magnetometer's performance, is discussed in Sec. 8.2.2.

5.2   Vapour Cell

At the heart of many optical magnetometers lies an atomic vapour cell which contains an alkali vapour (generally potassium, caesium, or rubidium). In a standard vapour cell containing a low-pressure vapour with no anti-relaxation coatings, the transverse spin-relaxation time is limited to the transit time of the optically pumped atoms as they cross the probe beam (cf. Sec. 5.2.2.3). This leads to broad magneto-optical resonances — due to the finite interaction time of the probe beam with the atoms — and therefore poor sensitivity when using them to perform magnetometry (cf. Sec. 6.4.2).

Improvements to the interaction time results in increased magnetic sensitivity, and can be accomplished in at least two ways: the use of buffer gas, or the use of anti-relaxation coatings. By inserting a relatively high-pressure buffer gas into the vapour cell, the interaction time — and therefore transverse spin-relaxation time — is increased as a result of reduced mean free path and subsequently increased diffusion time of the alkali atoms. Optically pumped atoms spend a longer time in the probe beam, thereby increasing the interaction time and reducing the resonance width.

Using buffer gas to extend the ground-state coherence time of the atomic vapour does have its drawbacks, namely: the pump and probe beams must be spatially overlapped due to the reduced mean free path of the alkali atoms, and pressure broadening from alkali-buffer-gas collisions (cf. Sec. 3.2.1). For these reasons, many optical magnetometers use anti-relaxation coatings instead. These coatings — typically composed of alkane (paraffin) or alkene molecules [239] — provide an inert barrier between the atomic vapour and the cell, which prevents spin-destruction collisions with the cell walls (cf. Sec. 5.2.2.4). Atoms which collide with the paraffin-coated walls of the cell tend to retain their atomic polarisation, and some fraction of them will subsequently re-enter the probe beam — effectively increasing their interaction time. The fraction of atoms which re-enter the probe beam depends, amongst other things, on the quality of the anti-relaxation coating — characterised by the number of wall collisions which can occur before the atoms lose their spin coherence. Transverse spin-relaxation times of $\mathcal{O}\left(100\,\mathrm{ms}\right)$ are commonplace when using paraffin anti-relaxation coatings, while minute-long relaxation times have been reported in alkene-coated cells [239].

The vapour cells used in the experiments detailed within this thesis are cylindrical (40 mm diameter and 40 mm length) and are made from borosilicate glass which has a low coefficient of thermal expansion. Inside the cell is an isotopically pure ^{87}Rb vapour at low pressure. The walls of the cell are anti-relaxation coated with paraffin, which extends the coherence time between ground-state Zeeman sublevels to $\mathcal{O}\left(40\,\mathrm{ms}\right)$ (cf. App. F).

5.2.1   Cell Mounting

Within the innermost magnetic shield resides a solenoid that creates a relatively homogeneous magnetic field across the vapour cell (cf. Sec. 5.4). The solenoid was handmade by carefully winding copper wire around a plastic former. In order to situate the vapour cell

Figure 5.4: Horizontal cross-section of a scale model of the magnetic shielding, vapour cell and solenoid. Depending on the requirements of the measurement, either all six layers of shielding were used or, alternatively, only the innermost three layers.

within the geometric centre of both the shield and solenoid, a mount was designed and subsequently 3D printed from polylactic acid (PLA). This material was chosen as it is *non-magnetic*, in order to minimise the ambient magnetic-field noise present within the shields. The mount is designed to fix the vapour cell within the centre of the shields, and connects to the plastic former using two end caps which slot onto the ends of the former. A horizontal cross-section of a scale model of the vapour cell mounting and magnetic shields used in the experiment is presented in Fig. 5.4.

In the case where we wanted to perform magnetic-gradiometer measurements, it was necessary to hold two cells within the magnetic shielding. In order to do this, a horizontal displacement was introduced between the two cells, leaving them with some degree of cross-sectional overlap. This enabled both cells to be optically pumped with the same pump beam, while being individually addressable using separate probe beams.

5.2.2   Relaxation Mechanisms and Resonance Broadening

Perhaps the most important parameter in regards to the sensitivity of an optical magnetometer is the transverse spin-relaxation time T_2 (cf. Sec. 6.4.1). The transverse spin-relaxation time is a measure of the time for which an optically pumped atomic vapour remains spin-polarised, and depends upon various relaxation processes which will be discussed in this section.

5.2.2.1 Spin-Exchange Collisions

A spin-exchange collision is a collision between two alkali atoms in which the *total* angular momentum $F_1 + F_2$ is conserved; though the direction of the individual spins may change, and one or both of the atoms may change hyperfine states as a result of the collision. This is represented symbolically as [162,288]

$$A\,(\uparrow) + B\,(\downarrow) \to A\,(\downarrow) + B\,(\uparrow)\,. \tag{5.1}$$

The effect of spin-exchange collisions is to destroy coherence between the ground-state Zeeman sublevels, and is therefore a relaxation mechanism. The rate at which spin-exchange collisions occur, Γ_{SE}, is given by $\Gamma_{\mathrm{SE}} = n_{\mathrm{v}}\sigma_{\mathrm{SE}}\bar{v}_{\mathrm{th}}$, where n_{v} is the atomic number density, $\sigma_{\mathrm{SE}} = 2 \times 10^{-18}\,\mathrm{m}^2$ is the alkali-alkali spin-exchange cross-section [164,167,170,240,289–291], and $\bar{v}_{\mathrm{th}}$ is the mean thermal velocity. At $T = 23°\mathrm{C}$, the number density of $^{87}\mathrm{Rb}$ is approximately $n_{\mathrm{v}} = 1.04 \times 10^{16}\,\mathrm{m}^{-3}$, and the mean thermal velocity is $\bar{v}_{\mathrm{th}} = 238\,\mathrm{ms}^{-1}$. Using these parameters, the spin-exchange collision rate is $\Gamma_{\mathrm{SE}} = 4.95\,\mathrm{s}^{-1}$. This contributes a FWHM of $\Gamma/2\pi = 2\Gamma_{\mathrm{SE}}/2\pi \approx 1.57\,\mathrm{Hz}$ to the resonance width, which is just over a factor of 8 less than the observed width under optimal conditions (cf. Ch. 10). Therefore, the experiments detailed within this thesis are *not* operating in a regime dominated by spin-exchange collisons.

5.2.2.2 Spin-Destruction Collisions

In addition to spin-exchange collisions which conserve the total angular momentum, there are also *non-conservative* spin-destruction collisions which may occur. Rather than transferring spin angular momentum between the two colliding atoms, spin angular momentum is transferred to *rotational* angular momentum. These collisions can either occur between two alkali atoms, or between alkali atoms and noble buffer gases (e.g. helium, neon, nitrogen, etc.). Analogous to the case with spin-exchange collisions, spin-destruction collisions can be written symbolically as [241]

$$A\,(\uparrow) + B\,(\downarrow) \to A\,(\downarrow) + B\,(\downarrow)\,. \tag{5.2}$$

Since the vapour cell used in the experiments detailed within this thesis contains $^{87}\mathrm{Rb}$ exclusively, with no buffer gas, only Rb-Rb spin-destruction collisions need to be considered. The spin-destruction cross-section for Rb-Rb collisions is $\sigma_{\mathrm{SD}} = 1.6 \times 10^{-21}\,\mathrm{m}^2$ [241], which is over three orders of magnitude *smaller* than the spin-exchange cross-section. Using the parameters discussed in Sec. 5.2.2.1, the spin-destruction collision rate for Rb-Rb collisions at $T = 23°\mathrm{C}$ is $\Gamma_{\mathrm{SD}} = n_{\mathrm{v}}\sigma_{\mathrm{SD}}\bar{v}_{\mathrm{th}} = 3.96\,\mathrm{ks}^{-1}$, which contributes a FWHM of $\Gamma/2\pi = 2\Gamma_{\mathrm{SD}}/2\pi \approx 1.26\,\mathrm{mHz}$. Given that this contribution is *many* orders of magnitude smaller than the observed resonance width of $\mathcal{O}\,(10\,\mathrm{Hz})$, spin-destruction collisions can be safely ignored as a dominant relaxation mechanism in experiments detailed within this thesis.

5.2.2.3   Transit-Time Broadening

Transit-time broadening is the broadening of a resonance due to the finite interaction time between the atoms and the light [165,200,202,235–238]. This broadening comes as a result of the time-energy uncertainty principle (cf. Sec. 3.2.1), in which the frequency uncertainty σ_ω, is related to the finite interaction time σ_t via $\sigma_\omega \geq 1/2\sigma_t$. Consider a beam of atoms moving orthogonally across a beam of light, with a mean thermal velocity of $\bar{v}_{\text{th}}$. Assuming that the atoms interact with the light across its entire $1/e^2$ diameter d, the atoms are subjected to a pulse of radiation with a duration $\sigma_t = d/\bar{v}_{\text{th}}$. Therefore, the expected transit broadening of the resonance is given approximately by

$$\sigma_\omega \geq \frac{\bar{v}_{\text{th}}}{2d}. \tag{5.3}$$

For room-temperature ^{87}Rb atoms with a mean velocity of $\bar{v} = 238\,\text{ms}^{-1}$ travelling across a Gaussian laser beam with a $1/e^2$ diameter of 1.5 mm, the expected resonance width due to the transit effect is $\Gamma/2\pi = 2\sigma_\omega/2\pi \gtrsim 25\,\text{kHz}$. This resonance width is *extremely* large and would severely inhibit the performance of a magnetometer operating on transit-broadened resonances. For this reason, and the reason to be discussed in Sec. 5.2.2.4, it is *necessary* to use anti-relaxation coatings on the cell walls in order to increase the *effective* interaction time of the atoms with the light, and thereby decrease the resonance width.

5.2.2.4   Wall Collisions

When an alkali atom collides with the bare glass wall of a cell, it is adsorbed into the surface for a finite period of time (on the order of $1\,\mu s$ – $10\,\mu s$ [292]), after which it is desorbed back into the cell volume. During the time in which the atom is adsorbed, it experiences large local electric and magnetic fields produced by the ions and molecules within the glass, causing the spin direction to be completely randomised [293]. Wall collisions therefore generally *completely dominate* all other relaxation processes. For this reason, wall collisions *must* be actively suppressed through the use of either buffer gas (to reduce atomic mean free path), or anti-relaxation coatings. These two scenarios will be considered below.

Wall collisions affect both the longitudinal *and* tranverse spin-relaxation times (T_1 and T_2, respectively) equally, with a relaxation rate γ_{wall}. The wall-collision-induced relaxation rate will depend drastically on the vapour cell, i.e. depending on whether it is filled with buffer gas, or coated with an anti-relaxation coating. In a buffer gas cell, the decay of polarisation is governed by the diffusion equation [297,298]:

$$\frac{\partial P}{\partial t} = \mathcal{D}\nabla^2 P, \tag{5.4}$$

where $\mathcal{D} = \mathcal{D}_0 P_A/P_v$ is the diffusion coefficient at a gas pressure of P_v, and $\mathcal{D}_0$ is the species-specific diffusion coefficient at atmospheric pressure P_A. For a spherical cell of radius R_{cell}, the lowest-order diffusion mode has a time constant given by (cf. App E.1) [298,299]
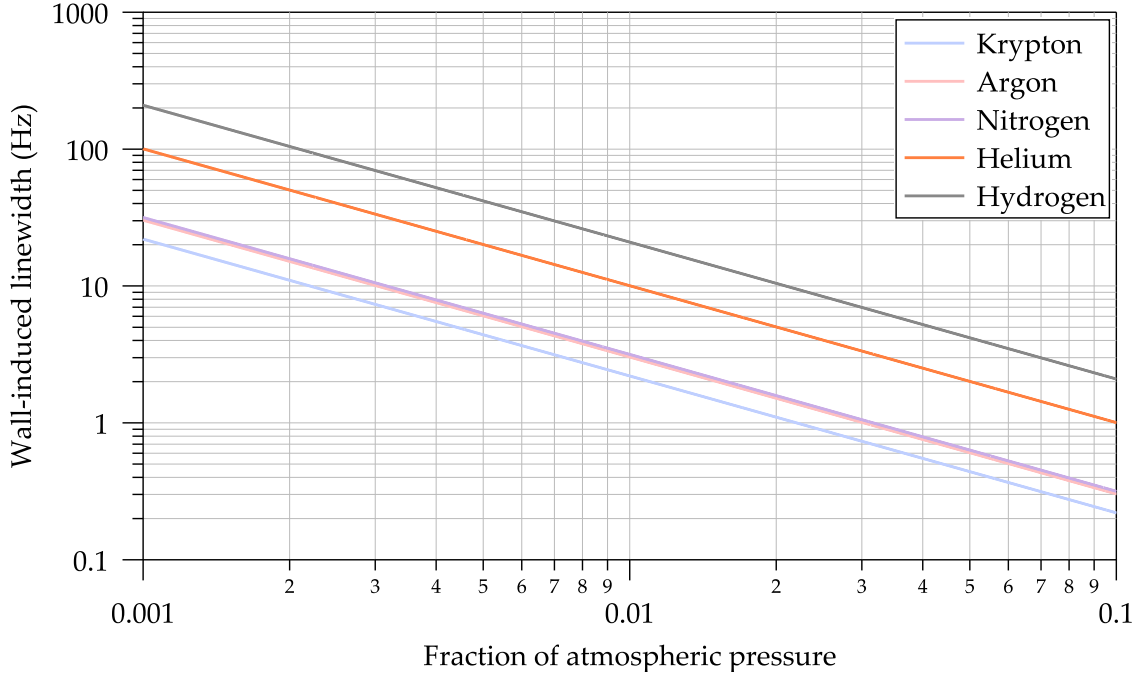
Figure 5.5: Predicted resonance linewidth (FWHM) due to wall collisions in a cylindrical cell in the presence of various noble buffer gases, calculated via $\Gamma/2\pi = 2/2\pi T_{\text{wall}}$ using Eq. (5.6). The radius and length of the cell used in the calculation is $R_{\text{cell}} = 20\,\text{mm}$ and $L_{\text{cell}} = 40\,\text{mm}$, respectively. Diffusion coefficients for the noble gases are presented in Refs. [294–296].

$$T_{\text{wall}} = \frac{1}{\mathcal{D}}\left(\frac{R_{\text{cell}}^2}{\pi^2}\right), \tag{5.5}$$

whereas for a cylindrical cell of radius R_{cell} and length L_{cell}, the lowest-order diffusion mode has a time constant given by (cf. App E.2) [298,300]

$$T_{\text{wall}} \approx \frac{1}{\mathcal{D}}\left(\frac{R_{\text{cell}}^2}{(2.40483)^2} + \frac{L_{\text{cell}}^2}{\pi^2}\right). \tag{5.6}$$

Although the resonance linewidth predicted by Eq. (5.6), and shown in Fig. 5.5, is sufficiently low to perform highly sensitive magnetometry, it does not tell the full story. When using buffer gas, there is also associated pressure broadening of the magneto-optical resonance (cf. Sec. 3.2.1), and the resonance is especially susceptible to broadening as a result of magnetic-field gradients (cf. Sec. 5.4). For these two reasons, optical magnetometers — at least ones which operate near room temperature — are typically constructed using cells which have been coated with paraffin, an inert alkene anti-relaxation coating which suppresses the wall-induced depolarisation. Paraffin coatings exhibit significantly reduced pressure broadening due to the low vapour pressure and, for a fixed set of magnetic gradients, are less susceptible to gradient broadening due to motional narrowing (cf. Sec. 5.4).

In an anti-relaxation-coated cell without any buffer gas, atoms follow a near-ballistic

trajectory[1], with a relaxation rate γ_{wall} that is proportional to the reciprocal of the average time between collisions T_{wall}^{-1}. For a spherical cell, this is given analytically by [299]

$$\gamma_{\text{wall}} = \frac{\bar{v}_{\text{th}} A}{4V}, \tag{5.7}$$

where A/V is the surface area to volume ratio of the cell, and $\bar{v}_{\text{th}}$ is the mean thermal velocity of the atoms. In the case of a cylindrical cell — as is used in the experiments detailed within this thesis — the dependence of γ_{wall} on the cell dimensions is nonlinear. In order to estimate the wall relaxation rate in cylindrical vapour cells with different dimensions, a Monte Carlo simulation was developed (cf. Sec. L.6).

The wall-collision rate γ_{wall} constitutes a FWHM resonance linewidth of $2\gamma_{\text{wall}}$. Using the results from the Monte Carlo simulation — presented in Fig. L.5 — the mean time between wall collisions in the cylindrical vapour cell used in this work is $T_{\text{wall}} = 106\,\mu\text{s}$, corresponding to a FWHM resonance linewidth of around $\Gamma/2\pi = 2/2\pi T_{\text{wall}} \approx 3\,\text{kHz}$. This is over two orders of magnitude larger than the resonance width of $\mathcal{O}\,(10\,\text{Hz})$ observed in the paraffin-coated cell (cf. App. F). Fortunately, however, anti-relaxation coatings enable atoms to undergo as many as 10^6 wall collisions before depolarising [239], yielding ultra-narrow resonance widths observed to be as low as $\Gamma/2\pi \approx 1\,\text{Hz}$ [236]. These ultranarrow resonances — typically referred to as 'wall-induced Ramsey resonances' [200] — appear as a narrow feature on top of the broad background pedestal which arises due to the transit effect (cf. Sec. 5.2.2.3).

5.3 Magnetic Shielding

Earth's geomagnetic field is large in amplitude — relative to the fields generated by small objects of interest — with typical magnitudes ranging anywhere from $25\,\mu\text{T}$ to $65\,\mu\text{T}$. Additionally, it also exhibits *strong* temporal dependence, with ambient magnetic-field noise ranging from $\sim 2\,\text{nT}_{\text{rms}}/\sqrt{\text{Hz}}$ at $0.01\,\text{Hz}$, to $\sim 2\,\text{pT}_{\text{rms}}/\sqrt{\text{Hz}}$ at $1\,\text{Hz}$ [302] (cf. Fig. 8.5). These two factors necessitate the use of magnetic shielding in order to create a quiet magnetic environment in which to perform sensitive measurements.

The temporal dependence of the ambient magnetic field in the laboratory is actually *significantly worse* than that of Earth's geomagnetic field. A measurement of the ambient magnetic field in the laboratory over $35\,\text{h}$ — in close proximity to the magnetometer — is presented in Fig. 5.6. The amplitude spectral density of the laboratory field ranges from around $\sim 10\,\text{nT}_{\text{rms}}/\sqrt{\text{Hz}}$ at $0.01\,\text{Hz}$, dipping down to $\sim 4\,\text{nT}_{\text{rms}}/\sqrt{\text{Hz}}$ at $0.1\,\text{Hz}$, and rising back up to $\sim 10\,\text{nT}_{\text{rms}}/\sqrt{\text{Hz}}$ at $1\,\text{Hz}$ — between one to four orders of magnitude larger than that of Earth's geomagnetic field over the same range. Given such large fluctuations, it was necessary to use either three layers or six layers of μ-metal magnetic shielding in order to adequately suppress the magnetic-field fluctuations to an acceptable level — i.e. low enough

---

[1]Despite being in near-vacuum, there is still a relatively high rate of velocity-changing collisions between alkali atoms and background gas within the cell at room temperature. Recent work has measured this to be on the order of $1 \times 10^6\,\text{s}^{-1}$ for a background gas pressure of $3.0\,\text{Pa}$, which corresponds to a mean free path of about $1\,\text{mm}$ for rubidium atoms [301].
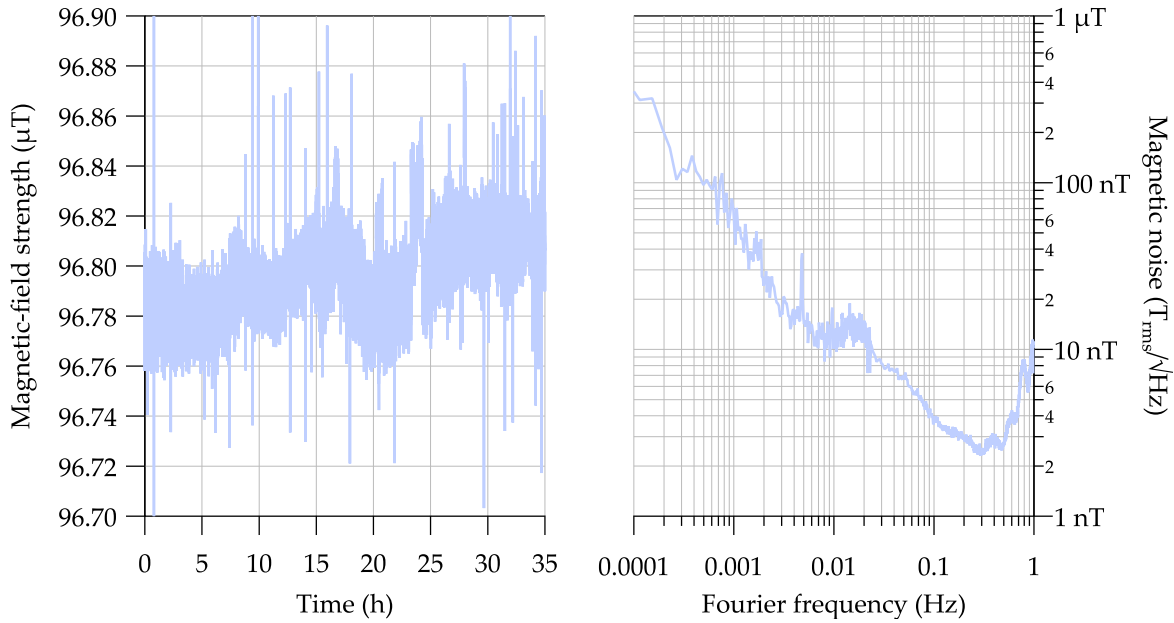
Figure 5.6: Time trace (left) and amplitude spectral density (right) of the magnetic-field strength in the laboratory, measured over a period of 35 h using a three-axis flux-gate magnetometer.

to measure the magnetometer noise floor. A schematic of the magnetic shields used in the experiment is shown in Fig. 5.4.

5.3.1 Efficiency of Magnetic Shielding

The efficiency of a magnetic shield is quantified by the factor to which an external magnetic field is attenuated inside the shields. That is, the ratio of the field measured inside the shields, B_{int}, to the field measured outside the shields, B_{ext}:

$$A_{\text{s}}^{(\text{tot})} \equiv \frac{B_{\text{int}}}{B_{\text{ext}}}. \tag{5.8}$$

The physical mechanism(s) responsible for shielding magnetic fields depends upon the frequency of the field. For static and low-frequency fields, the dominant mechanism is flux shunting due to the high magnetic permeability of the material [150,303]. The reluctance, $\mathcal{R}$, of a *magnetic* circuit is analogous to the *resistance* of an *electric* circuit. Just as electric current will tend to follow the path of least electrical resistance, magnetic flux will tend to follow the path of least magnetic reluctance. Given that $\mathcal{R} \propto \mu_{\text{P}}^{-1}$ [150], where μ_{P} is the magnetic permeability of the medium, magnetic flux will be guided along high-permeability paths (cf. Fig. 5.7). Magnetic shielding of static fields can therefore be achieved by *shunting* the flux *around* the shielded region via high-permeability materials (such as μ-metal). In this case, the shielding efficiency is given by an approximate formula [150,203,303]:

$$A_{\text{s}}^{(\text{tot})} \simeq A_{\text{s}}^{(n)} \prod_{i=1}^{n-1} A_{\text{s}}^{(i)} \left[1 - \left(\frac{X_{i+1}}{X_i} \right)^k \right]^{-1}, \tag{5.9}$$

Figure 5.7: Illustration of the field distribution around a single-layer cylindrical shield in the transverse (left) and longitudinal (right) planes due to flux shunting [303].

where X_i is the radius or length of the i^th layer (depending on the relative orientation of the magnetic field and the layer), $A_\text{s}^{(i)}$ is the shielding efficiency of the i^th layer, and k is a power which accounts for the geometry of the shielding. For a spherical shield, $k \simeq 3$, while for a cylindrical shield, $k \simeq 2$ for the transverse axis, and $k \simeq 1$ for the longitudinal axis. The shielding efficiency of the i^th layer, $A_\text{s}^{(i)}$, can be calculated via [150,203,303]

$$A_\text{s}^{(i)} \simeq \frac{X_i}{\mu_r^{(i)} W_i},\tag{5.10}$$

where $\mu_r^{(i)}$ and W_i are the relative permeability and thickness of the i^th layer respectively. Note that this formula assumes that $X_i > X_{i+1}$, and that $\mu_r^{(i)} \gg X_i/W_i$. Therefore, $A_\text{s}^{(n)}$ corresponds to the innermost shield, with the value i incrementing from the outside shield towards the inside shield.

The three-layer magnetic shields used in some experiments had a measured shielding factor of $\left(A_\text{s}^{(\text{tot})}\right)^{-1} \approx 1800$ — consistent with the manufacturer's specifications. This was sufficient for high-frequency measurements such as those described in Ch. 7 and Ch. 11, but not suitable for the sensitive, long-term measurements described in Ch. 10. In this case, the three-layer magnetic shield was nested inside a larger three-layer shield — providing a total of six layers of shielding (cf. Fig. 5.4). Using the material properties and dimensions of the six-layer magnetic shield employed in our experiment, Eq. (5.9) predicts a shielding efficiency of $A_\text{s}^{(\text{tot})} \approx 4.5 \times 10^{-6}$ for the longitudinal axis, and $A_\text{s}^{(\text{tot})} \approx 1.9 \times 10^{-7}$ for the transverse axis. This corresponds to a shielding factor of $\left(A_\text{s}^{(\text{tot})}\right)^{-1} \approx 2 \times 10^5$ and $\left(A_\text{s}^{(\text{tot})}\right)^{-1} \approx 5 \times 10^6$ for the longitudinal and transverse axes, respectively. The residual magnetic field measured inside the innermost layer of shielding is therefore predicted to be around a million times smaller than that of the external field.

5.3.2   Demagnetisation

Provided that the shielding factor is sufficiently large relative to the ambient magnetic-field strength, the main factor in determining the residual magnetic field present within the shields is the magnetisation of the innermost shielding layer. Magnetisation of this layer can occur in a number of ways, but perhaps the two most prominent causes are: the applied magnetic field (due to hysteresis), and physical handling of the shield. Simply disassembling and reassembling the shields has been shown to increase the residual magnetic field within the innermost layer by an order of magnitude when compared to the typical value after demagnetisation [150].

The small residual magnetic fields that might occur may not necessarily be an issue *per se* for NMOR-based magnetometers as they operate at non-zero magnetic field; however, it is naive to assume that the magnetisation will have *no* effect. In practice, the magnetisation *will* be spatially dependent, which will give rise to a magnetic gradient within the shielded volume. This is problematic, as the presence of a magnetic-field gradient will reduce the transverse spin-relaxation time and therefore result in decreased sensitivity of the magnetometer (cf. Sec. 5.4).

In order to reduce the residual magnetic field (and/or gradient) inside the shield, a special procedure is performed which decreases the magnetisation of the shields. This procedure is known as demagnetisation[2], and involves subjecting the shield to a strong, slowly decaying alternating magnetic field. The applied field should be *purely* ac in order to randomly orient the magnetic domains, as the presence of a dc component will produce a net magnetisation during the process. The field should be strong enough to completely saturate the material, which is typically recommended to be five times the material's coercivity [304].

The simplest method of demagnetising a shield is to wind a thick-gauge[3] insulated copper wire around the innermost shield in a solenoidal configuration. Not only is a solenoid relatively easy to wind (due to its simple geometry), it also provides a relatively large magnetic field per unit current. The coil is then supplied current at an industrial frequency (i.e. 50–60 Hz), which is most easily achieved by connecting it to the electricity supply of the building. In Australia, the electrical grid is 230 V at 50 Hz, and hence the demagnetisation coil *must* be loaded in series in order to limit the current. This can be achieved by using any passive electrical component which has high thermal dissipation power. The series resistance of the circuit should be chosen such that, at maximum voltage, the current passing through the demagnetisation coil produces a magnetic field which is large enough to saturate the material. The material should be held at saturation for a small period of time, before gradually reducing the applied field to zero. This is readily achieved by using a variable transformer, which enables the supply voltage to be manually controlled.

When demagnetising the innermost layer of magnetic shielding in the experiment de-

---

[2]Demagnetisation is quite often referred to as degaussing, since the old unit of magnetic-field strength — now deprecated — was the Gauss.

[3]The resistance of a copper wire is inversely proportional to its diameter, so by using thick-gauge copper, the wire can sustain large currents without dissipating excessive power in the form of heat. In a similar manner, multi-strand wire is also useful for carrying large currents.

Figure 5.8: Measured temperature fluctuations of the laboratory (left), as well as inside six layers of μ-metal shielding (right). The periodic structure, observed roughly every 3–4 hours, is due to the cyclic operation of the air conditioner in the laboratory. The temperature fluctuations inside the six layers of shielding are attenuated by approximately a factor of 200 relative to the fluctuations of the laboratory.

scribed here, a variable transformer was used to ramp the voltage from 110% of mains ($\sim$250 V), down to zero, over a period of about 40 s. The series resistance of the circuit was $\sim$49 Ω, which yielded a maximum current of approximately 5 A, and a power dissipation of $\mathcal{O}(1\,\text{kW})$ into the load resistor.

5.3.3 Temperature Stability

Given that both the amount of absorption and optical rotation are directly proportional to the atomic number density (cf. Sec. 4.8), it is desirable that the temperature of the atomic vapour remain relatively constant. Fortunately, one of the intrinsic benefits of using multiple layers of magnetic shielding is the enhanced thermal isolation from temperature fluctuations in the laboratory. In order to assess the temperature stability in both the lab and inside the shields over long timescales, the resistance of two negative-temperature-coefficient thermistors — one placed inside the innermost shield, and one placed outside the shields — were monitored using a multimeter. The resulting temperature fluctuations in both the laboratory and within the shields are presented in Fig. 5.8. Evident in Fig. 5.8 is a periodic temperature dependence with a period of about 3–4 hours associated with the cyclic operation of the air conditioner in the laboratory. Fortunately however, the temperature fluctuations inside the six layers of shielding are attenuated by approximately a factor of 200 relative to those outside of the shields.

5.4   Magnetic Field Generation and Gradients

Due to the fact that optical magnetometers based on NMOR rely on measuring the Larmor frequency of a spin-polarised alkali vapour, the magnetic-field strength distribution over the cell volume *must* be as homogenous as possible in order to achieve a narrow resonance linewidth and subsequently high sensitivity. If there is a magnetic-field strength *gradient* present within the cell volume, the Larmor frequency becomes *spatially dependent* and hence atoms in various regions of the cell undergo spin precession at different frequencies. Each atom will in turn acquire a different phase between pumping cycles — depending upon their trajectory — which reduces the ground-state coherence time and therefore leads to a broadened resonance linewidth and subsequently reduced magnetic sensitivity.

It has been shown that the spin-relaxation rate depends quadratically on the first-order magnetic field gradient [234]. For a spherical cell of radius R_{cell}, the *longitudinal* spin-relaxation rate, T_1, is given explicitly by [305]

$$\frac{1}{T_1} = \mathcal{D}\frac{|\nabla B_x|^2 + |\nabla B_y|^2}{\bar{B}^2} \, , \tag{5.11}$$

where $\mathcal{D}$ is the diffusion coefficient (which is inversely proportional to the gas pressure), $\bar{B}$ is the mean magnetic-field strength[4] (which is assumed to lie along the z-axis), and ∇B_i is the gradient of the i^{th} component of the magnetic field. At low pressures, when the reciprocal of the Larmor frequency (i.e. the Larmor period) is much larger than the characteristic time required for the atoms to diffuse through the cell, Eq. (5.12) simplifies to [305]

$$\frac{1}{T_1} = \frac{8R_{\text{cell}}^4\gamma_{\text{g}}^2}{175\mathcal{D}}\left(|\nabla B_x|^2 + |\nabla B_y|^2\right) \, , \tag{5.13}$$

where γ_{g} is the ground-state gyromagnetic ratio. In the same low-pressure regime, the *transverse* spin-relaxation time, T_2, is given by [305]

$$\frac{1}{T_2} = \frac{4R_{\text{cell}}^4\gamma_{\text{g}}^2}{175\mathcal{D}}\left(|\nabla B_x|^2 + |\nabla B_y|^2 + 2|\nabla B_z|^2\right) \, . \tag{5.14}$$

In the high-pressure regime, the atoms have a relatively small mean free path, and hence the temporal evolution of the spins is determined by the magnetic field in their local position. However, in the low-pressure regime, the mean free path is large and hence the atoms can sample the cell volume over one Larmor precession cycle. In this case, the atomic Larmor precession frequency is governed by the *average* magnetic field in the cell, rather than at the instantaneous position of the atom. This phenomenon is known as *motional narrowing*, due to the fact that Eqs. (5.13) and (5.14) reveal that $T_1 \propto \mathcal{D}$ and $T_2 \propto \mathcal{D}$. Since $\mathcal{D}$ is inversely proportional to the gas pressure, this implies that when the pressure is low (i.e. when the atoms have a large mean free path), the relaxation time is *increased* for a fixed set of gradi-

---

[4]The mean magnetic-field strength is simply the average value of the field over the cell volume V_{cell}:

$$\bar{B} = \frac{1}{V_{\text{cell}}} \iiint_{\text{cell}} B\left(r\right)\mathrm{d}V \, . \tag{5.12}$$

ents ∇B_i. Another interesting observation to note is that, if the gradient of the longitudinal magnetic field is zero ($\nabla B_z = 0$), then $T_2 = 2T_1$. This is true for all pressures, not just in the low-pressure regime. If however $\nabla B_z \neq 0$, then $T_2 < 2T_1$.

Due to the quadratic dependence of the ground-state relaxation rate on the first-order magnetic-field gradient, significant effort must be taken to both *reduce* existing gradients, and *prevent* creating new ones in the process[5]. The former can be achieved through a combination of magnetic shielding (to provide a quiet magnetic environment), demagnetisation of said shielding, and the use of gradient-compensation coils (such as an anti-Helmholtz coil or saddle coil [306]). The latter is generally achieved through a combination of well-designed field-generation coils, as well as gradient-compensation coils[6].

Perhaps the two most common types of coil used to generate homogeneous magnetic fields are the Helmholtz coil and the solenoid. The solenoid has the benefit of being the easiest coil to make (as it has a simple geometry), as well as a large current-to-field ratio. However; for a solenoid to provide a sufficient region of homogeneous field, the ratio of its length to its radius must be quite large. If coil complexity is not much of an issue, a greater field homogeneity can be achieved — for comparable coil dimensions — by using alternate, more complex coil geometries (e.g. a Lee-Whiting coil [307]). Ultimately, the achievable field uniformity will depend upon the size of the coil, and the minimisation of winding imperfections. Generally speaking, larger coils yield superior field uniformity.

The coil geometry chosen in the experimental work detailed here is a solenoid, with a large (almost 6:1) length-to-radius ratio. The analytic solutions for the off-axis magnetic-field components of a solenoid are non-trivial, and involve either elliptic integrals [308] or infinite series of Legendre polynomials [309]. Therefore, in order to calculate the magnetic-field strength distribution, it is *simpler* to solve the Biot-Savart law numerically. The Biot-Savart law states that the magnetic field, $\boldsymbol{B}\left(\boldsymbol{r}\right)$, produced by a steady-state current I and measured at a position $\boldsymbol{r}$ in three-dimensional space, is given by

$$\boldsymbol{B}\left(\boldsymbol{r}\right) = \frac{\mu_0 I}{4\pi} \int_{\mathcal{P}} \frac{\mathrm{d}\boldsymbol{r}' \times \left(\boldsymbol{r} - \boldsymbol{r}'\right)}{\left|\boldsymbol{r} - \boldsymbol{r}'\right|^{3/2}}, \tag{5.15}$$

where μ_0 is the vacuum permeability, and $\boldsymbol{r}'$ is a vector along the path $\mathcal{P}$ which points in the direction of current flow. The Biot-Savart law is incredibly powerful, as it allows the magnetic field of a coil with any arbitrary shape to be numerically calculated, provided that we can write a parametric equation for the path $\mathcal{P}$. In general, we write the parametric equation as a function of t in Cartesian coordinates, which takes the form

$$\boldsymbol{r}' = \left\langle x', y', z' \right\rangle = \left\langle f\left(t\right), g\left(t\right), h\left(t\right) \right\rangle. \tag{5.16}$$

---

[5]One of Maxwell's equations states that the divergence of the magnetic-field vector is zero. Mathematically, this is described by $\nabla \cdot \boldsymbol{B} = \frac{\partial B_x}{\partial x} + \frac{\partial B_y}{\partial y} + \frac{\partial B_z}{\partial z} = 0$. This implies that, if a magnetic-field gradient is created along one axis, there *must necessarily* be a gradient created along the other axes to satisfy the condition that the divergence is zero.

[6]It can be relatively straightforward to compensate for linear gradients using compensation coils; however, if there are higher-order gradients (cf. Fig. 5.10), this process can be complex — if not impossible.

Figure 5.9: A scale model of the solenoid and former around which it is wrapped. The
solenoid has a radius of 56.5 mm, a length of 317 mm, and approximately 370
turns. Refer to Fig. 5.4 for a detailed cross-section of the cell, solenoid and μ-
metal shields.

The differential $\mathrm{d}\boldsymbol{r}'$ is thus given by $(\mathrm{d}\boldsymbol{r}/\mathrm{d}t)\,\mathrm{d}t$ in the parametric coordinate t. The full nu-
merical Biot-Savart law is therefore given by

$$\boldsymbol{B}\left(\boldsymbol{r}\right)=\frac{\mu_{0}I}{4\pi}\int_{a}^{b}\frac{\langle\frac{\mathrm{d}f}{\mathrm{d}t},\frac{\mathrm{d}g}{\mathrm{d}t},\frac{\mathrm{d}h}{\mathrm{d}t}\rangle\times\left(\langle x,y,z\rangle-\langle f\left(t\right),g\left(t\right),h\left(t\right)\rangle\right)}{\left|\langle x,y,z\rangle-\langle f\left(t\right),g\left(t\right),h\left(t\right)\rangle\right|^{3/2}}\,\mathrm{d}t\,. \tag{5.17}$$

where a and b are points which characterise the beginning and end of the coil.

A solenoid is simply a circular helix which, given a longitudinal axis along $\hat{\boldsymbol{z}}$ and a trans-
verse plane lying in the xy-plane, is described by the following parametric equation:

$$\langle f\left(t\right),g\left(t\right),h\left(t\right)\rangle=\left\langle r_{\mathrm{c}}\cos\left(t\right),r_{\mathrm{c}}\sin\left(t\right),\left(\frac{\Delta W}{2\pi}\right)t\right\rangle, \tag{5.18}$$

where r_{c} is the radius of the coil, and ΔW is the spacing between consecutive winds of the
coil. The Cartesian coordinates of the magnetic-field vector can be subsequently calculated
by computing the inner product with the $\hat{\boldsymbol{x}}$, $\hat{\boldsymbol{y}}$, or $\hat{\boldsymbol{z}}$ unit vectors (e.g. $B_{x}\left(\boldsymbol{r}\right)=\boldsymbol{B}\left(\boldsymbol{r}\right)\cdot\hat{\boldsymbol{x}}$).
Note that, due to the sinusoidal nature of Eq. (5.18), one turn of coil corresponds to an
integration of t through 2π. The integration variable therefore needs to be calculated from
$a=-2\pi L/2\Delta W$ to $b=2\pi L/2\Delta W$, where L is the length of the solenoid.

The solenoid used in this work — a scale model of which is presented in Fig. 5.9 —
has a radius of 56.5 mm and a length of 317 mm, yielding approximately 370 turns. Using

Figure 5.10: Calculated deviation of the magnetic-field strength along the longitudinal (blue) and radial (red) axes, over the region containing the vapour cell, calculated using Eqs. (5.17) and (5.18). The longitudinal and radial magnetic-field strength deviations were calculated at the centre of the transverse and longitudinal planes, respectively. Note that there is a minor numerical instability in the solution of the magnetic-field strength along the radial axis.

Eqs. (5.17) and (5.18), the magnetic-field strength deviation across the region containing the vapour cell is shown in Fig. 5.10. Through differentiation of the data presented in Fig. 5.10, it is possible to calculate the longitudinal and radial magnetic-field strength gradients. These gradients, evaluated at a dc magnetic-field strength of $2.5\,\mu$T, are presented in Fig. 5.11. Residuals of a linear fit indicate the presence of higher-order gradients, which are arduous to compensate.

5.4.1   Power Supply

Given that the photon shot-noise limit of the magnetometer described in this thesis is a fractional sensitivity of $\mathcal{O}\left(10\,\mathrm{ppb}/\sqrt{\mathrm{Hz}}\right)$, it is impossible — at least at the time of writing — to produce a voltage/current supply that will perform at the same level at low Fourier frequencies. This necessitates the use of either active stabilisation (e.g. through the use of a transconductance amplifier), or passive stabilisation via filtering, in order to suppress the voltage/current fluctuations to a level that is acceptable over the frequency range of interest.

At the heart of the power supply used to drive the solenoid lies an ultra-precision voltage reference — the Analog Devices LTZ1000 — which produces a stable output voltage using a temperature-stabilised subsurface Zener diode. The output voltage of the LTZ1000 is filtered using a Butterworth filter with a measured cut-off frequency of $\nu_{\mathrm{co}} = 0.30\,\mathrm{mHz}$ (cf. Sec. 8.2), in order to suppress any residual voltage/current fluctuations. The performance of the power supply, and its effective magnetic-field noise, is discussed in detail in Sec. 8.1.1.

Figure 5.11: Calculated gradients of the magnetic-field strength along the longitudinal (blue) and radial (red) axes, over the region containing the vapour cell, at a dc magnetic-field strength of 2.5 µT. The gradients were calculated via differentiation of a cubic spline interpolation of the data presented in Fig. 5.10. The residuals of a linear fit are shown at the top, indicating the presence of higher-order gradients.

5.5 Balanced Polarimetry

Performing sensitive magnetometry with an NMOR magnetometer relies on making exquisitely precise measurements of the polarisation state of the light. More specifically, the *frequency* of the optical rotation of the linearly polarised probe light must be accurately and precisely measured. Given that typical optical-rotation amplitudes are $\mathcal{O}\,(10\,\text{mrad})$, this can be an arduous task. However, the use of a *balanced polarimeter* can yield angular sensitivity of better than $10^{-8}\,\text{rad}_{\text{rms}}/\sqrt{\text{Hz}}$ [299].

At the output of the vapour cell, the pump beam is discarded while the probe beam is collected and analysed by a balanced polarimeter consisting of a Wollaston prism and two photodetectors. The Wollaston prism splits the probe beam into its two orthogonal linear polarisation components, and each component is measured using a separate photodetector. The polarimeter is arranged such that the optical power on each photodetector is equal; in which case, if the polarisation axis of the probe light rotates, a misbalance of optical powers is measured at the photodetectors, which can then be used to determine the polarisation-rotation angle.

In order to derive an expression for the output of a balanced polarimeter, consider defining P_0 as the sum of the two photodetector signals P_1 and P_2. This allows us to write $P_1 = P_0 \cos^2 \phi$ and $P_2 = P_0 \sin^2 \phi$. If we now take the difference of the two photodetector signals and apply standard double-angle trigonometric identities, we have

$$P_1 - P_2 = P_0 \cos 2\phi. \tag{5.19}$$

Provided the optical powers are balanced at the output of the Wollaston prism, then the polarisation angle ϕ is given by

$$\phi = \frac{\pi}{4} + \delta\phi, \tag{5.20}$$

where $\delta\phi$ is some small change in polarisation angle. Now if we substitute this expression for ϕ into our expression for $P_1 - P_2$, we have

$$P_1 - P_2 = P_0 \cos\left(\frac{\pi}{2} + 2\delta\phi\right)$$
$$= -P_0 \sin(2\delta\phi),$$

Now making use of the fact that $P_0 = P_1 + P_2$, and taking the magnitude of the rotation angle, we have

$$|\delta\phi| = \frac{1}{2} \arcsin\left(\frac{P_1 - P_2}{P_1 + P_2}\right). \tag{5.21}$$

Equation (5.21) enables the optical-rotation angle to be calculated using the optical powers measured by the photodetectors, and is valid in *both* the time domain *or* the modulation-frequency domain.

Experimental Techniques

This chapter provides an in-depth overview of key experimental techniques used in the work detailed within this thesis, beginning with time- and frequency-domain measurement and analysis techniques. This subsequently leads into a discussion about magnetometer performance, including topics such as: noise characterisation, performance metrics, fundamental sensitivity, parameter optimisation, and bandwidth.

6.1 Data Acquisition and Analysis

Depending on the particular experimental quantity or phenomenon wished to be observed, experimental data can be acquired and analysed in various measurement domains, using a myriad of acquisition techniques. In some cases, time-domain measurements are performed by recording the two photodetector signals (from the balanced polarimeter) on an oscilloscope, and subsequently calculating the time-dependent optical-rotation signal in post-processing. This is particularly useful when observing *transient* signals, or signals with rich spectral content. However, in the case that the observed signal is a steady-state oscillation at a single known frequency — as is typically the case when performing magnetometry via synchronous optical pumping on a resonance condition — it becomes more convenient to measure the signal in the frequency domain using a phase-sensitive detector such as a lock-in amplifier. These two classes of measurement techniques are discussed below.

6.1.1 Time-Domain Measurements

Perhaps the most general class of measurement techniques used for the purpose of investigating nonlinear magneto-optical effects — especially transient signals — are time-domain measurements. As the name suggests, a time-domain measurement is performed by observing the optical rotation (i.e. the output of the balanced polarimeter), as a function of time. This is typically either done via a free-induction decay measurement in which the atomic vapour is optically pumped for some period of time and then allowed to freely decay, or via synchronous optical pumping on a resonance condition.

6.1.1.1 Free-Induction Decay

A free-induction decay measurement is one in which the atomic vapour is optically pumped for some period of time[1] and, when the pump is switched off, the free temporal evolution of the optical-rotation signal is observed with a weak probe. As discussed in Sec. 4.8, the

---

[1]This period of time is generally the time in which it takes the optical pumping to reach steady state.

optical-rotation observable — under the experimental conditions detailed within this thesis — is proportional to a sum over the coherences between ground-state Zeeman sublevels which satisfy $|\Delta m_F| = 2$. As such, the temporal evolution is sinusoidal at a frequency of $2\Omega_L$, as governed by the time-dependent Schrödinger equation.

Due to relaxation processes (cf. App. 5.2.2), the optical-rotation signal is *transient*, and decays exponentially with a time constant characterised by T_2, the transverse spin-relaxation time[2]. Assuming a static magnetic field (i.e. $dB/dt = 0$), this gives rise to an optical-rotation signal which takes the form of a damped, single-frequency sinusoid. This is explicitly given by

$$\phi(t) = \phi_0 \sin(2\Omega_L t + \varphi) \exp\left(-\frac{t}{T_2}\right), \tag{6.1}$$

where ϕ_0 is the peak optical-rotation amplitude at $t = 0$, Ω_L is the Larmor frequency, φ is an arbitrary phase offset, and T_2 is the transverse spin-relaxation time.

An example of an FID measurement is shown in Fig. 6.1. Although the FID signal is transient, it is still useful for at least two applications; transient measurements of the magnetic-field strength (cf. Sec. 2.2.1), and characterisation of the quality of anti-relaxation-coated cells (cf. App. F). In the former case, generally a *train* of FID measurements is performed [140], which has been demonstrated to enable quasi-static (or slowly varying) magnetic fields to be resolved [144]. However, this technique is limited to low frequencies due to aliasing [144], and often requires *intentional destruction* of ground-state coherence in order to increase *both* the amplitude and frequency response[3].

6.1.1.2   Synchronous Optical Pumping

Rather than pumping once (or for some short period of time) and allowing the signal to *freely* evolve/decay, a quasi-static signal can be obtained by pumping approximately once *every* oscillation cycle, i.e. $\Omega_m \approx 2\Omega_L$. This technique is known as *synchronous optical pumping*, and gives rise to an output signal which behaves like a damped, driven harmonic oscillator. This gives rise to a resonant behaviour in which the signal amplitude is maximised *exactly* when the pump modulation frequency equals the oscillation frequency (i.e. $\Omega_m = 2\Omega_L$), thereby enabling the magnetic-field strength to be determined by finding the pumping frequency at which the signal is maximised. An example of a typical synchronously pumped optical-rotation signal is presented in Fig. 6.2.

However, performing magnetometry in the time domain by synchronously optically pumping and monitoring the signal amplitude is *only possible* provided that the modulation frequency is *exactly* on resonance. If this is not the case, there will be an error in the predicted magnetic-field strength, relative to its true value. Although this is true for frequency-domain

---

[2]The optical-rotation signal is affected by the transverse spin-relaxation time T_2, rather than the longitudinal spin-relaxation time T_1, because the atoms are spin-polarised and subsequently undergo Larmor precession in the plane orthogonal to both the light propagation vector and magnetic-field direction (cf. Sec. 6.4.1).

[3]This is generally unwanted as, in accordance with Eq. (6.18), the magnetic sensitivity is negatively affected as a result of reduced ground-state coherence time through the relation $\delta B_{SQL} \propto 1/\sqrt{T_2}$ (cf. Sec. 6.4.1).

Figure 6.1: Left: Free-induction decay of the optical-rotation signal (blue) measured in a single shot after the pump beam (grey) has been switched off. Right: First 200 µs of the FID data, with a fit using Eq. (6.1). The fit returns a transverse spin-relaxation time of $T_2 = 43.4$ ms (cf. App. F) and a Larmor frequency of $\Omega_L/2\pi = 15.31$ kHz, which corresponds to a magnetic-field strength of $B_{dc} = 2.188$ µT.

measurements as well, this issue is *further compounded* in the time domain by the fact that the signal amplitude is *symmetric* about the resonance centre — i.e. it is not possible to determine whether the magnetic field has *increased* or *decreased*. Unfortunately, it is difficult to ensure that the resonance condition is met without being able to observe the complete resonance profile. For this reason, it is generally recommended to perform synchronous-optical-pumping measurements in the frequency domain using frequency-demodulation techniques.

6.1.1.3   Subharmonic Optical Pumping

In fact, not only is there a resonance observed when $\Omega_m = 2\Omega_L$, but resonances can also occur *any time* the condition $\Omega_m = 2\Omega_L/n$ where $n \in \mathbb{Z}^* \setminus \{0\}$ is satisfied. Simply put, a resonance can occur when the pump modulation frequency is any subharmonic of the quantum-beat frequency[4]. An example of the optical-rotation signal observed when pumping at the 10th subharmonic of the quantum-beat frequency is presented in Fig. 6.3.

---

[4]Resonances can also be observed when pumping at higher-order harmonics of the Larmor frequency, i.e. $\Omega_m = p2\Omega_L$ for $p \in \mathbb{Z}^* \setminus \{0\}$; however, these resonances are associated with higher-order polarisation moments [245,285,286,310] (hexadecapole, hexacontatetrapole, etc.) and require multi-photon interactions in a ground state which can support their associated $|\Delta m_F| = 2p$ coherences.

Figure 6.2: Synchronously pumped optical-rotation signal measured after 2,048 averages (blue), with the pump beam's square-wave modulation shaded in grey. The data has been fit using a single-frequency sinusoid, yielding a Larmor frequency of $\Omega_L/2\pi = 20.00\,\text{kHz}$, which corresponds to a magnetic-field strength of $B = 2.859\,\mu\text{T}$. The residuals to a sinusoidal fit contain a small amount of second-harmonic contamination due to detection nonlinearity.

6.1.2 Frequency-Domain Measurements

In the case where the signal is quasi-static (e.g. a single-frequency sinusoid that occurs when performing synchronous optical pumping), it becomes advantageous to use phase-sensitive detection techniques (cf. Sec. 4.11 and App. K). That is, rather than measuring the time-dependent signal across a broad range of Fourier frequencies, the signal is demodulated at a single reference frequency[5]. The immediate benefit of measuring in this way is that it is a *narrowband* measurement, i.e. only the signal at one particular Fourier frequency of interest is observed. This significantly enhances the SNR, as only the noise within a small bandwidth around the demodulation frequency is measured. If sensitive magnetometry is to be performed, this generally becomes the measurement technique of choice.

There are two ways in which to observe the magneto-optical resonance in the modulation-frequency domain: either by holding the pump modulation frequency fixed and scanning the magnetic-field strength across the resonance condition, or by holding the magnetic-field strength fixed and scanning the pump modulation frequency[6]. During the sweep, the

---

[5]This reference (or demodulation) frequency is typically chosen to be either the known carrier frequency of the oscillatory signal, or the frequency of an applied modulation (or a harmonic thereof).

[6]Care must be taken when doing this because, if either the modulation frequency or magnetic-field strength are swept too rapidly, there appear transient oscillatory features in the resonance [311–313]. These oscillations arise due to fact that the atomic medium does not reach equilibrium before significant changes in the physical parameters occur [311–313].

Figure 6.3: Optical-rotation signal (blue) measured after 1,024 averages when pumping at the 10^{th} subharmonic of quantum-beat frequency, $\Omega_{\text{m}}/2\pi \approx (2\Omega_{\text{L}}/10)/2\pi = 2.71\,\text{kHz}$. The pump beam modulation is shaded in grey. Note that there is negligible relaxation between pulses due to a large tranverse spin-relaxation time of $T_2 \approx 40\,\text{ms}$ under the experimental conditions under which this data was measured.

optical-rotation amplitude is recorded, which exhibits a resonant enhancement of the amplitude when the resonance criteria, $\Omega_{\text{m}} = 2\Omega_{\text{L}}$, is met. An example of a typical frequency-domain resonance measurement is shown in Fig. 6.4.

As discussed in App. K, demodulation with a lock-in amplifier yields two outputs: one that is *in phase* with the reference frequency, and one that is shifted $90°$ *out of phase*. As their names suggest, these two channels are called the 'in-phase' (ϕ_P) and 'quadrature' (ϕ_Q) components, respectively. The in-phase and quadrature components, ϕ_P and ϕ_Q, of the NMOR resonance studied here are well approximated by the *real* and *imaginary* components respectively of a complex-Lorentzian lineshape $\mathcal{L}_{\text{c}}$ defined by

$$\mathcal{L}_{\text{c}}\left(\omega - \omega_0\right) = \mathcal{A}\left[\frac{\left(\frac{\Gamma}{2}\right)^2}{\left(\omega - \omega_0\right)^2 + \left(\frac{\Gamma}{2}\right)^2} + i\,\frac{\frac{\Gamma}{2}\left(\omega - \omega_0\right)}{\left(\omega - \omega_0\right)^2 + \left(\frac{\Gamma}{2}\right)^2}\right], \tag{6.2}$$

where $\mathcal{A}$ is the resonance amplitude, Γ is the resonance width (FWHM), and ω_0 is the resonant angular frequency.

A measurement of the dc magnetic-field strength using frequency-domain techniques simply corresponds to locating the resonant frequency ω_0, as this resonant frequency coincides with the satisfaction of the resonance condition $\Omega_{\text{m}} = 2\Omega_{\text{L}} = 2\gamma_{\text{g}}B$. In principle, this can be achieved in one of two ways: by either measuring the modulation frequency which maximises the in-phase component, or by measuring the frequency of the zero-crossing of

Figure 6.4: Typical NMOR resonances as measured in the frequency domain using a lock-in amplifier, at low (left) and high (right) magnetic-field strength. The measured data is shown in grey markers, while fits using Eq. (6.2) are shown in the coloured traces. Fits to the in-phase and quadrature components, ϕ_P and ϕ_Q, are shown in blue and red respectively, while a fit to the resonance magnitude — given by the quadrature sum of the two components — is shown in purple. Note that the fits on the right-hand side are simply a sum of three individual complex-Lorentzian profiles. The residuals of the fits to the resonance magnitudes (purple markers) show residual structure, indicating that the complex-Lorentzian function is only an approximation to the measured lineshape. The data on the right was reproduced, with permission, from Ref. [314].

the quadrature component. However, to first-order the peak of the in-phase component has *no* sensitivity to small changes in magnetic-field strength. Additionally, the in-phase component is *symmetric*, which means that both increases *or* decreases in magnetic-field strength are *indistinguishable*. The quadrature component on the other hand has an amplitude which changes linearly (and asymmetrically) in response to small changes in magnetic-field strength, with a sensitivity which is determined, in part, by the slope[7] (cf. Eq. 6.24). The quadrature component therefore enables the resonant frequency to be measured with high precision, yielding a sensitive magnetic-field strength measurement.

Generally speaking, if short-term magnetic-field measurements are to be performed, this is typically best achieved by tuning the pump modulation frequency to the resonance condition by locating the zero-crossing of the quadrature component. Once this has been located, the voltage fluctuations of this component are measured, and the resonance slope is used to

---

[7]The sensitivity of the magnetometer will therefore be, in part, determined by the resonance slope. It is therefore paramount to ensure that the parameters of the experiment are set to their optimal configuration, which corresponds to the optimal trade-off between resonance slope and resultant measurement noise. This is discussed in further detail in Sec. 6.4.3.

convert the voltage fluctuations to effective magnetic-field fluctuations. However, if one is to perform long-term measurements, or measurements of ac magnetic fields with amplitudes or frequencies which exceed the resonance bandwidth, it is beneficial to use a 'resonance locking' technique.

6.1.3 Resonance Locking

As discussed in previous sections, rapid magnetic-field fluctuations (e.g. above 1 Hz) are readily observed by tuning to the resonance condition, $\Omega_m = 2\Omega_L$, and measuring the zero-crossing voltage fluctuations of the quadrature component of the lock-in amplifier output. This works sufficiently well in the short term; however, over long timescales the dc value of the magnetic field may wander significantly such that the resonance condition is no longer met. In this circumstance, it is beneficial to implement a locking system which ensures that the resonant optical pumping condition is *always* satisfied.

In the experiments detailed within this thesis, resonance locking has been achieved by using the apparatus presented in Fig. 6.5. Rather than measuring the voltage noise of the quadrature output directly, this signal is instead fed to a locking servomechanism. By setting the locking point such that the quadrature component is zero, the pump-modulation frequency is adjusted by feeding the locking-servo control signal to the signal generator that provides the amplitude modulation for the pump beam. If the quadrature output drifts away from zero, the error signal — defined as the difference between the desired set point and the measured value — becomes non-zero. A proportional-integral controller computes the control signal from this error signal, which is fed to the modulation port[8] of the signal generator and subsequently adjusts the modulation frequency in order to compensate for the field change. When locked, rather than measuring the magnetic-field strength via voltage fluctuations of the quadrature component, the resonant modulation frequency is measured directly instead. This measurement technique has the added benefit of being intrinsically calibration-free, and hence does not require an experimentally determined calibration slope in order to convert the magnetometer output into effective magnetic-field units.

6.2 Noise Characterisation

In order to assess the performance of a measurement device, it is necessary to characterise both the magnitude of noise present within the system, as well as its temporal and/or frequency dependence. Combined with the signal amplitude, the noise level determines the signal-to-noise-ratio of the device and, subsequently, its performance.

There are two broad classes of techniques for characterising the noise of a device: frequency-domain techniques, and time-domain techniques. Both of these classes of techniques have advantages and disadvantages, and one can be more advantageous than the other depending on the situation. Frequency-domain techniques are generally useful for determining the frequency components present within the noise spectrum of a system, which

---

[8]The sensitivity of the modulation port can be set as required, but in most measurements is around 4 Hz/V.
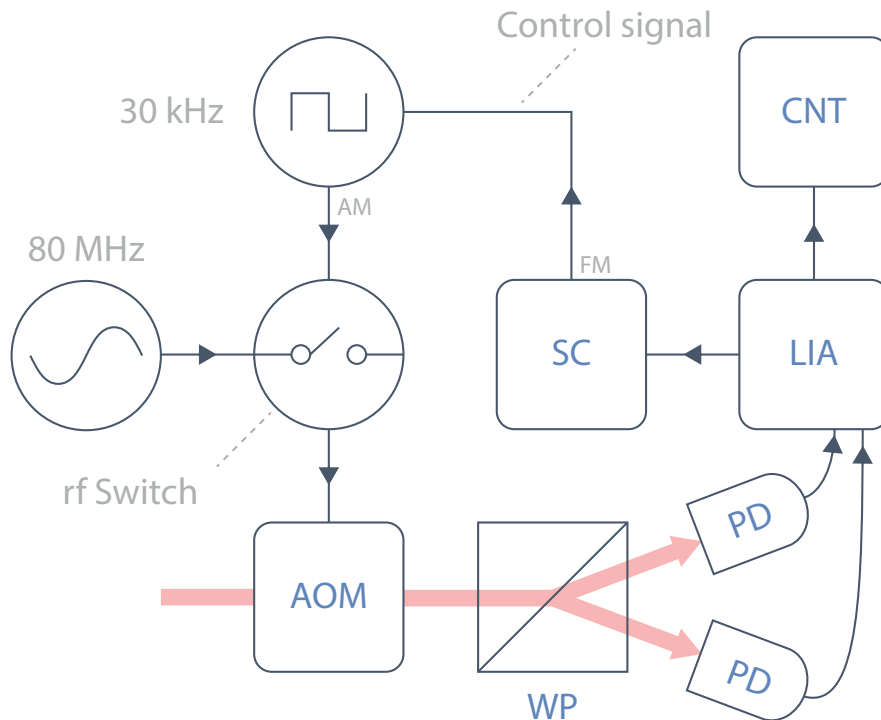
Figure 6.5: Key components of an NMOR resonance lock: AOM — acousto-optic modulator, PD — photodetector, SC — servo controller, LIA — lock-in amplifier, CNT — frequency counter, WP — Wollaston prism. The AM and FM labels indicate amplitude modulation and frequency modulation, respectively.

is critical when attempting to isolate the origin of performance-limiting noise. On the other hand, time-domain techniques are useful for determining the expected drift of the output signal as a function of time.

This chapter introduces two of the main noise-characterisation techniques — one frequency-domain technique, and one time-domain technique — that are used frequently throughout this thesis. These techniques are the amplitude/power spectral density, and the Allan deviation/variance, respectively. The Cramér-Rao lower bound and its application to magnetometry is also discussed.

6.2.1   Amplitude and Power Spectral Density

The power spectral density is a frequency-domain measurement which quantifies the amount of signal *power* contained in each frequency component within the signal of interest. Similarly, the amplitude spectral density quantifies the amount of signal *amplitude* contained in each frequency component. These are useful in the context of noise characterisation, as they can be used to quantify both the *magnitude* and *frequency dependence* of noise present within the system. The former is necessary when calculating the SNR — and by proxy, the sensitivity — of the magnetometer (cf. Sec. 6.3.1), while the latter is specifically useful for locating any performance-limiting noise contributions via the fingerprint that their frequency dependence provides (cf. App. H.3).

In practice, computing the amplitude/power spectral density of a signal requires an ex-

Figure 6.6: Amplitude spectral density of the magnetometer noise floor (blue) calculated via Welch's method (cf. App. H.1.1.3), using a 26-h continuous measurement. All performance-limiting technical noise sources discussed in Sec. 8 are presented here, as well as fundamental noise sources: power-supply noise (green), shield noise (grey), Johnson noise (purple), photon shot noise (red), and spin-projection noise (orange). Note that the Johnson noise is low-pass filtered due to the bandwidth of the magnetometer under the experimental conditions (cf. Sec. 6.6.2).

perimental measurement of the signal that has been sampled at a known frequency. Given the known sample rate, the amplitude/power spectral density of the data can then be computed using Fourier transforms via the methods discussed in Apps. G and H.1. An example of an amplitude spectral density — which is the square root of the power spectral density — for the magnetometer over a 26-h measurement is presented in Fig. 6.6.

6.2.2 Allan Deviation and Allan Variance

The Allan variance is a time-domain measurement which quantifies the variance in the value of a signal as a function of integration time. Simply put, given two measurements spaced apart by some period of time, the Allan variance determines the expected variance (or spread) of these measurements. This essentially constitutes a stability measurement and, in the context of magnetometry, this is useful for characterising the long-term performance of the magnetometer.

In practice, the requirements for computing an Allan variance are the same as per the power spectral density — an experimental measurement of the signal that has been sampled at a known frequency. Given this information, the Allan variance can be computed using the method discussed in App. H.2. An example of an Allan deviation — the square root of the Allan variance — for the magnetometer over a 26-h measurement is shown in Fig. 6.7.

Figure 6.7: Allan deviation of the magnetometer measured over a 26-h period (blue trace). The Allan deviation increases monotonically, approximately in proportion to τ^1 (black, dashed), for integration times $\tau \gtrsim 10\,\text{s}$ — indicative of long-term drift.

6.2.3   Cramér-Rao Lower Bound

The Cramér-Rao lower bound (CRLB) is extremely useful for estimating the lower bound on the variance (and therefore standard deviation) of an unbiased estimator of a deterministic parameter [315–318]. Simply put, for a given set of experimental parameters, the CRLB estimates the maximum precision to which one can measure an unknown parameter — subject to Gaussian white noise — in a finite measurement time. This can subsequently be used as a noise metric in terms of calculating the *absolute* and/or *fractional uncertainty* of a given measurand — for example, the Larmor frequency.

A derivation of the CRLB for a general signal $s[n]$ is given in App. H.5. Of particular interest in the context of magnetometry, is the measurement of the frequency of a sinusoid in the presence of Gaussian white noise. This is described mathematically as $s[n] = a\cos(2\pi\nu\Delta tn + \varphi) + w[n]$. As derived in App. H.5.3.2, the CRLBs in this case are given by [315]

$$\text{var}(\hat{a}) \geq \frac{2\sigma^2}{N}, \tag{6.3}$$

$$\text{var}(\hat{\nu}) \geq \frac{24\sigma^2}{(2\pi a\Delta t)^2\, N\,(N^2 - 1)}, \tag{6.4}$$

$$\text{var}(\hat{\varphi}) \geq \frac{4\sigma^2\,(2N - 1)}{a^2 N\,(N + 1)}, \tag{6.5}$$

where a is the sinusoidal amplitude, ν is the linear frequency, Δt is the sample time of the

signal, φ is an arbitrary phase offset, N is the total number of data points, and σ^2 is the variance of the Gaussian white noise $w[n]$.

As discussed in Secs. 3.6 and 3.7, an NMOR magnetometer determines the local magnetic-field strength by measuring a sinusoid with an angular frequency of $2\pi\nu = 2\Omega_{\rm L} = 2\gamma_{\rm g}B$. By expanding the denominator of Eq. (6.4), and noting that $N^3 \gg N$, the following simplification can be made:

$$\text{var}(\hat{\nu}) \gtrsim \frac{12\sigma^2}{(2\pi)^2 a^2 \nu_{\rm BW} T^3}, \tag{6.6}$$

where $T = \Delta t N$ is the measurement time, and $\nu_{\rm BW} = (2\Delta t)^{-1}$ is the measurement bandwidth (Nyquist frequency). In the case of white noise, Eq. (H.10) can be used to write $\sigma^2 = S_{\rm PSD}(\nu)\,\nu_{\rm BW}$. Substituting this into Eq. (6.6), and converting into magnetic-field units, results in the following standard deviation:

$$\sigma_B \gtrsim \frac{\sqrt{3 S_{\rm PSD}(\nu)}}{a\gamma_{\rm g} T^{3/2}}. \tag{6.7}$$

Equation (6.7) is a rather profound result. It claims that for a dc magnetic field which gives rise to a single-frequency sinusoidal signal, the uncertainty in the measured value of the field scales with measurement time as $T^{3/2}$.

6.3 Performance Metrics

Given that there are numerous different types of magnetometer — each with their own characteristic advantages and disadvantages — it becomes necessary to use *multiple* performance metrics in order to draw fair comparisons between them. This section aims to introduce some of the main performance metrics, to then calculate these quantities for the magnetometer described within this thesis, and subsequently make comparisons with state-of-the-art devices found in the literature.

6.3.1 Absolute Sensitivity

The most commonly used measure of magnetometer performance is the absolute sensitvity. That is, the minimum detectable magnetic field per unit bandwidth, denoted δB. This quantity describes the amplitude spectral density (cf. Sec. H.1) of the magnetometer in magnetic units and, for an NMOR magnetometer using linearly polarised light, is calculated from the minimum detectable optical rotation per unit bandwidth, $\delta\phi$, via [206] (cf. Sec. 6.4.2)

$$\delta B = \frac{\delta\phi}{2\mathcal{S}\gamma_{\rm g}}, \tag{6.8}$$

where $\mathcal{S}$ is the slope of the resonance and $\gamma_{\rm g}$ is the ground-state gyromagnetic ratio. For the $F = 2$ ground state of ^{87}Rb — the ground state used in all of the measurements performed throughout this thesis — the gyromagnetic ratio is given by $\gamma_{\rm g}/2\pi = 6.9958\,\text{GHz/T}$ [229].

The photon shot-noise-limited absolute sensitivity of the magnetometer detailed within this thesis is $15\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ [206] (cf. Sec. 6.4.2), which, at the time of writing, is the highest demonstrated absolute sensitivity of an NMOR magnetometer [206]. For comparison, the highest absolute sensitivity achieved by a SERF magnetometer (in a gradiometer configuration) is $160\,\mathrm{aT_{rms}}/\sqrt{\mathrm{Hz}}$ [168], which is two orders of magnitude superior to the NMOR magnetometer presented here. However; one of the trade-offs for this increase in absolute sensitivity is a significant reduction in dynamic range. Spin-exchange relaxation-free magnetometers must be operated in near-zero magnetic field ($B_{\mathrm{dc}} \lesssim 1\,\mathrm{nT}$), which is *at least* 3–4 orders of magnitude less than the ambient fields permitted by the NMOR technique. This necessitates the use of an ambient-field-normalised sensitivity, the *fractional* sensitivity, which takes into account the difference in operational ambient fields.

6.3.2   Fractional Sensitivity

As alluded to in Sec 6.3.1, absolute sensitivity is *not necessarily* the best measure of performance. In many applications, sensitive measurements of magnetic fields *must* be performed within large ambient dc fields, which prevents the use of magnetometers that are restricted to operating in near-zero ambient fields (e.g. SERF magnetometers). For this reason, it is useful to define a fractional sensitivity, which normalises the absolute sensitivity by the ambient field strength. The mathematical definition of the fractional sensitivity is therefore given by

$$\delta B_{\mathrm{frac}} = \sqrt{2}\left(\frac{\delta B}{B_{\mathrm{dc}}}\right), \qquad (6.9)$$

where the factor of $\sqrt{2}$ converts the magnetic sensitivity from root mean square (rms) to amplitude, in order to be compared with the dc field B_{dc}. The magnetometer described in this thesis achieves fractional sensitivity of $\delta B_{\mathrm{frac}} = 9\,\mathrm{ppb}/\sqrt{\mathrm{Hz}}$ in an ambient field of $B_{\mathrm{dc}} = 2.5\,\mathrm{\mu T}$ [206]. At the time of writing, this is the highest demonstrated fractional sensitivity for a room-temperature optical magnetometer of any kind. For comparison, this fractional sensitivity is about 2 orders of magnitude superior to state-of-the-art SERF magnetometers which generally operate in dc fields of $B_{\mathrm{dc}} \lesssim 1\,\mathrm{nT}$.

6.3.3   Spatiotemporal Resolution

In situations where spatial resolution is important (e.g. magnetic resonance imaging), it is necessary to normalise the absolute sensitivity by the active measurement volume, in order to quantify the spatiotemporal resolution of the device. The sensitivity per spatiotemporal bandwidth is given by $\delta B \sqrt{V}$ [153,319], where δB is the absolute sensitivity and V is the active measurement volume. There is some ambiguity here as to whether to use the volume of the cell, or the active measurement volume defined by the probe beam. For the NMOR magnetometer detailed within this thesis, the spatiotemporal resolution is $\delta B \sqrt{V_{\mathrm{cell}}} = 106\,\mathrm{fT\,cm^{3/2}}/\sqrt{\mathrm{Hz}}$ using the volume of the cell, whereas using the active measurement volume defined by the probe beam yields $\delta B \sqrt{V_{\mathrm{probe}}} = 4\,\mathrm{fT\,cm^{3/2}}/\sqrt{\mathrm{Hz}}$.

Given that NMOR magnetometers conventionally use room-temperature alkali vapour which has a low optical depth, one must typically either multipass the probe beam, or use a cell with a relatively long path length in order to achieve high absolute sensitivity. The latter of these two solutions, unfortunately, reduces the spatiotemporal resolution. In the literature, the optical magnetometers which frequently achieve high spatiotemporal resolution are SERF magnetometers and diamond magnetometers. Spin-exchange relaxation-free magnetometers require high vapour temperatures ($\sim 150°$C) in order to suppress spin-exchange collisions (cf. Sec. 2.2.3), which results in large optical depth and therefore permits small measurement volumes. On the other hand, diamond magnetometers — despite their comparatively poor *absolute* sensitivities — have *extremely small* measurement volumes of around $(13 \times 200 \times 2000)\,\mu m^3$ [73]. For these reasons, SERF magnetometers have demonstrated spatiotemporal sensitivities of around $0.1\,fT\,cm^{3/2}/\sqrt{Hz}$ [168,169], while diamond magnetometers have demonstrated $34.2\,fT\,cm^{3/2}/\sqrt{Hz}$ [73] — with a projected quantum limit of $22\,aT\,cm^{3/2}/\sqrt{Hz}$ [73].

6.3.4   Energy Resolution

Building upon the concept of spatiotemporal resolution, it is also possible to quantify the *energy* resolution of the device. This idea was first proposed in Ref. [320], in which the energy per unit bandwidth is defined by [169,319,321]

$$E_R = \frac{\left(\delta B \sqrt{V}\right)^2}{2\mu_0}\,, \tag{6.10}$$

where $\delta B \sqrt{V}$ is the minimum detectable magnetic field per unit spatiotemporal bandwidth, and μ_0 is the vacuum permeability. Again, there is some ambiguity as to which volume to use here: the volume of the cell, or the active measurement volume defined by the probe beam. For the NMOR magnetometer described in this thesis, the energy resolution calculated using the volume of the cell is $E_R = 4.27 \times 10^7 \hbar$, while the same calculation using the active measurement volume yields $E_R = 6 \times 10^4 \hbar$. Since the energy resolution scales linearly with the measurement volume, and quadratically with the absolute sensitivity, it should come as no surprise that SERF magnetometers achieve the highest energy resolution out of any optical magnetometer. The best reported result in the literature is $E_R = 44\hbar$ [168], which was achieved through an absolute sensitivity of $160\,aT_{rms}/\sqrt{Hz}$ in an active measurement volume of $0.45\,cm^3$ [168].

6.4   Fundamental Sensitivity and Parameter Optimisation

When performing magnetometry using an optical magnetometer, there are fundamental noise sources which limit the sensitivity. These fundamental noise sources can be split into two categories: classical noise and quantum noise. Generally, one will optimise various experimental parameters in order to minimise the contribution from both of these sources.

Figure 6.8: Illustration of Larmor precession in the plane orthogonal to the magnetic-field direction, due to the torque $\boldsymbol{\tau} = \boldsymbol{\mu} \times \boldsymbol{B}$ exerted on the magnetic dipole moment $\boldsymbol{\mu}$ by the magnetic field $\boldsymbol{B}$. The transverse angle spanned in a time $\mathrm{d}t$ is given by $\mathrm{d}\varphi = \Omega_{\mathrm{L}}\mathrm{d}t$, where Ω_{L} is the Larmor frequency.

This section aims to describe the two main fundamental noise sources, and how they can be minimised through optimisation of experimental parameters.

6.4.1 Standard Quantum Limit

When measuring magnetic fields using atoms, quantum mechanics sets the ultimate limit on the sensitivity that can be achieved. One of the quantum sensitivity limits for an optical magnetometer is spin-projection noise[9], which arises from the fact that spin projections along different axes do not commute.

In order to derive a quantitative expression for the standard quantum limit due to spin-projection noise, consider an atomic ensemble which has been optically pumped into a fully spin-polarised state along the $\hat{x}$ direction (i.e. $\langle F_x \rangle = F$). If there is a magnetic field aligned along $\hat{y}$, i.e. $\boldsymbol{B} = B_y\hat{y}$, then the magnetic field will exert a torque on the magnetic moment of the atom, which subsequently causes the spin to precess about the field in the xz-plane[10]

---

[9]Spin-projection noise is also commonly referred to as atomic shot noise, due to its analogue with classical photon or electron shot noise

[10]The torque exerted on the magnetic moment by a magnetic field is given by $\boldsymbol{\tau} = \boldsymbol{\mu} \times \boldsymbol{B}$. Given the non-zero total torque applied to the system, the system's total angular momentum $\boldsymbol{F}$ must change in accordance with $\mathrm{d}\boldsymbol{F}/\mathrm{d}t = \boldsymbol{\tau}$. Since the magnetic moment is related to the total angular momentum by $\boldsymbol{\mu} = \gamma_{\mathrm{g}}\boldsymbol{F}$, the equation of

(cf. Fig. 6.8). In a scalar atomic magnetometer, e.g. an NMOR magnetometer, the Larmor frequency is best measured by finding the frequency of an oscillating $\langle F_z \rangle$ component [150]. One therefore measures the component of the spin perpendicular to the direction of the magnetic field.

Given a fully polarised spin state with $\langle F_x \rangle = F$, the spin wavefunction $|\psi\rangle$ with quantisation along $\hat{z}$ can be found by rotating the spin state by $\pi/2$ around the y-axis using Wigner's (small) d-function (cf. App. M.2) [150]:

$$|\psi\rangle = \sum_{m_F=-F}^{F} d^{(F)}_{m_F,F}\left(\frac{\pi}{2}\right) |F,F\rangle . \tag{6.11}$$

Using Wigner's formula given by Eq. (M.4), this simplifies to [150]

$$\begin{aligned}
|\psi\rangle &= \sum_{m_F=-F}^{F} \sum_{n=F-m_F}^{F-m_F} \left(\frac{\sqrt{(2F)!\,(F-m_F)!\,(F+m_F)!}}{(2F-n)!\,(F-m_F-n)!\,n!\,(m_F+n-F)!} \right) \frac{(-1)^{3n+m_F-F}}{2^F} |F,m_F\rangle \\
&= \sum_{m_F=-F}^{F} \left(\sqrt{\frac{(2F)!}{(F+m_F)!\,(F-m_F)!}} \right) \frac{(-1)^{F-m_F}}{2^F} |F,m_F\rangle .
\end{aligned} \tag{6.12}$$

In a measurement of F_z, the probability of obtaining each value of m_F is proportional to $\left[d^{(F)}_{m_F,F}(\pi/2) \right]^2$ [150], such that the standard deviation of many measurements on different atoms will then be given by [150]

$$\delta F_z = \sqrt{\text{var}(F_z)} = \sqrt{\sum_{m_F=-F}^{F} m_F^2 \left[d^{(F)}_{m_F,F}\left(\frac{\pi}{2}\right) \right]^2} = \sqrt{\frac{F}{2}}. \tag{6.13}$$

The uncertainty in the transverse spin projection, σ_{F_z}, gives rise to an uncertainty in both the tranverse polarisation, $\sigma_{P_z} = \sigma_{F_z}/F$, and the azimuthal phase $\sigma_\varphi = \sigma_{P_z}$. If N uncorrelated measurements of the phase φ are performed, the uncertainty σ_φ is decreased by a factor of $\sqrt{N}$ (cf. Eq. H.28), yielding the following:

$$\sigma_\varphi = \frac{1}{\sqrt{2FN}}. \tag{6.14}$$

The azimuthal phase is the 'Larmor phase', as it is the precession angle of the spin during coherent evolution time T:

$$\varphi = \int_0^T \Omega_L(t)\,dt = \gamma_g \int_0^T B(t)\,dt, \tag{6.15}$$

where γ_g is the gyromagnetic ratio. Assuming that the field is static (i.e. $d\Omega_L/dt = dB/dt = 0$), then Eq. (6.15) simplifies to $\varphi = \Omega_L T = \gamma_g B T$. Using this relation, the magnetic-field uncertainty σ_B for a single coherent evolution during time T is given by

---

motion becomes $d\boldsymbol{\mu}/dt = \gamma_g \boldsymbol{\mu} \times \boldsymbol{B}$ [322]. The change in magnetic moment is therefore *orthogonal* to the magnetic field, and hence Larmor precession occurs in the plane perpendicular to the applied field as shown in Fig. 6.8.

Figure 6.9: Illustration of performing M repeated uncorrelated measurements of length T of the coherent evolution (Larmor precession), over an integration time of τ with a duty cycle D.

$$\sigma_B = \frac{\sigma_\varphi}{\gamma_g T} = \frac{1}{\gamma_g T \sqrt{2FN}} . \tag{6.16}$$

Equation (6.16) describes the magnetic-field uncertainty in a single-shot measurement using N atoms. However, for M repeated *uncorrelated* measurements of the coherent evolution, the standard deviation (i.e. the uncertainty) decreases by a factor of $\sqrt{M}$ (cf. App. H.2.2, specifically Eq. H.28). For an integration time τ, and a single-shot measurement time of T_{shot}, the number of possible repeated measurements is $M = \tau/T_{\text{shot}}$ (cf. Fig. 6.9). The single-shot measurement time depends upon the measurement duty cycle, D, via $T_{\text{shot}} = T/D$. Substituting these expressions into Eq. (6.16) yields

$$\sigma_B = \frac{1}{\gamma_g \sqrt{2FND\tau T}} . \tag{6.17}$$

Now, assuming that the measurement has zero dead time (i.e. $D = T/T_{\text{shot}} = 1$), and after dividing by the square root of the measurement bandwidth $\sqrt{\nu_{\text{BW}}} = (2\tau)^{-1/2}$, we obtain the expression for the minimum detectable magnetic field per unit bandwidth due to spin-projection noise:

$$\delta B_{\text{SQL}} = \frac{1}{\gamma_g \sqrt{FNT_2}} , \tag{6.18}$$

where we have made the substitution $T \to T_2$, where T_2 is the transverse spin-relaxation time. Equation (6.18) is known as the 'standard quantum limit' due to spin-projection noise.

For a ^{87}Rb vapour at 296 K (23°C), the expected number density using Eq. (3.5) is $n_v = 1.04 \times 10^{16}$ atoms/m^3. For a probe volume of 70 mm^3, this corresponds to $N = 7.38 \times 10^8$ atoms. For the $F = 2$ ground state, the gyromagnetic ratio is $\gamma_g/2\pi = 6.9958\,\text{GHz/T}$ [229].

Combining all of this with a transverse spin-relaxation time of $T_2 = 36.5\,\text{ms}$, yields a spin-projection noise of $\delta B_{\text{SQL}} = 3.1\,\text{fT}_{\text{rms}}/\sqrt{\text{Hz}}$, which is less than a factor of 5 below the demonstrated photon shot-noise limit of $15\,\text{fT}_{\text{rms}}/\sqrt{\text{Hz}}$ [206].

The spin-projection noise derived above is referred to as the *standard* quantum limit, as it defines the maximum absolute sensitivity achievable under 'standard' conditions. However, by using some quantum trickery which exploits the Heisenberg uncertainty principle, it *is* possible to surpass the standard quantum limit. This is achieved by using squeezed light, in which case $\delta F_z \leq \sqrt{F/2N}$ and hence $\delta B \leq \delta B_{\text{SQL}}$. In the literature, magnetometers based on NMOR have demonstrated sensitivity enhancements on the order of 2–5 dB using polarisation-squeezed light [323–325].

6.4.2 Classical Limit

Although spin-projection noise sets the ultimate *quantum* limit of the sensitivity of an atomic magnetometer, the demonstrated sensitivity of most atomic magnetometers is typically limited by *classical* noise. For *optical* atomic magnetometers, this classical limit is due to photon shot noise.

Generally speaking, the sensitivity of magnetic field measurements, δB, using an optical magnetometer, is given by [150]

$$\delta B = \left(\frac{\text{d}\phi}{\text{d}B}\right)^{-1} \delta\phi, \tag{6.19}$$

where ϕ is the optical-rotation angle, and $\delta\phi$ is the sensitivity of the rotation of the light polarisation. When performing magnetic field measurements using NMOR signals observed in the modulation-frequency domain (i.e. phase-sensitive detection via a lock-in amplifier), Eq. (6.19) becomes

$$\delta B = \left(\frac{\text{d}\phi}{\text{d}\Omega_{\text{m}}}\frac{\text{d}\Omega_{\text{m}}}{\text{d}B}\right)^{-1} \delta\phi, \tag{6.20}$$

where Ω_{m} is the modulation frequency. When using linearly polarised light to generate ground-state coherence between Zeeman sublevels with $|\Delta m_F| = 2$, the resonance condition is satisfied when $\Omega_{\text{m}} = 2\Omega_{\text{L}}$, and hence

$$\frac{\text{d}\Omega_{\text{m}}}{\text{d}B} = 2\frac{\text{d}\Omega_{\text{L}}}{\text{d}B} = \frac{2g_F\mu_{\text{B}}}{\hbar}, \tag{6.21}$$

where g_F is the Landé g-factor, μ_{B} is the Bohr magneton, and $\hbar$ is the reduced Planck constant. The sensitivity of the optical rotation measurements in the modulation-frequency domain is given by the slope, $\mathcal{S}$, of the quadrature component of the resonance near the zero-crossing:

$$\mathcal{S} = \lim_{\Omega_{\text{m}}\to 2\Omega_{\text{L}}} \frac{\text{d}\phi}{\text{d}\Omega_{\text{m}}} = \frac{2\mathcal{A}}{\Gamma}, \tag{6.22}$$

where $\mathcal{A}$ is the amplitude of the resonance, and Γ is the resonance width (FWHM). Substi-

tuting Eqs. (6.21) and (6.22) into Eq. (6.20) yields the following formula for the sensitivity of the magnetic field measurements using NMOR:

$$\delta B = \frac{\Gamma \hbar}{4 \mathcal{A} g_F \mu_B} \delta \phi \,. \tag{6.23}$$

Further simplification can be made by noting that Eq. (6.23) contains terms which constitute the gyromagnetic ratio, $\gamma_g = g_F \mu_B / \hbar$. After making this substitution, Eq. (6.23) simplifies to [206]

$$\delta B = \frac{\delta \phi}{2 \mathcal{S} \gamma_g} \,. \tag{6.24}$$

When using an optical magnetometer, the fundamental (classical) limit of the optical-rotation noise, $\delta \phi$, is set by photon shot noise. For a beam of photons with time-averaged power $\overline{P}$, the power spectral density as measured at the photodetector is given by $S_{\mathrm{PSD}}(\nu) = 2h\nu \overline{P} \eta(\nu)$, where $h\nu$ is the energy per photon, and $\eta(\nu)$ is the quantum efficiency of the photodiode. Given equal detected powers, $\overline{P}$, on each photodetector in the balanced polarimeter, the optical-rotation noise per unit bandwidth is given by [206]

$$\delta \phi \approx \frac{1}{2} \sqrt{\frac{h\nu}{\eta(\nu) \overline{P}}} \,. \tag{6.25}$$

For the work detailed in this thesis, the optical-rotation noise per unit bandwidth, due to photon shot noise, was $\delta \phi = 0.24 \, \mu\mathrm{rad_{rms}}/\sqrt{\mathrm{Hz}}$. Combined with a resonance slope of $2\pi \mathcal{S} = 1.13 \, \mathrm{mrad/Hz}$ — which in magnetic units is equivalent to $2\mathcal{S}\gamma_g = 15.78 \, \mathrm{mrad/nT}$ — this optical-rotation noise corresponds to an effective magnetic noise per unit bandwidth of $15 \, \mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ [206].

6.4.3 Parameter Optimisation

Optimisation is one of the most important aspects of building a highly sensitive optical magnetometer as there are a multitude of experimental parameters and, if they are not all set to their optimal value, the performance of the magnetometer can be relatively poor.

By inspection of Eq. (6.24), it is clear that the magnetic noise is inversely proportional to the slope $\mathcal{S}$ of the quadrature component of the resonance, and directly proportional to the optical-rotation noise $\delta \phi$. It is therefore apparent that, in order to achieve high sensitivity, one must maximise the resonance slope whilst *simultaneously* minimising the optical-rotation noise[11]. Generally speaking, the resonance slope has a complicated dependence on a myriad of experimental parameters. Some of these parameters are either difficult or infeasible to change and, for all intents and purposes, must be considered fixed. However, some parameters are readily changed in real time, and can be optimised with relative ease.

In order to optimise any experimental parameter $\mathcal{P}$, one will measure a resonance profile

---

[11]It should be noted that some of the experimental parameters which affect $\mathcal{S}$ also affect $\delta \phi$. Therefore, there is quite often a trade-off between multiple experimental parameters, which results in an optimal configuration in which the sensitivity is maximised.

at each value of the parameter in question, calculate the resonance slope $\mathcal{S}(\mathcal{P})$, measure the optical-rotation noise $\delta\phi(\mathcal{P})$, and subsequently determine the magnetic noise $\delta B(\mathcal{P})$. The optimal configuration for that particular parameter can then be located, which corresponds to the maximum in sensitivity — i.e. the optimal value(s) of $\mathcal{P}$ is $\arg\min_{\mathcal{P}} \delta B(\mathcal{P})$. Note however that many of the experimental parameters may not be entirely orthogonal, and hence a multi-dimensional parameter space must be mapped out in order to find the *global* maximum in sensitivity.

In this work, two of the major parameters to explore in terms of importance were the optical powers of the pump and probe beams incident on the vapour cell. As the name suggests, nonlinear magneto-optical rotation is highly nonlinear in the optical power used to both pump and probe the atoms. In order to optimise the optical powers, one will typically measure a resonance signal in the modulation-frequency domain (i.e. the in-phase and quadrature components of a lock-in amplifier) as a function of both pump and probe powers. By using least-squares fitting with complex-Lorentzian profiles, it is possible to extract the amplitude, width and slope of the resonances. The photon shot noise $\delta\phi$ is then measured for each probe power, from which a shot-noise-limited sensitivity can be calculated using Eq. (6.24) and the resonance slope at each point.

By mapping out the resonance profile as a function of optical powers four contour plots can be created, one of which will reveal the optimal optical power configuration that produces the highest shot-noise-limited performance. An example of optical-power optimisation can be seen in Fig. 6.10, and is described in further detail in both Ch. 10 and Ref. [206]. Note that the optical-power optimisation will need to be performed whenever other important parameters are changed significantly (e.g. magnetic-field strength or atomic number density).

6.5 Technical Noise

In addition to the fundamental noise sources discussed in Sec. 6.4 — photon shot noise and spin-projection noise — there are also a myriad of technical noise sources that can limit the performance of an optical magnetometer. In fact, *all* magnetometers will be limited by technical noise over *some* range of Fourier frequencies. Due to the typical $\nu^{-\alpha}$ (for $\alpha > 0$) power-law frequency-dependence of *most* technical noise sources, this will generally occur at low Fourier frequencies.

Many technical noise sources do not produce *real* magnetic noise — i.e. fluctuations in the local magnetic-field strength — but rather, they produce *fictitious* magnetic noise. That is, fluctuations in the magnetometer output which appear *indistinguishable* from real fluctuations in the measured magnetic-field strength, but do *not necessarily* correlate with any variation of the field. Given the fact that fictitious magnetic noise and real magnetic noise are *indistinguishable* from each other, they are *both* equally as important when attempting to improve the performance of the magnetometer.
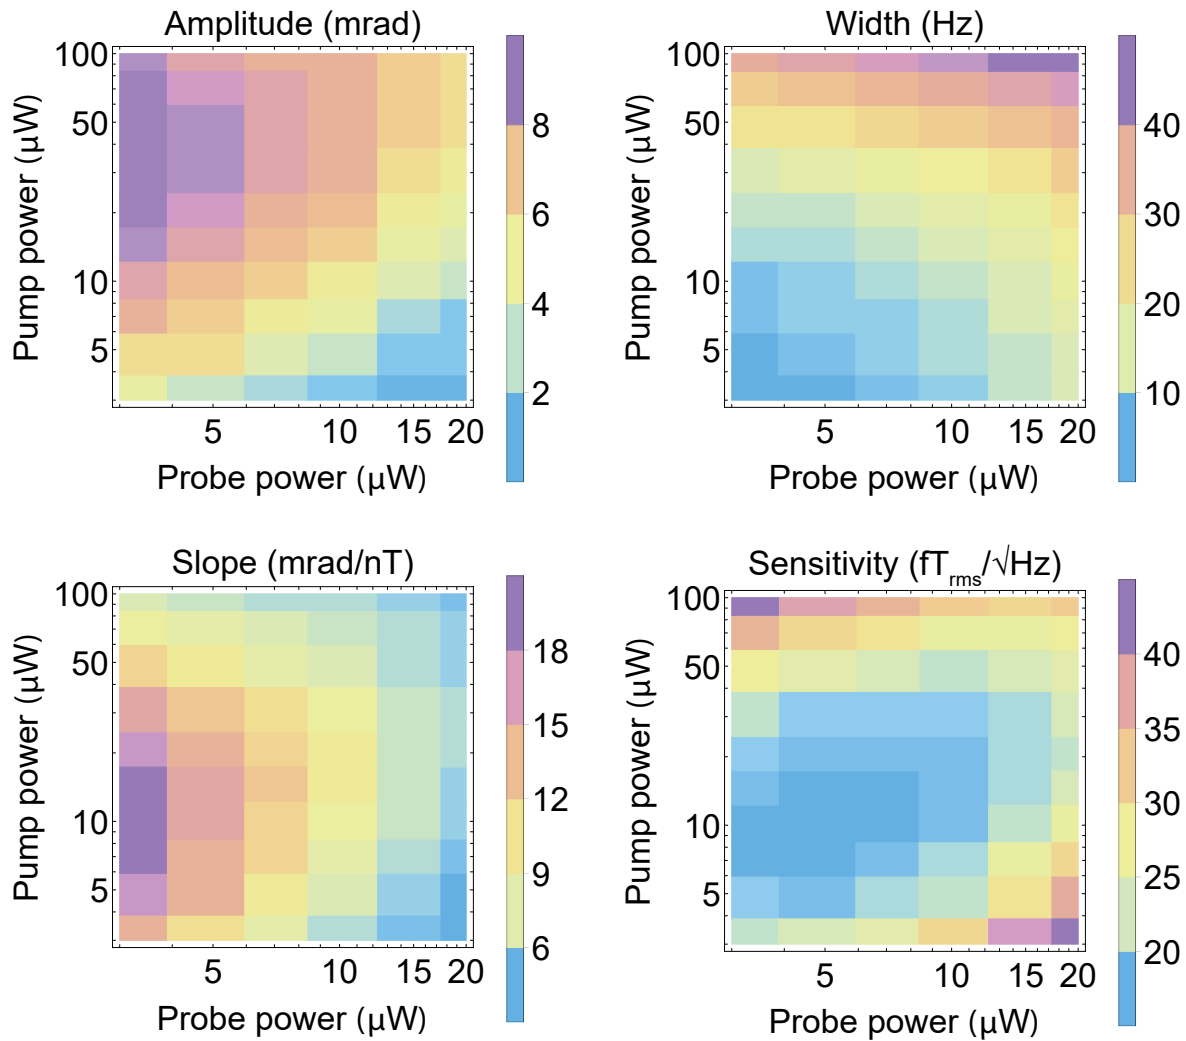
Figure 6.10: Optimisation contour plots of the resonance parameters and subsequent sensitivity of the magnetometer as a function of pump and probe optical powers.

Characterisation of technical noise sources can quite often be an arduous task, but generally involves the same step-by-step process. The initial step typically involves measuring a calibration slope. This is achieved by modulating an experimental parameter, and measuring the response of the magnetometer to this modulation. In doing so, a calibration slope is obtained, which enables fluctuations in the parameter to be mapped to an 'effective magnetic noise'. Once the calibration slope has been obtained, the stability of the experimental parameter in question is measured directly, and characterised using either an amplitude/power spectral density or Allan deviation/variance (cf. Apps. H.1 and H.2, respectively). The measured fluctuations of the experimental parameter are then scaled by the calibration slope, revealing its equivalent magnetic noise. If the equivalent magnetic noise for the parameter matches the magnetic noise floor measured by the magnetometer, and it is above *both* of the fundamental limits discussed in Sec. 6.4, then the instability of this parameter is a limitation of the magnetometer and its influence must somehow be reduced in order to improve performance. However, if the equivalent magnetic noise is *below* the observed detector noise floor, then it is not a limitation at present.

Although there are only two fundamental noise sources for an NMOR magnetometer, there are a *vast* number of technical noise sources. Many of these noise sources have been measured and quantified for the magnetometer described within this thesis. A detailed discussion of technical noise sources, and their impact on the performance of the magnetometer, is presented in Ch. 8.

6.6 Bandwidth

This section discusses the bandwidth limitations of the magnetometer, in terms of both amplitude and frequency response, when performing conventional phase-sensitive detection of a synchronously optically pumped resonance in the modulation-frequency domain. A resonance-locking technique, which overcomes these two limitations, is also described. Another techniques to overcome these limitations relies on performing measurements in the temporal domain, and is discussed in further detail in Ch. 7.

6.6.1 Amplitude Response

In a conventional phase-sensitive detection measurement, a lock-in amplifier is used to demodulate the optical-rotation signal at the pump-modulation frequency (cf. Sec. 4.11, Sec. 6.1.2 and App. K). When synchronously optically pumping, the system behaves as a damped driven harmonic oscillator, and therefore exhibits a resonant response which has a finite width (FWHM) Γ that is inversely proportional to the tranverse spin-relaxation time T_2. This width can be converted to magnetic units, Γ_B, via $\Gamma_B = \Gamma/2\gamma_g$, where γ_g is the ground-state gyromagnetic ratio. If the magnetic-field is changed by more than $\pm\Gamma_B/2$ — whether due to long-term drift or rapid sinusoidal (or otherwise) modulation — it is no longer possible to measure the field change[12]. This restricts the amplitude response to $|\Delta B| < \Gamma_B/2 = \Gamma/4\gamma_g \approx 0.5\,\text{nT}$ for the magnetometer described in this thesis under conditions optimised for high sensitivity (cf. Ch. 10). When measuring in this fashion, the longer the transverse spin-relaxation time, the smaller the amplitude response.

6.6.2 Frequency Response

Given that Larmor precession has no inertia [150,238], one might naively conclude that, provided the amplitude of the modulation is less than half the resonance width, the frequency-response bandwidth should be limited *only* by electronics. However, since the dc magnetic-field strength gives rise to a sinusoid at an angular frequency $2\Omega_L$, a sinusoidal modulation of the field gives rise to sinusoidal modulation of the carrier frequency. As discussed in further detail in Sec. 7, this sinusoidal modulation gives rise to sidebands around the central carrier frequency $\omega_c = 2\Omega_L$, appearing at $\omega = \omega_c + n\omega_m$ for $n \in \mathbb{Z} \setminus \{0\}$, where ω_m is the angular modulation frequency. The first-order sidebands therefore appear at $\omega = \omega_c \pm \omega_m$, and hence the signal of interest (i.e. the modulation) appears offset from the resonance peak.

---

[12]In fact, the FWHM Γ_B is *not* linear over the entire range, and hence the *quantitatively* measurable magnetic-field deviation is actually much less than this.

Figure 6.11: Left: Illustration depicting the first-order sideband positions on the complex-Lorentzian resonance as the modulation frequency is increased. Right: Transfer function of the sideband amplitude calculated using Eq. (6.26). The dashed grey lines indicate the $-3\,\mathrm{dB}$ bandwidth, which occurs at a modulation frequency of $\omega_\mathrm{m} = \Gamma/2$, where Γ is the resonance FWHM.

The lineshape of the resonance is well-approximated by a complex-Lorentzian profile, and is given by Eq. (6.2). The magnitude of the resonance is given by the quadrature sum of its real and imaginary components — the in-phase and quadrature components, respectively — which, after algebraic manipulation, simplifies to

$$|\mathcal{L}_\mathrm{c}\left(\omega - \omega_0\right)| = \frac{\mathcal{A}}{\sqrt{1 + 4\left(\frac{\omega-\omega_0}{\Gamma}\right)^2}}. \tag{6.26}$$

Equation (6.26) has the form of a first-order Butterworth filter, with a bandwidth given by $\Gamma/2$, the resonance HWHM. This means that, regardless of the modulation amplitude, the sideband maps out the complex-Lorentzian lineshape as the modulation frequency is increased (cf. Fig. 6.11). The frequency response is therefore limited to a bandwidth given by the resonance HWHM. Since the width of the resonance is inversely proportional to the transverse spin-relaxation time, there is a trade-off between sensitivity and bandwidth — at least when performing conventional phase-sensitive detection of a modulated resonance.

If additional bandwidth is required, and sensitivity is not of immediate concern, this can be achieved through the use of spin-decoherence mechanisms. By intentionally reducing the transverse spin-relaxation time, T_2, the resonance width is broadened and the bandwidth is subsequently increased. Examples of spin-decoherence mechanisms include gradient broadening [234] (cf. Sec. 5.4), power broadening [150,199,201,238] (cf. Sec. 3.2.2) and spin-exchange broadening [240] (cf. Sec. 5.2.2.1). In fact, by heating the atomic vapour to

Figure 6.12: Diagram of the three-stage process in the NMOR resonance lock. The pump-modulation frequency is initially tuned and locked to the resonance condition $\Omega_m = 2\Omega_L$. Any subsequent change in magnetic-field strength shifts the resonant frequency, causing a change in the measured optical-rotation amplitude. An error signal is created from the change in optical-rotation amplitude, which is used to adjust the modulation-frequency back to the resonance condition.

induce spin-exchange broadening, the magnetometer bandwidth can be increased without significant reduction in sensitivity due to the corresponding increase in optical depth [240] — though there will be a plateau (and possible turning point) at which further increases in optical depth do not yield improvements to (or actually *decrease*) the sensitivity.

6.6.3 Resonance Locking

The resonance-locking technique described in Sec. 6.1.3 can overcome both of the afore-mentioned limitations — namely the amplitude and frequency response — by using active feedback to ensure that the pump-modulation frequency (and therefore demodulation frequency) is *always* centred on resonance [326]. If the magnetic-field strength deviates from its initial value, a voltage offset is generated in the quadrature output of the lock-in amplifier. This voltage offset results in a non-zero error signal, from which the locking servo generates a control signal that is fed back to the signal generator which controls the pump-modulation frequency. The signal generator subsequently adjusts the modulation frequency in order to return the quadrature output back to zero voltage (i.e. the resonance condition $\Omega_m = 2\Omega_L$). The result is that, rather than the pump-modulation frequency sitting passively in one position and measuring reduced optical-rotation amplitude, the modulation frequency tracks the magnetic-field resonance. This is shown diagrammatically in Fig. 6.12. When using the resonance-locking technique, rather than measuring voltage fluctuations, the modulation frequency is measured instead. This has the added benefit of being a calibration-free technique, as the magnetometer output no longer requires scaling by the resonance slope.

The amplitude and frequency response when using the resonance-locking technique are, in principle, limited only by the locking servomechanism; with previous results demonstrating a frequency-response bandwidth of $\mathcal{O}\,(100\,\text{kHz})$ [326]. However, it should be noted that, due to the reasons described in Sec. 6.6.2, the SNR of this technique decreases linearly with increasing frequency. Therefore, the magnetic sensitivity scales as ω^{-1}.

Another technique that can extend the frequency response of atomic magnetometers is so-called 'self-oscillation' [327–330]. This technique is essentially the same as resonance locking — insomuch as they both lock the pump-modulation frequency to the resonance condition — however, self-oscillation involves *directly* feeding the output-signal frequency back to the input pump-modulation frequency [327–330]. In this way, there is no need for a lock-in amplifier or locking servomechanism, and the system will 'self-oscillate' [327–330]. Provided sufficient amplification, the oscillation builds up spontaneously upon being seeded by noise [330]. Previous results using the self-oscillation technique have demonstrated frequency-response bandwidths of $\mathcal{O}\,(1\,\text{kHz})$ [328–330].

6.6.4 Instantaneous-Phase Retrieval

An alternative technique for measuring high-frequency ac magnetic fields, which relies on observing the instantaneous phase of a freely precessing atomic polarisation, is known as 'instantaneous-phase retrieval' [326,331,332]. This technique has been shown to be able to measure ac magnetic fields with amplitudes and frequencies more than two and three orders of magnitude greater than is permissible using conventional phase-sensitive detection techniques [326,331,332], respectively. This novel technique is discussed in detail in Ch. 7, with the results of the technique demonstrated in Ch. 11 and Ref. [331].

High-Bandwidth Magnetometry via Instantaneous-Phase Retrieval

The amplitude and frequency-response limitations described in Sec. 6.6 arise due to synchronous optical pumping at the resonance condition $\Omega_m = 2\Omega_L$. Under these conditions the system behaves as a damped, driven harmonic oscillator, with a narrow response bandwidth centred around the resonance condition. Typical exquisitely sensitive magnetometers exhibit *extremely* narrow resonances [236] (due to long transverse spin-relaxation times), and will therefore exhibit *severely* restricted amplitude and frequency response when operated in this manner. Improvements to the response bandwidth can be achieved by intentionally degrading the ground-state coherence via spin-decoherence mechanisms (e.g. power– [150,199,201,238], spin-exchange– [240], or gradient broadening [234]); though this comes at the expense of sensitivity.

One pathway to measuring high-frequency ac magnetic fields using optical magnetometers, is by tuning the dc magnetic-field strength such that the ground-state Zeeman splitting is equal to the frequency of the ac magnetic field of inerest (cf. Sec. 2.2.6). By doing so, it is possible to measure high-frequency magnetic fields with high sensitivity. These devices are known as radio-frequency optical atomic magnetometers (rf OAMs) [85,88–94,209–217]. Unfortunately however, these devices also suffer from a narrow response bandwidth, and are therefore only suitable when the signal has narrowband spectral content, and is at a frequency which is known *a priori*. Moreover, these devices require amplitude calibration through a convoluted process which involves artificial broadening of the magneto-optical resonance [212]. Again, improvements to the response bandwidth of these devices can be achieved through intentional spin-decoherence mechanisms at the cost of sensitivity.

However, by instead allowing the atomic ensemble to undergo free-induction decay and monitoring the instantaneous phase of the Larmor precession, it is possible to perform sensitive magnetometry with *exceedingly large* amplitude and frequency response, *without* intentionally degrading the ground-state coherence time. This chapter describes, in detail, the application of instantaneous-phase retrieval to sensitive magnetometry. Explicit topics discussed in this chapter include: the theory of frequency modulation, Bessel functions, the Hilbert transform, and the instantaneous phase and its numerical computation. Limitations of this technique are also discussed, with quantitative bounds derived for both the amplitude and frequency response.

7.1   Frequency Modulation and Bessel Functions

As has been well established throughout this thesis, an NMOR magnetometer works by measuring the sinusoidal modulation of the plane of polarisation of a linearly polarised probe beam, at a 'carrier' frequency ν_c equal to twice the Larmor frequency (i.e. $2\pi\nu_c = 2\Omega_L$). The measurement of this frequency is then used to infer the volume average of the magnetic-field strength, B_{dc}, within the atomic vapour. However, if the magnetic field has an ac component, i.e. $B(t) = B_{dc} + \Delta B(t)$, the optical-rotation signal takes the form of a frequency-modulated sinusoid. In general, the frequency-modulated optical-rotation signal is given by

$$
\begin{aligned}
\phi(t) &= a_c \cos\left(2\pi \int_0^t \nu_I(\tau)\,d\tau + \varphi_c\right) \\
&= a_c \cos\left(2\pi \int_0^t [\nu_c + \delta\nu x_m(\tau)]\,d\tau + \varphi_c\right) \\
&= a_c \cos\left(2\pi\nu_c t + 2\pi\delta\nu \int_0^t x_m(\tau)\,d\tau + \varphi_c\right),
\end{aligned}
\tag{7.1}
$$

where a_c is the carrier amplitude, $\nu_I(\tau)$ is the 'instantaneous frequency', ν_c is the carrier frequency, $\delta\nu$ is the frequency deviation, $x_m(\tau)$ is the modulation waveform, and φ_c is an arbitrary phase offset of the carrier. In the case of sinusoidal modulation $x_m(\tau) = \cos(2\pi\nu_m\tau + \varphi_m)$ for some modulation frequency ν_m and phase offset φ_m, and hence the optical-rotation signal becomes

$$
\phi(t) = a_c \cos(2\pi\nu_c t + \beta \sin[2\pi\nu_m t + \varphi_m] + \varphi_c),
\tag{7.2}
$$

where $\beta = \delta\nu/\nu_m$ is the modulation index. Using trigonometric and Bessel identities[1], it is possible to show that Eq. (7.2) can be written in the 'Jacobi-Anger expansion' [333]:

$$
a_c \cos(2\pi\nu_c t + \beta \sin[2\pi\nu_m t + \varphi_m] + \varphi_c) = a_c \sum_{n=-\infty}^{\infty} J_n(\beta) \cos(2\pi[\nu_c + n\nu_m]t + \varphi_c + n\varphi_m),
\tag{7.6}
$$

where $J_n(\beta)$ is the n^{th}-order Bessel function of the first kind, evaluated at the modulation index β. Equation (7.6) shows that sinusoidal frequency modulation of a sinusoidal carrier signal gives rise to an *infinite* series of sidebands in the frequency domain, each spaced from the carrier frequency by integer multiples of the modulation frequency, and with amplitudes $a_n = a_c J_n(\beta)$. The implication of this result is that, when measuring ac magnetic fields using

---

[1] The Bessel-function identities required to derive the Jacobi-Anger expansion from Eq. (7.2) are

$$
\cos(z \sin\theta) = J_0(z) + 2\sum_{k=1}^{\infty} J_{2k}(z) \cos(2k\theta),
\tag{7.3}
$$

$$
\sin(z \sin\theta) = 2\sum_{k=0}^{\infty} J_{2k+1}(z) \sin([2k+1]\theta),
\tag{7.4}
$$

$$
J_{-n}(z) = (-1)^n J_n(z).
\tag{7.5}
$$

an NMOR magnetometer in the temporal domain, the measured signal corresponding to the ac field will be spread among a large number of sidebands[2]. In order to calculate the true amplitude of the ac field in the Fourier-frequency domain, one must therefore compute the frequency deviation via the inverse Bessel function using numerical methods:

$$a_n = a_c J_n(\beta)$$
$$\Rightarrow \delta\nu = \nu_m J^{-1}\left(\frac{a_n}{a_c}\right).$$

(7.7)

However, if it is possible to measure the 'instantaneous phase' of the signal, $\varphi_I(t) = 2\pi\nu_c t + \beta\sin(2\pi\nu_m t + \varphi_m) + \varphi_c$, then the amplitude of the sinusoidal modulation can be measured directly without undoing any Bessel-function attenuation.

7.2   Instantaneous Phase and the Hilbert Transform

Extracting the instantaneous phase from a signal requires that its quadrature component — the version of the signal which is out-of-phase by $90°$ with respect to the original — is known. For a simple signal with a known closed form, the quadrature component is *generally* also known. An example of this is a simple trigonometric function, $\sin(x)$, the quadrature of which is given by $-\cos(x)$. However, for arbitrarily complicated signals with no closed form — e.g. experimentally measured signals — the quadrature component is *generally* non-trivial and a closed-form solution may not be permissible. In this case, the quadrature signal can be obtained numerically by shifting each Fourier component by $90°$; a task which is accomplished through the use of the Hilbert transform.

Given a real-valued function $x(t)$, the analytic representation of this function can be constructed from its quadrature component $y(t)$. That is, a function $x_a(t) = x(t) + iy(t)$ satisfying the Cauchy-Riemann equations can be found, where $y(t) = \mathcal{H}\{x(t)\}$ is the Hilbert transform of $x(t)$, and where the Hilbert transform is defined as the convolution of $(\pi t)^{-1}$ with the signal $x(t)$:

$$\mathcal{H}\{x(t)\} = \frac{1}{\pi t} * x(t)$$
$$= \frac{1}{\pi}\text{p.v.}\int_{-\infty}^{\infty}\frac{x(\tau)}{t - \tau}d\tau,$$

(7.8)

where p.v. denotes the Cauchy principal value[3]. Although the Hilbert transform may appear somewhat abstract by its definition, it has a simple effect; it imparts a $90°$ phase shift to every

---

[2]In principle, the signal will be spread amongst an *infinite* number of sidebands. However; due to the Bessel-function scaling of the sideband amplitudes, only a small number of sidebands will contain non-negligible signal power.

[3]The Cauchy principal value of the integral must be computed because the function $(\pi t)^{-1}$ is not integrable — due to the first-order pole at $t = 0$ — and hence the integral defining the Hilbert transform does *not necessarily* converge. Recall that for a singularity of the function $f(x)$ at the point $x = a$, the Cauchy principal value is defined by

$$\text{p.v.}\int_{-\infty}^{\infty} f(x)\,dx = \lim_{\epsilon\to 0^+}\left[\int_{-\infty}^{a-\epsilon} f(x)\,dx + \int_{a+\epsilon}^{\infty} f(x)\,dx\right].$$

(7.9)

Fourier component of the function $x(t)$. This can be seen explicitly through the use of the convolution theorem for Fourier transforms[4], which results in the spectrum of the Hilbert transform being given by

$$
\begin{aligned}
\mathscr{F}\left\{\mathcal{H}\left\{x(t)\right\}\right\} &= \mathscr{F}\left\{\frac{1}{\pi t}\right\}\mathscr{F}\left\{x(t)\right\} \\
&= -i\operatorname{sgn}(\omega)\,\mathscr{F}\left\{x(t)\right\} \\
&= \begin{cases} e^{\frac{i\pi}{2}}\mathscr{F}\left\{x(t)\right\} & \text{for } \omega < 0, \\ 0 & \text{for } \omega = 0, \\ e^{-\frac{i\pi}{2}}\mathscr{F}\left\{x(t)\right\} & \text{for } \omega > 0. \end{cases}
\end{aligned}
\tag{7.11}
$$

Therefore, negative frequency components of $x(t)$ are shifted by $+90°$, while positive frequency components are shifted by $-90°$. After constructing the analytic representation of the signal, $x_a(t) = x(t) + i\mathcal{H}\{x(t)\}$, it can then be written in polar form via Euler's formula[5]:

$$
x_a(t) = a(t)\,e^{i\varphi_I(t)},
\tag{7.13}
$$

where $a(t)$ is the time-dependent amplitude of the signal, and $\varphi_I(t)$ is the instantaneous phase, which are given explicitly by

$$
a(t) = \sqrt{[x(t)]^2 + [\mathcal{H}\{x(t)\}]^2},
\tag{7.14}
$$

$$
\varphi_I(t) = \arg\left[x(t) + i\mathcal{H}\{x(t)\}\right].
\tag{7.15}
$$

Equation (7.15) reveals a method by which to calculate the instantaneous phase $\varphi_I(t)$ of the signal $x(t)$ using the Hilbert transform. This can be demonstrated explicitly by considering the case of a simple sinusoid, i.e. $x(t) = \sin(2\pi\nu_c t + \varphi_c)$, in which case the instantaneous phase is given by

$$
\begin{aligned}
\varphi_I(t) &= \arctan\left(\frac{-\operatorname{sgn}(2\pi\nu_c)\cos(2\pi\nu_c t + \varphi_c)}{\sin(2\pi\nu_c t + \varphi_c)}\right) \\
&= -\arctan\left[\cot(2\pi\nu_c t + \varphi_c)\right].
\end{aligned}
\tag{7.16}
$$

Provided that the quadrant in which the point $(\cos(2\pi\nu_c t + \varphi_c), \sin(2\pi\nu_c t + \varphi_c))$ lies is taken into account when the arctangent is calculated, Eq. (7.16) simplifies to $\varphi_I(t) = 2\pi\nu_c t +$

---

[4]Given two functions $f(x)$ and $g(x)$, the Fourier transform of the convolution $f(x) * g(x)$ is given by the product of the Fourier transforms of $f(x)$ and $g(x)$. Mathematically, the convolution theorem is given by

$$
\mathscr{F}\left\{f(x) * g(x)\right\} = \mathscr{F}\left\{f(x)\right\}\mathscr{F}\left\{g(x)\right\}.
\tag{7.10}
$$

[5]Recall that Euler's formula is the fundamental relationship between trigonometric functions and the complex exponential function:

$$
e^{ix} = \cos(x) + i\sin(x).
\tag{7.12}
$$

$\varphi_c - \pi/2$, with the caveat that the phase is *wrapped* — i.e. that it is constrained to its principal value $\varphi_I(t) \in (-\pi, \pi]$. However, provided that the SNR is sufficiently high, it is possible to unwrap the phase (i.e. by removing the discontinuities) using simple post-processing techniques, such that the phase is a continuous function of time. By extension, for a sinusoidally modulated oscillatory signal at a frequency ν_m, the instantaneous phase becomes $\varphi_I(t) = 2\pi\nu_c t + \beta \sin(2\pi\nu_m t + \varphi_m) + \varphi_c - \pi/2$.

7.3 Instantaneous Frequency and Signal-to-Noise Ratio

Once the instantaneous phase has been obtained, it is possible to calculate the instantaneous frequency via

$$\nu_I = \frac{1}{2\pi} \frac{d\varphi_I(t)}{dt}, \tag{7.17}$$

which, for a sinusoidally modulated oscillatory signal with an instantaneous phase given by $\varphi_I(t) = 2\pi\nu_c t + \beta \sin(2\pi\nu_m t + \varphi_m) + \varphi_c - \pi/2$, gives rise to an instantaneous frequency of

$$\begin{aligned}
\nu_I &= \frac{2\pi\nu_c}{2\pi} + \left(\frac{2\pi\nu_m}{2\pi}\right) \beta \cos(2\pi\nu_m t + \varphi_m) \\
&= \nu_c + \nu_m \left(\frac{\delta\nu}{\nu_m}\right) \cos(2\pi\nu_m t + \varphi_m) \\
&= \nu_c + \delta\nu \cos(2\pi\nu_m t + \varphi_m).
\end{aligned} \tag{7.18}$$

From Eq. (7.18) it is apparent that, by calculating the instantaneous frequency from the instantaneous phase, it is possible to *directly* measure the frequency deviation, $\delta\nu$, with no unwanted scaling factors. However, it should be noted that amplitude of the oscillatory component of the instantaneous phase is given by β, which is proportional to ν_m^{-1}. Given that this is essentially the 'signal', the SNR therefore decreases in proportion to ν_m^{-1}. This is also apparent in the fact that, given white phase noise, the derivative of this noise results in violet/purple noise, which has a power spectral density that scales as ν^2 and hence an amplitude spectral density which scales as ν.

This definition of instantaneous frequency is suitable when dealing with generated data with an SNR that is limited only by machine precision. However, it is less suitable when using experimentally measured data, as noise is amplified via numerical differentiation[6]. Fortunately however, numerical differentiation can be avoided by considering each Fourier

---

[6]One way to reduce the noise amplification of the numerical differentiation process is to use a multi-point derivative. Through the use of a Taylor-series expansion, it is possible to show that the first derivative with a sixth-order accuracy is given by

$$\begin{aligned}
\frac{df(t_n)}{dt} \approx &-\frac{49}{20h} f(t_n) + 6f(t_{n+1}) - \frac{15}{2h} f(t_{n+2}) + \frac{20}{3h} f(t_{n+3}) \\
&- \frac{15}{4h} f(t_{n+4}) + \frac{6}{5h} f(t_{n+5}) - \frac{1}{6h} f(t_{n+6}) + \mathcal{O}(h^6).
\end{aligned} \tag{7.19}$$

Equation (7.19) enables the derivative to be approximated using multiple points. To reduce noise amplification, it is advisable to not use directly adjacent points, but rather, use points that are spread by some distance. This reduces the amount of noise amplification, but comes at the cost of decreased temporal resolution.

component individually. For the case of a simple single-frequency sine wave (i.e. a single Fourier-frequency component), Eq. (7.17) enables the frequency deviation to be calculated via the simple relation $\delta \nu = \beta \nu_m$. This can be extended to arbitrarily complicated signals that admit a Fourier series, by multiplying each Fourier-frequency component by its corresponding Fourier frequency, thereby avoiding the need for numerical differentiation.

7.4 Instantaneous Magnetic-Field Strength

As discussed previously in Sec. 3.3, the Larmor frequency is directly related to the magnetic-field strength via quantum mechanics — specifically, the Zeeman effect. Modulation of the Larmor frequency therefore corresponds to modulation of the magnetic-field strength. Once the instantaneous frequency has been obtained, this can then be used to calculate the instantaneous magnetic-field strength, $B(t)$, via

$$B(t) = \frac{\pi \nu_I}{\gamma_g},\tag{7.20}$$

where γ_g is the gyromagnetic ratio of the ground-state hyperfine level. Using the example of a single-frequency sine wave, conversion from frequency deviation to ac magnetic-field amplitude is possible via $\Delta B = \pi \delta \nu / \gamma_g = \pi \beta \nu_m / \gamma_g$.

7.5 Performance Comparison with Conventional Phase-Sensitive Detection Techniques

Perhaps one of the most important questions regarding the instantaneous-phase-retrieval technique is: how does it compare with the conventional phase-sensitive detection technique in terms of standard performance metrics? What is the sensitivity like in comparison to lock-in amplifier demodulation when synchronously optically pumping on the resonance condition $\Omega_m = 2\Omega_L$? There must be *some* trade-off for the substantial increase in amplitude and frequency response.

As discussed in Sec. 6.4.2, the photon shot-noise limit of the magnetometer when using the conventional phase-sensitive detection technique is $15\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$, with a frequency-response bandwidth of about $7\,\mathrm{Hz}$. Using the instantaneous-phase-retrieval technique, the photon shot-noise-limited noise floor has a magnitude of about $44\,\mathrm{\mu rad_{rms}}/\sqrt{\mathrm{Hz}}$. Using Eqs. (7.17) and (7.20), this phase noise corresponds to a violet magnetic noise of about $3.1\nu\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$, with a demonstrably flat frequency response that extends orders of magnitude beyond that of the conventional technique [331] (cf. Ch. 11). A graph showing the absolute sensitivity comparison between the two techniques is presented in Fig. 7.1.

The discrepancy in sensitivity ($\sim$40%) between the two techniques as shown in Fig. 7.1 is attributed to slight differences in experimental conditions between the two measurements, rather than any inherent limitation of the instantaneous-phase-retrieval technique. However, there is one notable drawback to the instantaneous-phase-retrieval technique. Because

Figure 7.1: Absolute sensitivity of the magnetometer using two different measurement techniques: the conventional phase-sensitive detection technique (blue), and the instantaneous-phase-retrieval technique (red). The vertical dashed line corresponds to the lowest Fourier frequency that the instantaneous-phase-retrieval technique can measure in the experiments detailed within this thesis — which is about 22 Hz. A visible sensitivity discrepancy of about ∼40% is attributed to slight differences in experimental conditions between the two measurements, rather than any inherent limitation of the instantaneous-phase-retrieval technique.

the signal is transient, with an exponential decay governed by the transverse spin-relaxation time T_2, there exists a lower bound on the modulation frequency which can be measured in a single shot. Explicitly, this lower bound is given by $\nu_m \gtrsim T_2^{-1}$ which, under the experimental conditions detailed within this thesis, corresponds to approximately 22 Hz. However, with an increase in transverse spin-relaxation time, the modulation-frequency lower bound can be extended. Alternatively, the measurement of a train of FID sequences would enable resolution of low-frequency magnetic-field modulations. Such a technique has previously been demonstrated in Ref. [144], for example, and is described in further detail in Sec. 2.2.1.

7.6 Validity of the Hilbert Transform for Quadrature Estimates

Although the analytic Hilbert transform *exactly* calculates the quadrature of the signal, there are limitations of the *numerical* Hilbert transform which arise due to the way in which it is computed. The origin of these limitations, as well as their impact on quadrature estimation, are discussed in this section.

7.6.1   Modulation Beyond the Carrier Frequency

To lowest order, most of the power of the modulated signal is contained within the two first-order sidebands appearing at $\nu = \nu_c \pm \nu_m$. When frequency-modulating faster than the carrier frequency (i.e. $\nu_m > \nu_c$), the low-frequency first-order sideband at $\nu = \nu_c - \nu_m$ appears at a *negative* frequency. This has ramifications for the numerical calculation of the instantaneous phase, which are discussed below.

7.6.1.1   Splitting of First-Order Sideband into Two Components

There are a number of different ways in which the Hilbert transform $\mathcal{H}\{x(t)\}$ of a signal $x(t)$ can be computed numerically; some of which are performed in the time domain, while others are performed in the frequency domain. Frequency-domain techniques — which make use of discrete Fourier transforms — are most commonly used, and typically involve discarding the negative-frequency Fourier components of $x(t)$ (cf. App. I).

    When $\nu_m > \nu_c$, discarding the negative-frequency components results in half of the modulated signal power contained within the two first-order sidebands being discarded. Fortunately however, due to the Hermitian symmetry of the Fourier transform of a real signal, there is also a carrier frequency and corresponding sidebands in *negative* frequencies, i.e. at $\nu = -\nu_c \pm \nu_m$ (cf. Fig. 7.2). When $\nu_m > \nu_c$, the first-order *negative* sideband at $\nu = -\nu_c + \nu_m$ becomes *positive* and, after subsequently calculating the Hilbert transform, *still* remains within the frequency spectrum. It is therefore possible to retrieve the 'discarded' signal power; however, the signal power *will not* be found at a single frequency component in this case.

    In the case of weak modulation (i.e. $\beta \lesssim 1$, which is generally satisfied when $\nu_m > \nu_c$), the majority of the modulated power of the analytic signal is contained within the two positive frequencies $\nu = \nu_c + \nu_m$ and $\nu = -\nu_c + \nu_m$. The spacing of these two frequency components from the positive-frequency carrier is ν_m and $|2\nu_c - \nu_m|$, respectively. The asymmetry of these two dominant sidebands results in a modification of the instantaneous phase retrieved using Eq. (7.15), taking the form of double-sideband suppressed-carrier amplitude modulation:

$$
\begin{aligned}
\varphi_I(t) &= 2\pi\nu_c t + \frac{\beta}{2}\sin\left(2\pi\nu_m t + \varphi_m\right) + \frac{\beta}{2}\sin\left(2\pi\left(2\nu_c - \nu_m\right)t - \varphi_m\right) + \varphi_c \\
&= 2\pi\nu_c t + \beta\cos\left(2\pi(\nu_m - \nu_c)t + \varphi_m\right)\sin\left(2\pi\nu_c t\right) + \varphi_c.
\end{aligned}
\tag{7.21}
$$

    In accordance with Eq. (7.21), the resulting frequency spectrum of $\varphi_I(t)$ contains two tones: one at ν_m, and one at $|2\nu_c - \nu_m|$. However, in the small β approximation (valid for $\nu_m > \nu_c$ under *realistic* experimental conditions), each of these two frequency components has an amplitude of $\beta/2$, and hence the amplitude of ac magnetic fields can still be imputed unambiguously in the supra-carrier modulation regime. Furthermore, since linear regression to the instantaneous phase can still be used to determine ν_c in this regime, the unambiguous determination of the frequency of ac magnetic fields is also permissible.

Figure 7.2: Diagram of the truncated frequency spectrum of the Jacobi-Anger expansion, Eq. (7.6), when $\nu_\mathrm{m} > \nu_\mathrm{c}$. In this case, the low-frequency first-order sideband crosses from positive to negative frequencies; however, due to the Hermitian symmetry of the Fourier transform of a real signal, there is a corresponding negative-frequency mirror component which crosses over into positive frequencies.

7.6.1.2 Retrieving Amplitude, Phase and Frequency Information

Given that the upper first-order sideband at $\nu = \nu_\mathrm{c} + \nu_\mathrm{m}$ *always* remains positive, this sideband can *always* be used to retrieve the correct modulation frequency ν_m. However, this sideband alone is *not necessarily* sufficient to accurately retrieve the frequency deviation $\delta\nu$.

  In theory, it should be possible to retrieve $\delta\nu$ by simply multiplying the signal amplitude obtained from the $\nu = \nu_\mathrm{c} + \nu_\mathrm{m}$ sideband by a factor of 2 to account for the lost signal power; however, in practice there may be electronic (or otherwise) transfer functions present which yield erroneous results when taking into account only a single sideband. In practice, it is therefore advisable to take into account the signal amplitude in *both* the $\nu = \nu_\mathrm{c} + \nu_\mathrm{m}$ and $\nu = -\nu_\mathrm{c} + \nu_\mathrm{m}$ sidebands. The amplitudes of these components should be added *linearly*, from which the frequency deviation $\delta\nu$ can be retrieved with greater accuracy.

7.6.2 Maximum Modulation Index

Although to lowest order the modulation signal is contained within the first-order sidebands, there are, in principle, an *infinite* number of sidebands. If the modulation index is large enough, considerable signal power will be contained within higher-order ($|n| \geq 2$) sidebands, and hence they can no longer be safely ignored. Depending on the ratio of the modulation frequency to the carrier frequency, these higher-order sidebands *may* be discarded when computing the Hilbert transform numerically — even if the modulation frequency is *below* the carrier frequency. This sets an upper limit on the allowed modulation

index before significant errors are accrued when computing the Hilbert transform numerically.

In this section, quantitative bounds are derived for the maximum permissible modulation index, such that minimal signal power is discarded when computing the Hilbert transform via numerical methods.

7.6.2.1 Power in Negative-Frequency Sidebands

As discussed previously in Sec. 7.1, the sideband amplitudes scale as $J_n(\beta)$, and the power in these sidebands therefore scales as $J_n(\beta)^2$. In order to calculate the number of sidebands n over which the signal power extends appreciably, one can calculate the power-weighted rms harmonic:

$$n_{\text{rms}} = \sqrt{\langle n^2 \rangle} = \sqrt{\frac{\sum\limits_{n=-\infty}^{\infty} n^2 J_n(\beta)^2}{\sum\limits_{n=-\infty}^{\infty} J_n(\beta)^2}} = \sqrt{\frac{2\sum\limits_{n=1}^{\infty} n^2 J_n(\beta)^2}{J_0(\beta)^2 + 2\sum\limits_{n=1}^{\infty} J_n(\beta)^2}} = \frac{\beta}{\sqrt{2}}. \tag{7.22}$$

The result of Eq. (7.22) suggests that the power in the sidebands extends appreciably to $\nu_c \pm \beta \nu_m$. One method for estimating the upper limit of the modulation index, β_{max}, for a given modulation-frequency ratio, $\alpha = \nu_m/\nu_c$, is to set an acceptable power limit, P_{max}, which is permitted to reside within the discarded higher-order sidebands. This discarded signal power, if large enough, can introduce significant errors into the computation of the Hilbert transform, and the subsequent estimation of magnetic-field modulations.

For any ν_c and ν_m, the largest value of n such that $\nu_c - n\nu_m > 0$, is given by

$$n_{\text{max}} = \text{ceil}\left(\alpha^{-1}\right) - 1. \tag{7.23}$$

In order to calculate β_{max} for a given α, the total signal power in the discarded higher-order sidebands must be less than or equal to the acceptable power limit P_{max}. Mathematically, this is given by

$$\sum_{n=n_{\text{max}}+1}^{\infty} J_n(\beta_{\text{max}})^2 \lesssim P_{\text{max}}. \tag{7.24}$$

Note that Eq. (7.24) is an *infinite* sum with no closed-form simplification. This summation *does* converge quite rapidly and hence a good approximation to this sum can be acquired by evaluating only a small number of terms; however, it is *desirable* to obtain a closed-form solution which does not require an infinite summation (or an approximation thereof). Note that instead of summing over an infinite number of negative sidebands, one can instead sum over an infinite number of *positive* sidebands. In this case, Eq. (7.24) is equivalent to

$$\sum_{n=-\infty}^{n_{\text{max}}} J_n(\beta_{\text{max}})^2 \gtrsim 1 - P_{\text{max}}. \tag{7.25}$$

By exploiting the symmetry of the sidebands and using Bessel-function summation identi-

Figure 7.3: Diagram illustrating the derivation of Carson's bandwidth rule. To lowest order, the majority of the frequency modulation is contained within the first-order sidebands. A good approximation to the required bandwidth for a given frequency modulation is determined by the spacing of the two first-order sidebands, relative to the carrier frequency. The maximum possible spacing of these two sidebands corresponds to Carson's bandwidth rule.

ties[7], it is possible to write Eq. (7.25) in a form which does not require infinite summation:

$$\sum_{n=-n_{\text{max}}}^{n_{\text{max}}} J_n \left(\beta_{\text{max}}\right)^2 \gtrsim 1 - 2P_{\text{max}} . \tag{7.27}$$

A conservative upper limit to the lost power fraction in higher-order sidebands is $P_{\text{max}} \approx$ 0.01. The maximum modulation index, β_{max}, calculated using Eq. (7.27) for a fractional power of $P_{\text{max}} = 0.01$, as a function of α, is shown in Fig. 7.4.

7.6.2.2 Carson's Bandwidth Rule

Carson's bandwidth rule is an approximate bandwidth requirement for the frequency modulation of a carrier signal [334]. As discussed in Sec. 7.1, a frequency-modulated sinusoidal carrier signal has an infinite number of sidebands, each spaced by an integer multiple of the modulation frequency. To lowest order, the majority of the frequency modulation is contained within the two first-order sidebands. Carson's bandwidth rule derives the ap-

---

[7]The required Bessel-function identities and symmetry relations make use of Eq. (7.5), and are given explicitly by

$$\sum_{n=-\infty}^{0} J_n \left(\beta\right)^2 = \sum_{n=0}^{\infty} J_n \left(\beta\right)^2 = \frac{1}{2} + \frac{1}{2} J_0 \left(\beta\right)^2 . \tag{7.26}$$

proximate bandwidth required by considering *only* these two sidebands. As illustrated in Fig. 7.3, the maximum spacing of the two first-order sidebands from the carrier frequency, is given by

$$\Delta \nu_{\text{CBR}} = 2 \left(\nu_{\text{m}} + \delta \nu \right) , \tag{7.28}$$

where ν_{m} is the modulation frequency, and $\delta \nu$ is the frequency deviation. Equation (7.28) is known as Carson's bandwidth rule [334]. In the limit of small modulation index (i.e. $\beta \ll 1$, or $\delta \nu \ll \nu_{\text{m}}$), then the required bandwidth is $\Delta \nu_{\text{CBR}} \approx 2\nu_{\text{m}}$. This regime is called narrowband frequency modulation. On the other hand, if the modulation index is large (i.e. wideband frequency modulation with $\beta \gg 1$, or $\delta \nu \gg \nu_{\text{m}}$), then the required bandwidth is $\Delta \nu_{\text{CBR}} \approx 2\delta \nu$.

In the context of the Jacobi-Anger expansion detailed in Sec. 7.1, Carson's bandwidth rule is set by the requirement that the low-frequency first-order sideband *does not* appear at a negative frequency. Mathematically, this is written as

$$\nu_{\text{m}} + \delta \nu = \frac{\Delta \nu_{\text{CBR}}}{2} < \nu_{\text{c}} . \tag{7.29}$$

Re-writing Eq. (7.29) in terms of β using the relation $\delta \nu = \beta \nu_{\text{m}}$, gives a maximum modulation index of

$$\beta_{\text{max}} \approx \alpha^{-1} - 1 . \tag{7.30}$$

A comparison between the expected maximum modulation index using the aforementioned estimation methods, as a function of fractional modulation frequency α, is shown in Fig. 7.4.

7.6.2.3   Accounting for Signal Power Lost in Higher-Order Sidebands

The calculations of the maximum modulation index, β_{max}, performed in Secs. 7.6.2.1 and 7.6.2.2, assume that the signal power in higher-order sidebands (i.e. $n \leq -2$) is *irretrievably lost* in negative-frequency space. However, since the signal is real (i.e. experimentally measured data), the Fourier transform exhibits Hermitian symmetry. That is, the spectrum observed in positive frequencies is mirrored in negative frequencies. Therefore, when the higher-order *positive*-frequency sidebands become negative, the corresponding *negative*-frequency sidebands become positive. Because of this fact, it is actually possible to retrieve the signal power which is 'lost' due to discarding negative-frequency components when computing the Hilbert transform numerically. This would, at least in principle, enable the modulation index to be increased *beyond* the conservative limits calculated in Secs. 7.6.2.1 and 7.6.2.2 — *provided* that the SNR is sufficiently high to resolve the higher-order sidebands. However, this would come at the cost of increased signal-processing complexity.

Figure 7.4: Maximum modulation index permissible, $\beta_{\max}$, for given modulation-frequency ratios, $\alpha = \nu_{\mathrm{m}}/\nu_{\mathrm{c}}$, when numerically calculating the Hilbert transform. The blue trace was calculated using Eq. (7.27) for a fractional power loss of 1%, the red trace corresponds to Carson's bandwidth rule given by Eq. (7.30), the purple trace corresponds to $n_{\max}$ given by Eq. (7.23), and the grey trace corresponds to the condition that $\Delta B < B_{\mathrm{dc}}$ (cf. Sec. 7.7.1).

7.7   Additional Considerations

In addition to the sideband-power considerations discussed in Sec. 7.6.2, there are also a number of other potential issues which must be considered when using the instantaneous-phase-retrieval technique to measure rapidly varying magnetic fields. This section contains a detailed discussion of such issues, with quantitative bounds derived in order to characterise when they become problematic. Where possible, mitigation techniques which possess the potential to overcome these limitations are also discussed.

7.7.1   Changing Direction of the Field

Equation (7.27) was derived with the aim of determining an upper bound for β, which prevents significant errors from being introduced into the instantaneous-phase estimate when numerically calculating the Hilbert transform of the analytic signal. However, if one wishes to use $\beta_{\max}$ to estimate the maximum ac magnetic-field amplitude which can be measured, $\Delta B_{\max}$, there is an additional constraint. This constraint is that, for a given $\beta_{\max}$, the sign of $B(t)$ *cannot* change. This limits the maximum measurable ac magnetic-field amplitude to $\Delta B < B_{\mathrm{dc}}$, which is equivalent to $\delta\nu < \nu_{\mathrm{c}}$ and thus $\beta < \nu_{\mathrm{c}}/\nu_{\mathrm{m}} = \alpha^{-1}$. This constraint becomes stricter than the sideband-power considerations discussed in Sec. 7.6.2.1 for $\alpha \gtrsim (4P_{\max})^{-1/2}$.

7.7.2 Arbitrarily Oriented Magnetic Fields

The magnetometry technique described in this chapter is primarily applicable to longitudinal ac magnetic fields (i.e. oscillating fields along the axis of the dc magnetic field). For arbitrarily oriented ac fields there will be a dead band of $\delta v_\perp = \gamma_g \Delta B_\perp / 2\pi$ about the Larmor frequency, where a transverse component of the oscillating field with amplitude $\Delta B_\perp$ can drive Zeeman transitions[8]. Driving of these Zeeman transitions would result in amplitude modulation of the polarisation rotation [335], and render the instantaneous-phase-retrieval technique inapplicable. Fortunately, however, this dead band decreases with decreasing ac field strength.

7.7.3 Nonlinear Zeeman Effect

The instantaneous-phase-retrieval technique, as discussed in this chapter, is based on the presence of *one* carrier frequency v_c. At low magnetic-field strength this condition is satisfied, and the carrier frequency corresponds to the quantum-beat frequency of the $|\Delta m_F| = 2$ ground-state coherences. However, as the magnetic-field strength is increased, the frequencies of these coherences diverge quadratically in accordance with Eq. (3.51), giving rise to *three distinct* 'carrier frequencies'. This results in an optical-rotation signal that is no longer described by Eq. (7.6), subsequently rendering the instantaneous-phase-retrieval technique, as formulated here, inapplicable.

As it stands, this caveat restricts the application of the instantaneous-phase-retrieval technique to situations in which the ambient magnetic field is small. Unfortunately, many field-based applications require operation at relatively large ambient fields — such as that produced by Earth. However, there do exist techniques which can be employed in order to suppress the nonlinear Zeeman effect.

Two examples of cancellation techniques for the nonlinear Zeeman effect include light shifts [256,336] and spin-locking [255]. In the former case, as its name suggests, high-intensity light is used to create ac Stark shifts (cf. Sec. 3.2.2.2) which *counteract* the shifts produced by the nonlinear Zeeman effect [256,336]. In the latter case, a radio-frequency magnetic field — oriented transverse to the dc magnetic field — is used to prevent the atomic spins from undergoing alignment-to-orientation conversion [255]. In both of these cases the result is the same; only one oscillation frequency is present within the system — the Larmor frequency (or a harmonic thereof) corresponding to the dc field. The use of either of these two suppression methods — or indeed, any alternatives which yield the same outcome — would enable application of the instantaneous-phase-retrieval technique to situations in which the ambient dc magnetic-field strength is large.

---

[8]This dead band is analogous to that observed in conventional rf OAMs which sense *transverse* ac magnetic fields, in which case an oscillating longitudinal component near the Larmor frequency dithers the Zeeman resonance.

7.8 Arbitrarily Modulated Magnetic Fields

Thus far, only simple sinusoidal modulation of the magnetic field has been considered. However, the instantaneous-phase-retrieval technique can resolve arbitrarily complicated modulations that admit a Fourier series. Consider a general time-dependent magnetic field described by the following Fourier cosine series:

$$B\left(t\right) = B_{\mathrm{dc}} + \sum_{j=1}^{N} \Delta B_j\left(t\right) \cos\left(2\pi \nu_j t + \varphi_j\right), \tag{7.31}$$

where $\Delta B_j\left(t\right)$, ν_j, and φ_j are the time-dependent ac field amplitude, linear oscillation frequency, and arbitrary phase of the j^{th} Fourier component, respectively. The resulting optical-rotation signal takes the form

$$\phi\left(t\right) = a_{\mathrm{c}} \cos\left(2\pi \nu_{\mathrm{c}} t + \sum_{j=1}^{N} \beta_j\left(t\right) \sin\left[2\pi \nu_j t + \varphi_j\right] + \varphi_{\mathrm{c}}\right), \tag{7.32}$$

where $\beta_j\left(t\right) = \delta \nu_j\left(t\right) / \nu_j = \gamma_{\mathrm{g}} \Delta B_j\left(t\right) / \pi \nu_j$ is the time-dependent modulation index of the j^{th} Fourier component.

Comparison of Eq. (7.32) with Eq. (7.2) reveals essentially the *same* functional form of the optical-rotation signal, except that in the case of an arbitrarily modulated magnetic field, the phase now contains a *sum* of frequency components rather than just one. However, as has been explicitly demonstrated in Ch. 11 and Ref. [331], the instantaneous phase of atomic spins, subject to a magnetic field which takes the form of Eq. (7.31), can still be accurately retrieved using the technique derived in this chapter.

7.9 Demonstration

An experimental demonstration of the novel instantaneous-phase-retrieval technique, applied to the measurement of rapidly varying ac magnetic fields, is presented both in Ch. 11 and Ref. [331]. The manuscript presented in Ch. 11 explicitly demonstrates the enhanced amplitude and frequency response of the technique in comparison to conventional detection methods, as well as the ability to measure arbitrarily complex magnetic fields with rich and dynamical spectral content.

Sources of Technical Noise

Technical noise sources can be split into two distinct categories: real magnetic noise, and fictitious magnetic noise. Both of these categories of noise are *indistinguishable*, and hence they are equally as important when it comes to eliminating performance-limiting noise sources. This chapter provides an in-depth, quantitative analysis of a myriad of technical noise sources and, where possible, techniques to overcome these limitations are presented and discussed.

8.1 Real Magnetic Noise

Real magnetic noise is defined as unwanted fluctuations in the output of the magnetometer that arise from changes in the local magnetic-field strength within the vapour cell. This section discusses a few of the major sources of real magnetic noise, as well as ways in which their contribution may be reduced.

8.1.1 Power-Supply Noise

Perhaps the first thing that comes to mind when considering technical noise sources of an optical magnetometer, is the real magnetic noise generated by the current instability of the power supply which produces the bias field within the shields. Depending on the voltage/current supply being used to produce the bias field, the resulting noise generated from this source can be enormous and may completely dominate all other noise sources. This is especially true when using a typical 'laboratory dc power supply' or a signal generator.

If a modest current is all that is required (e.g. in the case of small bias fields), an efficient method of generating a bias field with minimal magnetic noise is through the use of an ultra-precision voltage reference such as the Analog Devices LTZ1000. Typical voltage references are able to supply $\mathcal{O}\left(10\,\mathrm{mA}\right)$ of current, with some being able to supply up to $50\,\mathrm{mA}$ or even $100\,\mathrm{mA}$ at the cost of slightly increased voltage noise. If more current is needed than this, transistors must be incorporated into the output stage.

Due to the exceptional fractional stability of the LTZ1000, it was impossible to measure its performance using a standard 6.5-digit multimeter. For this reason, the data presented in Fig. 8.1 was necessarily measured using an 8.5-digit multimeter (HP 3458a). A 16-hour time trace of the supply voltage produced by the LTZ1000, and its corresponding Allan deviation, is presented in Fig. 8.1. The Allan deviation has a demonstrated fractional stability of about 30 ppb at 1 s, which slowly degrades to 1 ppm for integration times of 10^4 s. Given a dc magnetic-field strength of about $2.5\,\mu\mathrm{T}$, this corresponds to an effective magnetic-field noise of $7.5\,\mathrm{fT}$ at 1 s, and $2.5\,\mathrm{pT}$ at 10^4 s.

Figure 8.1: Time trace (left) and fractional Allan deviation (right) of the LTZ1000 ultra-precision voltage reference which drives the solenoid in the experiment. The measured voltage noise occurs prior to any filtering, and exhibits a fractional Allan deviation of about 30 ppb at 1 s, slowly degrading to 1 ppm for integration times of 10^4 s.

Even when using the best commercially available voltage references, the fractional stability of these devices may still not be sufficient for many optical magnetometers — specifically those based on nonlinear magneto-optical rotation, as they exhibit exceptional dynamic range. For this reason, filtering the output voltage using a Butterworth filter with a sufficiently low cut-off frequency can be beneficial. The Butterworth filter employed in this work is comprised of an RC circuit — that is, a resistance and capacitance in series. Such a circuit gives rise to a first-order low-pass filter with a cut-off frequency of $\nu_{co} = (2\pi RC)^{-1}$, resulting in an amplitude transfer function that is described by

$$\mathcal{T}(\nu) = 10 \log_{10} \left(\frac{1}{\sqrt{1 + \left(\frac{\nu}{\nu_{co}}\right)^2}} \right). \tag{8.1}$$

The Butterworth filter has a resistance of 2.5 kΩ, and a capacitance of 200 mF (achieved through the use of two 100 mF aluminium electrolytic capacitors arranged in parallel). The resistor implemented here is an ultrahigh-precision foil resistor, with an ultralow temperature coefficient of 200 ppb/$^\circ$C in order to minimise thermal fluctuations in the supply current. The calculated cut-off frequency for this Butterworth filter is $\nu_{co} = 0.32$ mHz, which is in close agreement with the measured value of $\nu_{co} = 0.30$ mHz (cf. Fig. 8.2). There is also another 2.5 kΩ resistor which appears after the Butterworth filter — placed between the solenoid and ground — which gives a total series resistance of 5 kΩ. This yields a typical current of around 1.4 mA flowing through the solenoid.

Figure 8.2: Measured impulse response of the Butterworth filter to a 10 V input voltage (left), and amplitude transfer function (right) with a fit using Eq. (8.1) (red trace). The $-3\,\mathrm{dB}$ power bandwidth of the transfer function is measured to be $\nu_{\mathrm{co}} = 0.30\,\mathrm{mHz}$, which is in close agreement with the calculated value of $\nu_{\mathrm{co}} = 0.32\,\mathrm{mHz}$. A plateau in attenuation is seen above 10 Hz, and is associated with the equivalent series resistance of the electrolytic capacitors.

8.1.2   Johnson Noise

As discussed in Sec. 8.1.1, a low-noise magnetic field is generated within the shields using a solenoid that is driven by a custom-made, ultra-low-noise power supply. In most experiments, the resistance of the coil is very small — typically $\mathcal{O}(10\,\Omega)$. Rather than applying a *low* voltage to the coils in order to produce a *sensible* current, a high voltage is applied and attenuated appropriately using low-temperature-coefficient resistors[1]. Although this is fine in principle, Johnson-current noise is inadvertently added into the system due to the thermal motion of charge carriers within the low-resistance conductor. Given a conductor with a frequency-dependent impedance $Z(\nu)$, the Johnson-current noise is given by [337]

$$\delta I(\nu) = \sqrt{\frac{4k_{\mathrm{B}}T\,\mathrm{Re}\left\{Z(\nu)\right\}}{\left|Z(\nu)\right|^2}}\,,\tag{8.2}$$

where k_{B} is the Boltzmann constant and T is the absolute temperature of the resistor. For a solenoid with in-series resistors, the Johnson-current noise scales as $\mathrm{Re}\{Z(\nu)\}^{-1/2}$ at low frequencies[2], and hence the contribution from the two $2.5\,\mathrm{k\Omega}$ resistors is negligible. Therefore, only the solenoid will be considered from hereafter unless explicitly stated otherwise.

---

[1]It is generally recommended to drive a load at a relatively high voltage and using passive components to set the desired current, rather than driving with a low voltage. This is due to most electronic devices exhibiting higher fractional stability at the upper end of their range.

[2]This is due to the impedance of a solenoid being given by $Z(\nu) = R + i2\pi\nu L$, where R is the resistance and L is the inductance. At low frequencies $\mathrm{Im}\{Z\} \to 0$, resulting in $|Z| \to \mathrm{Re}\{Z\}$ and hence $\delta I \propto \mathrm{Re}\{Z\}^{-1/2}$.

Figure 8.3: Effective magnetic noise due to Johnson current arising from the thermal motion of charge carriers in the solenoid. The blue trace was calculated using Eq. (8.2) and the finite impedance of the solenoid, while the red trace is the Johnson noise predicted from a SPICE model which incorporates the entire circuit.

In the measurements detailed within this thesis, a voltage of 7.1513 V is applied to the coils which are in series with a total resistance of 5 kΩ. The solenoid has a complex impedance of $Z(\nu) = R + i2\pi\nu L$, where $R = 6\,\Omega$ is the measured resistance and $L = 5.5\,\text{mH}$ is the calculated inductance. The corresponding Johnson-current noise arising from the solenoid, using Eq. (8.2), is $50\,\text{pA}_{\text{rms}}/\sqrt{\text{Hz}}$ for frequencies below the self-imposed filter frequency. Using the solenoid calibration factor of around $1.75\,\mu\text{T/mA}$, this results in an effective magnetic noise of approximately $90\,\text{fT}_{\text{rms}}/\sqrt{\text{Hz}}$. A graph of the Johnson-current noise of the solenoid calculated using Eq. (8.2) is shown in Fig. 8.3, along with that predicted by a SPICE model which takes into account the entire circuit.

8.1.3 Current Shot Noise

Another source of 'real' magnetic noise arises due to the discrete nature of electric charge. That is, the shot noise associated with the random arrival times of electrons that comprise the current flowing through the solenoid, which generates the magnetic field in the experiment. The current noise per unit bandwidth due to shot noise, δI, is given by

$$\delta I = \sqrt{2eI}\,, \tag{8.3}$$

where $e = 1.6021766208(98) \times 10^{-19}\,\text{C}$ [338] is the elementary charge, and I is the electric current. Given an applied voltage of about 7.1513 V into a resistance of 5 kΩ, the current flowing through the solenoid was about 1.43 mA. This yields a current shot noise of about

$21.4\,\mathrm{pA_{rms}}/\sqrt{\mathrm{Hz}}$; which, using a solenoid calibration factor of around $1.75\,\mathrm{\mu T/mA}$, corresponds to a magnetic noise of approximately $37.5\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$. This is just over a factor of 2 above the demonstrated photon shot-noise limit.

8.1.4 Ambient Field Leakage

Another major noise source which comes to mind when considering technical noise limitations — at least in the context of 'real' magnetic noise — is the leakage of the ambient magnetic field through the shields. This ambient field may be dominated by either the magnetic noise of the laboratory environment, or that of the geomagnetic field. The amplitude spectral density of Earth's geomagnetic field is presented in Fig. 8.5, while that of the laboratory field is presented in Fig. 5.6.

Most of the commercially available magnetic shields are typically made of μ-metal and, for a cylindrical shield with removal end caps, will have a three-layer shielding factor of $\mathcal{O}\left(10^3\right)$. Depending on the level of the ambient magnetic noise in the measurement environment, and the noise requirements within the shielded volume, a three-layer shield *may* be sufficient. However, given that most ambient noise sources exhibit a $\nu^{-\alpha}$ dependence on frequency (where $\alpha > 0$), additional layers of shielding may become necessary if sensitive low-frequency measurements need to be performed. As shown in Fig. 8.4, there was *significant* leakage of the ambient magnetic field through the three-layer cylindrical shield — which had a measured shielding factor of 1800 — that was initially implemented in this work. Therefore, in order to further suppress the ambient field, a six-layer cylindrical μ-metal shield was utilised in the low-frequency magnetic-field measurements performed within this thesis. A 3D render of the six-layer shield, as well as the solenoid and cell mounting, is presented in Fig. 5.4.

From the data presented in Fig. 8.5, the expected magnetic-field noise of Earth's ambient field ranges from $30\,\mathrm{nT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$, to about $2\,\mathrm{pT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{Hz}$. Given a calculated shielding factor of 2×10^5 for the six-layer μ-metal shield described in Sec. 5.3, this corresponds to an expected magnetic-field noise of $150\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$ and $10\,\mathrm{aT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{Hz}$ within the innermost shield, respectively. Likewise, the magnetic-field noise of the laboratory ranges from $70\,\mathrm{nT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$, dropping as low as $4\,\mathrm{nT_{rms}}/\sqrt{\mathrm{Hz}}$ at $0.1\,\mathrm{Hz}$, and increasing to about $10\,\mathrm{nT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{Hz}$ (cf. Fig. 5.6). Within the shielded volume, this level of noise is expected to be attenuated to $350\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$ and $50\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{Hz}$. By inspection of Fig. 6.6, it is clear that ambient field leakage — either from Earth's geomagnetic field, or the laboratory field — is *not* a performance limitation at present.

Although the ambient magnetic field (and its associated noise) can be significantly attenuated through the use of multi-layer shields, there *is* a limit to the level of magnetic-field noise that will exist within the shielded volume. This limit arises from the thermal magnetic noise of the shields themselves, and is discussed in detail in the next section.

Figure 8.4: Time trace demonstrating the ambient magnetic field leaking through three layers of cylindrical μ-metal shields with a measured shielding factor of 1800. All highlighted features were correlated with timestamps of known events. The green and purple shaded regions correspond to the movement of an elevator (and corresponding iron counterweight) within the building, travelling from the basement to the fourth floor and back down again, respectively. The five spikes (circled) within the yellow shaded region correspond to the electromagnetic lock on the laboratory door. Note that all other noise sources described in this chapter would also be included in this measurement.

8.1.5 Shield Noise

In spite of high-permeability shields being used to attenuate external magnetic fields, the *shields themselves* are a source of magnetic noise, albeit small. However, if a magnetometer is sufficiently sensitive, then the overall noise floor may in fact be limited by the shields. There are two sources of magnetic noise in high-permeability magnetic shields: magnetic noise due to the thermal agitation of charge carriers within the shield material (i.e. Johnson noise), and thermal fluctuations in the magnetisation of the shield.

The magnetic-field noise due to Johnson noise, δB_J, in a cylindrical magnetic shield of finite length, L, is given by [339]

$$\delta B_J = \frac{\mu_0 \sqrt{k_B T \sigma_c W}}{a} \sqrt{\frac{2}{3\pi}} G \,, \tag{8.4}$$

where μ_0 is the permeability of free space, k_B is the Boltzmann constant, σ_c is the electrical conductivity of the shield material, W is the shield thickness, a is the shield radius, and G is a factor which depends upon the shield geometry. For a cylindrical shield of finite length, the factor G is given by [339]

Figure 8.5: Amplitude spectral density of Earth's geomagnetic field, reproduced from Ref. [302].

$$G = F_1\left(\frac{L}{a}\right) + 2F_2\left(\frac{L}{a}\right). \tag{8.5}$$

The functions F_1 and F_2 are given by [339]

$$F_1(x) = \sum_\alpha \frac{1}{J_1^2(\alpha)\sinh^2\left(\frac{\alpha x}{2}\right)}, \tag{8.6}$$

$$F_2(x) = \int_0^{1/2} x\left[\sum_\alpha \frac{1}{J_1(\alpha)}\frac{\cosh(\alpha x z)}{\sinh\left(\frac{\alpha x}{2}\right)}\right]^2 dz, \tag{8.7}$$

where J_1 is the first-order Bessel function of the first kind, and the summation values, α, are the roots of the zeroth-order Bessel function of the first kind — i.e. they are the solutions of the equation $J_0(\alpha) = 0$. Since the Bessel functions of the first kind are highly oscillatory, so too are the functions F_1 and F_2. These functions do *eventually* converge; however, as shown in Fig. 8.6, the summation must be truncated at a large number of terms.

Calculating the G factor for our shields, using experimentally measured values and $n = 3000$ summation terms, yields a value of $G \approx 0.43563$. This is within 0.01% of the value $G \approx 0.435345$ that is obtained for an infinitely long cylindrical shield [339]. The G factor as a function of the length-to-radius ratio of a cylindrical shield is shown in Fig. 8.6. Using the experimental measurements, and the calculated value of G for our shield geometry, we obtain a Johnson (white) magnetic noise of $\delta B_J \approx 7.75\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$.

In addition to white Johnson magnetic noise, there is also magnetic noise that arises due to thermal fluctuations in the magnetisation of the shields, which typically dominates at low

Figure 8.6: Left: Convergence of the G factor for a finite cylindrical shield with length-to-radius ratio $L/a = 2$. Calculated G factors as a function of the number of terms are shown in blue markers, with linear interpolation between points to guide the eye. Due to the highly oscillatory nature of the Bessel function of the first kind, as well as the slow damping of the oscillation, the infinite series in Eqs. (8.6) and (8.7) must be evaluated for a large number of terms. Right: Calculated G factor for a finite cylindrical shield, as a function of the length-to-radius ratio. The grey dashed line indicates the G factor for an infitely long cylindrical shield.

Fourier frequencies. The noise due to magnetisation fluctuations is related to the Johnson magnetic noise via [339]

$$\frac{\delta B_{\mathrm{M}}}{\delta B_{\mathrm{J}}} = \sqrt{\frac{3}{2}} \frac{\delta_{\mathrm{skin}}(\nu)}{W} \sqrt{\tan \delta_{\mathrm{loss}}}, \tag{8.8}$$

where $\delta_{\mathrm{skin}}(\nu)$ is the frequency-dependent skin depth, and $\tan \delta_{\mathrm{loss}}$ is the loss tangent, where each of these are defined by [339]

$$\delta_{\mathrm{skin}}(\nu) = \frac{1}{\sqrt{\pi \operatorname{Re}(\mu_{\mathrm{P}}) \sigma_{\mathrm{c}} \nu}}, \tag{8.9}$$

$$\tan \delta_{\mathrm{loss}} = \frac{\operatorname{Im}(\mu_{\mathrm{P}})}{\operatorname{Re}(\mu_{\mathrm{P}})}, \tag{8.10}$$

where μ_{P} is the complex permeability of the shield material.

The total magnetic noise arising from the shields is then given by the quadrature sum of the Johnson noise and magnetisation noise components — that is, they are assumed to be *uncorrelated*. Calculating the total magnetic noise arising from the shield in our experiment gives a noise floor shown in Fig. 8.7. As demonstrated in Sec. 10 and Ref. [206], the magnetic noise arising from the μ-metal shields was a performance limitation in the 1 mHz — 30 mHz range. This is due to the high conductivity (equivalently, low resistivity) of μ-metal,

Figure 8.7: Calculated magnetic noise arising from the μ-metal shields in the experiment. The Johnson (white) noise component, δB_J, is shown in the blue dashed line, whereas the magnetisation noise component, δB_M, is shown in the red dashed line. The total noise floor is the quadrature sum of the two uncorrelated components, and is shown in the grey curve. Note that the magnetisation noise scales as $\nu^{-1/2}$.

which results in large Johnson noise due to the scaling $\delta B_J \propto \sqrt{\sigma_c}$. However, this limitation could be remedied through the use of an innermost layer of ferrite shielding. Ferrite shields have similar permeabilities — and therefore shielding performance — as μ-metal shields; however, they have significantly reduced thermal magnetic-field noise owing to an electrical resistivity which is as much as 6 orders of magnitude larger than μ-metal [340]. Ferrite shields have been demonstrated to generate up to 25 times less Johnson noise, and less than half the thermal magnetisation noise, of μ-metal shields of similar dimensions [340].

8.2 Fictitious Magnetic Noise

Fictitious magnetic noise is defined as unwanted fluctuations in the output of the magnetometer, despite there being no correlated changes in the local magnetic-field strength within the vapour cell. There are a slew of noise sources of this type, and unfortunately many of them are difficult to characterise. However, a general overview of the method used to characterise these types of noise sources is discussed in Sec. 6.5.

This section discusses some of these noise sources, as well as strategies in which to mitigate their effect. The list of sources discussed in this section may not be exhaustive, and the magnitude of their effect may vary wildly between different experimental setups; however, all dominant sources for the experiment detailed within this thesis are discussed here.

8.2.1   Optical-Power Noise

Changes in optical power, whether it be the pump or probe beam, will affect the demodulated NMOR signal in two ways: by changing the resonance slope (cf. Sec. 6.1.2), and by changing the phase of the optical-rotation signal. Typically one will perform experiments by locking to the zero-crossing of the quadrature component of the resonance using a locking system such as the one described in Sec. 6.1.3. In principle this should result in immunity to changes in the resonance slope, as the offset should be zero. In practice however, there are generally small voltage offsets present in the electronics, which will result in a non-zero sensitivity to changes in resonance slope and hence to optical-power fluctuations. Additionally, even if the electronics were ideal components with zero offsets present, changes in optical power will affect the phase of the demodulated signal, which will cause the locking system to adjust the modulation frequency in order to satisfy the condition that the quadrature output of the lock-in amplifier is zero.

By intentionally modulating the optical powers of the pump and probe beams separately, and monitoring the output of the magnetometer, the sensitivity of the fictitious magnetic-field fluctuations to input optical powers are $0.44\,\mathrm{fT/nW}$ and $13.35\,\mathrm{fT/nW}$ for the pump and probe beams, respectively. Using the measured optical-power stability of each beam, these calibration factors can be used to convert the optical-power noise into an effective magnetic-field noise. As shown in Fig. 8.8, the measured fictitious magnetic noise for Fourier frequencies above $1\,\mathrm{mHz}$ is below $1\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ for the pump beam; however, for the probe beam the fictitious magnetic noise is significantly higher, with a maximum of $100\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$. In any case, the fictitious magnetic noise generated by either beam is below the measured magnetic noise floor of the magnetometer, and therefore does not contribute significantly to the magnetic-field measurements.

8.2.2   Polarisation Drift

When setting up the magnetometer to perform a measurement, the optimal condition for the balanced polarimeter is to have an exact 1:1 ratio of the optical powers measured at each photodetector. If there are long-term drifts in the polarisation — typically the result of thermal effects — whether before the cell or after, the balance of the optical powers is disrupted. A mismatch in powers on the photodetectors does two different things: it changes the phase of the optical-rotation signal, and it reduces the slope of the quadrature component. Both of these effects — especially when coupled with the locking system described in Sec. 6.1.3 — give rise to fictitious magnetic-field fluctuations.

By intentionally modulating the angle of the linear polarisation of the probe beam, and monitoring the output of the magnetometer, the sensitivity of the fictitious magnetic-field fluctuations to polarisation drift was observed to be approximately $197\,\mathrm{fT/mrad}$. The polarisation noise of the probe beam measured after the cell, has a flicker-noise characteristic with a magnitude of $700\,\mathrm{\mu rad_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$. This corresponds to an effective magnetic noise due to thermal polarisation drift of the probe beam of about $100\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$.

Figure 8.8: Effective magnetic noise due to optical-power fluctuations of the pump (red) and probe (blue) beams, calculated using experimentally determined calibration slopes of 0.44 fT/nW and 13.35 fT/nW, respectively. Note that the red trace is limited by electronic (white) noise of the photodetector above 0.1 Hz.

Note that the polarisation drift of the pump beam will also introduce an effective magnetic noise, though the magnitude is presumed to be significantly less than that of the probe beam. Given that the polarisation drift of the probe beam was not a performance limitation of the magnetometer under the circumstances presented in this thesis, the effective magnetic noise introduced by the pump beam was not explicitly investigated.

8.2.3 Modulation-Frequency Instability

As explained in previous sections, in order to measure the Larmor frequency (and hence the magnetic-field strength), the amplitude of the pump beam is modulated and a resonance in the optical rotation of the probe beam is observed when the condition $\Omega_m = 2\Omega_L$ is satisfied. By using frequency-demodulation techniques, an asymmetric error signal is created, which enables small changes in Ω_L to be precisely measured. However, when performing an open-loop measurement[3], fluctuations in the modulation frequency give rise to a detector response which is indistinguishable from fluctuations in the magnetic-field strength

In most experiments, the amplitude modulation of the pump beam will be controlled using a signal generator. Signal generators commonly use a temperature-controlled internal Quartz oscillator as a frequency reference, which typically have a fractional frequency stability of around 2 ppb at 1 s, 3 ppb at 10 s, and that degrades to 100 ppb at around 10^3 s.

---

[3]Open loop in this context refers to a measurement in which the pump-modulation frequency remains fixed. This is in contrast to a closed-loop measurement, in which the pump-modulation frequency is locked to, and subsequently tracks, the resonance (cf. Sec 6.1.3).

In the experiments detailed within this thesis, the magnetic-field strength is approximately $2.5\,\mu T$, which corresponds to a carrier frequency of $2\Omega_L/2\pi \approx 35\,kHz$. A fractional frequency stability of 3 ppb at this field strength corresponds to 7.5 fT at 10 s, which is within a small factor of the demonstrated noise floor (cf. Fig. 6.7). For this reason, the signal generator which controls the modulation frequency — and all other electronics, where possible — is referenced to the 10 MHz output of a caesium atomic clock.

The caesium atomic clock used in the experiment is the Datum 4310A, which has a fractional Allan deviation (cf. App. H.2) given in Table. 8.1. The Allan deviation of the atomic clock averages down as white noise with increasing integration time, in accordance with Eq. (H.20). Using Eq. (H.19), this corresponds to an amplitude spectral density of about $600\,nHz_{rms}/\sqrt{Hz}$ at a modulation frequency of $\Omega_m/2\pi = 34.7\,kHz$. In magnetic-field units, this frequency noise corresponds to about $40\,aT_{rms}/\sqrt{Hz}$, which is around two orders of magnitude below the photon shot-noise limit of $15\,fT_{rms}/\sqrt{Hz}$ (cf. Sec. 6.4.2).

Integration time (s)	Fractional Allan deviation
10^0	1.2×10^{-11}
10^1	8.5×10^{-12}
10^2	2.7×10^{-12}
10^3	8.5×10^{-13}
10^4	2.7×10^{-13}
10^5	8.5×10^{-14}
Noise floor	5×10^{-14}

Table 8.1: Fractional Allan deviation of the Datum 4310A caesium atomic clock.

If both the photon shot-noise limit and spin-projection noise were to *somehow* drop below $40\,aT_{rms}/\sqrt{Hz}$ — perhaps through the use of cell heating or by multi-passing the probe to increase the optical depth — the instability of the modulation frequency *may* become a limitation. In this case, it would become necessary to use an optical clock with superior frequency stability. The current state-of-the-art fractional frequency stability of an optical clock is that of a strontium optical lattice clock, which achieves 4.8×10^{-17} at 1 s and 6×10^{-19} at 1 h [341]. In principle, using an optical clock with this level of precision could yield an effective open-loop magnetic noise of $2 \times 10^{-22}\,T_{rms}/\sqrt{Hz}$ due to modulation-frequency instability — over 5 orders of magnitude below the photon shot-noise limit of the current state-of-the-art SERF magnetometer [168].

8.2.4 Photodetector Noise

In the absence of incident light (and hence photon shot noise) on the photodetectors, there will still be some finite electronic noise. This electronic noise is known as 'dark-current noise' and, depending on the experimental circumstances, may be larger in amplitude than

the shot noise of the incident light. Therefore, dark-current noise can become a performance limitation under certain experimental conditions[4].

Fortunately, it is relatively easy to quantify the level of dark-current noise. By blocking any incident light from the photodetectors and measuring the resultant voltage noise of the output, one obtains the dark-current noise multiplied by the gain resistor of the photodetectors[5]. This voltage noise can then be converted to an effective magnetic noise using the slope of the resonance (in voltage units), and the ground-state gyromagnetic ratio.

The combined dark-current noise of the two photodetectors in this work is about $20\,\text{nV}_{\text{rms}}/\sqrt{\text{Hz}}$ which, after scaling by the resonance slope and ground-state gyromagnetic ratio, corresponds to an effective magnetic noise of approximately $1\,\text{fT}_{\text{rms}}/\sqrt{\text{Hz}}$. This is an order of magnitude less than the photon shot-noise limit demonstrated in Sec. 6.4.2.

8.2.5 Pump Scatter

As discussed in Sec. 6.1.2, the local magnetic-field strength is determined by demodulating the balanced-polarimeter signal at the pump-modulation frequency, when it is tuned (or locked) to the resonance condition $\Omega_{\text{m}} = 2\Omega_{\text{L}}$. Due to the pump light being amplitude modulated at the *same frequency* as the measurement frequency, if there is pump light scattering onto the photodetectors, this will yield a dc offset in the output of the lock-in amplifier[6]. In principle this does not affect the magnetometer provided that the dc offset (i.e. the pump-power level) remains constant — since the locking point is permitted to be at a non-zero voltage — however; as discussed in Sec. 8.2.1, the optical powers have a finite noise associated with them despite being locked. Fluctuations in pump power can therefore yield noise in the quadrature output of the lock-in amplifier and, regardless of whether the measurement is performed in open or closed loop, this results in a fictitious magnetic-field noise.

The presence of pump light on the polarimeter photodetectors is especially problematic when using either spatially overlapped pump and probe beams[7], or co-propagating beams at a small angle. In order to remove, or at least minimise, the amount of pump light scattered onto the photodetectors, the pump and probe beams were counter-propagated through the vapour cell. By doing so, it was possible to reduce the presence of any pump light on the photodetectors to below measurable levels, thereby removing this as a source of performance-limiting fictitious magnetic-field noise in the experiment.

---

[4]This is typically the case at low levels of incident optical power.

[5]The photodetectors utilised in these experiments contain a transimpedance amplifier, which sets the gain using a feedback resistor placed between the input and the output.

[6]The presence of pump light on the polarimeter photodetectors will manifest as a dc offset in the output of the lock-in amplifier because the pump beam is *always* modulated at the *same frequency* as the demodulation. The signal resulting from the pump light will remain constant as the modulation frequency is scanned, resulting in a dc offset.

[7]Spatially overlapped beams are necessary when using vapour cells which contain buffer gas. In these cells the transverse spin-relaxation time T_2 is extended by reducing the mean free path of the alkali atoms, thereby reducing the probability of spin-destruction collisions with the cell walls. It is the reduction in mean free path that renders spatial mode matching a necessity.

8.2.6   Counting Noise

As discussed in Sec. 6.1.3, when the pump-modulation frequency is locked to the magnetic resonance in a feedback loop, the magnetic-field fluctuations are measured by counting the modulation frequency rather than measuring voltage fluctuations in the lock-in amplifier output directly. The precision to which the modulation frequency can be determined has a complicated dependence on parameters such as the signal-to-noise-ratio of the input signal to the counter, as well as the integration time of the counter. The frequency-counter noise floor appears to scale linearly as a function of SNR under the experimental circumstances described here. Furthermore, an order of magnitude increase in integration time yields an order of magnitude reduction in observed noise floor.

Under the experimental circumstances in this work, using a gate time of $1\,\mathrm{s}$, the frequency-counter floor is measured to be equivalent to $2\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$, which is approximately an order of magnitude below the photon shot-noise limit. This rules out the frequency counter as a source of performance-limiting noise in the work detailed within this thesis.

8.2.7   Lock-in Amplifier Noise

Given that the lock-in amplifier is one of the major active components of the magnetometer, it has the potential to introduce performance-limiting noise. This section discusses the potential sources of lock-in-amplifier-induced fictitious magnetic noise, and quantifies their effect on the magnetometer under the experimental conditions detailed within this thesis.

8.2.7.1   Gain Noise

In principle, when the pump-modulation frequency is locked to the magnetic resonance via a feedback loop (cf. Sec. 6.1.3), the quadrature component of the lock-in amplifier *should always* yield an output voltage of zero. The magnetometer should therefore be immune to any processes which yield a change in the amplitude of the resonance. However; in practice, any slight offset of the lock point from the zero-crossing of the quadrature component will result in some finite sensitivity to changes in the resonance amplitude. One such mechanism by which the resonance amplitude can be influenced is the gain stage of the lock-in amplifier. If the gain of the lock-in amplifier fluctuates, and there is a finite offset from zero in the quadrature output, the resultant change in amplitude will manifest as a fictitious magnetic-field fluctuation.

When performing a magnetic-field measurement by locking the pump-modulation frequency to resonance, a change in gain by 250% yielded a fictitious magnetic-field deviation of about $8\,\mathrm{pT}$. This corresponds to a calibration factor of about $30\,\mathrm{fT}/\%$. The lock-in amplifier has a specified gain accuracy of $\pm 1\%$, which means that the maximum fictitious magnetic-field deviation expected is approximately $\pm 30\,\mathrm{fT}$. By inspection of the Allan deviation of the magnetometer (cf. Fig. 6.7), this value is below the measured noise floor for all integration times and is therefore not a performance limitation of the magnetometer.

8.2.7.2 Phase Noise

As discussed in App. K, the lock-in amplifier can demodulate the input signal at any arbitrary phase relative to the reference signal. Since the in-phase and quadrature channels of the lock-in amplifier are fixed at $90°$ out-of-phase, it is important to set the arbitrary phase such that the *demodulated signal* in the in-phase and quadrature channels are completely orthogonal. If this is not successfully accomplished, there will be some overlap between the two channels, which will give rise to asymmetry in the in-phase and quadrature outputs.

When fixing the arbitrary demodulation phase and subsequently locking the pump-modulation frequency to the magnetic resonance, if the demodulation phase changes and produces asymmetry in the quadrature output, the pump-modulation frequency will be adjusted to ensure that the quadrature output remains zero. This change in the pump modulation frequency will appear *indistinguishable* from a real magnetic-field strength fluctuation, and hence fluctuations in the demodulation phase constitutes an effective magnetic noise.

According to the manufacturer's specifications, the phase noise of the lock-in amplifier is about $60\,\mathrm{nrad_{rms}}/\sqrt{\mathrm{Hz}}$. By introducing a large change in the demodulation phase, and measuring the response of the magnetometer (while locked), it is possible to obtain a calibration factor by which to map this phase noise to effective magnetic noise. This calibration factor was measured to be $1.6\,\mathrm{nT/rad}$, which gives rise to an effective magnetic noise of approximately $90\,\mathrm{aT_{rms}}/\sqrt{\mathrm{Hz}}$. This is around two orders of magnitude below the photon shot-noise limit of $15\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ (cf. Sec. 6.4.2), and is therefore not a performance limitation for the magnetometer under any realistic circumstances.

8.2.7.3 Input Noise

The two input stages of the lock-in amplifier will add a finite amount of electronic noise onto the input signals. This noise is indistinguishable from real magnetic-field-induced fluctuations in the signal. For the lock-in amplifier used in the experiment (SRS SR830), the input noise is specified to be $6\,\mathrm{nV_{rms}}/\sqrt{\mathrm{Hz}}$ per channel. Assuming that the input noise on each channel is uncorrelated, the combined noise is summed in quadrature. After performing the quadrature sum and scaling by the gain of the lock-in amplifier, the expected voltage noise at the output is approximately $1.7\,\mathrm{\mu V_{rms}}/\sqrt{\mathrm{Hz}}$. Diving by the resonance slope and the ground-state gyromagnetic ratio yields an effective magnetic noise of $0.4\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$, which is over an order of magnitude below the demonstrated photon shot-noise limit.

8.2.7.4 Output Noise

In principle, in the same way that input noise can yield fictitious magnetic-field noise, so too can the output noise. However, under any sensible circumstances, the output noise of the lock-in amplifier is many orders of magnitude smaller than the amplified input noise, and can therefore be disregarded. The only circumstance in which the output noise may be comparable to the input noise, is if the signal amplitude and gain of the lock-in amplifier are

both small. This is generally never the case, at least not under the experimental conditions detailed here, and hence this noise contribution can be safely ignored.

8.2.8   Optical-Frequency Drift

Another source of fictitious magnetic-field noise arises from instability in the optical frequency. Given that the role of the pump beam is only to spin-polarise the atomic vapour and generate ground-state coherences between Zeeman states with $|\Delta m_F| = 2$ (which is possible at non-zero detuning[8] $\Delta \omega \neq 0$), the magnetometer is significantly less sensitive to fluctuations in the optical frequency of the pump beam than of the probe beam. In the case of the probe beam however, since the measured optical rotation is dependent upon detuning (cf. Sec. 4.9), any fluctuations in the optical frequency of the probe beam can yield changes in the magnetometer's output which are indistinguishable from real magnetic-field noise.

Due to the nonlinearity of the observables, and the necessity to numerically solve the density matrix for non-trivial transitions, it is difficult to calculate the exact influence on the magnetometer due to drifts in optical frequency. However, by adjusting the carrier frequency of the AOMs of both the pump beam and probe beams by $\pm 300\,\text{kHz}$[9], there was no detectable effect on the output of the magnetometer. Given that the optical-frequency stability of the laser while frequency locked is less than $100\,\text{kHz}$ for all integration times under 6 hours (cf. App. D, specifically Fig. D.2), this effect can be safely ruled out as a source of performance-limiting technical noise under the experimental conditions detailed within this thesis.

8.2.9   Spectrum-Analyser Noise

When measuring rapid magnetic-field noise (e.g. $>1\,\text{Hz}$) by tuning the modulation frequency to resonance, $\Omega_\text{m} = 2\Omega_\text{L}$, and measuring the quadrature output directly with a spectrum analyser, the spectrum analyser itself may become a performance limitation. This is true in situations where the input noise of the spectrum analyser is larger than the photon shot noise.

The input noise of the spectrum analyser is readily determined by measuring the noise in the *output* of the spectrum analyser when the inputs are terminated with $50\,\Omega$ resistors. This can then be converted to an effective magnetic noise in the same way as per the photodetector noise (cf. Sec. 8.2.4), through division by the resonance slope (in voltage units)

---

[8]This is also true in regards to the detuning of the pump beam relative to the probe beam. In many experiments, as is the case here, the pump and probe beam interact with different velocity classes. In an uncoated cell, this would generally result in no measurable signal; however, in cells with anti-relaxation coatings, it is possible for optically pumped atoms to interact with the probe beam via velocity-changing collisions with the cell walls [342–344].

[9]As discussed in App. C, the carrier freqency of the AOM determines the frequency shift of the first-order diffraction. By adjusting the carrier frequency up or down, it is therefore possible to control the optical frequency of the beam. There will also be an undesirable change in optical power due to a change in driving efficiency (cf. App. C); however, provided the change in carrier frequency is sufficiently small (such as the case here), this effect should be negligible. Furthermore, if the optical powers are stabilised using the locking system described in Sec. 5.1.2, the optical power will remain constant despite changes in the AOM carrier frequency.

Figure 8.9: Comparison of the magnetometer noise floor — characterised using the amplitude spectral density — as measured at various stages along its development.
Between the first noise floor ever measured in early 2015 (red), and the final one
measured in late 2018 (blue), the performance has been increased between 3 and
4 orders of magnitude across the entire range of Fourier frequencies.

and the ground-state gyromagnetic ratio. For the measurements performed within this thesis, the effective magnetic noise of the spectrum analyser is consistently below $1\,\text{fT}_{\text{rms}}/\sqrt{\text{Hz}}$
at 1 Hz, dropping as low as $150\,\text{aT}_{\text{rms}}/\sqrt{\text{Hz}}$ at high Fourier frequencies.

8.3   Summary

The characterisation and subsequent reduction of the technical noise sources discussed in
this chapter, and the fundamental sources mentioned in Sec. 6.4, has been a long and difficult road. Years of hard work have been invested into increasing the performance of the
magnetometer, with improvements occurring incrementally over time. From the first noise
floor ever measured in early 2015, to the final noise floor measured in late 2018, the performance has been increased between 3 and 4 orders of magnitude across the entire range of
Fourier frequencies. A graph showing the amplitude spectral density of the magnetometer
at various stages along its development is presented in Fig. 8.9.

A summary of the various noise sources described in this chapter, as well as the fundamental noise sources detailed in Sec. 6.4, is listed in Table 8.2. The values in this table
correspond to those measured towards the end of the magnetometer's development, i.e.
after significant time had been invested into reducing them.

Noise source	Noise type	Magnitude
Photon shot noise	Fundamental	$15\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$
Spin-projection noise	Fundamental	$3.1\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$
Power supply (without filtering)	Technical	$7.5\,\mathrm{fT}$ at $1\,\mathrm{s}$
		$2.5\,\mathrm{pT}$ at $10^4\,\mathrm{s}$
Power supply (with filtering)	Technical	Indeterminate
Johnson noise	Technical	$90\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$
Current shot noise	Technical	$37.5\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$
Earth-field leakage (six layers)	Technical	$150\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$
		$10\,\mathrm{aT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{Hz}$
Lab-field leakage (six layers)	Technical	$350\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$
		$50\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{Hz}$
Shield noise	Technical	$1.4\,\mathrm{pT_{rms}}/\sqrt{\mathrm{Hz}}$ at $100\,\mathrm{\mu Hz}$
		$7.75\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ above $100\,\mathrm{Hz}$
Optical-power noise (probe)	Technical	$100\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$
Optical-power noise (pump)	Technical	$1\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$
Polarisation drift (probe)	Technical	$100\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{mHz}$
Polarisation drift (pump)	Technical	Indeterminate
Modulation-frequency instability	Technical	$40\,\mathrm{aT_{rms}}/\sqrt{\mathrm{Hz}}$
Photodetector dark-current noise	Technical	$1\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$
Pump scatter	Technical	Indeterminate
Counting noise	Technical	$2\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$
Lock-in gain noise	Technical	Indeterminate
Lock-in phase noise	Technical	$90\,\mathrm{aT_{rms}}/\sqrt{\mathrm{Hz}}$
Lock-in input noise	Technical	$0.4\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$
Lock-in output noise	Technical	Indeterminate
Optical-frequency drift	Technical	Indeterminate
Spectrum-analyser input noise	Technical	$1\,\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at $1\,\mathrm{Hz}$
		$150\,\mathrm{aT_{rms}}/\sqrt{\mathrm{Hz}}$ above $100\,\mathrm{Hz}$

Table 8.2: Fundamental and technical noise sources of the magnetometer — measured towards the end of the magnetometer's development — and their corresponding magnitude.

Part IV

Results

Nonlinear Magneto-Optical Rotation with an Electro-Optic Frequency Comb

This chapter is based on the published article:

9.1 Overview and Motivation

Nonlinear magneto-optical effects are conventionally generated and observed using continuous-wave lasers, which are monochromatic and therefore interact with the medium at a single optical frequency. This optical frequency is generally chosen to be near the centre of a Doppler-broadened transition in order to obtain high interaction strength. One of the drawbacks to using continuous-wave lasers, especially as a probe beam, is that atomic dynamics which exhibit *both* transient temporal *and* spectral dependence are difficult — if not impossible — to observe in real time.

   This chapter describes experimental work in which nonlinear magneto-optical rotation was observed, for the first time, in both the temporal and spectral domains simultaneously. This was achieved by probing the atomic vapour with an electro-optic frequency comb, which was generated via psuedo-random phase modulation of a cw laser. Simultaneous observation of both the temporal and spectral domains enabled polarisation dynamics which had both transient temporal and spectral dependence to be measured in real time. This measurement technique was shown to be able to accurately predict the fraction of polarised atoms within the vapour, as well as the quality of the anti-relaxation coating on the cell walls. Furthermore, this measurement technique has the potential to revolutionise our understanding of transient nonlinear magneto-optical effects, and has immediate applications such as investigating collisional transfer of higher-order polarisation moments between alkali atoms [345].

9.2 Statement of Contribution

9.2.1 Conceptualisation

The idea of using a frequency comb as a probe beam, in order to investigate transient nonlinear magneto-optical effects with both temporal and spectral dependence, is attributed to Szymon Pustelny and André Luiten.

9.2.2 Realisation

The experimental apparatus required to produce the results presented in this manuscript essentially consisted of two independent experiments: an optical magnetometer, and an electro-optic frequency comb. The bulk amount of design and construction of the optical magnetometer was performed by Nathanial Wilson, with assistance from Christopher Perrella and Philip Light. The electro-optic frequency comb was constructed by Nicolas Bourbeau Hébert, based on previous work with Jérôme Genest. The experimental data was acquired by both Nathanial Wilson and Nicolas Bourbeau Hébert, under the supervision of Christopher Perrella, Philip Light, and André Luiten. Analysis of the experimental data was performed by Nathanial Wilson, Nicolas Bourbeau Hébert, Christopher Perrella, Philip Light and André Luiten. Complimentary theoretical modelling was performed by Nathanial Wilson (Monte Carlo) and Christopher Perrella (density matrix), in collaboration with Szymon Pustelny.

9.2.3 Documentation

This paper was written primarily by Nathanial Wilson, Christopher Perrella and André Luiten, with editing by all other authors.

Candidate overall percentage contribution: 50%.

Nathanial Wilson

9.2.4   Co-author Contributions

By signing below, each co-author certifies that:

- The candidate's stated contribution to the publication is accurate (as stated above)

- Permission is granted for the candidate to include the publication in the thesis

Co-author certifications are in order of published authorship unless otherwise specified.

_____

Nicolas Bourbeau Hébert

_____

Christopher Perrella

_____

Philip Light

_____

Jérôme Genest

_____

Szymon Pustelny

_____

André Luiten

Simultaneous Observation of Nonlinear Magneto-Optical Rotation in the Temporal and Spectral Domains with an Electro-Optic Frequency Comb

Nathanial Wilson,[1,][*] Nicolas Bourbeau Hébert,[1,2] Christopher Perrella,[1] Philip Light,[1] Jérôme Genest,[1,2] Szymon Pustelny,[3] and André Luiten[1]

[1]*Institute for Photonics and Advanced Sensing (IPAS), and School of Physical Sciences, The University of Adelaide, South Australia 5005, Australia*
[2]*Centre d'Optique, Photonique et Laser, Université Laval, Québec, QC GIV 0A6, Canada*
[3]*Marian Smoluchowski Institute of Physics, Jagiellonian University, Łojasiewicza 11, Kraków 30-348, Poland*

We simultaneously observe the Larmor precession of rubidium atoms in both the temporal and spectral domains using an electro-optically generated frequency comb. Rubidium vapor is optically pumped on the $F = 2 \rightarrow F' = 1$ hyperfine transition of the ^{87}Rb D$_1$ manifold, while its response is observed with a frequency comb that spans 8 GHz of the spectrum with a temporal resolution of 9.78 μs. The frequency comb modes experience optical rotation by interacting with the $F = 2 \rightarrow F' = 1, 2$ transitions. The spectral and temporal resolution of the comb allows us to observe that there are two separate channels for polarized atoms to appear in the probe beam: one pathway where atoms travel directly from the pump to the probe region, and a secondary pathway that involves interaction of the polarized atoms with the cell walls. The unique features of the comb allow a direct estimate of the relative density of polarized atoms in the vapor cell, as well as a measurement of the quality of the antirelaxation coating on the cell walls. We show that rotation measurements with the comb approach the limits set by photon shot noise.

I. INTRODUCTION

The ability to measure magnetic fields with high sensitivity is important for a myriad of applications, including medical diagnoses and imaging [1–5], measurements of geomagnetism [6], and fundamental physics research [7–10]. One approach to highly sensitive magnetometry exploits nonlinear magneto-optical effects in which a polarized medium undergoes temporal evolution due to interactions with external electromagnetic fields. This temporal evolution results in measurable modulation of the optical parameters of light transmitted through the medium [11]. Through the use of antirelaxation coatings, ultranarrow magneto-optical resonances on the order of approximately 1 Hz can be observed [12], which enables ultrasensitive magnetometry to be performed.

Nonlinear magneto-optical effects are conventionally observed using a continuous-wave (CW) probe laser, which naturally interacts with the medium at just a single optical frequency. This observation frequency is usually chosen to be at the center of a Doppler-broadened transition in order to obtain the highest interaction strength

[13]. However, in this paper we use an electro-optic frequency comb (EOFC) [14–21] as the probe, which allows real-time observations at multiple optical frequencies. This alternative approach has the capability of exploring physical processes such as collisional transfer of higher-order polarization moments between alkali atoms [22], as well as the processes that lead to polarization of the atomic population.

The comb is generated using an electro-optic modulator and enables spectroscopy with temporal and spectral resolutions that approach the limits set by the Fourier transform [16,17]. The ability to simultaneously resolve both temporal and spectral domains provides deep insight into physical processes and can allow observation of previously unobserved effects such as the role of radiation reabsorption in nonlinear spectroscopy [17].

Here, the unique features of the comb are used to show that there are multiple pathways to produce polarized atoms in the medium, that these pathways have different temporal evolutions, and that the fraction of pumped atoms within the vapor cell is approximately 4.4%. Furthermore, a comparison is made between the noise of optical rotation measurements made with a frequency comb and the noise that arises in the conventional approach of a continuous-wave laser. We show that both techniques are compatible with photon shot-noise-limited performance.

_____
[*]nathanial.wilson@adelaide.edu.au

II. EXPERIMENTAL SETUP

The experimental setup is shown in Fig. 1. The magneto-optical effects are generated in a 50-mm-long rubidium vapor cell that has a paraffin coating on the cell walls to extend the spin-relaxation time to approximately 100 μs. The vapor density is set to approximately 7×10^{10} atoms/cm³ [23] by heating the cell to 315 K (42 °C), where the temperature is accurately measured using the Doppler width of the optical transitions [24]. The cell is housed within a three-layer cylindrical μ-metal magnetic shield with a measured shielding factor of approximately 1800 to provide a quiet magnetic environment. Within the innermost shield, a constant-bias magnetic field of approximately 0.8 μT is generated along the longitudinal axis using a Helmholtz coil.

Both pump and probe light propagate through the rubidium vapor parallel to the magnetic field in a conventional Faraday geometry [13], with a beam displacement of approximately 10 mm. The pump beam originates from an external cavity diode laser (ECDL), which is stabilized to a frequency that is 175 MHz red-detuned from the ^{87}Rb $F = 2 \rightarrow F' = 1$ hyperfine transition using saturated absorption spectroscopy. The pump is amplitude modulated via an acousto-optic modulator (AOM) to produce either a single pulse for the free-induction decay (FID) measurements or a continuous train of pulses for the synchronously pumped measurements. The pump beam has a $1/e^2$ diameter of 1.5 mm with a peak power of 15 mW input to the cell and is linearly polarized along the vertical using a Glan-Thompson (GT) prism.

The atomic vapor is probed with an EOFC, which is generated by passing 795-nm radiation from a titanium-sapphire laser through a high-bandwidth, high-index electro-optic modulator (EOM). The EOM is driven between two phase states in accordance with a pseudorandom binary sequence (PRBS). The modulation generates an optical spectrum consisting of a comb of sidebands with a frequency separation of 39.37 MHz, set by the repetition rate of the PRBS. The comb has an approximate sinc² power envelope with its first zeros ±5 GHz away from the seed laser's frequency. The generation of the EOFC, as well as its temporal and spectral profile, is presented in further detail in Ref. [16].

At the input to the vapor cell, the comb is linearly polarized at 45° from vertical by a GT prism. A total comb power of 500 μW reaches the cell, corresponding to approximately 4 μW per mode near the center of the comb. Analysis of the comb's polarization rotation through the vapor cell is achieved using a vertically aligned Wollaston prism (WP), which separates the two orthogonal polarization components. Each of these components is then coupled into single-mode optical fibers.

In parallel, a fraction of the titanium-sapphire light is frequency shifted by an AOM to generate a local oscillator

(LO) signal. This LO is then combined with the two outputs of the Wollaston prism, and the resulting mixing products are detected by separate high-bandwidth photodetectors. The photodetector signals are simultaneously recorded on two channels of a high-speed oscilloscope (2.5-GHz bandwidth, with a sample rate of 20 GS/s).

A total of 500 μs of data is recorded in order to observe the temporal evolution of the rubidium spectrum. The data is sliced into 9.78-μs-long segments by integrating 385 consecutive binary sequences. Each segment is individually Fourier transformed, producing time-resolved spectra with both a high temporal ($\Delta t = 9.78$ μs) and spectral ($\Delta \nu = 39.37$ MHz) resolution. To improve the signal-to-noise ratio, the data presented here consist of 400 averages. The spectral shape of the comb is measured using a reference path (cf. Fig. 1) that bypasses the vapor cell, enabling calibration of the optical and electrical transfer functions.

The transmission $\mathcal{T}_{\mathrm{exp}}(\omega, t)$ and optical rotation $\phi_{\mathrm{exp}}(\omega, t)$ experienced by each mode of the EOFC can be calculated using the powers measured on the two photodetectors P_1 and P_2, together with the total incident comb power P_0, via

$$\mathcal{T}_{\mathrm{exp}}(\omega, t) = \frac{P_1 + P_2}{P_0}, \tag{1}$$

$$\phi_{\mathrm{exp}}(\omega, t) = \frac{1}{2} \arcsin\left(\frac{P_1 - P_2}{P_1 + P_2}\right), \tag{2}$$

where the incident power is measured using the reference path. Figure 2 displays the measured rubidium transmission spectrum $\mathcal{T}_{\mathrm{exp}}(\omega)$, along with a nonlinear fit of the
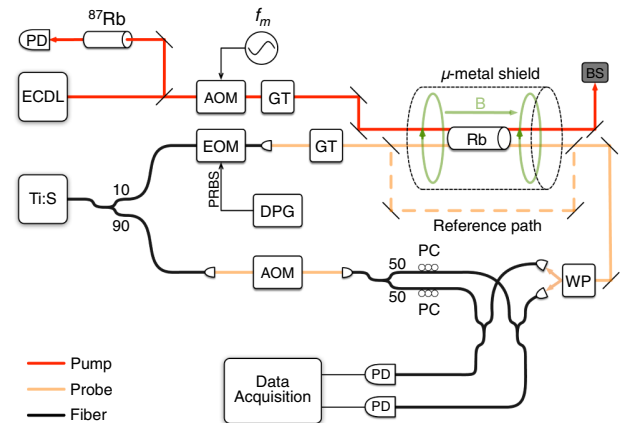

FIG. 1. Simplified experimental setup, showing the external cavity diode laser (ECDL), titanium-sapphire laser (Ti:S), acousto-optic modulator (AOM), Glan-Thompson prism (GT), Wollaston prism (WP), photodetector (PD), beam stopper (BS), polarization controller (PC), electro-optic modulator (EOM), and digital pattern generator (DPG).

FIG. 2.   Rubidium D_1 spectrum as measured by the EOFC with an integration time of 400×9.78 μs (dots), with a fit using Eq. (3). The inset shows the Zeeman sublevels of ^{87}Rb involved in the nonlinear magneto-optical effects observed with the comb, with one possible transition that is connected by the probe radiation shown using red arrows. The shaded features in the transmission spectrum correspond to the optical transitions between the levels shown in the inset, while the unshaded features are associated with the ^{85}Rb isotope ($F = 3 \rightarrow F' = 2, 3$).

form

$$\mathcal{T}_{\text{fit}}(\omega) = \exp\left(-\sum_i \alpha_i \exp\left[\frac{(\omega - \omega_i)^2}{2\sigma^2} \right] \right), \quad (3)$$

where the sum is performed over all hyperfine transitions. Here, σ is the Doppler width, while α_i and ω_i are the optical depth and frequency detuning of the ith allowed transition, respectively.

III. THEORETICAL ANALYSIS

The observation of the atomic vapor with a frequency comb provides insight into the production and relaxation of optically polarized atoms. This section of the paper aims to provide a clear theoretical framework for discussion of the experimental observations.

We measure the temporally and spectrally dependent polarization changes that are imposed on the probe beam in order to track the response of the medium to both optical pumping and the influence of an external magnetic field. The optical rotation of the transmitted probe, $\phi_{\text{th}}(\omega, t)$, is calculated by decomposing the incident optical probe into right- and left-circularly polarized components (σ_+ and σ_-, respectively) and determining the influence on each of these by their associated complex refractive indices, $\eta_+(\omega, t)$ and $\eta_-(\omega, t)$ [25]:

$$\phi_{\text{th}}(\omega, t) = \frac{\omega l}{2c} \text{Re}\left[\eta_+(\omega, t) - \eta_-(\omega, t) \right], \quad (4)$$

while the optical depth of the probe, under the condition that $\text{Im}[\eta_+(\omega, t)] = \text{Im}[\eta_-(\omega, t)]$ (valid for symmetric

ground-state distributions), is given by [26]

$$\alpha_{\text{th}}(\omega, t) = \frac{2\omega l}{c} \text{Im}\left[\eta_\pm(\omega, t) \right], \quad (5)$$

where l is the sample length, ω is the angular frequency of the probe radiation, and c is the speed of light.

Within an atomic vapor at thermal equilibrium, the population distribution over the Zeeman sublevels leads to $\eta_+ = \eta_-$ and hence $\phi_{\text{th}}(\omega, t) = 0$, while $\alpha_{\text{th}}(\omega, t)$ is proportional to the number density. In an atomic population that has been optically pumped by linearly polarized light, coherences appear between the Zeeman sublevels with $\Delta m_F = 2$, which leads to $\eta_+ \neq \eta_-$ (see Appendix A). In this case, the absorption is modified and a rotation of the polarization angle of the transmitted probe is observed.

In the presence of a longitudinal magnetic field, the energies of the ground-state Zeeman sublevels are split by the Larmor frequency, $\Omega_L = m_F \mu_B g_F B / \hbar$, where m_F is the magnetic quantum number, μ_B is the Bohr magneton, g_F is the hyperfine Landé g factor, B is the magnetic field strength, and $\hbar$ is the reduced Planck constant. If an initially polarized medium is allowed to freely relax (i.e., free-induction decay), then the coherences between states with $\Delta m_F = 2$ will evolve at a frequency of $2\Omega_L$ [27,28]. In this case, $\text{Re}[\eta_+(\omega, t) - \eta_-(\omega, t)]$ contains a component that oscillates at a frequency of $2\Omega_L$ (see Appendix A), resulting in modulation of the plane of polarization of the probe laser at the same frequency.

In order to maximize the polarization of the medium at a nonzero magnetic field, one can optically pump the medium with a laser that is amplitude modulated at a frequency $\Omega_m \sim 2\Omega_L$. In this case, the probe beam will show a modulation of the plane of polarization at the driving frequency Ω_m and will be resonantly enhanced if $\Omega_m = 2\Omega_L$ [29].

The narrow spectral width of the pump means that it can only directly optically pump atoms that reside in a relatively narrow spectral window (set by the power-broadened width of the pump-atom interaction). At first sight, one might expect that the probe will only observe a relatively narrow peaked spectrum of polarized atoms. However, in a cell with an antirelaxation coating, atoms may bounce from the wall while maintaining their ground-state coherence. This process produces a thermalization of their longitudinal velocity so that typical experiments observe polarization across the full Doppler-broadened spectrum. However, these temporally and spectrally resolved frequency-comb observations allow us to directly observe two different populations of polarized atoms contributing to the optical rotation of the probe—the atoms that have taken a direct pathway from the pump to the probe, as well as those that have bounced from the cell wall in the time between those two interactions.

These two atomic populations are distinguishable by their spectral distributions. Polarized atoms that have taken the direct pathway between pump and probe regions show a narrow spectral peak that rapidly disappears, while atoms that have bounced from the cell wall exhibit a longer-lived Gaussian-shaped spectrum:

$$\phi_{\text{th}}(\omega) = \sum_i \phi_i \exp\left[-\frac{(\omega - \omega_i)^2}{2\sigma^2}\right], \qquad (6)$$

characterized by the Doppler width σ and a center frequency ω_i and where ϕ_i is the peak optical rotation of the ith transition. The sum in this case is performed over all hyperfine transitions. The probe absorption also shows a typical thermal velocity distribution because the atomic vapor is in thermal equilibrium with the cell walls:

$$\alpha_{\text{th}}(\omega) = \sum_i \alpha_i \exp\left[-\frac{(\omega - \omega_i)^2}{2\sigma^2}\right], \qquad (7)$$

where α_i is the peak optical depth on the ith resonance and all other terms are defined as per Eq. (6). Our measurements show that the Doppler width of the optical rotation spectrum in Eq. (6) and the width in Eq. (7) are identical within experimental error.

The unique ability of the frequency comb to isolate the two polarized atomic populations enables us to make two independent estimates of the fraction of polarized atoms within the cell. The first approach recognizes that unpolarized atoms make no contribution to the optical rotation, while making the predominant contribution to the optical absorption. This approach gives us the opportunity to directly measure the fraction of the polarized atoms by taking the ratio of Eqs. (6) and (7). Employing the three-level density matrix model detailed in Appendix A, we can simplify the ratio of the optical rotation amplitude to the optical depth for the two ^{87}Rb hyperfine transitions spanned by the frequency comb:

$$\frac{\phi_i(\omega)}{\alpha_i(\omega)} = 0.63\epsilon, \qquad (8)$$

where i corresponds to the ith hyperfine transition and $\epsilon = N_\phi/N$ is the fractional density of polarized atoms, where N_ϕ is the density of optically pumped atoms which generate optical rotation and N is the total atomic number density. The prefactor originates from averaging the ratio over a given observation time based on the experimental parameters used below.

IV. RESULTS

A. Free-induction decay

The capability of the EOFC to resolve the spectral and temporal evolution of the optical rotation is demonstrated by observing the free-induction decay of the atomic polarization. A CW laser optically pumps the $F = 2 \rightarrow F' = 1$ hyperfine transition of the ^{87}Rb D$_1$ manifold to create a steady-state atomic polarization. The evolution of the decay is then observed with the EOFC after the pump light is switched off (cf. Fig. 3).

A number of strong features are evident in Fig. 3. In the time domain, the optical rotation shows an exponentially damped sinusoidal modulation. This modulation is a result of the applied approximately 0.8-μT magnetic field producing a Larmor frequency of $\Omega_L/2\pi = 5.7$ kHz. The exponential decay has a time constant of 100 μs, which reflects the relatively poor quality of the antirelaxation coating in this particular cell [30–32]. Additionally, the optical rotation for the two hyperfine transitions, $F = 2 \rightarrow F' = 1, 2$, is observed to evolve out of phase due to a sign difference between their respective Clebsch-Gordan coefficients (cf. Appendix B).

The pump laser is red detuned 175 MHz from the center of the Doppler profile of the $F = 2 \rightarrow F' = 1$ transition and hence can only polarize atoms within a narrow velocity class. However, as described earlier, the line shape of the optical rotation spectrum (cf. Fig. 3(b) for detailed cross sections) is much broader than the spectral width of the pump and exhibits predominantly a Gaussian line shape, particularly for measurements greater than 20 μs after the pump beam is switched off. As described earlier, the Gaussian line shape arises due to the atomic velocities of the polarized atoms that have been thermalized by one or more wall collisions prior to entering the probe region [33]. This is in line with previous experiments, which show that only a small number of collisions with the cell walls is sufficient to reproduce the full Doppler profile for a gas [34].

A closer examination of the initial time steps in Fig. 3(b) reveals a peaked structure that differs from the Gaussian line shape but decays within approximately 20 μs. We postulate that this anomaly is associated with a small number of atoms that are optically polarized by the pump and then travel directly into the probe beam without interacting with the cell walls [35]. The timescale for the rapid decay is consistent with the typical time for an atom to transit from the pump to the probe beam ($\tau_{\text{tr}} \sim 33 \pm 14$ μs, where the errors are from the uncertainty in the pump and probe separation). To analyze this feature in more detail, we generate a significantly enhanced atomic polarization using a synchronous-optical-pumping technique.

B. Synchronously pumped optical rotation

The synchronous-optical-pumping technique achieves a large steady-state atomic polarization by modulating the pump amplitude at the second harmonic of the Larmor frequency, $\Omega_m = 2\Omega_L$. In this experiment, we generate a 100% amplitude modulation of the pump with a 25% duty cycle at a rate of $\Omega_m/2\pi = 10$ kHz. The atomic
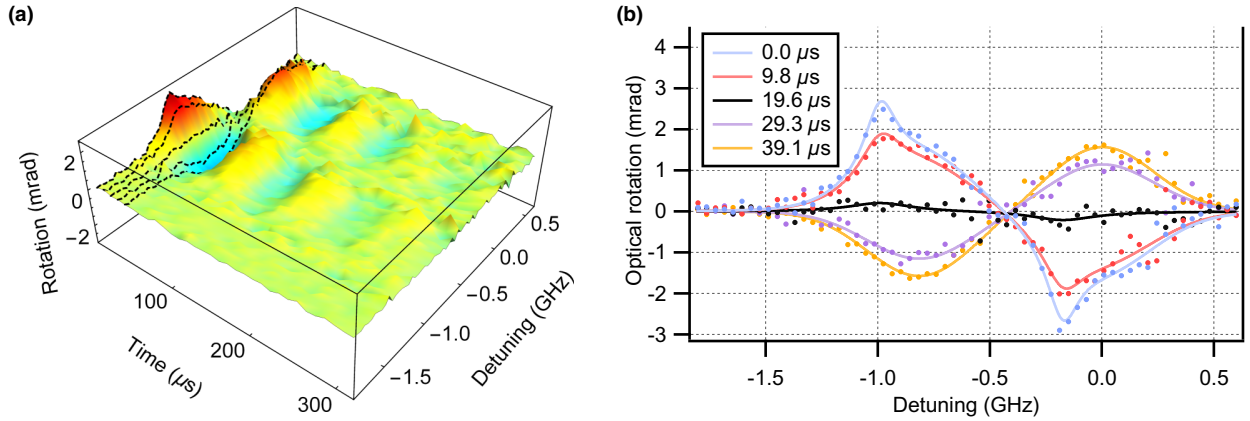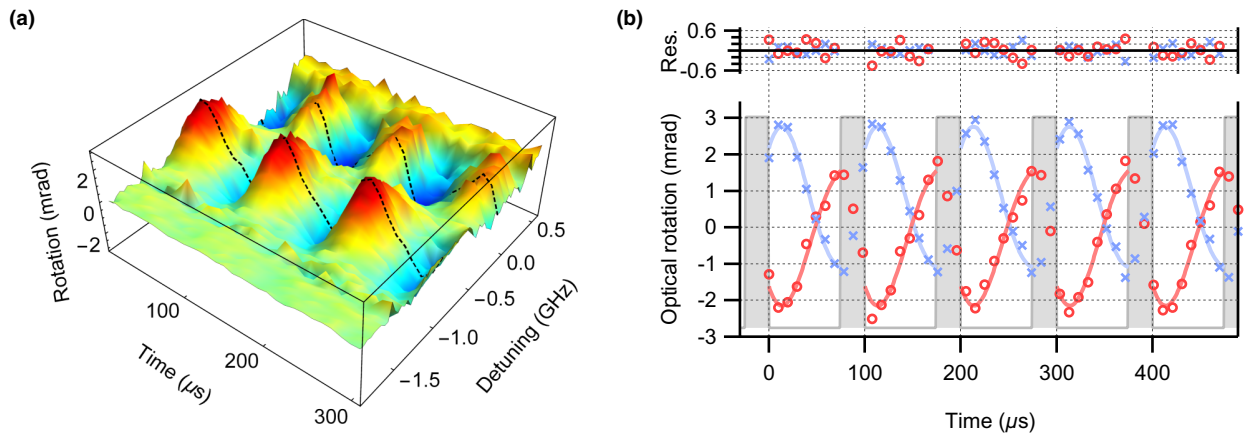
FIG. 3. (a) Three-dimensional surface plot of the optical rotation of the electro-optic frequency comb using the free-induction decay technique for the ^{87}Rb $F = 2 \rightarrow F' = 1, 2$ transitions at a relative detuning of -0.82 and 0.0 GHz, respectively. The temporal and spectral resolutions are $\Delta t = 9.78$ μs and $\Delta \nu = 39.37$ MHz, respectively. The data are averaged 400 times, with the atomic-linear-polarization response removed in postprocessing. (b) Spectra for the first five time steps indicated by the dashed lines on the surface plot. Each spectrum is fit using Eq. (10), where the fitting parameters have been held fixed using those obtained from fitting Eq. (3) to Fig. 2.

response in the temporal and spectral domains (cf. Fig. 4) exhibits similar features to the FID data; however, the synchronous optical pumping allows us to average over multiple pumping cycles.

The time-domain optical rotation $\phi_{\exp}(t)$ for the central comb modes of the two hyperfine transitions is displayed in Fig. 4(b). A relatively nonsinusoidal optical rotation is observed in these data due to a combination of inadvertently detuning from the resonance condition ($\Omega_m = 2\Omega_L$) by one resonance width (approximately 1.5 kHz)

and ground-state relaxation between pumping pulses. As a result, each comb mode's optical rotation is only fitted during the period in which the pump beam is switched off. In this case, the optical rotation signals exhibit the following form:

$$\phi_{\text{fit}}(t) = a \sin(\omega t + \varphi) \exp\left(-\frac{t}{\tau}\right) + c, \qquad (9)$$

where ω is the angular frequency; a is the optical rotation amplitude; φ is the phase; τ is the spin-relaxation time,

FIG. 4. (a) Three-dimensional surface plot of the optical rotation of the electro-optic frequency comb using the synchronous-optical-pumping technique for the ^{87}Rb $F = 2 \rightarrow F' = 1, 2$ transitions at a relative detuning of -0.86 and 0.0 GHz, respectively. Dashed lines indicate the central comb mode for each hyperfine transition. The temporal and spectral resolutions are $\Delta t = 9.78$ μs and $\Delta \nu = 39.37$ MHz, respectively. The data are averaged 400 times with the atomic-linear-polarization response removed in postprocessing. (b) Optical rotation measured by the center comb modes of the hyperfine transitions (dashed lines on the surface plot) when pumped at a modulation frequency of $\Omega_m \sim 2\Omega_L$. Blue markers (crosses) correspond to the $F = 2 \rightarrow F' = 1$ hyperfine transition, while the red markers (circles) correspond to the $F = 2 \rightarrow F' = 2$ transition. A sinusoidal fitting function, Eq. (9), is applied to the comb data, with the pump-beam modulation shaded in gray.

which is held fixed at $\tau = 100\ \mu$s, obtained from fitting the FID data presented in Fig. 3; and c is an offset that accounts for the nonuniform optical rotation due to the optics over the span of the broadband comb.

We plot the optical rotation amplitude a for each of the comb modes in Fig. 5, which shows an improved signal-to-noise ratio when compared to the FID data presented in Fig. 3(b). The optical rotation spectrum $\phi(\omega)$ is modeled as

$$\phi_{\text{fit}}(\omega) = \sum_i \mathcal{L}_i(\omega - \omega_i - \omega_p) + \mathcal{G}_i(\omega - \omega_i), \quad (10)$$

where ω is angular frequency; ω_i is detuning relative to the $F = 2 \rightarrow F' = 2$ transition; ω_p is the pump detuning from resonance; and the sum is performed over the $F = 2 \rightarrow F' = 1, 2$ hyperfine transitions. As described in Sec. III, atoms that travel directly from the pump to the probe region should reflect the velocity-selective nature of the atom-light interaction and should therefore give rise to a Lorentzian spectral signature, $\mathcal{L}_i(\omega - \omega_i - \omega_p)$. A second group of atoms enters the probe region following one or more wall collisions and produces a Gaussian spectral profile $\mathcal{G}_i(\omega - \omega_i)$ due to the thermalization of the atoms' velocity along the probe-beam propagation vector.

The optical rotation spectrum presented in Fig. 5 has been fit using Eq. (10), with the Doppler width fixed at the value extracted from the transmission spectrum (cf. Fig. 2). Only three free parameters are allowed in the fitting algorithm: the amplitude of the two Gaussian components

FIG. 5.   Spectral dependence of the optical rotation amplitude a (blue markers) and average optical depth of the two polarization components (red dots) of each comb mode. The red line is a fit to the red dots using Eq. (3), while the blue line is a fit to the blue markers using Eq. (10). Also presented are the spectral contributions to the optical depth from each hyperfine transition (gray lines), along with Lorentzian contributions to the optical rotation amplitude (dashed gray line) that have been scaled by a factor of 2 for visibility. Two ^{85}Rb hyperfine transitions ($F = 3 \rightarrow F' = 2, 3$) extend off the graph at positive detuning.

and the amplitude and width of the two Lorentzian features (which are assumed to be the same for the two resonances). The Lorentzian detuning ω_p is found to be consistent with the known pump frequency, while the fitted width of approximately 100 MHz is consistent with the power broadening expected from a pump beam with an input intensity that is around 200 times above the saturation intensity.

The extracted amplitude of the Gaussian component of the rotation spectrum in Fig. 5 is $\phi = 3.06$ mrad, while the optical depth in Fig. 2 is $\alpha = 0.11$. As outlined in Sec. III, these two values, along with Eq. (8), can be used to estimate that the fraction of polarized atoms within the vapor is approximately 4.4%. An approximate estimate of the expected number of polarized atoms can be derived from the knowledge that the pump can only interact with about 20% of the atoms in the thermal distribution (as shown in Fig. 5), that the pump has an effective pumping volume of around 36% of the cell (2-cm diameter) in a single pump interaction of 25 μs duration with a typical thermal velocity of rubidium (245 ms^{-1} at 42 °C), and that the spin-relaxation time is approximately 100 μs. This combination of effects leads to an expected polarized population averaged over the pump cycle of around 5%, which is in good agreement with the experimental observation.

The comb also allows an estimate of the quality of the antirelaxation coating by considering the relative fraction of optically polarized atoms that enter the probe via the two possible routes: either directly from the pump to the probe regions or through one or more wall collisions. The amplitude of the Lorentzian component shown in Fig. 5 is about 30% of the Gaussian component, which indicates the relative density of atoms arriving via the two pathways. On the other hand, Monte Carlo modeling using the size of the two beams (1.5- and 3.5-mm $1/e^2$ diameter) and their separation (6–10 mm) suggests that only about 3%–9% of the atoms leaving the pump beam should directly pass through the probe region. This result suggests that about 75% of atoms bouncing from the cell wall maintain their ground-state coherence, which is consistent with the spin-relaxation time of this cell (approximately 100 μs, corresponding to a few bounces before relaxation).

C. Noise comparison between EOFC and CW probing

A key determination is whether the additional information that is extracted by the EOFC comes at the cost of a noise penalty. For these purposes, we made a careful comparison to a conventional single-frequency probe beam on the same experiment.

A total comb power of 500 μW is incident on the vapor cell, which corresponded to around 4 μW of optical power per mode for modes near the center of the comb. The cell transmission, polarization splitting at the Wollaston prism, and coupling to the photodetectors reduces the detected

power per mode to around 0.3 μW. The signals from each comb mode are detected using a heterodyne read-out with a local oscillator of much higher power, in which case the shot noise of each comb mode is the dominant noise source [36]. The calculated shot-noise-limited sensitivity for a single comb mode is approximately 26 μrad (cf. Appendix C). The fit residuals presented in Fig. 4(b) have a standard deviation of approximately 150 μrad, around a factor of 6 above the shot-noise limit. The current performance limit is set by electronic noise of the photodetectors as confirmed by a dark-current measurement.

We compare the EOFC measurements with a CW probe measurement. In this case, the magnetic field and corresponding pump modulation frequency are 2.9 μT and $\Omega_m/2\pi = 40$ kHz, respectively. The probe beam had an optical power of 100 μW and is tuned to the center of the $F = 2 \rightarrow F' = 1$ hyperfine transition of the ^{87}Rb D$_1$ manifold. The optical rotation signal measured by this probe beam is presented in Fig. 6, along with a fit using Eq. (9). The fit residuals have a standard deviation of 3 μrad. Using a similar approach to that detailed in Appendix C, the shot-noise-limited optical rotation is readily calculated from the power on the photodetector along with Eq. (2) to be 0.5 μrad in the measurement bandwidth. This measurement is therefore a factor of 6 from the shot-noise limit and is consistent with the electronic noise of the photodetectors.

In the particular conditions here, photodetector noise set the limit because we chose to use relatively low probe powers to minimize the perturbation of the ground-state

FIG. 6. Optical rotation measurements using a conventional CW laser tuned to the center of the $F = 2 \rightarrow F' = 1$ hyperfine transition of the ^{87}Rb D$_1$ manifold. The data presented (blue markers) are averaged 2048 times and fitted (blue line) using Eq. (9). The pump-beam modulation is shaded in gray. An offset due to misbalancing the polarimeter has been removed in postprocessing.

atomic coherence. A change in experimental conditions is likely to yield shot-noise-limited EOFC measurements. The perhaps surprising observation here is that there is no significant introduction of noise arising from the use of a comb-based probe despite the additional information that is acquired from the atomic system.

V. CONCLUSION

We demonstrate the ability of an electro-optic frequency comb to simultaneously observe the Larmor precession of rubidium atoms in the temporal and spectral domains. The free-induction decay of the atomic vapor is observed on both the $F = 2 \rightarrow F' = 1, 2$ transitions, which shows the optical rotation of the two manifolds oscillating out of phase due to sign differences of the Clebsch-Gordan coefficients between the two transitions. We also observe the exponential decay of the atomic polarization due to collisions with the cell walls. The comb observations allow us to distinguish two groups of polarized atoms: those that have undergone one or more wall collisions, which exhibit a Doppler-broadened spectral profile, and those that have followed a direct path between the pump and probe regions, which show a much narrower spectral feature that rapidly relaxes.

We compare the rotational sensitivity performance of the comb measurement with that of a continuous-wave measurement and show that both techniques can approach the shot-noise limit. The comb technique thus allows rapid temporally and spectrally resolved measurements to be performed without introducing a significant amount of noise with respect to a continuous-wave measurement.

APPENDIX A: DENSITY MATRIX THEORY

The dynamics of the ^{87}Rb atoms in the pump and probe beams are independent as there is no cross-sectional overlap between the two beams. The interaction between the pump and probe is therefore mediated by atoms traversing these two beams. In order to model the observed optical rotation signals, it is necessary to incorporate both the pumping and probing processes, as well as the kinetics that lead to atoms traversing through two separate optical fields.

The ^{87}Rb atoms are initially optically pumped by the pump beam, which is tuned to the $F = 2 \rightarrow F' = 1$

hyperfine transition with a Rabi frequency of $\Omega \approx 56$ MHz—much larger than the decay rate of the excited state, $\Gamma = 5.75$ MHz. A strongly driven optical transition such as this, $\Omega \gg \Gamma$, results in the $F = 2$ ground state being rapidly optically pumped into a population distribution, which forms a dark state, established on a timescale of Γ^{-1}. This dark state exhibits negligible absorption and thus slows the optical pumping of atoms into the uncoupled $F = 1$ ground state. Since the atomic population is "trapped" in the $F = 2$ ground state during the pumping process, with little population dynamics, this enables the system to be simplified into a three-level model to better understand the pumping and probing processes.

The simplified energy-level diagram used in these calculations is shown as red arrows in the inset of Fig. 2. The quantization axis is defined to be in the direction of the magnetic field, which is parallel to the pump and probe beams. The linearly polarized pump and probe beams can be decomposed into two circularly polarized components, leading to our simple model using only three energy levels: two ground states $|m_1\rangle$ and $|m_{-1}\rangle$ and a single excited state $|e\rangle$. The excited-state decay rate is given by Γ, while the relaxation of the ground states due to dephasing collisions is given by γ. A first-order approximation to the full model of the ^{87}Rb atoms can be reproduced by summing a number of these simplified models together that cover all of the ground and excited states.

The temporal evolution of the density matrix ρ, representing the simple model in both beams, is governed by the Liouville equations [37]. The Liouville equations governing the dynamics of the simple model atom take the general form

$$\frac{d\rho}{dt} = -\frac{i}{\hbar}[\mathscr{H}_{\text{atom}} + \mathscr{H}_{\text{int}}, \rho] + \Lambda(\rho), \qquad (A1)$$

where the atomic and light-atom interaction Hamiltonians under the rotating wave approximation are given by

$$\mathscr{H}_{\text{atom}} + \mathscr{H}_{\text{int}} = \hbar \begin{pmatrix} \Omega_L & 0 & \Omega/2 \\ 0 & -\Omega_L & -\Omega/2 \\ \Omega/2 & -\Omega/2 & -\Delta \end{pmatrix}, \quad (A2)$$

where Ω_L is the Zeeman splitting of the ground states relative to the $|m_0\rangle$ ground state, Δ is the optical detuning of the laser field from the atomic transition frequency, and Ω is the Rabi frequency of the atom-light interaction given by

$$\Omega = \frac{\mu E_0}{\hbar}, \qquad (A3)$$

where E_0 is the amplitude of the driving electric field and μ is the electric dipole moment for the $|m_{\pm 1}\rangle \rightarrow |e\rangle$ transition.

Relaxation is included via the Λ term which takes the form

$$\Lambda(\rho) = \frac{\Gamma}{2} \begin{pmatrix} \rho_{e,e} & 0 & -\rho_{m_1,e} \\ 0 & \rho_{e,e} & -\rho_{m_{-1},e} \\ -\rho_{e,m_1} & -\rho_{e,m_{-1}} & -2\rho_{e,e} \end{pmatrix}$$

$$+ \frac{\gamma}{2} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} - \gamma\rho. \qquad (A4)$$

The first line is the contribution from spontaneous decay, while the second line contains the contributions from ground-state repopulation and relaxation, respectively.

The atomic system approaches a steady state on a timescale of Γ^{-1}, which is much shorter than both the time taken for the atoms to cross the pump beam and the time in which the pump beam is switched on. As a result, the simplified atomic system leaves the pump beam having been optically pumped to a steady state, with a population distribution and coherences given by

$$\rho_{m_1,m_1} = \rho_{m_{-1},m_{-1}} = 1/2, \qquad (A5)$$

$$\rho_{m_1,m_{-1}} = \rho_{m_{-1},m_1} = -1/2, \qquad (A6)$$

with all other terms equal to zero when $\Omega_{m_{\pm 1},e} \gg \Omega_L, \Gamma, \gamma$. After the pump light is switched off, the atomic system undergoes free evolution, which results in an oscillatory response of the ground-state coherences given by

$$\rho_{m_1,m_{-1}} = \rho_{m_{-1},m_1}^* = \frac{1}{2}e^{-\gamma t}e^{i2\Omega_L t}. \qquad (A7)$$

It can be seen that the coherences evolve at a frequency of $2\Omega_L$, which is the frequency separation between the two ground-state Zeeman sublevels, $|m_1\rangle$ and $|m_{-1}\rangle$, and decay at a rate of γ due to ground-state relaxation.

The atomic interaction with the probe beam can be greatly simplified under the assumption that the probe does not significantly perturb the atomic populations. This is a valid assumption given the relatively low optical power of the probe beam. When the steady-state populations and the freely evolving ground state coherences are substituted into Eq. (A1), the resulting time-dependent coherences between states $|m_{\pm 1}\rangle$ and $|e\rangle$ are given by

$$\frac{d\rho_{m_{\pm 1},e}}{dt} = \pm \frac{i\Omega}{4}\left(1 - e^{-\gamma t}e^{\mp 2it\Omega_L}\right)$$

$$- \left[\gamma + \frac{\Gamma}{2} + i(\Delta \pm \Omega_L)\right]\rho_{m_{\pm 1},e}, \qquad (A8)$$

where Ω is the Rabi frequency of the probe light. Solving this equation for $t \gg \Gamma^{-1}$ yields the solution:

$$\rho_{\pm 1,e} \approx A_{\pm} + B_{\pm}e^{-\gamma t}e^{\mp i(2\Omega_L t + \theta_{\pm})}. \qquad (A9)$$

The first term, A, represents the dc atomic response, which is present even when the pump beam is switched off, and

corresponds to the typical two-level absorption. The second term, which oscillates at a frequency of $2\Omega_L$, is the nonlinear magneto-optical rotation. This term is induced by the pump beam generating coherence between the two ground-state Zeeman sublevels and is observed experimentally. The coefficients $A_\pm$ and $B_\pm$ and the phase of the coherences $\theta_\pm$ are given by

$$A_\pm = \pm \frac{\Omega_{m_{\pm 1},e}}{2} \frac{i(\Gamma + 2\gamma) + 2(\Delta \pm \Omega_L)}{(\Gamma + 2\gamma)^2 + 4(\Delta \pm \Omega_L)^2}, \quad (A10)$$

$$B_\pm = \frac{\Omega_{m_{\pm 1},e}}{2} \frac{1}{\sqrt{\Gamma^2 + 4(\Delta \mp \Omega_L)^2}}, \quad (A11)$$

$$\theta_\pm = \arctan\left[\mp 2(\Delta \mp \Omega_L), \Gamma\right], \quad (A12)$$

The optical depth and optical rotation for an optically thin media can be calculated from [37]

$$\alpha = \frac{N\omega\hbar l}{\sqrt{2}} \frac{\Omega}{I} \text{Im}\left(\rho_{m_{-1},e} - \rho_{m_1,e}\right), \quad (A13)$$

$$\phi = -\frac{N_\phi \omega\hbar l}{\sqrt{2}} \frac{\Omega}{I} \text{Re}\left(\rho_{m_{-1},e} + \rho_{m_1,e}\right), \quad (A14)$$

where N_ϕ is the density of optically pumped atoms that generate optical rotation; N is the total atomic number density; l is the optical path length through the ^{87}Rb vapor; and I is the optical intensity, which is given by

$$I = \frac{c\epsilon_0 |E_0|^2}{2}, \quad (A15)$$

where ϵ_0 is the vacuum permittivity and c is the speed of light. Taking the ratio of the oscillating optical rotation to the steady-state optical depth when on resonance (i.e., $\Delta = 0$) using the solution to the density matrix, Eq. (A9), and noting that $\Omega_L \ll \Gamma$ leads to

$$\begin{aligned}
\frac{\phi}{\alpha} &= \frac{N_\phi}{N} \frac{\text{Re}\left(\rho_{m_{-1},e} + \rho_{m_1,e}\right)}{\text{Im}\left(\rho_{m_{-1},e} - \rho_{m_1,e}\right)}, \\
&= \frac{N_\phi}{N} \frac{\text{Re}(B_- + B_+)}{\text{Im}(A_- - A_+)}, \\
&= \frac{N_\phi}{N} \frac{(2\gamma + \Gamma)^2 + 4\Omega_L^2}{(2\gamma + \Gamma)\sqrt{\Gamma^2 + 4\Omega_L^2}} e^{-\gamma t}, \\
&\approx \frac{N_\phi}{N} e^{-\gamma t}.
\end{aligned} \quad (A16)$$

The second line follows from the fact that we are considering the amplitude of the oscillating optical rotation

at the resonant frequency, $2\Omega_L$, while the measured optical depth is the dc component. The final line follows from the excited-state decay rate being much faster than both the Larmor frequency and the ground-state relaxation rate, i.e., $\Gamma \gg \Omega_L, \gamma$. As the optical rotation $\phi(t)$ decays due to ground-state relaxation, the ratio between the optical rotation and optical depth also exponentially decays with time.

Atoms that contribute to the optical rotation signal are those that possess coherence between the ground-state Zeeman sublevels upon being probed. On the other hand, all available atoms contribute to the optical depth measured by the probe beam. Thus, the relative density of polarized atoms, $\epsilon = N_\phi/N$, can be directly related to the ratio of the optical rotation to the optical depth using Eq. (A16). For the experimental parameters used here, i.e., a ground-state relaxation rate of $\gamma \approx 10^4$ Hz and a pumping period of $t_0 = 100$ μs, the fraction of polarized atoms is given by

$$\begin{aligned}
\frac{\phi}{\alpha} &= \frac{1}{t_0} \int_0^{t_0} \frac{N_\phi}{N} e^{-\gamma t} dt, \\
&= 0.63\epsilon,
\end{aligned} \quad (A17)$$

which is in good agreement with experimental observations. The validity of this simplified model is checked against a numerical model that incorporated the full ^{87}Rb energy-level structure, with both models predicting a similar result to Eq. (A17).

APPENDIX B: HYPERFINE PHASE SHIFT

The π phase shift between the two hyperfine transitions observed in Fig. 4 is the result of a sign difference of the Clebsch-Gordan coefficients between right- and left-circularly polarized light for the two transitions. This can clearly be seen when evaluating the simplified model shown in Fig. 2. In order to model the $F = 2 \rightarrow F' = 1$ hyperfine transition of the ^{87}Rb D$_1$ manifold, two of the five ground states, $|m_\pm\rangle = |F = 1, m_f = \pm 1\rangle$, and one of the three excited states, $|e\rangle = |F' = 1, m_f = 0\rangle$, are considered (cf. Appendix A). After implementing the rotating-wave approximation, the light-atom interaction Hamiltonian for this transition is given by

$$\mathcal{H}_{2,1} = \frac{\hbar\Omega_{2,1}}{2} \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 1 \\ -1 & -1 & 0 \end{pmatrix}, \quad (B1)$$

where $\Omega_{2,1}$ is the Rabi frequency for the $F = 2 \rightarrow F' = 1$ transition. However, using the same approximations, the interaction Hamiltonian for the $F = 2 \rightarrow F' = 2$ hyperfine

transition exhibits a sign change of two of the elements:

$$\mathscr{H}_{2,2} = \frac{\hbar\Omega_{2,2}}{2} \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}. \qquad \text{(B2)}$$

The elements in question correspond to one of the circular components of the linearly polarized light. Changing the sign of these elements produces a π phase shift between the observed optical rotation in the two hyperfine transitions.

APPENDIX C: SHOT-NOISE CALCULATION FOR THE EOFC

Provided that the power of the local oscillator (LO) is significantly greater than the comb power (as is the case here), the effective shot noise for each mode is independent of the LO and is instead proportional to the power in the mode [36,38]. To calculate the resulting shot-noise-limited rotation noise, we substitute the shot noise of each of the photodetectors into Eq. (2):

$$\sigma_\phi = \sqrt{\left(\frac{\partial\phi}{\partial P_1}\sigma_{P_1}\right)^2 + \left(\frac{\partial\phi}{\partial P_2}\sigma_{P_2}\right)^2}, \qquad \text{(C1)}$$

where P_1 and P_2 are the optical powers measured at the two photodetectors, σ_{P_1} and σ_{P_2} are the shot-noise-induced power fluctuations of the two photodetectors, and σ_ϕ is the optical rotation noise. For a balanced polarimeter, the optical power on each photodetector is approximately equal ($P_1 \approx P_2$) and hence so is the shot noise $\sigma_{P_1} \approx \sigma_{P_2}$. This simplifies Eq. (C1) to

$$\sigma_\phi = \sqrt{2}\frac{\partial\phi}{\partial P_1}\sigma_{P_1}. \qquad \text{(C2)}$$

Calculation of the partial derivative and subsequent simplification yields

$$\sigma_\phi = \frac{1}{2\sqrt{2}}\frac{\sigma_{P_1}}{P_1}. \qquad \text{(C3)}$$

Since the optical power on the photodetector is linearly proportional to the photocurrent I, Eq. (C3) can be rewritten in terms of current as

$$\sigma_\phi = \frac{1}{2\sqrt{2}}\frac{\sigma_{I_1}}{I_1}, \qquad \text{(C4)}$$

where I_1 is the photocurrent of one of the photodetectors and σ_{I_1} is the root-mean-square current due to shot noise.

The current shot noise can be calculated from

$$\sigma_I = \sqrt{\frac{2eI}{\Delta t N}}, \qquad \text{(C5)}$$

where e is the charge of an electron, I is the photocurrent, Δt is the measurement time, and N is the number

of averages. For a detected power per mode of around 0.3 μW, we calculate that the photocurrent produced by the photodetector (responsivity of 0.1 A/W at 795 nm) is $I_1 = 0.03$ μA. The total measurement time for a single comb mode shown in Fig. 4 is 400 averages of a 9.78-μs measurement, yielding an estimated rotational shot noise of 20 μrad. In addition, we use a calibration spectrum that increases this noise by a factor of $\sqrt{2}$, giving rise to a root-mean-square rotation of $\sigma_\phi = 26.1$ μrad on each comb mode.

—————

[1] G. Bison, N. Castagna, A. Hofer, P. Knowles, J.-L. Schenker, M. Kasprzak, H. Saudan, and A. Weis, A room temperature 19-channel magnetic field mapping device for cardiac signals, Appl. Phys. Lett. **95**, 173701 (2009).

[2] Svenja Knappe, Tilmann H. Sander, Olaf Kosch, Frank Wiekhorst, John Kitching, and Lutz Trahms, Cross-validation of microfabricated atomic magnetometers with superconducting quantum interference devices for biomagnetic applications, Appl. Phys. Lett. **97**, 133703 (2010).

[3] Cort Johnson, Peter D. D. Schwindt, and Michael Weisend, Magnetoencephalography with a two-color pump-probe, fiber-coupled atomic magnetometer, Appl. Phys. Lett. **97**, 243703 (2010).

[4] Shoujun Xu, Valeriy V. Yashchuk, Marcus H. Donaldson, Simon M. Rochester, Dmitry Budker, and Alexander Pines, Magnetic resonance imaging with an optical atomic magnetometer, Proc. Natl. Acad. Sci. U.S.A. **103**, 12668 (2006).

[5] G. Bison, R. Wynands, and A. Weis, A laser-pumped magnetometer for the mapping of human cardiomagnetic fields, Appl. Phys. B **76**, 325 (2003).

[6] H. B. Dang, A. C. Maloof, and M. V. Romalis, Ultrahigh sensitivity magnetic field and magnetization measurements with an atomic magnetometer, Appl. Phys. Lett. **97**, 151110 (2010).

[7] M. Smiciklas, J. M. Brown, L. W. Cheuk, S. J. Smullin, and M. V. Romalis, New Test of Local Lorentz Invariance using a 21Ne-Rb-K Comagnetometer, Phys. Rev. Lett. **107**, 171604 (2011).

[8] C. J. Berglund, L. R. Hunter, D. Krause Jr, E. O. Prigge, M. S. Ronfeldt, and S. K. Lamoreaux, New Limits on Local Lorentz Invariance from Hg and Cs Magnetometers, Phys. Rev. Lett. **75**, 1879 (1995).

[9] Szymon Pustelny *et al.*, The global network of optical magnetometers for exotic physics (GNOME): A novel scheme to search for physics beyond the standard model, Ann. Phys. **525**, 659 (2013).

[10] Derek F. Jackson Kimball, Ian Lacey, Julian Valdez, Jerlyn Swiatlowski, Cesar Rios, Rodrigo Peregrina-Ramirez, Caitlin Montcrieffe, Jackie Kremer, Jordan Dudley, and C. Sanchez, A dual-isotope rubidium comagnetometer to search for anomalous long-range spin-mass (spin-gravity) couplings of the proton, Ann. Phys. **525**, 514 (2013).

[11] Marcis Auzinsh and Ruvin Ferber, *Optical Polarization of Molecules* (Cambridge University Press, New York, 2005), Vol. 4.

[12] Dmitry Budker, Valeriy Yashchuk, and Max Zolotorev, Nonlinear Magneto-optic Effects with Ultranarrow Widths, Phys. Rev. Lett. **81**, 5788 (1998).

[13] Dmitry Budker and Derek F. Jackson Kimball, *Optical Magnetometry* (Cambridge University Press, New York, 2013).

[14] David A. Long, Adam J. Fleisher, David F. Plusquellic, and Joseph T. Hodges, Multiplexed sub-Doppler spectroscopy with an optical frequency comb, Phys. Rev. A **94**, 061801 (2016).

[15] Vincente Duran, Santiago Tainta, and Victor Torres-Company, Ultrafast electrooptic dual-comb interferometry, Opt. Express **23**, 30557 (2015).

[16] Nicolas Bourbeau Hébert, Vincent Michaud-Belleau, James D. Anstie, Jean-Daniel Deschênes, Andre N. Luiten, and Jérôme Genest, Self-heterodyne interference spectroscopy using a comb generated by pseudo-random modulation, Opt. Express **23**, 27806 (2015).

[17] Nicolas Bourbeau Hébert, Vincent Michaud-Belleau, Christopher Perrella, Gar-Wing Truong, James D. Anstie, Thomas M. Stace, Jérôme Genest, and Andre N. Luiten, Real-time Dynamic Atomic Spectroscopy using Electro-optic Frequency Combs, Phys. Rev. Applied **6**, 044012 (2016).

[18] David A. Long, Adam J. Fleisher, Kevin O. Douglass, Stephen E. Maxwell, K. Bielska, Joseph T. Hodges, and David F. Plusquellic, Multiheterodyne spectroscopy with optical frequency combs generated from a continuous-wave laser, Opt. Lett. **39**, 2688 (2014).

[19] Yuan Bao, Xingwen Yi, Zhaohui Li, Qingming Chen, Jianping Li, Xudong Fan, and Xuming Zhang, A digitally generated ultrafine optical frequency comb for spectral measurements with 0.01-pm resolution and 0.7-μs response time, Light Sci. Appl. **4**, e300 (2015).

[20] Guy Millot, Stéphane Pitois, Ming Yan, Tatevik Hovhannisyan, Abdelkrim Bendahmane, Theodor W. Hänsch, and Nathalie Picqué, Frequency-agile dual-comb spectroscopy, Nat. Photonics **10**, 27 (2016).

[21] Pedro Martín-Mateos, Borja Jerez, and Pablo Acedo, Dual electro-optic optical frequency combs for multiheterodyne molecular dispersion spectroscopy, Opt. Express **23**, 21149 (2015).

[22] B. Patton, O. Neitzke, S. Rochester, E. Bahr, S. Guttikonda, D. F. Kimball, B. Coste, I. Novikova, and D. Budker, in *APS Division of Atomic, Molecular and Optical Physics Meeting Abstracts* (2012).

[23] Daniel Adam Steck, Rubidium 87 D Line Data, http://steck.us/alkalidata (revision 2.1.5, 23 October 2017).

[24] G.-W. Truong, J. D. Anstie, E. F. May, T. M. Stace, and A. N. Luiten, Accurate lineshape spectroscopy and the Boltzmann constant, Nat. Commun. **6**, 8345 (2015).

[25] Wojciech Gawlik and Szymon Pustelny, in *High Sensitivity Magnetometers*, edited by Asaf Grosz, Michael J.

Haji-Sheikh, and Subhas C. Mukhopadhyay (Springer, Switzerland, 2017), p. 425.

[26] Gordon W. F. Drake, *Springer Handbook of Atomic, Molecular, and Optical Physics* (Springer Science & Business Media, New York, 2006).

[27] Evgeniy B. Alexandrov, Marcis Auzinsh, Dmitry Budker, Derek F. Kimball, Simon M. Rochester, and Valeriy V. Yashchuk, Dynamic effects in nonlinear magneto-optics of atoms and molecules: Review, J. Opt. Soc. Am. B **22**, 7 (2005).

[28] V. Acosta, M. P. Ledbetter, S. M. Rochester, D. Budker, D. F. Jackson Kimball, D. C. Hovde, W. Gawlik, S. Pustelny, J. Zachorowski, and V. V. Yashchuk, Nonlinear magneto-optical rotation with frequency-modulated light in the geophysical field range, Phys. Rev. A **73**, 053404 (2006).

[29] W. Gawlik and S. Pustelny, *New Trends in Quantum Coherence and Nonlinear Optics* (Nova Publishers, New York, 2009), p. 47.

[30] Miriam T. Graf, Derek F. Kimball, Simon M. Rochester, K. Kerner, C. Wong, D. Budker, E. B. Alexandrov, M. V. Balabas, and V. V. Yashchuk, Relaxation of atomic polarization in paraffin-coated cesium vapor cells, Phys. Rev. A **72**, 023401 (2005).

[31] W. Franzen, Spin relaxation of optically aligned rubidium vapor, Phys. Rev. **115**, 850 (1959).

[32] S. J. Seltzer, D. M. Rampulla, S. Rivillon-Amy, Y. J. Chabal, S. L. Bernasek, and M. V. Romalis, Testing the effect of surface coatings on alkali atom polarization lifetimes, J. Appl. Phys. **104**, 103116 (2008).

[33] S. I. Kanorsky, A. Weis, and J. Skalla, A wall-collision-induced Ramsey resonance, Appl. Phys. B **60**, S165 (1995).

[34] C. H. Holbrow, A. P. Ghosh, D. Heinzen, X. Zhu, W. W. Quivers Jr, G. Shimkaveg, P. G. Pappas, J. E. Thomas, and M. S. Feld, Complete doppler coverage in laser optical pumping by wall-induced velocity-changing collisions, Phys. Rev. A **34**, 2477 (1986).

[35] D. Budker, D. F. Kimball, S. M. Rochester, and V. V. Yashchuk, Nonlinear electro- and magneto-optical effects related to Bennett structures, Phys. Rev. A **65**, 033401 (2002).

[36] J. Reichert, R. Holzwarth, T. Udem, and T. W. Hänsch, Measuring the frequency of light with mode-locked lasers, Opt. Commun. **172**, 59 (1999).

[37] Marcis Auzinsh, Dmitry Budker, and Simon M. Rochester, *Optically Polarized Atoms* (Oxford University Press, New York, 2010).

[38] Hugo Bergeron, Jean-Daniel Deschênes, and Jérôme Genest, Improving the signal-to-noise ratio of the beat note between a frequency comb and a tunable laser using a dynamically tracking optical filter, Opt. Lett. **41**, 4253 (2016).

Ultrastable Optical Magnetometry

This chapter is based on the published article:

Nathanial Wilson, Philip Light, André Luiten, and Christopher Perrella, *Ultrastable Optical Magnetometry*, Phys. Rev. Appl. **11**, 044034 (2019).

10.1   Overview and Motivation

Optical magnetometers have demonstrated exceptional performance, with absolute sensitivities which rival, or even surpass that of superconducting quantum interference devices. One of the major limitations of these devices, however, is their relatively poor performance over long timescales. This reduction in performance is typically the result of contributions from a myriad of technical noise sources, rather than due to fundamental limitations of the device itself. However, regardless of the origin, reduced performance over long timescales prohibits the applicability of these devices in situations where small, slowly varying magnetic fields must be measured. Examples of such applications include magnetic anomaly detection, earthquake monitoring, and the search for plant biomagnetism.

This chapter describes experimental work in which an optical magnetometer — based on nonlinear magneto-optical rotation — was used to perform sensitive magnetometry over a 26-h period. This was achieved through optimisation of a number of experimental parameters, and active stabilisation of these parameters at their optimal values. Additionally, a multitude of noise sources — both fundamental and technical — were identified, and their relative contributions to the noise floor of the device were characterised and subsequently minimised where possible.

To the best of our knowledge, and at the time of publishing, the demonstrated performance of the device was the highest absolute sensitivity of a nonlinear magneto-optical rotation based magnetometer across 8 orders of magnitude in Fourier frequency, as well as the highest dynamic-range magnetic-field measurement performed by an optical magnetometer of any kind. The observed low-frequency performance was vastly superior to previously demonstrated work, and paves the way for the use of optical magnetometers in applications which demand high sensitivity over long timescales.

10.2   Statement of Contribution

10.2.1   Conceptualisation

The idea of performing ultrastable optical magnetometry over long timescales, using a magnetometer based on nonlinear magneto-optical rotation , is attributed to André Luiten.

10.2.2   Realisation

The experimental apparatus used to produce the results presented in this manuscript was constructed by Nathanial Wilson, under the supervision of Christopher Perrella and Philip Light. The power supply used to generate the stable magnetic field in this work was built by Philip Light and Nathanial Wilson, using a design created by Philip Light. Optimisation and characterisation of the device was performed by Nathanial Wilson, with some assistance from André Luiten and Philip Light. The experimental data was acquired by Nathanial Wilson, with some assistance from Christopher Perrella. Analysis of the experimental data was performed by Nathanial Wilson. Theoretical modelling used in support of the experimental results was performed by Nathanial Wilson, with some assistance from Philip Light.

10.2.3   Documentation

This manuscript was written primarily by Nathanial Wilson and André Luiten, with editing by Christopher Perrella and Philip Light. Additional editorial support was provided by Russell Anderson.

Candidate overall percentage contribution: 70%.

_____

Nathanial Wilson

10.2.4 Co-author Contributions

By signing below, each co-author certifies that:

- The candidate's stated contribution to the publication is accurate (as stated above)

- Permission is granted for the candidate to include the publication in the thesis

Co-author certifications are in order of published authorship unless otherwise specified.

_____

Philip Light

_____

André Luiten

_____

Christopher Perrella

Ultrastable Optical Magnetometry

Nathanial Wilson,[*] Philip Light, André Luiten, and Christopher Perrella

Institute for Photonics and Advanced Sensing (IPAS), and School of Physical Sciences,The University of Adelaide,
South Australia 5005, Australia

We report on an ultrastable optical magnetometer based on nonlinear magneto-optical rotation in ^{87}Rb vapor. The atomic vapor is both optically pumped and probed on the $F = 2 \rightarrow F' = 1$ hyperfine transition of the D_1 manifold. A measurement over 26 h quantifies the magnetometer's performance across 8 orders of magnitude in the Fourier-frequency domain, allowing us to measure the magnetic response into the microhertz domain. We demonstrate a room-temperature sensitivity floor of 15 $\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ at a magnetic-field strength of 2.5 μT, which corresponds to a record 9 $\mathrm{ppb}/\sqrt{\mathrm{Hz}}$ fractional sensitivity. The magnetometer's performance is photon shot-noise limited from 40 Hz to 10 kHz, while below 40 Hz it slowly degrades as approximately $f^{-1/4}$. At 1 mHz the performance is still better than 1 $\mathrm{pT_{rms}}/\sqrt{\mathrm{Hz}}$. We show that this worsening performance is not a characteristic of the sensor, but instead associated with minute fluctuations of the magnetic field at the sensor from various identified sources.

I. INTRODUCTION

Optical magnetometers are presently among the most sensitive devices for measuring magnetic fields; their sensitivities can rival [1] or even surpass [2–4] that of superconducting quantum interference devices (SQUIDs). With their intrinsic accuracy, room-temperature operability, and portability, these devices are finding applications in numerous fields ranging from medical diagnostics and imaging [5–15], to geomagnetism [4], and fundamental physics research [16–19].

Optical magnetometers come in a number of flavors, each with characteristic advantages and disadvantages. Perhaps the two most frequently encountered flavors make use of the nonlinear magneto-optical rotation (NMOR) [20] or spin-exchange relaxation-free (SERF) [21] techniques. SERF magnetometers have demonstrated magnetic sensitivities below 1 $\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$ [2–4]; however, they can only operate in near-zero magnetic field ($\lesssim$ 1 nT) and have a relatively low bandwidth. These aspects limit their practicality in many real-world applications. In contrast, NMOR magnetometers, utilizing frequency [22,23], amplitude [1,24], or polarization [25] modulated-pumping techniques, can operate in magnetic fields higher than that of Earth [22,26]. The absolute sensitivity of NMOR magnetometers is generally 1 to 2 orders of magnitude worse than that of their SERF counterparts.

The sensitivity of an optical magnetometer is subject to two fundamental limits [27–29]: spin-projection noise

and photon shot noise. For NMOR magnetometers, photon shot noise always dominates over spin-projection noise, and sets the performance limit of a well-designed magnetometer at relatively high frequencies ($\gtrsim$ 10 Hz). However, at low frequencies a number of technical-noise sources typically limit the magnetometer's performance, or the ability to measure that performance, such that one cannot achieve this fundamental limit. These technical-noise sources include fluctuations in optical power (both pump and probe), optical-polarization drift, ambient-field fluctuations, electronic noise in the detection, magnetic-shield noise etc. This can prevent the application of a magnetometer in circumstances that demand high performance at low frequencies. Examples would include magnetic anomaly detection [30], earthquake monitoring [31–33], and the search for plant biomagnetism [34].

In this work, we demonstrate an optical magnetometer, which operates at room temperature and exhibits high sensitivity over a broad range of Fourier frequencies. We generate a magnetic field of 2.5 μT within a magnetic shield and subsequently measure that field with a shot-noise-limited sensitivity of 15 $\mathrm{fT_{rms}}/\sqrt{\mathrm{Hz}}$. This corresponds to a record fractional sensitivity of 9 $\mathrm{ppb}/\sqrt{\mathrm{Hz}}$. Over a 26-h measurement, the measured magnetic field drifts by only 100 ppm, and this instability arises in the applied field rather than from the sensor itself. The result presented here is, to the best of our knowledge, the most sensitive magnetic-field measurement performed by an NMOR magnetometer over an exceptionally broad range of Fourier frequencies from 100 μHz to 10 kHz [1,35], as well as the highest dynamic-range measurement performed by an optical magnetometer of any kind. The fractional

*[*]nathanial.wilson@adelaide.edu.au

sensitivity reported here is around 2 orders of magnitude superior to the best reported measurements with a SERF magnetometer [2–4].

When compared to state-of-the-art superconducting quantum-interference devices, the reported performance is only a factor of 40 inferior to the best-ever reported result at high frequencies [36], and less than 2 orders of magnitude worse at ultralow frequencies around 100 μHz [37]. However, the optical approach reported here benefits from avoiding cryogenics and giving an intrinsically accurate measurement. This demonstration of low-frequency performance of an optical magnetometer opens the potential for many new applications, which are currently dominated by the use of SQUIDs.

II. MAGNETOMETER CONCEPT

The magnetometer utilizes amplitude-modulated nonlinear magneto-optical rotation (AMOR) [20,24], in which resonant, linearly polarized light produces spin coherence in an alkali vapor through periodic optical pumping. Pumping with linearly polarized light generates atomic alignment (a rank-two polarization moment) [38]. Using an atomic description in which the quantization axis is aligned with the applied magnetic field, this corresponds to a symmetric distribution of population between Zeeman sublevels with respect to the $|m_F = 0\rangle$ state, and coherences between ground-state Zeeman sublevels of $\Delta m_F = 2$ [22].

When exposed to a weak longitudinal magnetic field, the energy levels of the Zeeman states are shifted by the Larmor frequency, $\Omega_L = \mu_B g_F B / \hbar$, where μ_B is the Bohr magneton, g_F is the hyperfine Landé g-factor, B is the magnetic-field strength and $\hbar$ is the reduced Planck constant. This energy splitting induces temporal evolution of the $\Delta m_F = 2$ coherences at a frequency of $2\Omega_L$ [39], yielding a modulation of the complex refractive indices for the two polarization eigenmodes [28] (circular birefringence). This consequently modulates the polarization plane of a linearly polarized probe beam at the same frequency [27,40]. One finds a strong resonant pumping condition, which maximizes the atomic alignment within the vapor, if the pump beam is modulated at a frequency $\Omega_m = 2\Omega_L$. A measurement of the frequency of this resonant condition allows accurate estimation of the local magnetic-field strength.

III. EXPERIMENTAL SETUP

The experimental setup is presented in Fig. 1. Isotopically pure ^{87}Rb is contained in a cylindrical vapor cell with a 40-mm diameter and 40-mm length. The walls of the cell are coated with paraffin in order to preserve coherences between ground-state Zeeman sublevels, with a measured transverse spin-relaxation time of $T_2 = 36.5$ ms. The cell remains at room temperature ($23.10 \pm 0.04\,^\circ$C), and is
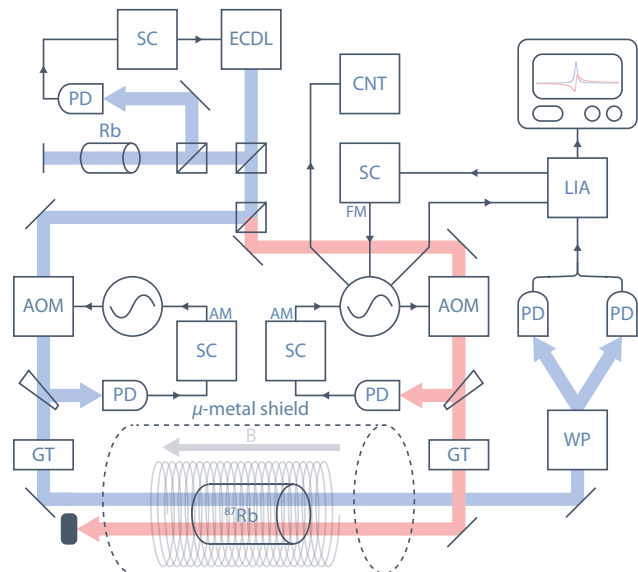
FIG. 1. Simplified experimental setup, showing the external cavity diode laser (ECDL), acousto-optic modulator (AOM), Glan-Thompson prism (GT), Wollaston prism (WP), photodetector (PD), lock-in amplifier (LIA), servo controller (SC), and frequency counter (CNT). The probe beam is shown in blue, while the pump beam is shown in red. All relevant electronic components are referenced to the 10-MHz output of a cesium atomic clock. The AM and FM labels denote amplitude and frequency modulation, respectively.

housed within a six-layer cylindrical μ-metal magnetic shield that has a calculated shielding factor of 2×10^5 [41]. A constant-bias magnetic field of 2.5 μT is generated along the longitudinal axis of the cell using a solenoid installed within the innermost shield. This value is chosen as it is the strongest magnetic field that could be applied without deteriorating the magnetometer's performance via gradient broadening [42] due to the inhomogeneity of the field generated by the solenoid.

The solenoid is driven by an ultralow-noise power supply, which produces a fixed voltage with a square-root Allan variance [43] of 3×10^{-8} at 1 s, slowly degrading to 10^{-6} for integration times of 10^4 s. In order to obtain fractional current fluctuations in the part-per-billion range at high Fourier frequencies, the output voltage of the power supply is filtered prior to driving the solenoid using a first-order Butterworth filter with a cut-off frequency of 0.3 mHz. The solenoid current is set by series resistors with an ultralow temperature coefficient of 200 ppb/$^\circ$C. The current noise after the filter is below the level that can be measured by any conventional instrumentation and is seen not to limit the magnetometer performance above 10 mHz. Below this frequency, we observe a temperature dependence in the Butterworth filter that causes unwanted current fluctuations that can be observed by the magnetometer.

The atomic vapor is optically pumped and probed on the $5^2S_{1/2} \rightarrow 5^2P_{1/2}$ (D_1) transition of ^{87}Rb by an external cavity diode laser tuned near 795 nm. The laser is frequency locked to the $F = 2 \rightarrow F' = 1$ hyperfine transition using saturated absorption spectroscopy in a reference cell. Both pump and probe beams pass through acousto-optic modulators (AOMs) that downshift their frequencies by 80 MHz, and are also used to control the optical power incident on the vapor cell. The AOM for the pump beam also applies a square-wave amplitude modulation at $\Omega_m/2\pi \approx 34.7$ kHz, with a modulation depth of 100% and a duty cycle of 20%. This amplitude modulation resonantly drives the polarization of the medium.

Prior to the vapor cell, a small fraction of the pump and probe beams is detected by two separate photodetectors. These signals are used to keep the power of each beam stabilized by actively controlling the rf driving power of the AOMs. As the pump beam is amplitude modulated, the photodetector signal is demodulated at $\Omega_m/2\pi$ using a lock-in amplifier referenced to the pump modulation frequency prior to being fed into the power servo. Following stabilization, the square-root Allan variance of the pump and probe powers is relatively flat (i.e., independent of integration time) and below 10^{-4} for integration times ranging between 0.2 and 400 s.

We note that optical-power fluctuations can result in small changes in the observed Larmor frequency through optical saturation and other nonlinear optical effects. Without sufficient power stabilization, this has the effect of giving rise to a fictitious magnetic field. We measure the sensitivity of the magnetometer to these effects by intentionally inducing a large power modulation, and measuring the magnetometer response. Using this sensitivity, we calculate that the stabilized pump and probe beams induce fictitious magnetic fluctuations with a flicker-noise characteristic, the magnitude of which is 1 fT$_{\mathrm{rms}}/\sqrt{\mathrm{Hz}}$ and 100 fT$_{\mathrm{rms}}/\sqrt{\mathrm{Hz}}$ at 1 mHz for the pump and probe beams, respectively. These unwanted fluctuations are not a dominant limitation of the magnetometer presented here.

Immediately prior to entering the cell, the pump beam is vertically polarized by a Glan-Thompson (GT) prism and set to propagate parallel to the magnetic field. The probe beam is linearly polarized at 45° from the vertical by a GT prism, and propagates antiparallel to the magnetic field. Both pump and probe beams have a $1/e^2$ diameter of 1.5 mm, and are horizontally displaced by approximately 10 mm relative to each other.

The stability of the polarization of the probe beam is an important quantity, as fluctuations in the polarization of the probe can give rise to asymmetry in the resonance signals, which is indistinguishable from real magnetic noise. The stability of the polarization of the probe beam, as measured after the cell, has a flicker-noise characteristic with a magnitude of 700 μrad$_{\mathrm{rms}}/\sqrt{\mathrm{Hz}}$ at 1 mHz. By intentionally inducing a large polarization change in the probe beam and

subsequently measuring the magnetometer response, we conclude that the effective magnetic noise associated with probe polarization drift is about 100 fT$_{\mathrm{rms}}/\sqrt{\mathrm{Hz}}$ at 1 mHz. As with the optical-power fluctuations described previously, the polarization instability of the probe beam is not a dominant limitation for the magnetometer presented here.

After traversing the vapor cell, the probe beam passes through a Wollaston prism, which separates the beam into its orthogonal polarization components. These orthogonal components are measured on separate photodetectors, forming a balanced polarimeter. The optical power on each photodetector can be converted into a polarization-rotation angle via [27,40]

$$\phi = \frac{1}{2} \arcsin\left(\frac{P_1 - P_2}{P_1 + P_2}\right), \tag{1}$$

where P_1 and P_2 are the optical powers on the two photodetectors and ϕ is the polarization-rotation angle. The optical-rotation signal is measured in the frequency domain by demodulating the balanced polarimeter output at $\Omega_m/2\pi$.

IV. OPTIMIZATION

The magnetic-detection limit of the magnetometer is determined by its signal-to-noise ratio; thus, for ultimate performance, one needs to understand the dependence of both signal and noise on key operational settings.

The two-channel demodulated polarimeter signal, containing the AMOR resonance (cf. Fig. 2), has in-phase, $\phi_P(v)$, and quadrature, $\phi_Q(v)$, components. These components are well approximated as the real and imaginary components of a complex Lorentzian lineshape defined by

$$\mathcal{L}(v) = \mathcal{A}\left[\frac{\left(\frac{\Gamma}{2}\right)^2}{(v-v_0)^2 + \left(\frac{\Gamma}{2}\right)^2} + i\frac{\frac{\Gamma}{2}(v-v_0)}{(v-v_0)^2 + \left(\frac{\Gamma}{2}\right)^2}\right], \tag{2}$$

where $\mathcal{A}$ is the amplitude, Γ is the full width at half maximum, and v_0 is the resonant frequency. The maximum slope $\mathcal{S}$ of the quadrature component, ϕ_Q, is given by

$$\mathcal{S} = \lim_{v \to v_0} \frac{d\phi_Q}{dv} = \frac{2\mathcal{A}}{\Gamma}. \tag{3}$$

Equation (3) is effectively the sensitivity of the magnetometer. This sensitivity is seen to vary as certain operational parameters are changed. Here, we explore how this sensitivity varies with the time-averaged pump and probe power. To understand the origin of this sensitivity variation, we measure the resonant response in each operational setting (cf. Fig. 2) and fit to it an equation of the form of Eq. (2). This allows us to obtain the optical-power dependence of the resonant amplitude $\mathcal{A}$, resonant width

FIG. 2. In-phase and quadrature components of the AMOR resonance measured using time-averaged pump and probe powers of $P_{pump} = 15\,\mu W$ and $P_{probe} = 5\,\mu W$ respectively, with a complex Lorentzian fit using Eq. (2). The quadrature component has an amplitude of $\mathcal{A} = 7.57$ mrad, a width of $\Gamma = 13.4$ Hz $= 0.96$ nT, and a calculated slope of $\mathcal{S} = 1.13$ mrad/Hz using Eq. (3). The resonance is centered around a modulation frequency of $\Omega_m/2\pi = 34.672$ kHz, corresponding to a magnetic-field strength of $B = 2.4781\,\mu T$.

Γ, and consequently $\mathcal{S}$ through Eq. (3). These outcomes are shown across the first three panels of Fig. 3.

Figure 3(a) demonstrates that for low pump power there is a monotonic increase in resonance amplitude, $\mathcal{A}$, as the pump power is increased. This is reasonable as in this limit the pump power determines the maximum achievable spin polarization. However, for the highest pump powers shown, the atomic response is saturated and we observe a plateau and subsequent decrease of the resonance amplitude [44]. On the other hand, as the probe power increases we observe a monotonic decrease in the resonance amplitude, which is associated with the destruction of ground-state coherence via unwanted optical pumping by the probe [27]. An optimum amplitude is observed at $(P_{pump}, P_{probe}) = (30\,\mu W, 3\,\mu W)$. By contrast, the resonance width, Γ, presented in

Fig. 3(b), shows a broadening that scales approximately as $\Gamma \propto \sqrt{P_{pump}^2 + 3P_{probe}^2}$. Figure 3(c) gives the resulting slope of the resonance in magnetic units, $\mathcal{S}\gamma/\pi$, calculated via Eq. (3). The steepest slope (and hence maximum sensitivity) is found at $(P_{pump}, P_{probe}) = (7\,\mu W, 3\,\mu W)$.

We now turn our attention to the noise of the detection process, which determines the minimum detectable optical rotation per unit bandwidth, $\delta\phi$. Combining this with the slope, $\mathcal{S}$, we obtain the minimum detectable magnetic signal per unit bandwidth, δB:

$$\delta B = \frac{\pi\,\delta\phi}{\mathcal{S}\gamma}, \tag{4}$$

where γ is the gyromagnetic ratio of the ground state. For the $F = 2$ ground state of ^{87}Rb, the gyromagnetic ratio is $\gamma/2\pi = 6.9958$ GHz/T [45].

For all measurements presented here the photodetectors are shot-noise limited. In this case, a beam of light with time-averaged power, P, has a power spectral density as measured at the photodetector of $S(\nu) = 2h\nu P\eta(\nu)$, where $h\nu$ is the energy per photon and $\eta(\nu)$ is the quantum efficiency of the photodiode. Therefore, given approximately equal detected powers, P, on the two photodetectors in the polarimeter, the minimum detectable optical rotation per unit bandwidth in the case of shot noise can be calculated from Eq. (1) as $\delta\phi \approx \frac{1}{2}\sqrt{\frac{h\nu}{\eta(\nu)P}}$. In Fig. 3(d) we combine the signal and noise results to estimate the magnetic-noise floor of the instrument. This indicates an optimum operation point, $(P_{pump}, P_{probe}) = (15\,\mu W, 5\,\mu W)$, that is not identical to the optimal sensitivity (slope) point seen in Fig. 3(c). The atomic response around the magnetic resonance is shown under these optimal conditions in Fig. 2.

V. PERFORMANCE

The magnetometer performance is examined over a wide range of Fourier frequencies in Fig. 4. For fast fluctuations (> 0.5 Hz) we characterize this performance by
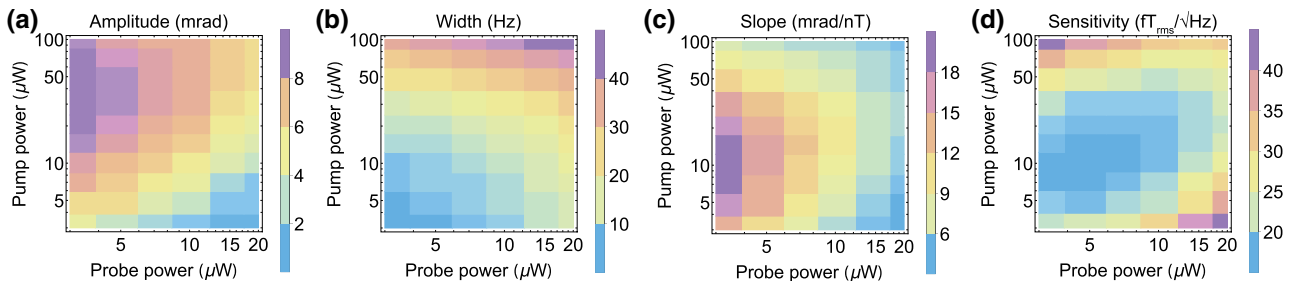
FIG. 3. Optical-power-optimization contour plots for the magnetometer, showing the resonance amplitude $\mathcal{A}$ (a), resonance width Γ (b), resonance slope in magnetic units $\mathcal{S}\gamma/\pi$ (c) and the shot-noise-limited sensitivity δB (d). The optimum operating point occurs for time-averaged pump and probe powers of $P_{pump} = 15\,\mu W$ and $P_{probe} = 5\,\mu W$, respectively. Using this combination of optical powers results in a shot-noise-limited sensitivity of 15 fT$_{rms}/\sqrt{Hz}$.

FIG. 4. Amplitude spectral density of the magnetometer noise floor (blue trace) calculated via Welch's method [46], using a 26-h measurement. The noise floor is shot-noise limited (red trace) above 40 Hz, but exhibits excess noise at low Fourier frequencies arising from unwanted fluctuations in the magnetic field at the sensor. In the intermediate frequency range (30 mHz to 40 Hz), Johnson current noise in the solenoid generates magnetic noise (solid purple trace, with a 5-Hz-filtered version shown in the dashed purple trace). Between 1 and 30 mHz, magnetic noise arising from the innermost shield dominates [47] (black trace), while at the lowest frequencies, temperature sensitivity of the solenoid current modulates the magnetic field (green trace). The measured temperature fluctuations of the current source are scaled by an experimentally-determined factor of 100 pT/°C.

directly measuring ϕ_Q with a spectrum analyzer when the pump modulation frequency is tuned to the resonance condition $\Omega_m = 2\Omega_L$. We subsequently convert this signal to an equivalent magnetic fluctuation using Eq. (4). Slow fluctuations (< 0.5 Hz) are measured by actively locking the pump modulation frequency to the magnetic resonance, and then using a frequency counter to measure the variations in pump modulation frequency. Both approaches are shown in Fig. 4 and are in excellent agreement where they overlap.

The key determinant of the read-out system noise is photodetection noise. The bespoke photodetectors used here implement a transimpedance amplifier with a gain of 10^6 Ω, and show an effective magnetic noise of about 2 $fT_{rms}/\sqrt{Hz}$ due to the combined dark-current noise of both photodetectors.

The photodetector output signals are then sent to a commercial lock-in amplifier with a measured input noise of 6 $nV_{rms}/\sqrt{Hz}$. Using the experimental parameters reported here, the input noise of the lock-in amplifier gives rise to an effective magnetic noise of 0.4 $fT_{rms}/\sqrt{Hz}$. We note that changes in the phase of the lock-in amplifier can give rise to an asymmetry in the resonance signals, which is indistinguishable from real magnetic fluctuations. The specified phase noise of the lock-in amplifier is about 60 $nrad_{rms}/\sqrt{Hz}$, which corresponds to an effective magnetic noise of about 90 $aT_{rms}/\sqrt{Hz}$.

When measuring high-frequency magnetic-field noise using a spectrum analyzer, the input noise of the spectrum analyzer will give rise to an effective magnetic noise. For

the measurements performed here, this effective noise is consistently below 1 $fT_{rms}/\sqrt{Hz}$ at 1 Hz, dropping as low as 150 $aT_{rms}/\sqrt{Hz}$ at high Fourier frequencies (i.e., above 100 Hz).

When measuring low-frequency magnetic-field noise by actively locking the pump modulation frequency to the magnetic resonance and subsequently observing variations in the modulation frequency, the effective magnetic noise introduced by the frequency counter must also be considered. The noise of the frequency counter is dependent upon both the signal-to-noise ratio of the signal, which it is counting, as well as the gate time. Under the experimental conditions presented here, the measured effective magnetic noise due to the frequency counter is about 2 $fT_{rms}/\sqrt{Hz}$.

The bandwidth of the magnetometer is determined by introducing a small sinusoidal modulation to the applied magnetic field using a high-bandwidth (> 100 kHz) coil wrapped around the center of the solenoid. The magnetometer's response to this modulation is observed by performing a free-induction decay measurement [48], in which the spin-polarized atomic vapor is allowed to relax in the presence of the modulated magnetic field. An example free-induction decay time trace and its corresponding Fourier transform is shown in Fig. 5, for an applied ac

FIG. 5. (a) Free-induction decay measurement performed in the temporal domain, with a carrier frequency of $\Omega_L/\pi = 29.567$ kHz and a small sinusoidal frequency modulation at 500 Hz. Dashed black lines highlight the damped exponential envelope, which has a time constant of 14.5 ms. (b) Fourier transform of the free-induction decay data presented in (a), showing a peak at the carrier frequency of 29.567 kHz (indicated by the dashed vertical line), and both first- and second-order sidebands arising from the sinusoidal modulation at 500 Hz.

magnetic-field strength of 10 nT at a frequency of 500 Hz. The measured transfer function shows that the bandwidth of the magnetometer is greater than 10 kHz.

The photon shot noise, calculated from the measured current from the photodetectors, corresponds to a magnetic noise of 15 $fT_{rms}/\sqrt{Hz}$, and is shown in red in Fig. 4. We see that the measured magnetometer performance is limited by this noise source over a frequency range of 40 Hz to 10 kHz. This corresponds to a dynamic range of the instrument of over $10^8/\sqrt{Hz}$, which, to the best of our knowledge, is the highest ever reported for an optical magnetometer of any kind.

Evident in the high-frequency region of Fig. 4 is a peak in the noise at 50 Hz and harmonics thereof. The amplitude of this noise component is observed to be largely independent of the number of layers of shielding. Since the high-permeability shields are appropriate for the attenuation of low-frequency magnetic fields via flux shunting, the noise must be originating from within the shields. Given that the measured transfer function of the Butterworth filter is around -100 dB at 50 Hz, the electronics, which drive the solenoid, cannot be the source. We therefore believe that these peaks originate from electromagnetic coupling of ambient magnetic fields through the conductors that connect the solenoid to the driving electronics.

In the intermediate region between 30 mHz and 40 Hz, the performance is limited by Johnson current noise arising in the finite impedance of the solenoid. A SPICE model including the voltage reference, current-setting resistors, and the solenoid (both inductance and resistance) predicts a current noise of 53 $pA_{rms}/\sqrt{Hz}$, corresponding to 92 $fT_{rms}/\sqrt{Hz}$ of magnetic noise (cf. solid purple trace in Fig. 4). The SPICE model also predicts that this Johnson noise should have a low-pass-filter characteristic with a cut-off frequency of 170 Hz. We see excellent agreement between the observed magnetic noise and the model prediction at low frequency; however, the experiment suggests a cut-off frequency close to 5 Hz. The dashed purple trace in Fig. 4 shows the Johnson noise filtered using a 5-Hz low-pass filter. The origin of this difference is unknown; however, we speculate that the distributed nature of the solenoid noise may lead to some additional filtering over that predicted by the lumped-element SPICE model.

For frequencies below 30 mHz, the performance is limited by a combination of two factors: thermal fluctuations of the magnetization of the innermost layer of magnetic shielding [47] (shown by the black trace in Fig. 4), and a temperature sensitivity of the current source that drives the solenoid. We noted a strong correlation between the observed magnetic field and the temperature of both the Butterworth filter and voltage reference. The measured temperature fluctuations are shown in the green trace in Fig. 4, after conversion to magnetic-field units using an experimentally-determined scaling factor of 100 pT/°C

($\equiv$ 57 nA/°C), that was measured independently by deliberately heating the Butterworth filter.

We note that in the region of elevated noise (i.e., below 40 Hz), the observed fluctuations arise not from the sensor itself, but rather from low-level residual fluctuations in the magnetic field generated inside the magnetic shields.

VI. DISCUSSION

A. Measurement limitations

As discussed in Sec. V of the manuscript, the sensitivity across the majority of the Fourier-frequency range presented here is limited by the stability of the measured magnetic field, rather than by the magnetometer itself. All of the present limitations have been identified, and there are a number of avenues remaining to be explored, which may yield increased performance in these areas.

At high Fourier frequencies (i.e., $\gtrsim$ 40 Hz), the shield noise is only a factor of 2 below the shot-noise limit and hence an improvement in sensitivity of the magnetometer by only a small factor would result in the shield noise limiting the measurement. In this frequency region, the shield noise is associated with Johnson currents arising from the relatively low resistivity of μ-metal. The magnitude of this noise contribution can be reduced through the use of ferrite as the inner layer of shielding, as ferrite shields have been demonstrated to generate up to 25 times less Johnson magnetic noise than that of μ-metal shields of similar dimensions, owing to an electrical resistivity that is as much as 6 orders of magnitude larger than μ-metal [49]. Furthermore, ferrite shields also generate less than half the thermal magnetization noise than their μ-metal counterparts [49], which would reduce the noise that is presently limiting the measurement in the 1 to 30 mHz range.

In the intermediate range between 30 mHz and 40 Hz, the measurement is limited by Johnson current noise arising from the solenoid. SPICE modeling indicates that this noise could be reduced by using a higher-impedance conductor, which would improve the performance over this frequency range.

In the ultralow-frequency range below 1 mHz, the measurement is limited by magnetic noise associated with temperature drift of the coil-driving electronics, as determined by a strong correlation between the observed magnetic field and the temperature of the Butterworth filter and voltage reference. The noise in this region can be effectively suppressed via active temperature stabilization of these components.

B. Magnetometer limitations

The ultimate performance limit of the magnetometer is determined by various noise sources, which are intrinsic to

the device. Many of these sources have been identified and have immediate pathways for improvement.

At high Fourier frequencies (i.e., $\gtrsim 40$ Hz), the magnetometer is limited by photon shot noise. There are two ways in which to improve the performance in this frequency region: by increasing the sensitivity of the magnetometer to changes in magnetic-field strength, and by reducing the effect of shot noise.

The optical rotation experienced by the probe beam may be improved by increasing the optical depth of the vapor via heating the cell. Due to the tradeoff between increasing vapor density, spin-exchange collisions [50], and shot noise, the optimal optical depth is around unity [24]. This optical depth could be achieved with a temperature increase of about $30\,^{\circ}$C [45], which would yield approximately an order of magnitude increased magnetic sensitivity.

The shot-noise limit could be improved through the use of polarization-squeezed light, which has previously been shown to improve the sensitivity of NMOR magnetometers by 2–3 dB [51,52].

At low frequencies, i.e., below 1 Hz, the two main limitations intrinsic to the magnetometer are the optical-power stability and polarization stability of the probe beam.

The fractional optical-power stability of the probe beam is presently about 100 ppm, which corresponds to an effective magnetic noise of 100 fT$_{\mathrm{rms}}/\sqrt{\mathrm{Hz}}$ at 1 mHz. Fractional power stabilities on the order of 1 ppm have previously been demonstrated [53], which would reduce the effective magnetic noise due to optical power instability by 2 orders of magnitude if implemented here.

The polarization drift of the probe beam also introduces an effective noise of 100 fT$_{\mathrm{rms}}/\sqrt{\mathrm{Hz}}$ at 1 mHz, as discussed in Sec. V. This polarization drift is associated with mechanical or thermal alignment drift of the polarization optics with respect to the incident beam. This can, in principle, be improved through the use of mounting materials with low thermal-expansion coefficients [54].

C. Applications

The measurements presented here are performed in a 2.5-μT bias field due to gradient-broadening limitations described in Sec. III. In principle, if the field variations across the cell volume can be minimized, this demonstrated performance can be achieved at geophysical magnetic fields [22].

As it stands, the demonstrated performance of the magnetometer is sufficient for high-sensitivity, high-bandwidth applications such as magnetocardiography [13,55] and magnetoencephalography [55]. These two applications require the measurement of small magnetic fields of around 100 pT and 100 fT, respectively, within a bandwidth of up to 1 kHz [55]. Furthermore, due to the large dynamic range of this device, these applications may also be possible within Earth's field, thereby removing the need for magnetically shielded rooms. The demonstrated sensitivity on long timescales is also highly suited for low-frequency applications such as the search for plant biomagnetism, which requires resolution of fields less than 60 pT at frequencies around 1 mHz [34].

We envisage that future work will involve measuring small magnetic samples within the shielded environment. For larger samples, the techniques presented here can be readily transferred to an unshielded environment in a similar fashion to a previous demonstration using a SERF magnetometer [56]. However, depending on the magnitude and inhomogeneity of the ambient field, there may be some reduction in sensitivity arising from the nonlinear Zeeman effect [22] and gradient broadening [42], respectively.

VII. CONCLUSION

We develop a magnetometer based on amplitude-modulated nonlinear magneto-optical rotation in ^{87}Rb vapor. Through careful optimization and stabilization of experimental parameters, a 2.5-μT magnetic field is measured with a shot-noise-limited noise floor of 15 fT$_{\mathrm{rms}}/\sqrt{\mathrm{Hz}}$, corresponding to a fractional sensitivity of 9 ppb$/\sqrt{\mathrm{Hz}}$. We observe the performance of the magnetometer over 8 orders of magnitude in Fourier frequency, and note that the low-frequency performance is limited by the ability to produce a sufficiently quiet magnetic environment, rather than a limitation of the sensor itself.

[1] V. G. Lucivero, P. Anielski, W. Gawlik, and M. W. Mitchell, Shot-noise-limited magnetometer with sub-picotesla sensitivity at room temperature, Rev. Sci. Instrum. **85**, 113108 (2014).

[2] I. K. Kominis, T. W. Kornack, J. C. Allred, and M. V. Romalis, A subfemtotesla multichannel atomic magnetometer, Nature **422**, 596 (2003).

[3] D. Sheng, S. Li, N. Dural, and M. V. Romalis, Sub-femtotesla Scalar Atomic Magnetometry Using Multipass Cells, Phys. Rev. Lett. **110**, 160802 (2013).

[4] H. B. Dang, A. C. Maloof, and M. V. Romalis, Ultrahigh sensitivity magnetic field and magnetization measurements

with an atomic magnetometer, Appl. Phys. Lett. **97**, 151110 (2010).

[5] G. Bison, N. Castagna, A. Hofer, P. Knowles, J.-L. Schenker, M. Kasprzak, H. Saudan, and A. Weis, A room temperature 19-channel magnetic field mapping device for cardiac signals, Appl. Phys. Lett. **95**, 173701 (2009).

[6] G. Bison, R. Wynands, and A. Weis, A laser-pumped magnetometer for the mapping of human cardiomagnetic fields, Appl. Phys. B **76**, 325 (2003).

[7] S. Knappe, T. H. Sander, O. Kosch, F. Wiekhorst, J. Kitching, and L. Trahms, Cross-validation of microfabricated atomic magnetometers with superconducting quantum interference devices for biomagnetic applications, Appl. Phys. Lett. **97**, 133703 (2010).

[8] C. Johnson, P. D. D. Schwindt, and M. Weisend, Magnetoencephalography with a two-color pump-probe, fiber-coupled atomic magnetometer, Appl. Phys. Lett. **97**, 243703 (2010).

[9] S. Xu, V. V. Yashchuk, M. H. Donaldson, S. M. Rochester, D. Budker, and A. Pines, Magnetic resonance imaging with an optical atomic magnetometer, Proc. Natl. Acad. Sci. USA **103**, 12668 (2006).

[10] E. Boto, N. Holmes, J. Leggett, G. Roberts, V. Shah, S. S. Meyer, L. D. Muñoz, K. J. Mullinger, T. M. Tierney, S. Bestmann, G. R. Barnes, R. Bowtell, and M. J. Brookes, Moving magnetoencephalography towards real-world applications with a wearable system, Nature **555**, 657 (2018).

[11] K. Jensen, M. A. Skarsfeldt, H. Strkind, J. Arnbak, M. V. Balabas, S.-P. Olesen, B. H. Bentzen, and E. S. Polzik, Magnetocardiography on an isolated animal heart with a room-temperature optically pumped magnetometer, Sci. Rep. **8**, 16218 (2018).

[12] J. Belfi, G. Bevilacqua, V. Biancalana, S. Cartaleva, Y. Dancheva, and L. Moi, Cesium coherent population trapping magnetometer for cardiosignal detection in an unshielded environment, J. Opt. Soc. Am. B **24**, 2357 (2007).

[13] G. Bison, R. Wynands, and A. Weis, Dynamical mapping of the human cardiomagnetic field with a room-temperature, laser-optical sensor, Opt. Express **11**, 904 (2003).

[14] H. Xia, A. Ben-Amar Baranga, D. Hoffman, and M. V. Romalis, Magnetoencephalography with an atomic magnetometer, Appl. Phys. Lett. **89**, 211104 (2006).

[15] O. Alem, T. H. Sander, R. Mhaskar, J. LeBlanc, H. Eswaran, U. Steinhoff, Y. Okada, J. Kitching, L. Trahms, and S. Knappe, Fetal magnetocardiography measurements with an array of microfabricated optically pumped magnetometers, Phys. Med. Biol. **60**, 4797 (2015).

[16] M. Smiciklas, J. M. Brown, L. W. Cheuk, S. J. Smullin, and M. V. Romalis, New Test of Local Lorentz Invariance Using a ^{21}Ne-Rb-K Comagnetometer, Phys. Rev. Lett. **107**, 171604 (2011).

[17] C. J. Berglund, L. R. Hunter, D. Krause, Jr., E. O. Prigge, M. S. Ronfeldt, and S. K. Lamoreaux, New Limits on Local Lorentz Invariance from Hg and Cs Magnetometers, Phys. Rev. Lett. **75**, 1879 (1995).

[18] S. Pustelny, D. F. Jackson Kimball, C. Pankow, M. P. Ledbetter, P. Wlodarczyk, P. Wcislo, M. Pospelov, J. R. Smith, J. Read, W. Gawlik, and D. Budker, The global network of optical magnetometers for exotic physics (GNOME): A novel scheme to search for physics beyond the standard model, Ann. Phys. **525**, 659 (2013).

[19] D. F. Jackson Kimball, I. Lacey, J. Valdez, J. Swiatlowski, C. Rios, R. Peregrina-Ramirez, C. Montcrieffe, J. Kremer, J. Dudley, and C. Sanchez, A dual-isotope rubidium comagnetometer to search for anomalous long-range spin-mass (spin-gravity) couplings of the proton, Ann. Phys. **525**, 514 (2013).

[20] D. Budker, D. F. Jackson Kimball, S. M. Rochester, V. V. Yashchuk, and M. Zolotorev, Sensitive magnetometry based on nonlinear magneto-optical rotation, Phys. Rev. A **62**, 043403 (2000).

[21] J. C. Allred, R. N. Lyman, T. W. Kornack, and M. V. Romalis, High-Sensitivity Atomic Magnetometer Unaffected by Spin-Exchange Relaxation, Phys. Rev. Lett. **89**, 130801 (2002).

[22] V. Acosta, M. P. Ledbetter, S. M. Rochester, D. Budker, D. F. Jackson Kimball, D. C. Hovde, W. Gawlik, S. Pustelny, J. Zachorowski, and V. V. Yashchuk, Nonlinear magneto-optical rotation with frequency-modulated light in the geophysical field range, Phys. Rev. A **73**, 053404 (2006).

[23] D. Budker, D. F. Jackson Kimball, V. V. Yashchuk, and M. Zolotorev, Nonlinear magneto-optical rotation with frequency-modulated light, Phys. Rev. A **65**, 055403 (2002).

[24] S. Pustelny, A. Wojciechowski, M. Gring, M. Kotyrba, J. Zachorowski, and W. Gawlik, Magnetometry based on nonlinear magneto-optical rotation with amplitude-modulated light, J. Appl. Phys. **103**, 063108 (2008).

[25] E. Breschi, Z. D. Grujić, P. Knowles, and A. Weis, A high-sensitivity push-pull magnetometer, Appl. Phys. Lett. **104**, 023501 (2014).

[26] A. Wojciechowski, E. Corsini, J. Zachorowski, and W. Gawlik, Nonlinear Faraday rotation and detection of superposition states in cold atoms, Phys. Rev. A **81**, 053420 (2010).

[27] W. Gawlik and S. Pustelny, *High Sensitivity Magnetometers* (Springer, Switzerland, 2017), p. 425.

[28] D. Budker, W. Gawlik, D. F. Jackson Kimball, S. M. Rochester, V. V. Yashchuk, and A. Weis, Resonant nonlinear magneto-optical effects in atoms, Rev. Mod. Phys. **74**, 1153 (2002).

[29] D. Budker and M. Romalis, Optical magnetometry, Nat. Phys. **3**, 227 (2007).

[30] A. Sheinker, L. Frumkis, B. Ginzburg, N. Salomonski, and B.-Z. Kaplan, Magnetic anomaly detection using a three-axis magnetometer, IEEE Trans. Magn. **45**, 160 (2009).

[31] M. Hayakawa, K. Hattori, and K. Ohta, Monitoring of ULF (ultra-low-frequency) geomagnetic variations associated with earthquakes, Sensors **7**, 1108 (2007).

[32] M. Hayakawa, R. Kawate, O. A. Molchanov, and K. Yumoto, Results of ultra-low-frequency magnetic field measurements during the Guam Earthquake of 8 August 1993, Geophys. Res. Lett. **23**, 241 (1996).

[33] A. C. Fraser-Smith, A. Bernardi, P. R. McGill, M. E. Ladd, R. A. Helliwell, and O. G. Villard, Jr., Low-frequency magnetic field measurements near the epicenter of the M_s

7.1 Loma Prieta Earthquake, Geophys. Res. Lett. **17,** 1465 (1990).

[34] E. Corsini, V. Acosta, N. Baddour, J. Higbie, B. Lester, P. Licht, B. Patton, M. Prouty, and D. Budker, Search for plant biomagnetism with a sensitive atomic magnetometer, J. Appl. Phys. **109,** 074701 (2011).

[35] I. Mateos, B. Patton, E. Zhivun, D. Budker, D. Wurm, and J. Ramos-Castro, Noise characterization of an atomic magnetometerat sub-millihertz frequencies, Sens. Actuators, A **224,** 147 (2015).

[36] M. Schmelz, R. Stolz, V. Zakosarenko, T. Schönau, S. Anders, L. Fritzsch, M. Mück, and H.-G. Meyer, Field-stable SQUID magnetometer with sub-fT Hz$^{-1/2}$ resolution based on sub-micrometer cross-type Josephson tunnel junctions, Supercond. Sci.Technol. **24,** 065009 (2011).

[37] D. Drung, Low-frequency noise in low-T_c multiloop magnetometers with additional positive feedback, Appl. Phys. Lett. **67,** 1474 (1995).

[38] E. B. Alexandrov, M. Auzinsh, D. Budker, D. F. Jackson Kimball, S. M. Rochester, and V. V. Yashchuk, Dynamic effects in nonlinear magneto-optics of atoms and molecules, J.Opt. Soc. Am. B **22,** 7 (2005).

[39] N. Wilson, N. B. Hébert, C. Perrella, P. Light, J. Genest, S. Pustelny, and A. Luiten, Simultaneous Observation of Nonlinear Magneto-Optical Rotation in the Temporal and Spectral Domains with an Electro-Optic Frequency Comb, Phys. Rev. Appl. **10,** 034012 (2018).

[40] W. Gawlik and S. Pustelny, *New Trends in Quantum Coherence and Nonlinear Optics* (Nova Publishers, New York, 2009), p. 47.

[41] T. J. Sumner, J. M. Pendlebury, and K. F. Smith, Convectional magnetic shielding, J. Phys. D **20,** 1095 (1987).

[42] S. Pustelny, D. F. Jackson Kimball, S. M. Rochester, V. V. Yashchuk, and D. Budker, Influence of magnetic-field inhomogeneity on nonlinear magneto-optical resonances, Phys. Rev. A **74,** 063406 (2006).

[43] D. W. Allan, Statistics of atomic frequency standards, Proc. IEEE **54,** 221 (1966).

[44] M. Auzinsh, D. Budker, and S. M. Rochester, *Optically Polarized Atoms* (Oxford University Press, New York, 2010).

[45] D. A. Steck, Rubidium 87 D Line Data, http://steck. us/alkalidata, note (revision 2.1.5, 13 September 2018).

[46] P. Welch, The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms, IEEE Trans. Audio Electroacoust. **15,** 70 (1967).

[47] S.-K. Lee and M. V. Romalis, Calculation of magnetic field noise from high-permeability magnetic shields and conducting objects with simple geometry, J. Appl. Phys. **103,** 084904 (2008).

[48] D. Hunter, S. Piccolomo, J. D. Pritchard, N. L. Brockie, T. E. Dyer, and E. Riis, Free-Induction-Decay Magnetometer Based on a Microfabricated Cs Vapor Cell, Phys. Rev. Appl. **10,** 014002 (2018).

[49] T. W. Kornack, S. J. Smullin, S.-K. Lee, and M. V. Romalis, A low-noise ferrite magnetic shield, Appl. Phys. Lett. **90,** 223501 (2007).

[50] M. T. Graf, D. F. Jackson Kimball, S. M. Rochester, K. Kerner, C. Wong, D. Budker, E. B. Alexandrov, M. V. Balabas, and V. V. Yashchuk, Relaxation of atomic polarization in paraffin-coated cesium vapor cells, Phys. Rev. A **72,** 023401 (2005).

[51] F. Wolfgramm, A. Cerè, F. A. Beduini, A. Predojević, M. Koschorreck, and M. W. Mitchell, Squeezed-Light Optical Magnetometry, Phys. Rev. Lett. **105,** 053601 (2010).

[52] T. Horrom, R. Singh, J. P. Dowling, and E. E. Mikhailov, Quantum-enhanced magnetometer with low-frequency squeezing, Phys. Rev. A **86,** 023803 (2012).

[53] F. Tricot, D. H. Phung, M. Lours, S. Guérandel, and E. de Clercq, Power stabilization of a diode laser with an acousto-optic modulator, Rev. Sci. Instrum. **89,** 113112 (2018).

[54] H. Duncker, O. Hellmig, A. Wenzlawski, A. Grote, A. J. Rafipoor, M. Rafipoor, K. Sengstock, and P. Windpassinger, Ultrastable, Zerodur-based optical benches for quantum gas experiments, Appl. Opt. **53,** 4468 (2014).

[55] *Magnetism in Medicine*, edited by W. Andrä and H. Nowak (Akademie Verlag, Weinheim, Germany, 2007), 2nd ed.

[56] S. J. Seltzer and M. V. Romalis, Unshielded three-axis vector operation of a spin-exchange-relaxation-free atomic magnetometer, Appl. Phys. Lett. **85,** 4804 (2004).

Quantum Sensing of Rapidly Varying Magnetic Fields

This chapter is based on the published article:

11.1   Overview and Motivation

Although recent developments in optical magnetometry have seen these devices reign supreme as exquisitely sensitive detectors of quasi-static magnetic fields — with sensitivities which rival, or even surpass that of superconducting quantum interference devices — they suffer from two major drawbacks: their relatively low amplitude and frequency response. This is, in part, due to the conventional phase-sensitive detection method most of them employ, which relies on synchronous optical pumping of a magneto-optical resonance.

This chapter describes experimental work in which a novel, temporal-domain measurement technique is developed and demonstrated to extend the amplitude and frequency response of an optical magnetometer by many orders of magnitude in comparison with conventional techniques. This is achieved by measuring the free-induction decay of a transverse atomic polarisation, during which an oscillating magnetic field is applied to the vapour cell. The instantaneous phase of the free-induction decay is extracted and analysed, yielding accurate and precise calibration-free measurements of both the dc magnetic-field strength, as well as any present ac field components. It is demonstrated that both the amplitude and frequency response are increased by over two and three orders of magnitude respectively, in comparison to the traditional phase-sensitive detection method.

This new measurement technique paves the way for optical magnetometers to be used in applications which require measurements of large-amplitude, high-frequency magnetic fields — a regime which was previously inaccessible for optical magnetometers due to prohibition by conventional measurement techniques.

11.2 Statement of Contribution

11.2.1 Conceptualisation

The idea of performing high-bandwidth optical magnetometry by monitoring phase fluctuations in the temporal domain is attributed to André Luiten.

11.2.2 Realisation

Nathanial Wilson constructed the experimental apparatus and acquired the data presented in this manuscript. The code used to generate the Fourier-series magnetic field was written by Christopher Perrella. Analysis of the experimental data was performed by Nathanial Wilson, with assistance from Russell Anderson. The analysis technique/algorithm — including the application of Hilbert transforms in extracting the instantaneous phase — and the code used to implement it, was developed by Nathanial Wilson. Theoretical analysis of the instantaneous-phase extraction algorithm — including its limitations — was performed by Nathanial Wilson, Philip Light and Russell Anderson, based on useful discussions with Yvonne Stokes.

11.2.3 Documentation

This manuscript was written primarily by Nathanial Wilson and Russell Anderson, with subsequent editing performed by Christopher Perrella, Philip Light and André Luiten.

Candidate overall percentage contribution: 70%.

_____

Nathanial Wilson

11.2.4 Co-author Contributions

By signing below, each co-author certifies that:

- The candidate's stated contribution to the publication is accurate (as stated above)

- Permission is granted for the candidate to include the publication in the thesis

Co-author certifications are in order of published authorship unless otherwise specified.

———————————————————

Christopher Perrella

———————————————————

Russell Anderson

———————————————————

André Luiten

———————————————————

Philip Light

Wide-bandwidth atomic magnetometry via instantaneous-phase retrieval

Nathanial Wilson ●,[1] Christopher Perrella,[1,*] Russell Anderson ●,[2] André Luiten,[1,†] and Philip Light ●[1]

[1]*Institute for Photonics and Advanced Sensing (IPAS), and School of Physical Sciences, University of Adelaide,
South Australia 5005, Australia*

[2]*La Trobe Institute for Molecular Science, and School of Molecular Sciences, La Trobe University, Victoria 3552, Australia*

We develop and demonstrate a new protocol that allows sensing of magnetic fields in an extra-ordinary regime for atomic magnetometry. Until now, the demonstrated bandwidth for atomic magnetometry has been constrained to be slower than the natural precession of atomic spins in a magnetic field—the Larmor frequency. We demonstrate an approach that tracks the instantaneous phase of atomic spins to measure arbitrarily modulated magnetic fields with frequencies up to 50 times higher than the Larmor frequency. By accessing this regime, we demonstrate magnetic-field measurements across four decades in frequency up to 400 kHz, over three orders of magnitude wider than conventional atomic magnetometers. Furthermore, we demonstrate that our protocol can linearly detect transient fields 100-fold higher in amplitude than conventional methods. We highlight the bandwidth and dynamic range of the technique by measuring a magnetic field with a broad and dynamical spectrum.

I. INTRODUCTION

Measuring magnetic fields with high accuracy and precision is paramount in myriad applications including medical diagnostics and imaging [1–11], geomagnetism [12], and fundamental physics [13–16]. Superconducting quantum interference devices (SQUIDs) [17,18] and optical atomic magnetometers (OAMs) [12,19,20] reign as exquisitely sensitive detectors of static and slowly changing magnetic fields with $fT/\sqrt{Hz}$ precision. However, there are many applications, e.g., biomagnetic signals [21], magnetic communications [22], and improvised threat detection [23] where there is a desire to detect time-varying magnetic fields.

One pathway to high-speed magnetometry is so-called resonantly tuned magnetometers, whose Larmor frequency is proximate to an oscillating magnetic field of interest [24–29]. These devices are suitable when the signal is relatively narrow band, about a frequency that is known *a priori*, and provided that the dc field strength can be tuned accordingly. The bandwidth of these magnetometers is set by the spin resonance linewidth $\Gamma \propto 1/T_2$, of order 3 to 400 Hz [24]. This approach to ac magnetic sensing surrenders two key benefits of dc atomic magnetometry: (a) the output signal is an indirect measure of the oscillating field amplitude and is no longer calibration free [30], while (b) the sensor only responds

linearly to magnetic field fluctuations that have a magnitude much less than the resonance width [25]—typically of order 1 nT and less. Although it is possible to enhance the amplitude range and bandwidth of this type of magnetometer through intentionally decreasing the effective spin relaxation time, this necessarily comes at the expense of sensitivity [31–36].

As an alternative, quantum metrology protocols such as dynamical decoupling [37–39], compressive sensing [40,41], or Hamiltonian estimation [42] can provide retrospective insight into magnetic waveforms. The recent demonstrations of quantum lock-in detection measure the frequency of continuously oscillating fields with superb submillihertz precision [43–45]. Contemporary approaches have used entanglement to enhance rf field detection [46], and predictive filters to track time-dependent signals [47]. However, the challenge with these protocols is the need to have prior information about the waveform or a requirement that its spectrum be single frequency.

None of these aforementioned approaches fulfills an urgent need for real-time and accurate detection of broadband magnetic signals with a mixture of frequencies and amplitudes. We present a protocol for time-dependent magnetometry that retrieves the instantaneous spin-precession frequency. This permits the observation of magnetic fields that are varying much faster than the spin-precession frequency itself. This regime of supra-Larmor-frequency modulation has been hitherto unexplored and perhaps surprisingly lies outside the realm of conventional FM signal processing.

Our approach exploits the technique of free-induction decay (FID) in a fundamentally original way. The standard approach is to (a) observe the FID of an ensemble of spins [48–53], (b) sinusoidal regression to the observed signal to obtain a single estimate of the Larmor frequency, and then (c) repeat this process in a train of optical pumping and

*chris.perrella@adelaide.edu.au
†andre.luiten@adelaide.edu.au

free-induction decay cycles [50]. This naturally imposes a maximum bandwidth on the order of the reciprocal of the repetition rate; e.g., an FID train with a repetition rate of 1 kHz could track magnetic fields oscillating orthogonal to the dc field up to 100 Hz [54] and parallel to the dc field with a bandwidth up to 1 kHz [50].

We instead use an innovative protocol to measure the instantaneous phase of precessing spins, rather than their average frequency. In this way, we obtain an instantaneous measure of the time-dependent Larmor frequency, in direct proportion to the time-dependent magnetic-field strength. Our approach can thus deliver the key benefits of a low-frequency atomic magnetometer, i.e., calibration-free and linear [12,19,20,55]. Moreover, as spin precession has no inertia [31,33], the Larmor frequency will respond instantaneously to a change in the external magnetic field. Here we demonstrate this by way of measurement of single-tone modulations up to 50 times higher than the Larmor frequency and by measuring arbitrarily modulated magnetic signals with significant spectral content above the Larmor frequency. We note that, to our knowledge, no one has yet demonstrated a protocol that has allowed access to this extraordinary regime.

II. INSTANTANEOUS PHASE RETRIEVAL

We induce transverse magnetization in an atomic vapor by amplitude modulating an optical pumping beam near twice the Larmor frequency [56]. Upon extinction of the pump beam, Faraday polarimetry of an off-resonant optical probe constitutes a weak measurement of the freely precessing atomic spins [33,57]. For a dc magnetic-field strength B_{dc}, the polarization rotation of the probe beam $\phi(t) \propto \sin(2\pi f_c t)$ has a "carrier" frequency $f_c = 2f_L$, where $2\pi f_L = \gamma B_{dc}$ is the Larmor frequency and γ is the gyromagnetic ratio of the ground-state hyperfine level. In our configuration, the polarization-rotation oscillates at twice the Larmor frequency, owing to $|\Delta m_F| = 2$ ground-state coherences [58]. For a time-varying magnetic field, we will see a polarization rotation:

$$\phi(t) = a_c \cos\left(2\pi \int_0^t f_I(\tau)d\tau + \varphi_c\right), \quad (1)$$

where a_c is the carrier amplitude and φ_c is the arbitrary phase of the carrier. The instantaneous frequency $f_I(t)$ provides a direct measure of the instantaneous magnetic-field strength $B(t)$ via

$$B(t) = \frac{\pi f_I(t)}{\gamma}. \quad (2)$$

To calculate the instantaneous frequency, we retrieve the instantaneous phase of the measured polarization rotation using the analytic representation $\phi_a(t) = \phi(t) + i\mathcal{H}\{\phi(t)\}$, where $\mathcal{H}$ is the Hilbert transform [59]

$$\mathcal{H}\{\phi(t)\} = \frac{1}{\pi}\mathcal{P}\int_{-\infty}^{\infty}\frac{\phi(\tau)}{t-\tau}d\tau. \quad (3)$$

In the above, $\mathcal{P}$ denotes the Cauchy principal value, and the Hilbert transform imparts a 90° phase shift to every Fourier component of $\phi(t)$. The instantaneous phase of $\phi(t)$ is then

estimated via

$$\varphi_I(t) = \arg\{\phi_a(t)\}. \quad (4)$$

Upon unwrapping $\varphi_I(t)$ we obtain a continuous function of t and compute the instantaneous frequency via [60]

$$f_I(t) = \frac{1}{2\pi}\frac{d\varphi_I(t)}{dt}. \quad (5)$$

When the magnetic field is sinusoidally modulated, $B(t) = B_{dc} + \Delta B \cos(2\pi f_m t + \varphi_m)$, and the instantaneous frequency is given by $f_I(t) = f_c + \Delta f \cos(2\pi f_m t + \varphi_m)$, where $\Delta f = \gamma \Delta B/\pi$ is the frequency deviation, f_m is the modulation frequency, and φ_m is the arbitrary modulation phase offset. The polarization rotation is then

$$\phi(t) = a_c \cos(2\pi f_c t + \beta \sin[2\pi f_m t + \varphi_m] + \varphi_c), \quad (6)$$

where $\beta = \Delta f/f_m$ is the modulation index. Equation (6) makes plain the means to obtain the parameters of an applied magnetic field from the observed frequency-modulated polarization rotation. For single-tone modulation, the oscillating field amplitude can be determined directly from the oscillation amplitude of either the instantaneous-frequency ($\Delta B = \pi \Delta f/\gamma$) or the instantaneous phase ($\Delta B = \pi \beta f_m/\gamma$). This can be extended to more general time-dependent magnetic fields that admit a Fourier series, by multiplying the modulation index of each Fourier component by its corresponding frequency (see Sec. V).

The approach we have described above is one that appears to be unorthodox in signal processing since we have not imposed any restriction on the modulation frequency being below the carrier frequency. Below we show that we can successfully extract magnetic field modulation signals at higher frequency than the carrier.

III. EXPERIMENTAL SETUP

The experimental setup is presented in Fig. 1. Isotopically pure ^{87}Rb is contained in a cylindrical vapor cell with 40 mm diameter and 40 mm length. The walls are anti-relaxation-coated to extend the coherences between ground-state Zeeman sublevels, with a measured transverse spin-relaxation time of $T_2 = 45$ ms (cf. Fig. 2). The cell remains at room temperature and is housed within a three-layer cylindrical μ-metal magnetic shield that has a measured shielding factor of approximately 2×10^3. A constant-bias magnetic field of $B_{dc} \approx 2$ μT is generated along the longitudinal axis of the cell using a solenoid installed within the innermost shield. An oscillatory component of the magnetic field is generated using a separate high-bandwidth coil wrapped around the center of the solenoid.

The atomic vapor is optically pumped and probed using light from an external cavity diode laser tuned to 795 nm, near the $5^2S_{1/2} \rightarrow 5^2P_{1/2}$ (D_1) transition of ^{87}Rb. The laser is frequency-locked 80 MHz below the $F = 2 \rightarrow F' = 1$ hyperfine transition ($\gamma/2\pi = 6.9958$ GHz/T [61]) using saturated absorption spectroscopy in a separate reference cell. The optical pumping beam is amplitude modulated with a 20% duty-cycle square wave via an acousto-optic modulator at a frequency near $2f_L$, with a time-averaged power of $\overline{P}_{pump} \approx 25$ μW.
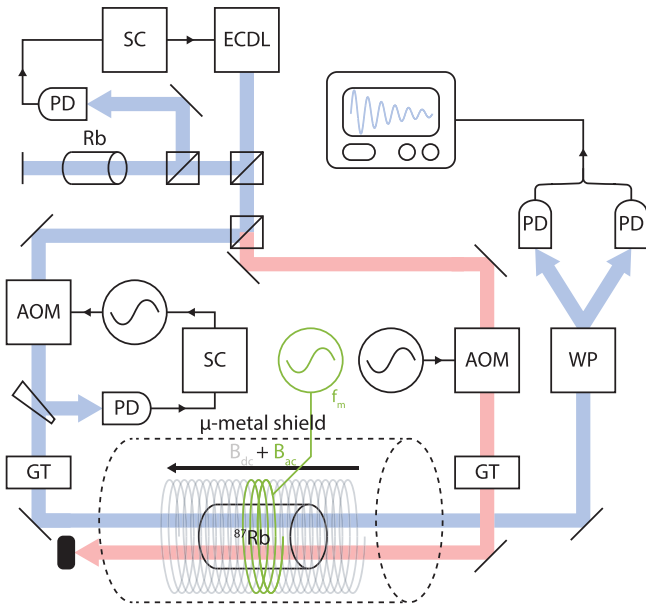
FIG. 1. Simplified experimental setup, showing the external cavity diode laser (ECDL), acousto-optic modulator (AOM), Glan-Thompson prism (GT), Wollaston prism (WP), photodetector (PD), and servo controller (SC). The probe beam is shown in blue, while the pump beam is shown in red. The modulation coil, shown in green, is wrapped around the center of the solenoid (gray).

The pump and probe beams are linearly polarized by Glan-Thompson (GT) prisms immediately prior to entering opposite sides of the vapor cell and propagate (anti)parallel to the magnetic field. Both beams have a $1/e^2$ diameter of 1.5 mm and are horizontally displaced from each other by approximately 10 mm.

After traversing the vapor cell, the probe beam ($P_{\mathrm{probe}} = 5\,\mu$W) passes through a Wollaston prism, which separates the beam into orthogonal linear-polarization components. These orthogonal components are measured on separate photodetectors, forming a balanced polarimeter. The optical power on each photodetector can be converted into a polarization-

FIG. 2. Polarization rotation $\phi(t)$ (blue trace) during free-induction decay of an atomic vapor, measured in a single shot. During the free evolution of the atomic spins, a transient, oscillating magnetic field with $f_{\mathrm{m}} = 500$ Hz was applied (black trace).

FIG. 3. Instantaneous phase $\varphi_{\mathrm{I}}(t)$ retrieved from the free-induction decay in Fig. 2. Linear regression to $\varphi_{\mathrm{I}}(t)$ yields a gradient $d\varphi_{\mathrm{I}}/dt = 192.49$ rad/ms, corresponding to a carrier frequency $f_{\mathrm{c}} = 2f_{\mathrm{L}} = 30.636$ kHz, which implies a dc magnetic-field strength $B_{\mathrm{dc}} = 2.1896\,\mu$T. The residual from the linear fit (top) shows the effect of the sinusoidal magnetic modulation, from which we derive the modulation frequency $f_{\mathrm{m}} = 500$ Hz, and modulation index $\beta = 1.73$ rad ($\Delta f = 865$ Hz) resulting from an ac magnetic-field strength $\Delta B = 61.8$ nT.

rotation angle $\phi(t)$, via [34]

$$\phi(t) = \frac{1}{2} \arcsin\left(\frac{P_1 - P_2}{P_1 + P_2}\right), \qquad (7)$$

where P_1 and P_2 are the optical powers on the two photodetectors.

IV. SINUSOIDAL MODULATION

Once the transverse spin coherence has been established, the pump beam is extinguished. During the FID, we modulate the magnetic field sinusoidally while measuring the resulting polarization rotation. We retrieve the instantaneous phase of that signal and—per Sec. II—are thus able to determine the average magnetic-field strength, as well as the frequency, amplitude, and phase of the magnetic-field modulation. An example of the recorded polarization rotation signal is presented in Fig. 2, with the process leading to the corresponding instantaneous phase retrieval shown in Fig. 3. The average gradient of the instantaneous phase can be used to calculate the carrier frequency of the signal using Eq. (5), thus delivering the Larmor frequency and hence the dc magnetic-field strength. The oscillating component of the instantaneous phase has an amplitude β, proportional to the ac magnetic-field amplitude.

Our protocol shows an outstanding amplitude linearity and dynamic range to transient fields when compared to traditional rf atomic magnetometers. In Fig. 4, we present the response of the new device to ac magnetic fields of different amplitudes at a fixed modulation frequency of $f_{\mathrm{m}} = 522$ Hz. The measured field amplitude ΔB across a ≈ 140 nT range was within about 3% of the value that was independently measured using a fluxgate magnetometer. This discrepancy can be accounted for by the slight difference in measurement volumes between the fluxgate magnetometer and the vapor cell, resulting in a different volume-averaged magnetic-field strength. Linear regression to the data in Fig. 4, with a fixed

FIG. 4. Linear amplitude response of the magnetometer to a sinusoidal modulation at a frequency of $f_m = 522$ Hz—measured amplitude ΔB compared with an independent measurement $(\Delta B)_{\text{induction}}$ using a fluxgate magnetometer. Bottom: Linear-regression (red, solid line) to the data (blue) with a fixed intercept of zero yielded a gradient of 0.9684(1), i.e., a $\approx 3\%$ deviation from the 1:1 line (gray, dashed). Single standard deviation error bars are indiscernible from the data points here. Top: Fractional residuals from linear regression to the data with a fixed, zero intercept. Shaded in gray is the typical linear amplitude response of conventional radio-frequency magnetometers. The maximum permissible modulation index under these conditions is $\beta_{\text{max}} \approx 57$ rad (cf. Appendix A).

offset of zero, yields a gradient of 0.9684(1), corresponding to a linearity of 126 ppm. The enhanced linearity and wide, calibration-free amplitude response augurs well for range tracking remote sensing targets of known characteristic size. A detailed estimation of the maximum permissible oscillating field amplitude is described in Appendix A.

By retrieving the instantaneous Larmor phase, it is possible to detect, in real time, magnetic field oscillations that are much faster than the Larmor frequency itself. Furthermore, the bandwidth of the sensor is also larger than the Larmor frequency. To demonstrate how our protocol allows access to this experimentally unprecedented regime, we detect $f_m \leqslant 50f_L$, over three orders of magnitude greater than the typical bandwidth of rf atomic magnetometers.

We characterized the frequency response and accuracy of the protocol by comparing the oscillating field amplitude ΔB imputed using our new instantaneous-phase retrieval to the value $(\Delta B)_{\text{induction}}$ determined from either a fluxgate ($f_m \leqslant 1$ kHz) or an induction-coil sensor ($f_m > 1$ kHz). The ratio of these independent measurements across four decades of modulation frequency is shown in Fig. 5 for two different Larmor/carrier frequencies, corresponding to dc fields 2.17 μT and 128 nT. This cross calibration of our FID magnetometer against inductive measurements demonstrates the ability to detect fields modulated well above the Larmor frequency using instantaneous phase retrieval. We attribute the small deviations from unity (≈ 2 dB) in the high frequency parts ($f_m > f_c$) of the cross calibration to mutual inductance between the modulation and induction coils, rather than any errant atomic response.

FIG. 5. Comparison of the ac field amplitude measured using instantaneous-phase retrieval to an independent inductive measurement vs modulation frequency. The ratio $\Delta B/(\Delta B)_{\text{induction}}$ is shown for two different carrier frequencies (dashed, vertical lines), $f_c = 2f_L = 1.8$ kHz (red) and 30.5 kHz (blue). Error bars correspond to the combined fractional uncertainty of the induction sensors (specified calibration accuracy of 5–25%, mauve) and the measurement uncertainty in ΔB (one standard error). The gray shaded region denotes the typical frequency response of conventional rf OAMs. Each datum was measured in a single shot, though averaging permits resolution of higher modulation frequencies; e.g., we observe up to $f_m = 400$ kHz for 20 shots. For these data, the modulation index spanned the range $\beta \in [0.003, 92]$ rad.

V. NONTRIVIAL MODULATION

To highlight the wide bandwidth and large dynamic range of the technique, we synthesized a magnetic field with a broad and dynamical spectrum that encodes the acronym NMOR (nonlinear magneto-optical resonance). We drove the modulation coil with a wide-band modulated current [see Fig. 6(a)], proportional to the resulting time-dependent field:

$$B(t) = B_{\text{dc}} + \sum_{j=1}^{N} \Delta B_j(t) \cos(2\pi f_j t + \varphi_j). \quad (8)$$

Here, we used the same dc magnetic field as displayed in Fig. 3 and $N = 9$ Fourier components, each spaced by $f_{j+1} - f_j = 1$ kHz, where $f_j \in [8, 16]$ kHz. The field amplitudes $\Delta B_j(t)$ were piecewise constant, varying abruptly at 1 ms intervals to encode the acronym in the time-frequency domain. The resulting polarization rotation was of the form

$$\phi(t) = a_c \cos \left[2\pi f_c t + \sum_{j=1}^{N} \beta_j(t) \sin(2\pi f_j t + \varphi_j) + \varphi_c \right]. \quad (9)$$

We retrieved the instantaneous phase from this measurand (Sec. II), and subtracted a linear fit (Sec. IV) to impute $\varphi_I(t) - 2\pi f_c t - \varphi_c$. The resulting oscillatory component of the instantaneous phase was analyzed in the time-frequency domain using the short-time Fourier transform (STFT). When the frequency binning is chosen to match $f_{j+1} - f_j$, the STFT amplitudes are the time-dependent modulation indices $\beta_j(t) \in [11.7, 36.2]$ mrad. The STFT amplitude was converted to magnetic field amplitudes $\Delta B_j = \pi \beta_j f_j / \gamma$ and is displayed

FIG. 6. Magnetometer response to a spectrally broad and dynamical magnetic-field modulation. (a) Real-time magnetic field measured using instantaneous phase retrieval (red) compared to the field predicted from the current used to drive the modulation coil (blue). Traces are vertically offset by ± 50 nT for clarity. Spectrograms of this prediction (b) and the measured magnetometer output (c) show the dynamical spectral components [ΔB_j in Eq. (8)] in close agreement.

as a spectrogram in Fig. 6(c). We compare this to the spectrogram of the intentionally applied field we predict at the sensor [Fig. 6(b)], based on the measured electrical current in the modulation coil and the Biot-Savart law. The two plots are in excellent agreement with only a small degradation in signal-to-noise ratio of the atomic magnetometer as t approaches $T_2 = 45$ ms.

VI. DISCUSSION

Under the experimental conditions detailed here, the instantaneous phase of the polarization rotation has an observed noise floor of $68\ \mu\mathrm{rad_{rms}}/\sqrt{\mathrm{Hz}}$ at frequencies above 100 Hz—within 60% of the shot-noise limit of $44\ \mu\mathrm{rad_{rms}}/\sqrt{\mathrm{Hz}}$ (cf. Appendix C). This level of phase noise corresponds to a magnetic-field noise of approximately $5(f/\mathrm{Hz})\,\mathrm{fT}/\sqrt{\mathrm{Hz}}$. All estimates of the modulation amplitude ΔB, modulation frequency f_m, and carrier frequency f_c pre-

sented here are within 15% of the Cramér-Rao lower bounds associated with this level of white phase noise [62].

For modulation frequencies below the carrier ($f_\mathrm{m} < f_\mathrm{c}$), there is a single frequency component present in the retrieved instantaneous phase, at f_m. However, when $f_\mathrm{m} > f_\mathrm{c}$ (supra-carrier modulation), the oscillating component of the instantaneous phase retrieved using Eq. (4) contains two tones (cf. Appendix B), at f_m and $2f_\mathrm{c} - f_\mathrm{m}$, with equal amplitudes $\beta/2$. The frequency and amplitude of ac magnetic fields can still be imputed in the supra-carrier regime, provided we have knowledge of f_c. This is guaranteed by the self-certifying estimate of f_c, which is naturally obtained during linear regression to $\varphi_\mathrm{I}(t)$ (Fig. 3). This is equally valid whether $f_\mathrm{m} > f_\mathrm{c}$ or $f_\mathrm{m} < f_\mathrm{c}$.

The ability to measure high-frequency magnetic-field fluctuations will, in practice, depend upon the signal-to-noise ratio of the measurement. This can be thought of in two equivalent ways: the measured instantaneous-phase noise is white (independent of frequency); however, the signal amplitude, when measured in rotation of the polarization, is given by the modulation index $\beta = \Delta f/f_\mathrm{m}$, which scales as f_m^{-1}. This results in a magnetic signal-to-noise ratio which also scales as f_m^{-1}. Alternatively, if we consider the signal amplitude in terms of the instantaneous frequency, then it is a constant value of Δf; however, the white instantaneous phase noise considered in terms of frequency will be violet (differentiated white noise). Once again, we see that the signal-to-noise ratio of an ac magnetic field measurement scales as f_m^{-1}.

The low-frequency response of this technique is limited to $f_\mathrm{m} \gtrsim T_2^{-1}$, below which the polarization-rotation signal will decay before a full modulation cycle has been observed. Of course, it would be possible to augment our technique with the traditional approach in which one uses repeated FID measurements: In this case, slow changes in f_L can be tracked [50].

We note that typical rf atomic magnetometers detect weak oscillating fields oriented transverse to the static background field [24–29], whereas our magnetometer senses longitudinally oscillating fields. For arbitrarily oriented rf fields, there will be a dead band of $\Delta f_\perp = \gamma \Delta B_\perp/2\pi$ about the Larmor frequency, where a transverse component of the oscillating field with amplitude $\Delta B_\perp$ can drive Zeeman transitions, resulting in amplitude modulation of the polarization rotation [39].

VII. CONCLUSION

We have developed a phase-retrieval technique that can extend the accuracy and applicability of dc optical atomic magnetometers so that they are now suitable for measuring ac magnetic fields. We explore a regime that was previously experimentally inaccessible, where the bandwidth of the sensor exceeds that of the Larmor frequency itself. We have demonstrated calibration-free measurement of oscillating fields in real time with amplitudes up to 150 nT and frequencies up to 400 kHz, an increase in the amplitude and frequency response of two and three orders of magnitude over conventional rf atomic magnetometers respectively. The instantaneous-phase retrieval can be applied more broadly to quantum sensors that employ weak continuous measurement of spin precession, including those exploiting synchronous detection and

feedback [63]. This capability augurs well for robust field-deployable magnetometers which, unlike laboratory-based apparatus, face inherently unpredictable operational environments and transient signals of interest.

APPENDIX A: MAXIMUM PERMISSIBLE MODULATION INDEX

The frequency-modulated polarization rotation deriving from an oscillating magnetic field [Eq. (6)] can be rewritten using the Jacobi-Anger expansion [64]:

$$\phi(t) = a_c \sum_{n=-\infty}^{\infty} J_n(\beta) \cos(2\pi[f_c + nf_m]t + \varphi_c + n\varphi_m),$$
(A1)

where J_n is the nth-order Bessel function of the first kind. Thus, the polarization rotation comprises an infinite series of sidebands, distributed symmetrically about the carrier at f_c, spaced by integer multiples of the modulation frequency f_m, with amplitudes $a_n = a_c J_n(\beta)$.

A conservative upper bound for the maximum permissible modulation index, β_{max}, is derived by considering the fractional power in sidebands about f_c which extend to negative frequencies. These sidebands are discarded when numerically computing the analytic signal, which can compromise the instantaneous-phase retrieval. For a given fractional modulation frequency, $\alpha = f_m/f_c$, the largest value of n such that $f_c - nf_m > 0$, is given by

$$n_{max} = \lceil (\alpha^{-1}) \rceil - 1.$$
(A2)

To calculate β_{max} for a given α, one must choose an acceptable fractional power P_{max} contained in the discarded sidebands. This constraint amounts to

$$\sum_{n=n_{max}+1}^{\infty} J_n(\beta_{max})^2 \leqslant P_{max}.$$
(A3)

Using Bessel-function identities and symmetry relations, Eq. (A3) simplifies to

$$\sum_{n=-n_{max}}^{n_{max}} J_n(\beta_{max})^2 \geqslant 1 - 2P_{max}.$$
(A4)

The maximum permissible modulation index calculated using Eq. (A4) is shown in Fig. 7, for two different power fractions: $P_{max} = 0.01$ and 0.001. Because of the nonlinearity of Eq. (A4), an increasingly stringent bound on P_{max} does not significantly reduce β_{max}.

FIG. 7. Maximum permissible modulation index, β_{max}, calculated using Eq. (A4) for $P_{max} = 0.01$ (blue trace) and $P_{max} = 0.001$ (red trace). For $\alpha > 1$, β_{max} remains constant until $\alpha \approx (4P_{max})^{-1/2}$, beyond which the constraint that $B(t) > 0$ dominates, and then $\beta_{max} = \alpha^{-1}$.

An additional constraint on β arises from requiring that the instantaneous magnetic field never be zero, which could result in nonadiabatic spin flips that compromise this measurement. This limits the maximum measurable ac magnetic-field amplitude to $\Delta B < B_{dc}$, equivalent to $\Delta f < f_c$ and thus $\beta < f_c/f_m = \alpha^{-1}$. This constraint becomes stricter than the sideband-power considerations for $\alpha \gtrsim (4P_{max})^{-1/2}$.

APPENDIX B: INSTANTANEOUS PHASE ESTIMATION WHEN MODULATING FASTER THAN THE CARRIER FREQUENCY

For supracarrier modulation ($f_m > f_c$), the lower sideband nearest the positive-frequency carrier ($f = f_c$) at $f = f_c - f_m$ becomes negative [$n = -1$ in Eq. (A1)]. When numerically computing the analytic signal in this regime, this sideband is discarded. However, due to the Hermitian symmetry of the Fourier transform of a real function, there is a corresponding upper sideband of the negative-frequency carrier ($f = -f_c$) at $f = -f_c + f_m$ which becomes positive. In the case of weak modulation (i.e., $\beta \lesssim 1$, which is generally the case when $f_m > f_c$), the majority of the modulated power of the analytic signal is contained in the two positive frequencies $f = f_c + f_m$ and $f = -f_c + f_m$. The spacing of these tones from the positive-frequency carrier f_c is f_m and $|2f_c - f_m|$, respectively. The asymmetry of these dominant sidebands results in a modification of the instantaneous phase retrieved using Eq. (4), in the form of double-sideband suppressed-carrier amplitude modulation:

$$\varphi_I(t) = 2\pi f_c t + \frac{\beta}{2} \sin(2\pi f_m t + \varphi_m)$$
$$+ \frac{\beta}{2} \sin[2\pi(2f_c - f_m)t - \varphi_m] + \varphi_c$$
$$= 2\pi f_c t + \beta \cos[2\pi(f_m - f_c)t + \varphi_m] \sin(2\pi f_c t) + \varphi_c.$$
(B1)

FIG. 8. Amplitude spectral density of the instantaneous phase for a carrier frequency $f_c = 2f_L = 30.636$ kHz, and with an applied magnetic-field modulation of amplitude $\Delta B = 102$ nT at a frequency $f_m = 10$ kHz.

The resulting spectrum of $\varphi_I(t) - 2\pi f_c t - \varphi_c$ contains two tones, at f_m and $2f_c - f_m$. However, as the power is evenly distributed between these two tones, and the method prescribed in Fig. 3 can still be used to determine f_c, the frequency and amplitude of ac magnetic fields can still be imputed unambiguously in the supracarrier regime.

APPENDIX C: SENSITIVITY

The magnetic sensitivity of this technique is ultimately determined by the ability to resolve small changes in the instantaneous phase. A representative amplitude spectral density of the instantaneous-phase is presented in Fig. 8, indicating a nearly white noise floor with magnitude 68 $\mu\mathrm{rad}_{rms}/\sqrt{\mathrm{Hz}}$. Using Eqs. (2) and (5), this level of instantaneous phase noise corresponds to a magnetic-field noise of $5.0(f/\mathrm{Hz})$ fT/$\sqrt{\mathrm{Hz}}$.

A photocurrent-referenced measurement of the optical power indicates the shot-noise limit for instantaneous-phase is 44 $\mu\mathrm{rad}_{rms}/\sqrt{\mathrm{Hz}}$, or about $3.1(f/\mathrm{Hz})$ fT/$\sqrt{\mathrm{Hz}}$ in magnetic units.

[1] G. Bison, N. Castagna, A. Hofer, P. Knowles, J.-L. Schenker, M. Kasprzak, H. Saudan, and A. Weis, A room temperature 19-channel magnetic field mapping device for cardiac signals, Appl. Phys. Lett. **95**, 173701 (2009).

[2] G. Bison, R. Wynands, and A. Weis, A laser-pumped magnetometer for the mapping of human cardiomagnetic fields, Appl. Phys. B **76**, 325 (2003).

[3] S. Knappe, T. H. Sander, O. Kosch, F. Wiekhorst, J. Kitching, and L. Trahms, Cross-validation of microfabricated atomic magnetometers with superconducting quantum interference devices for biomagnetic applications, Appl. Phys. Lett. **97**, 133703 (2010).

[4] C. Johnson, P. D. D. Schwindt, and M. Weisend, Magnetoencephalography with a two-color pump-probe, fiber-coupled atomic magnetometer, Appl. Phys. Lett. **97**, 243703 (2010).

[5] S. Xu, V. V. Yashchuk, M. H. Donaldson, S. M. Rochester, D. Budker, and A. Pines, Magnetic resonance imaging with an optical atomic magnetometer, Proc. Natl. Acad. Sci. USA **103**, 12668 (2006).

[6] E. Boto, N. Holmes, J. Leggett, G. Roberts, V. Shah, S. S. Meyer, L. Duque Muñoz, K. J. Mullinger, T. M. Tierney, S. Bestmann, G. R. Barnes, R. Bowtell, and M. J. Brookes, Moving magnetoencephalography towards real-world applications with a wearable system, Nature (London) **555**, 657 (2018).

[7] K. Jensen, M. A. Skarsfeldt, H. Stærkind, J. Arnbak, M. V. Balabas, S.-P. Olesen, B. H. Bentzen, and E. S. Polzik, Magnetocardiography on an isolated animal heart with a room-temperature optically pumped magnetometer, Sci. Rep. **8**, 16218 (2018).

[8] J. Belfi, G. Bevilacqua, V. Biancalana, S. Cartaleva, Y. Dancheva, and L. Moi, Cesium coherent population trapping magnetometer for cardiosignal detection in an unshielded environment, J. Opt. Soc. Am. B **24**, 2357 (2007).

[9] G. Bison, R. Wynands, and A. Weis, Dynamical mapping of the human cardiomagnetic field with a room-temperature, laser-optical sensor, Opt. Express **11**, 904 (2003).

[10] H. Xia, A. Ben-Amar Baranga, D. Hoffman, and M. V. Romalis, Magnetoencephalography with an atomic magnetometer, Appl. Phys. Lett. **89**, 211104 (2006).

[11] O. Alem, T. H. Sander, R. Mhaskar, J. LeBlanc, H. Eswaran, U. Steinhoff, Y. Okada, J. Kitching, L. Trahms, and S. Knappe, Fetal magnetocardiography measurements with an array of microfabricated optically pumped magnetometers, Phys. Med. Biol. **60**, 4797 (2015).

[12] H. B. Dang, A. C. Maloof, and M. V. Romalis, Ultrahigh sensitivity magnetic field and magnetization measurements with an atomic magnetometer, Appl. Phys. Lett. **97**, 151110 (2010).

[13] M. Smiciklas, J. M. Brown, L. W. Cheuk, S. J. Smullin, and M. V. Romalis, New Test of Local Lorentz Invariance Using a ^{21}Ne-Rb-K Comagnetometer, Phys. Rev. Lett. **107**, 171604 (2011).

[14] C. J. Berglund, L. R. Hunter, D. Krause Jr., E. O. Prigge, M. S. Ronfeldt, and S. K. Lamoreaux, New Limits on Local Lorentz Invariance from Hg and Cs Magnetometers, Phys. Rev. Lett. **75**, 1879 (1995).

[15] S. Pustelny, D. F. Jackson Kimball, C. Pankow, M. P. Ledbetter, P. Wlodarczyk, P. Wcislo, M. Pospelov, J. R. Smith, J. Read, W. Gawlik, and D. Budker, The global network of optical magnetometers for exotic physics (GNOME): A novel scheme to search for physics beyond the standard model, Ann. Phys. **525**, 659 (2013).

[16] D. F. Jackson Kimball, I. Lacey, J. Valdez, J. Swiatlowski, C. Rios, R. Peregrina-Ramirez, C. Montcrieffe, J. Kremer, J. Dudley, and C. Sanchez, A dual-isotope rubidium comagnetometer to search for anomalous long-range spin-mass (spin-gravity) couplings of the proton, Ann. Phys. **525**, 514 (2013).

[17] M. Schmelz, R. Stolz, V. Zakosarenko, T. Schönau, S. Anders, L. Fritzsch, M. Mück, and H.-G. Meyer, Field-stable SQUID magnetometer with sub-fT Hz$^{-1/2}$ resolution based on sub-micrometer cross-type Josephson tunnel junctions, Supercond. Sci. Technol. **24**, 065009 (2011).

[18] T. Schönau, M. Schmelz, V. Zakosarenko, R. Stolz, M. Meyer, S. Anders, L. Fritzsch, and H.-G. Meyer, SQUID-based setup for the absolute measurement of the Earth's magnetic field, Supercond. Sci. Technol. **26**, 035013 (2013).

[19] I. K. Kominis, T. W. Kornack, J. C. Allred, and M. V. Romalis, A subfemtotesla multichannel atomic magnetometer, Nature (London) **422**, 596 (2003).

[20] D. Sheng, S. Li, N. Dural, and M. V. Romalis, Subfemtotesla Scalar Atomic Magnetometry Using Multipass Cells, Phys. Rev. Lett. **110**, 160802 (2013).

[21] J. F. Barry, M. J. Turner, J. M. Schloss, D. R. Glenn, Y. Song, M. D. Lukin, H. Park, and R. L. Walsworth, Optical magnetic detection of single-neuron action potentials using quantum defects in diamond, Proc. Natl. Acad. Sci. USA **113**, 14133 (2016).

[22] V. Gerginov, F. C. S. da Silva, and D. Howe, Prospects for magnetic field communications and location using quantum sensors, Rev. Sci. Instrum. **88**, 125005 (2017).

[23] UK Defence Security Accelerator, Countering Drones: Finding and neutralising small uas threats, April 2019, https://www.gov.uk/government/publications/countering-drones-finding-and-neutralising-small-uas-threats.

[24] I. M. Savukov, S. J. Seltzer, M. V. Romalis, and K. L. Sauer, Tunable Atomic Magnetometer for Detection of Radio-Frequency Magnetic Fields, Phys. Rev. Lett. **95**, 063004 (2005).

[25] M. P. Ledbetter, V. M. Acosta, S. M. Rochester, D. Budker, S. Pustelny, and V. V. Yashchuk, Detection of radio-frequency magnetic fields using nonlinear magneto-optical rotation, Phys. Rev. A **75**, 023405 (2007).

[26] W. Chalupczak, R. M. Godun, S. Pustelny, and W. Gawlik, Room temperature femtotesla radio-frequency atomic magnetometer, Appl. Phys. Lett. **100**, 242401 (2012).

[27] S.-K. Lee, K. L. Sauer, S. J. Seltzer, O. Alem, and M. V. Romalis, Subfemtotesla radio-frequency atomic magnetometer for detection of nuclear quadrupole resonance, Appl. Phys. Lett. **89**, 214106 (2006).

[28] I. Savukov, T. Karaulanov, and M. G. Boshier, Ultra-sensitive high-density Rb-87 radio-frequency magnetometer, Appl. Phys. Lett. **104**, 023504 (2014).

[29] C. Deans, L. Marmugi, and F. Renzoni, Sub-picotesla widely tunable atomic magnetometer operating at room-temperature in unshielded environments, Rev. Sci. Instrum. **89**, 083111 (2018).

[30] J. Rajroop, Radio-frequency atomic magnetometers: An analysis of interrogation regimes, Ph.D. thesis, University College London, 2018.

[31] W. Gawlik and S. Pustelny, in *New Trends in Quantum Coherence and Nonlinear Optics* (Nova Publishers, Hauppauge, New York, 2009), pp. 47–84.

[32] W. Gawlik and S. Pustelny, in *High Sensitivity Magnetometers* (Springer, Berlin, 2017), pp. 425–450.

[33] D. Budker and M. Romalis, Optical magnetometry, Nat. Phys. **3**, 227 (2007).

[34] D. Budker and D. F. Jackson Kimball, *Optical Magnetometry* (Cambridge University Press, Cambridge, UK, 2013).

[35] R. Jiménez-Martínez, W. C. Griffith, S. Knappe, J. Kitching, and M. Prouty, High-bandwidth optical magnetometer, J. Opt. Soc. Am. B **29**, 3398 (2012).

[36] S. Pustelny, D. F. Jackson Kimball, S. M. Rochester, V. V. Yashchuk, and D. Budker, Influence of magnetic-field inhomogeneity on nonlinear magneto-optical resonances, Phys. Rev. A **74**, 063406 (2006).

[37] A. Z. Chaudhry, Utilizing nitrogen-vacancy centers to measure oscillating magnetic fields, Phys. Rev. A **90**, 042104 (2014).

[38] I. Baumgart, J.-M. Cai, A. Retzker, M. B. Plenio, and C. Wunderlich, Ultrasensitive Magnetometer Using a Single Atom, Phys. Rev. Lett. **116**, 240801 (2016).

[39] R. P. Anderson, M. J. Kewming, and L. D. Turner, Continuously observing a dynamically decoupled spin-1 quantum gas, Phys. Rev. A **97**, 013408 (2018).

[40] A. Cooper, E. Magesan, H. N. Yum, and P. Cappellaro, Time-resolved magnetic sensing with electronic spins in diamond, Nat. Commun. **5**, 3141 (2014).

[41] G. Puentes, G. Waldherr, P. Neumann, G. Balasubramanian, and J. Wrachtrup, Efficient route to high-bandwidth nanoscale magnetometry using single spins in diamond, Sci. Rep. **4**, 4677 (2014).

[42] L. E. de Clercq, R. Oswald, C. Flühmann, B. Keitch, D. Kienzler, H.-Y. Lo, M. Marinelli, D. Nadlinger, V. Negnevitsky, and J. P. Home, Estimation of a general time-dependent Hamiltonian for a single qubit, Nat. Commun. **7**, 11218 (2016).

[43] J. M. Boss, K. S. Cujia, J. Zopes, and C. L. Degen, Quantum sensing with arbitrary frequency resolution, Science **356**, 837 (2017).

[44] S. Schmitt, T. Gefen, F. M. Stürner, T. Unden, G. Wolff, C. Müller, J. Scheuer, B. Naydenov, M. Markham, S. Pezzagna, J. Meijer, I. Schwarz, M. Plenio, A. Retzker, L. P. McGuinness, and F. Jelezko, Submillihertz magnetic spectroscopy performed with a nanoscale quantum sensor, Science **356**, 832 (2017).

[45] D. R. Glenn, D. B. Bucher, J. Lee, M. D. Lukin, H. Park, and R. L. Walsworth, High-resolution magnetic resonance spectroscopy using a solid-state spin sensor, Nature (London) **555**, 351 (2018).

[46] F. Martin Ciurana, G. Colangelo, L. Slodička, R. J. Sewell, and M. W. Mitchell, Entanglement-Enhanced Radio-Frequency Field Detection and Waveform Sensing, Phys. Rev. Lett. **119**, 043603 (2017).

[47] R. Jiménez-Martínez, J. Kołodyński, C. Troullinou, V. G. Lucivero, J. Kong, and M. W. Mitchell, Signal Tracking Beyond the Time Resolution of an Atomic Sensor by Kalman Filtering, Phys. Rev. Lett. **120**, 040503 (2018).

[48] M. Jasperse, M. J. Kewming, S. N. Fischer, P. Pakkiam, R. P. Anderson, and L. D. Turner, Continuous Faraday measurement of spin precession without light shifts, Phys. Rev. A **96**, 063402 (2017).

[49] G. Bison, V. Bondar, P. Schmidt-Wellenburg, A. Schnabel, and J. Voigt, Sensitive and stable vector magnetometer for operation in zero and finite fields, Opt. Express **26**, 17350 (2018).

[50] D. Hunter, R. Jiménez-Martínez, J. Herbsommer, S. Ramaswamy, W. Li, and E. Riis, Waveform reconstruction with a Cs based free-induction-decay magnetometer, Opt. Express **26**, 30523 (2018).

[51] S. Afach, G. Ban, G. Bison, K. Bodek, Z. Chowdhuri, Z. D. Grujić, L. Hayen, V. Hélaine, M. Kasprzak, K. Kirch *et al.*, Highly stable atomic vector magnetometer based on free spin precession, Opt. Express **23**, 22108 (2015).

[52] Z. D. Grujić, P. A. Koss, G. Bison, and A. Weis, A sensitive and accurate atomic magnetometer based on free spin precession, Eur. Phys. J. D **69**, 135 (2015).

[53] D. Hunter, S. Piccolomo, J. D. Pritchard, N. L. Brockie, T. E. Dyer, and E. Riis, Free-Induction-Decay Magnetometer Based on a Microfabricated Cs Vapor Cell, Phys. Rev. Appl. **10**, 014002 (2018).

[54] P.-X. Miao, W.-Q. Zheng, S.-Y. Yang, B. Wu, B. Cheng, J.-H. Tu, H.-L. Ke, W. Yang, J. Wang, J.-Z. Cui, and L. Qiang, Wide-range and self-locking atomic magnetometer based on free spin precession, J. Opt. Soc. Am. B **36**, 819 (2019).

[55] N. Wilson, P. Light, A. Luiten, and C. Perrella, Ultrastable Optical Magnetometry, Phys. Rev. Appl. **11**, 044034 (2019).

[56] E. B. Alexandrov, M. Auzinsh, D. Budker, D. F. Jackson Kimball, S. M. Rochester, and V. V. Yashchuk, Dynamic effects in nonlinear magneto-optics of atoms and molecules, J. Opt. Soc. Am. B **22**, 7 (2005).

[57] Although we establish an atomic alignment [65] and subsequently track its precession, the technique described here would be equally applicable to extracting the instantaneous phase of an atomic orientation (oscillating about an $f_c = f_L$ carrier). We note that the use of an amplitude-modulated pump beam while establishing coherence and the relative orientation of the pump and probe beams are not critical elements of the technique reported here.

[58] N. Wilson, N. Bourbeau Hébert, C. Perrella, P. Light, J. Genest, S. Pustelny, and A. Luiten, Simultaneous Observation of Nonlinear Magneto-Optical Rotation in the Temporal and Spectral Domains with an Electro-Optic Frequency Comb, Phys. Rev. Appl. **10**, 034012 (2018).

[59] B. Boashash, Estimating and interpreting the instantaneous frequency of a signal. I. Fundamentals, Proc. IEEE **80**, 520 (1992).

[60] A. E. Barnes, The calculation of instantaneous frequency and instantaneous bandwidth, Geophysics **57**, 1520 (1992).

[61] D. A. Steck, Rubidium 87 D line data, rev. 2.1.5, http://steck.us/alkalidata.

[62] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory* (Prentice Hall, New York, 1993).

[63] R. Vijay, C. Macklin, D. H. Slichter, S. J. Weber, K. W. Murch, R. Naik, A. N. Korotkov, and I. Siddiqi, Stabilizing Rabi oscillations in a superconducting qubit using quantum feedback, Nature (London) **490**, 77 (2012).

[64] H. J. Korsch, A. Klumpp, and D. Witthaut, On two-dimensional bessel functions, J. Phys. A: Math. Gen. **39**, 14947 (2006).

[65] D. Budker, D. F. Jackson Kimball, S. M. Rochester, V. V. Yashchuk, and M. Zolotorev, Sensitive magnetometry based on nonlinear magneto-optical rotation, Phys. Rev. A **62**, 043403 (2000).

Conclusion and Future Work

The goal of this project was to develop an optical magnetometer for future applications in magnetic anomaly detection and medical diagnostics and imaging. In order to be suitable for these applications, the system must be: reliable, intrinsically accurate, able to operate at room temperature, highly sensitive over a broad range of Fourier frequencies, able to operate over a large bandwidth (in terms of *both* amplitude and frequency response), high dynamic range, absent of directional dead zones, low in heading errors, and able to be miniaturised. Though this is a long list of demanding criteria, *most* of them have been satisfied and demonstrated throughout the course of this thesis, with some necessitating the development of a novel measurement technique in order to be achieved.

A sensitive optical magnetometer based on nonlinear magneto-optical rotation — a process described theoretically in Sec. 3.6, Sec. 3.7 and Ch. 4 — was constructed, using experimental components described in Ch. 5. Chapter 6 saw the introduction of concepts such as noise characterisation, performance metrics, fundamental sensitivity, parameter optimisation and bandwidth. Each of these topics were used extensively in Ch. 10 to characterise and optimise the magnetometer's performance, resulting in a continuous measurement over 26 h which demonstrated state-of-the-art performance. This measurement satisfied a number of the strict requirements pertaining to the aforementioned applications, namely: reliability, intrinsic accuracy, room-temperature operability, high dynamic range, and high sensitivity over a broad range of Fourier frequencies. However, one of the criteria which was *not* met during this measurement was the need for a large amplitude and frequency response. This subsequently sparked further research into overcoming these limitations.

The bandwidth limitations of the device — along with their origin — were described towards the end of Ch. 6. This lead to the development of a novel measurement technique intended overcome these limitations, which is described in detail in Ch. 7. This new measurement technique relies upon observations made in the temporal domain, rather than the frequency domain. Specifically, instead of synchronous optical pumping of a magneto-optical resonance — which is the conventional measurement technique employed in most optical magnetometers, and that subsequently gives rise to amplitude and frequency-response limitations — the instantaneous phase of freely precessing atomic spins is monitored instead. In doing so, it is possible to enhance the amplitude and frequency response of the magnetometer by *many* orders of magnitude. An experimental demonstration of this novel technique is presented in Ch. 11, which shows an enhancement by two-to-three orders of magnitude — in comparison with the conventional technique — for both the amplitude *and* frequency response.

While drastic improvements to the magnetometer have been accomplished over the course of this endeavour, there is *always* additional work which can be done in order to further refine the system. There are a myriad of improvements and future developments which follow immediately from the research undertaken thus far, and these can be separated into two distinct categories: performance improvements, and application of the system to real-world scenarios. In relation to the former, there are two sub-streams: fundamental sensitivity improvements, and technical noise reduction.

Despite the magnetometer currently being limited by photon shot noise at high Fourier frequencies, and closely followed by spin-projection noise — both of which are fundamental in origin — there *are* avenues to explore which may result in improvements to *both* of these sources. As discussed in Sec. 6.4.2, the classical photon shot-noise limit is determined, in part, by the slope of the magneto-optical resonance. Assuming that the resonance *width* remains constant, the resonance slope can be improved by increasing the *amplitude* of the resonance. This can be achieved in a number of ways, though perhaps the most obvious two methods are: via heating the vapour cell (to increase vapour number density), and by multi-passing the probe beam through the cell. Both of these techniques would yield increased optical depth, and subsequently increase the resonance amplitude. Improvements via these two methods do have an upper bound though, as *too much* optical depth results in *reduced* performance. Furthermore, processes which give rise to increased optical depth may also *broaden* the resonance and reduce performance in that way. Balance is key here, and further exploration of parameter space will be required to find favourable operating configurations.

As discussed in Sec. 6.4.1, the magnetic-sensitivity limit due to spin-projection noise is inversely proportional to the square root of the number of interrogated spins. Therefore, by increasing the number of interacting spins — either by increasing the vapour number density or by multi-passing the probe beam through the cell — the standard quantum limit is subsequently improved. The two techniques discussed previously will therefore improve *both* the classical *and* quantum limits of the device simultaneously. Furthermore, polarisation squeezing can also reduce the spin-projection-noise limit to below that of the standard quantum limit [323–325].

Technical noise sources presently limiting the performance of the magnetometer over long timescales arise due to thermal effects. Specifically, these are: thermal fluctuations in the supply current which generates the dc magnetic field, thermal magnetisation noise of the μ-metal shields, and Johnson noise in the power-supply circuit. None of these limitations affect a real-world, field-deployed measurement (i.e. one that is performed *outside* of the shields); however, for applications in which a measurement target can be introduced within the shielded volume, these technical noise sources will place a lower bound on the sensitivity of the device.

Thermal noise sources are, unfortunately, notoriously difficult to control. In a perfect world, thermal fluctuations in the supply current can in principle be minimised through the use of active temperature stabilisation — via a Peltier thermoelectric device under control of a proportional-integral-derivative feedback loop, for example. On the other hand, ther-

mal magnetisation noise of the μ-metal shields can be reduced by placing a layer of ferrite between the vapour cell and the innermost layer of μ-metal [340]. Lastly, Johnson noise in the power-supply circuit can in principle be reduced via active cooling, though this may be prohibited by technical issues in practice.

In applications such as magnetic anomaly detection — where the device would typically be unshielded and operating in Earth's ambient magnetic field — none of the aforementioned technical noise sources would be present. Instead, the measurement would most likely be dominated by the natural fluctuations in Earth's magnetic field. In this case, it would be beneficial to implement an *array* of magnetometers and measure difference signals between them (i.e. gradiometry), in order to exploit the large-scale uniformity of Earth's magnetic field and attenuate ambient magnetic-field noise via common-mode rejection.

Also of importance in the context of magnetic anomaly detection are directional sensitivity and heading errors. As it stands, conventional NMOR magnetometers in the Faraday geometry — such as the one described in this thesis — exhibit directional sensitivity (and dead zones) to the dc magnetic field [259]; and, at geophysical fields where the nonlinear Zeeman effect becomes non-negligible, they also suffer from heading errors [255,256,258–260]. These two phenomenon were not explicitly investigated in this thesis, and further experimentation will be required in order to find ways in which to mitigate their effects.

In both of the flagged applications, miniaturisation of the detector is paramount. Fortunately, given that the 'sensor head' of the detector is entirely optical it is, at least in principle, possible to *significantly* reduce its physical dimensions to that of a portable device. Whether the sensor head is intended to operate at a fixed location in the field, or as a medical device in a clinical setting, optical fibres can be used to send and retrieve optical signals to and from the magnetometer.

Although not explicitly discussed in this thesis, a rudimentary fibre-coupled gradiometer was constructed using two vapour cells. Preliminary measurements demonstrated unshielded gradiometry in Earth's field with some degree of common-mode rejection. This augurs well for the use of the magnetometer, or an array thereof, in magnetic anomaly detection or other field-based applications. Further work is currently underway which involves the construction of a miniaturised sensor-head prototype for use in field-deployed magnetic anomaly detection. The opportunity to develop a functional, field-deployed sensor is an exciting prospect — albeit a challenging one.

Part V

Appendix

Density Matrix Derivations

The purpose of this appendix is to provide detailed mathematical derivations of the formulae presented in the discussion of light-atom interactions in Ch. 4. Derivations given here closely follow those presented in Ref. [242], though are repeated here in standard international units.

A.1 The Liouville Equation

Just as the time-dependent Schrödinger equation calculates the temporal evolution of a wavefunction, the temporal evolution of the density matrix can also be calculated given knowledge of its initial conditions. The equation which governs density-matrix dynamics is known as the Liouville equation, and is derived by considering the time derivative of Eq. (4.5), which yields

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} = \frac{\mathrm{d}}{\mathrm{d}t} \sum_j P_j \left|\psi_j\right\rangle \left\langle\psi_j\right| . \tag{A.1}$$

Using the product rule to evaluate the derivative in Eq. (A.1) yields

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} = \sum_j P_j \left[\left(\frac{\mathrm{d}}{\mathrm{d}t} \left|\psi_k\right\rangle \right) \left\langle\psi_j\right| + \left|\psi_j\right\rangle \left(\frac{\mathrm{d}}{\mathrm{d}t} \left\langle\psi_j\right| \right) \right] . \tag{A.2}$$

Equation (A.2) contains temporal derivatives of the wavefunctions $\left|\psi\right\rangle$, which using Eq. (4.1), can be evaluated using the Schrödinger equation, subsequently yielding

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} = \sum_j P_j \left[-\frac{i}{\hbar} \mathscr{H} \left|\psi_j\right\rangle \left\langle\psi_j\right| - \frac{i}{\hbar} \left|\psi_j\right\rangle \left\langle\psi_j\right| \mathscr{H} \right] . \tag{A.3}$$

Subsequent algebraic manipulation, and the use of Eq. (4.5), results in the Liouville equation:

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} = -\frac{i}{\hbar} \left[\mathscr{H}, \rho \right] . \tag{A.4}$$

A.2 The Spherical Basis

Rather than working in the Cartesian basis, it is quite often simpler to perform calculations in the spherical basis — especially when using circularly polarised light[1]. The spherical

---

[1]Given that linearly polarised light is a superposition of left- and right-circularly polarised components, it is still far simpler to perform calculations in the spherical basis *even* when using linearly polarised light.

basis vectors are defined in terms of the Cartesian basis vectors via

$$\hat{\epsilon}_1 = -\frac{1}{\sqrt{2}}(\hat{x} + i\hat{y}), \tag{A.5}$$

$$\hat{\epsilon}_0 = \hat{z}, \tag{A.6}$$

$$\hat{\epsilon}_{-1} = \frac{1}{\sqrt{2}}(\hat{x} - i\hat{y}), \tag{A.7}$$

where $\hat{\epsilon}_0$ corresponds to linear polarisation along the $\hat{z}$ direction, and $\hat{\epsilon}_1$ and $\hat{\epsilon}_{-1}$ correspond to left- and right-circularly polarised light, respectively[2]. Conveniently, there is a unitary transform which maps vectors from the Cartesian basis to the spherical basis, and vice versa, according to $\epsilon = U_t r$, where $\epsilon = (\epsilon_1, \epsilon_0, \epsilon_{-1})$ and $r = (x, y, z)$. This unitary transformation matrix is given by

$$U_t = \begin{pmatrix} -\frac{1}{\sqrt{2}} & -\frac{i}{\sqrt{2}} & 0 \\ 0 & 0 & 1 \\ \frac{1}{\sqrt{2}} & -\frac{i}{\sqrt{2}} & 0 \end{pmatrix}. \tag{A.8}$$

A.3 The Wigner-Eckart Theorem

The Wigner-Eckart theorem states that matrix elements of spherical-tensor operators can be written as the product of two factors: one that is *independent* of angular momentum orientation, and the other a Clebsch-Gordan coefficient [224]:

$$\langle \xi J m_J | \mathcal{T}_q^\kappa | \xi' J' m_J' \rangle = (-1)^{2\kappa} \langle J m_J \kappa q | J' m_J' \rangle \langle \xi J \| \mathcal{T}^\kappa \| \xi' J' \rangle, \tag{A.9}$$

where $\mathcal{T}^\kappa$ is a spherical tensor of rank κ, $\langle J m_J \kappa q | J' m_J' \rangle$ is a Clebsch-Gordan coefficient, and $\langle \xi J \| \mathcal{T}^\kappa \| \xi' J' \rangle$ is the reduced matrix element. Given that the dipole operator when evaluated in the spherical basis is a spherical tensor, the Wigner-Eckart theorem can be used to calculate its matrix elements. In doing this, however, it should be noted that reduced matrix elements with different ordering of states are related by [224]

$$\langle \xi' J' \| \mathcal{T}^\kappa \| \xi J \rangle = (-1)^{J'-J} \sqrt{\frac{2J+1}{2J'+1}} \langle \xi J \| \mathcal{T}^\kappa \| \xi' J' \rangle^*. \tag{A.10}$$

Given the reduced dipole matrix element is real, $\langle \xi J \| d \| \xi' J' \rangle^* = \langle \xi J \| d \| \xi' J' \rangle$. Since the dipole operator is proportional to the position operator, it transforms as a tensor of rank $\kappa = 1$. The Wigner-Eckart theorem for the dipole operator is therefore given by [224]

---

[2]This is easily remembered by considering that a left-circularly polarised photon has a spin of $+\hbar$, while a right-circularly polarised photon has a spin of $-\hbar$. This is true using the convention that the light is viewed from the frame of reference of the *receiver*, looking *anti-parallel* to the propagation vector of light.

$$\langle \xi J m_J | d_q | \xi' J' m_J' \rangle = (-1)^{J'-J+m_J'-m_J} \sqrt{\frac{2J+1}{2J'+1}} \langle J m_J 1 \, (-q) | J' m_J' \rangle \langle \xi J \| d_q \| \xi' J' \rangle . \quad (A.11)$$

A.4 The Rotating-Wave Approximation

As detailed in Sec. 4.5, the electric-field Hamiltonian (and therefore the total Hamiltonian — cf. Sec. 4.4.4, specifically Eq. 4.22) contains temporal dependence at the optical frequency. Solving the Liouville equation using these terms would require *incredibly* high temporal resolution, resulting in significant computational expense. However, by changing the frame of reference from the laboratory frame, to that of the light (i.e. the rotating frame), it is possible to dispense of the high-frequency oscillatory terms. This change of reference frame is described by a unitary transform, which for the $F = 1 \to F' = 0$ transition is given by the following diagonal matrix:

$$U_t = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{-i\omega t} \end{pmatrix} . \quad (A.12)$$

Given that the reference frame is being rotated, the unitary transformation is applied to the basis, such that its effect on a state vector is given by

$$|\widetilde{\psi}\rangle = U_t^\dagger |\psi\rangle . \quad (A.13)$$

Not only must the state vector be transformed, but so too must the Hamiltonian. The *effective* Hamiltonian can be derived starting with the time-dependent Schrödinger equation in the laboratory frame. In bra-ket notation, the time-dependent Schrödinger equation, Eq. (4.1), is given by

$$i\hbar \frac{\partial}{\partial t} |\psi\rangle = \mathscr{H} |\psi\rangle . \quad (A.14)$$

Multiplying both sides by $U_t^\dagger$, and inserting the identity operator, $U_t U_t^\dagger$, in front of the state ket on each sides, yields

$$i\hbar U_t^\dagger \frac{\partial}{\partial t} \left(U_t U_t^\dagger |\psi\rangle \right) = U_t^\dagger \mathscr{H} \left(U_t U_t^\dagger |\psi\rangle \right) . \quad (A.15)$$

Acting on the state vector, $|\psi\rangle$, with the unitary matrix transforms the state vector to the rotating reference frame, $|\widetilde{\psi}\rangle$, in accordance with Eq. (A.13). By using this transformation, Eq. (A.15) simplifies to

$$i\hbar U_t^\dagger \frac{\partial}{\partial t} \left(U_t |\widetilde{\psi}\rangle \right) = U_t^\dagger \mathscr{H} \left(U_t |\widetilde{\psi}\rangle \right) . \quad (A.16)$$

Evaluating the partial derivative on the left-hand side of Eq. (A.16) using the product rule,

and re-arranging, yields the following:

$$i\hbar \frac{\partial}{\partial t}|\widetilde{\psi}\rangle = \left(U_t^\dagger \mathscr{H} U_t - i\hbar U_t^\dagger \frac{\partial U_t}{\partial t} \right) |\widetilde{\psi}\rangle \tag{A.17}$$

$$= \widetilde{\mathscr{H}}|\widetilde{\psi}\rangle . \tag{A.18}$$

The form of the transformed Hamiltonian, $\widetilde{\mathscr{H}}$, has therefore been derived in Eq. (A.17). Calculating the effective Hamiltonian for the total Hamiltonian derived in Sec. 4.4 yields

$$\widetilde{\mathscr{H}} = \begin{pmatrix} \hbar\Omega_L & 0 & 0 & \frac{\hbar\Omega_R}{2\sqrt{2}}\left(1+e^{-2i\omega t}\right) \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -\hbar\Omega_L & -\frac{\hbar\Omega_R}{2\sqrt{2}}\left(1+e^{-2i\omega t}\right) \\ \frac{\hbar\Omega_R}{2\sqrt{2}}\left(1+e^{-2i\omega t}\right) & 0 & -\frac{\hbar\Omega_R}{2\sqrt{2}}\left(1+e^{-2i\omega t}\right) & -\hbar\Delta\omega \end{pmatrix}, \tag{A.19}$$

where $\Delta\omega = \omega_0 - \omega$ is the optical detuning. Notice that the oscillatory terms which were proportional to $\cos(\omega t)$ are now proportional to $(1 + \exp[-2i\omega t])$. This means that there is a now a static (temporally independent) term, and all of the temporal dependence now appears at *twice* the optical frequency. The rotating-wave approximation is completed by *discarding* the terms which oscillate at twice the optical frequency, leaving a purely static Hamiltonian:

$$\widetilde{\mathscr{H}} = \begin{pmatrix} \hbar\Omega_L & 0 & 0 & \frac{\hbar\Omega_R}{2\sqrt{2}} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -\hbar\Omega_L & -\frac{\hbar\Omega_R}{2\sqrt{2}} \\ \frac{\hbar\Omega_R}{2\sqrt{2}} & 0 & -\frac{\hbar\Omega_R}{2\sqrt{2}} & -\hbar\Delta\omega \end{pmatrix}. \tag{A.20}$$

A.5   The Repopulation Matrix

As discussed in Sec. 4.6.2, repopulation of ground-state populations ρ_{ii} and coherences ρ_{ij} must be incorporated phenomenologically via the repopulation matrix Λ. This repopulation matrix is comprised of two separate matrices: one describing transit repopulation (i.e. new atoms entering the beam), and the other describing spontaneous decay from the excited states to the ground states. Transit repopulation is described by a diagonal matrix, with each ground-state population ρ_{ii} increasing proportionally to the transit rate γ_t (cf. Sec. 4.6.2). Spontaneous decay, on the other hand, is described by a spontaneous-emission operator.

Given a density matrix ρ_{rs} in an excited state, the rate of change of the ground state density matrix element ρ_{mn}, due to spontaneous emission, is given by [242,283]

$$\frac{d\rho_{mn}}{dt} = \sum_{r,s} \frac{\omega_{rm}^3}{3\pi\epsilon_0\hbar c^3} \boldsymbol{d}_{mr} \cdot \boldsymbol{d}_{sn}\rho_{rs} = \sum_{r,s} F_{mn}^{sr}\rho_{rs}, \tag{A.21}$$

where F_{mn}^{sr} is the spontaneous-emission operator. The sum over the upper pair of indices

enables Eq. (A.21) to be written in a compact form by tracing over the excited-state density matrix ρ':

$$\frac{\mathrm{d}\rho_{mn}}{\mathrm{d}t} = \mathrm{tr}\left(\rho' F_{mn}\right) . \tag{A.22}$$

By explicitly evaluating the dipole operators in Eq. (A.21), it is possible to simplify the expression for the spontaneous-emission operator. To begin with, it is necessary to evaluate the inner product of the dipole operators, in the spherical basis, using the following identity [165]:

$$\boldsymbol{A} \cdot \boldsymbol{B} = \sum_q (-1)^q A_q B_{-q} , \tag{A.23}$$

where the sum over q runs over the values $\{-1, 0, 1\}$. Using Eq. (A.23) to evaluate the inner product in Eq. (A.21) yields

$$
\begin{aligned}
F_{m_1 m_2}^{m'_1 m'_2} &= \frac{\omega^3}{3\pi\epsilon_0\hbar c^3} \boldsymbol{d}_{m_1 m'_2} \cdot \boldsymbol{d}_{m'_1 m_2} \\
&= \frac{\omega^3}{3\pi\epsilon_0\hbar c^3} \sum_q (-1)^q \langle Jm_1 | d_q | J'm'_2 \rangle \langle J'm'_1 | d_{-q} | Jm_2 \rangle \\
&= \frac{\omega^3}{3\pi\epsilon_0\hbar c^3} \sum_q \langle Jm_1 | d_q | J'm'_2 \rangle \langle J'm'_1 | (-1)^q d_{-q} | Jm_2 \rangle .
\end{aligned}
$$

Further simplification is possible by using the relation $(-1)^q d_{-q} = d_q^\dagger$. This can then be used to switch the order of states via the relation $\langle\psi| A^\dagger |\phi\rangle = \langle\phi| A |\psi\rangle^*$:

$$
\begin{aligned}
F_{m_1 m_2}^{m'_1 m'_2} &= \frac{\omega^3}{3\pi\epsilon_0\hbar c^3} \sum_q \langle Jm_1 | d_q | J'm'_2 \rangle \langle J'm'_1 | d_q^\dagger | Jm_2 \rangle \\
&= \frac{\omega^3}{3\pi\epsilon_0\hbar c^3} \sum_q \langle Jm_1 | d_q | J'm'_2 \rangle \langle Jm_2 | d_q | J'm'_1 \rangle^* .
\end{aligned}
$$

At this point, it is now possible to invoke the Wigner-Eckart theorem, Eq. (A.9), to write the dipole-matrix elements in terms of the reduced dipole matrix element, and their respective Clebsch-Gordan coefficients:

$$
\begin{aligned}
F_{m_1 m_2}^{m'_1 m'_2} &= \frac{\omega^3}{3\pi\epsilon_0\hbar c^3} \sum_q (-1)^{J'-J+m'_J-m_J} \sqrt{\frac{2J+1}{2J'+1}} \langle Jm_1 1 (-q) | J'm'_2 \rangle \langle J \| d_q \| J' \rangle \tag{A.24} \\
&\quad \times \left((-1)^{J'-J+m'_J-m_J} \sqrt{\frac{2J+1}{2J'+1}} \langle Jm_2 1 (-q) | J'm'_1 \rangle \langle J \| d_q \| J' \rangle \right)^* \\
&= \frac{\omega^3}{3\pi\epsilon_0\hbar c^3} \frac{2J+1}{2J'+1} |\langle J\|d\|J'\rangle|^2 \sum_q \langle Jm_1 1 (-q) | J'm'_2 \rangle \langle Jm_2 1 (-q) | J'm'_1 \rangle^* , \tag{A.25}
\end{aligned}
$$

where we have made use of $(-1)^{J'-J+m'_J-m_J} = 1$ for the two circularly polarised components which couple the $F = 1 \to F' = 0$ transition.

The pre-factor in Eq. (A.25) simplifies to the spontaneous-decay rate through the relation [224]

$$\frac{\omega^3}{3\pi\epsilon_0 \hbar c^3} \frac{2J+1}{2J'+1} |\langle J\|d\|J'\rangle|^2 = \Gamma_{\rm d}.$$
(A.26)

Making the above substitution into Eq. (A.25), yields an expression for the spontaneous-emission operator which depends *only* upon two Clebsch-Gordan coefficients and the spontaneous-decay rate:

$$F_{m_1 m_2}^{m_1' m_2'} = \Gamma_{\rm d} \sum_q \langle Jm_1 1 \,(-q)|J'm_2'\rangle \langle Jm_2 1 \,(-q)|J'm_1'\rangle^* .$$
(A.27)

Equation (A.27) is readily applicable to hyperfine states where F and m_F are the relevant quantum numbers, simply by substituting $J \to F$ and $m_J \to m_F$. Using Eq. (A.27) to calculate the matrix elements of the spontaneous-emission operator, for the $F = 1 \to F' = 0$ transition, yields the following:

$$F = \begin{pmatrix} F_{1,1} & F_{1,0} & F_{1,-1} \\ F_{0,1} & F_{0,0} & F_{0,-1} \\ F_{-1,1} & F_{-1,0} & F_{-1,-1} \end{pmatrix}$$
(A.28)

$$= \Gamma_{\rm d} \sum_q \begin{pmatrix} \langle 111q|00\rangle \langle 111q|00\rangle^* & \langle 111q|00\rangle \langle 101q|00\rangle^* & \langle 111q|00\rangle \langle 1\,(-1)\,1q|00\rangle^* \\ \langle 101q|00\rangle \langle 111q|00\rangle^* & \langle 101q|00\rangle \langle 101q|00\rangle^* & \langle 101q|00\rangle \langle 1\,(-1)\,1q|00\rangle^* \\ \langle 1\,(-1)\,1q|00\rangle \langle 111q|00\rangle^* & \langle 1\,(-1)\,1q|00\rangle \langle 101q|00\rangle^* & \langle 1\,(-1)\,1q|00\rangle \langle 1\,(-1)\,1q|00\rangle^* \end{pmatrix}$$
(A.29)

$$= \Gamma_{\rm d} \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}.$$
(A.30)

Note that in general, each element $F_{m_1 m_2}$ is typically an $n \times n$ matrix, where n is the number of excited states. However, in the case of the $F = 1 \to F' = 0$ transition, there is only one excited state, hence $n = 1$ and therefore each element is a 1×1 matrix (i.e. a scalar).

The rate of change of the ground-state density-matrix elements due to spontaneous emission can now be calculated using Eq. (A.22); except that in this case, the trace reduces to a simple multiplication due to there only being one excited state.

A.6 Observables

Consider the electric field $\mathcal{E}(r, t)$ of a plane wave of angular frequency ω and wavevector k, which is described using the ϕ-ϵ parametrisation [242,283]:

$$\mathcal{E}(r, t) = \text{Re}\left\{\mathcal{E}_0 e^{i(k \cdot r - \omega t + \varphi)}\left[(\cos\phi\cos\epsilon - i\sin\phi\cos\epsilon)\,\hat{\mathbf{e}}_1 + (\sin\phi\cos\epsilon + i\cos\phi\sin\epsilon)\,\hat{\mathbf{e}}_2\right]\right\},$$
(A.31)

where $\hat{\mathbf{e}}_1$ and $\hat{\mathbf{e}}_2 = \hat{k} \times \hat{\mathbf{e}}_1$ are orthogonal unit vectors, $\mathcal{E}_0$ is the electric-field amplitude, φ is an overall phase, ϕ is the polarisation angle (azimuth) relative to the $\hat{\mathbf{e}}_1$ axis, and ϵ is the ellipticity. The propagation of this electric field through a medium is governed by the wave equation, which is derived from Maxwell's equations:

$$\nabla \cdot \mathbf{D} = \rho_f,$$

$$\nabla \cdot \mathbf{B} = 0,$$

$$\nabla \times \mathcal{E} = -\frac{\partial \mathbf{B}}{\partial t},$$

$$\nabla \times \mathbf{H} = \mathbf{J}_f + \frac{\partial \mathbf{D}}{\partial t},$$

where ρ_f is the free-electric-charge density, $\mathbf{J}_f$ is the free-current density, $\mathbf{D} = \epsilon_0\mathcal{E} + \mathbf{P}$ is the displacement field, $\mathbf{H} = \mathbf{B}/\mu_0 - \mathbf{M}$ is the magnetising field, $\mathbf{P}$ is the polarisation field, and $\mathbf{M}$ is the magnetisation field. In the definitions of the displacement and magnetising fields, ϵ_0 and μ_0 are the permittivity and permeability of free space, respectively. Of interest is the case where the medium is non-magnetic (i.e. $\mathbf{M} = 0$) and both the free-electric-charge and free-current densities are zero. This implies that $\mathbf{H} \propto \mathbf{B}$, and Maxwell's equations simplify to

$$\nabla \cdot \mathbf{D} = 0,$$

$$\nabla \cdot \mathbf{B} = 0,$$

$$\nabla \times \mathcal{E} = -\frac{\partial \mathbf{B}}{\partial t},$$
(A.32)

$$\nabla \times \mathbf{B} = \frac{1}{c^2}\frac{\partial \mathcal{E}}{\partial t} + \frac{1}{c^2\epsilon_0}\frac{\partial \mathbf{P}}{\partial t},$$

where $c = (\epsilon_0\mu_0)^{-1/2}$ is the speed of light. To derive the wave equation from Maxwell's equations, the curl of both sides of Eq. (A.32) is calculated, yielding

$$\nabla \times (\nabla \times \boldsymbol{\mathcal{E}}) = \nabla \times -\frac{\partial \boldsymbol{B}}{\partial t}$$

$$= -\frac{\partial}{\partial t} (\nabla \times \boldsymbol{B})$$

$$= -\frac{\partial}{\partial t} \left(\frac{1}{c^2} \frac{\partial \boldsymbol{\mathcal{E}}}{\partial t} + \frac{1}{c^2 \epsilon_0} \frac{\partial \boldsymbol{P}}{\partial t} \right)$$

$$= -\frac{1}{c^2} \frac{\partial^2 \boldsymbol{\mathcal{E}}}{\partial t^2} - \frac{1}{c^2 \epsilon_0} \frac{\partial^2 \boldsymbol{P}}{\partial t^2} .$$

Now, using the trig identity $\nabla \times \nabla \times \boldsymbol{\mathcal{E}} = \nabla (\nabla \cdot \boldsymbol{\mathcal{E}}) - \nabla^2 \boldsymbol{\mathcal{E}}$, the above can be written as

$$\nabla (\nabla \cdot \boldsymbol{\mathcal{E}}) - \nabla^2 \boldsymbol{\mathcal{E}} = -\frac{1}{c^2} \frac{\partial^2 \boldsymbol{\mathcal{E}}}{\partial t^2} - \frac{1}{c^2 \epsilon_0} \frac{\partial^2 \boldsymbol{P}}{\partial t^2} . \tag{A.33}$$

Under the assumption that $\boldsymbol{\mathcal{E}}$ is a transverse plane wave, the field never points along the direction in which it varies, and hence $\nabla \cdot \boldsymbol{\mathcal{E}} = 0$. This simplifies Equation (A.33) to the vector wave equation:

$$\nabla^2 \boldsymbol{\mathcal{E}} = -\frac{1}{c^2} \frac{\partial^2 \boldsymbol{\mathcal{E}}}{\partial t^2} - \frac{1}{c^2 \epsilon_0} \frac{\partial^2 \boldsymbol{P}}{\partial t^2} . \tag{A.34}$$

The polarisation of the medium, $\boldsymbol{P}$, is induced by the electric field and will therefore oscillate at the optical frequency, ω. It is assumed that the polarisation takes the same form as the incident electric field [242,283]:

$$\boldsymbol{P} = \mathrm{Re} \left\{ e^{i(k \cdot r - \omega t + \varphi)} \left[(P_1 - iP_2) \, \hat{\boldsymbol{e}}_1 + (P_3 - iP_4) \, \hat{\boldsymbol{e}}_2 \right] \right\} , \tag{A.35}$$

where P_1 and P_3 are the in-phase components of the polarisation, P_2 and P_4 are the quadrature components of the polarisation, and φ in this case is not an independent parameter — it instead references the overall phase relative to that of the light.

For a plane wave, the light-field parameters vary *only* along the propagation direction, $\hat{k}$. This means that $k \cdot r = |k|l$, where l is the distance along the propagation direction of the light, and hence $\nabla^2 = \partial^2 / \partial l^2$. This simplifies the *vector* wave equation to the *scalar* wave equation:

$$\frac{\partial^2 \boldsymbol{\mathcal{E}}}{\partial l^2} = -\frac{1}{c^2} \frac{\partial^2 \boldsymbol{\mathcal{E}}}{\partial t^2} - \frac{1}{c^2 \epsilon_0} \frac{\partial^2 \boldsymbol{P}}{\partial t^2} . \tag{A.36}$$

Using the general forms of the $\boldsymbol{\mathcal{E}}$ and $\boldsymbol{P}$ fields given by Equations (A.31) and (A.35) respectively, it is possible to simplify the wave equation even further by calculating their time derivatives. Since the time-dependence is held in the exponential term in both the $\boldsymbol{\mathcal{E}}$ and $\boldsymbol{P}$ fields, their time derivatives are given by

$$\frac{\partial^2 \boldsymbol{\mathcal{E}}}{\partial t^2} = -\omega^2 \boldsymbol{\mathcal{E}} ,$$

$$\frac{\partial^2 \boldsymbol{P}}{\partial t^2} = -\omega^2 \boldsymbol{P} . \tag{A.37}$$

This simplifies the scalar wave equation to

$$\frac{\partial^2 \mathcal{E}}{\partial l^2} = k^2 \mathcal{E} + \frac{k^2}{\epsilon_0} \mathbf{P} , \tag{A.38}$$

where $k = \omega/c$ is the wavenumber. We may now solve Equation (A.38) for the changes in the light-field parameters by substituting the forms of $\mathcal{E}$ and $\mathbf{P}$ (Eqs. A.31 and A.35, respectively). The second derivative on the left-hand side of Equation (A.38) results in 36 terms, including those with second-order derivative terms such as $d^2\phi/dl^2$ or $(d\phi/dl)(d\varphi/dl)$. Provided that the fractional changes of the light-field parameters are small over a distance equal to the wavelength of the light, the second-order terms can be neglected [242,283], reducing the second derivative to 10 terms. In the context of the electric-field amplitude, $\mathcal{E}_0$, this is known as the slowly varying envelope approximation. After discarding all terms with second-order derivatives, we obtain the following simultaneous equations [242,283]:

$$-\frac{k^2}{\epsilon_0}\mathrm{P}_1 = 2k\mathcal{E}_0 \left(\cos\phi\sin\epsilon\frac{d\phi}{dl} + \sin\phi\cos\epsilon\frac{d\epsilon}{dl} + \sin\phi\sin\epsilon\frac{1}{\mathcal{E}_0}\frac{d\mathcal{E}_0}{dl} - \cos\phi\cos\epsilon\frac{d\varphi}{dl} \right) ,$$

$$\frac{k^2}{\epsilon_0}\mathrm{P}_2 = 2k\mathcal{E}_0 \left(-\sin\phi\cos\epsilon\frac{d\phi}{dl} - \cos\phi\sin\epsilon\frac{d\epsilon}{dl} + \cos\phi\cos\epsilon\frac{1}{\mathcal{E}_0}\frac{d\mathcal{E}_0}{dl} + \sin\phi\sin\epsilon\frac{d\varphi}{dl} \right) ,$$

$$-\frac{k^2}{\epsilon_0}\mathrm{P}_3 = 2k\mathcal{E}_0 \left(\sin\phi\sin\epsilon\frac{d\phi}{dl} - \cos\phi\cos\epsilon\frac{d\epsilon}{dl} - \cos\phi\sin\epsilon\frac{1}{\mathcal{E}_0}\frac{d\mathcal{E}_0}{dl} - \sin\phi\cos\epsilon\frac{d\varphi}{dl} \right) ,$$

$$\frac{k^2}{\epsilon_0}\mathrm{P}_4 = 2k\mathcal{E}_0 \left(\cos\phi\cos\epsilon\frac{d\phi}{dl} - \sin\phi\sin\epsilon\frac{d\epsilon}{dl} + \sin\phi\cos\epsilon\frac{1}{\mathcal{E}_0}\frac{d\mathcal{E}_0}{dl} - \cos\phi\sin\epsilon\frac{d\varphi}{dl} \right) .$$

Solving these simultaneous equations with Mathematica, and using numerous trigonometric identities, yields the following differential equations for the light-field parameters [242,283]:

$$\frac{1}{\mathcal{E}_0}\frac{d\mathcal{E}_0}{dl} = \frac{k}{2\mathcal{E}_0\epsilon_0} \left[\sin\phi \left(\cos\epsilon\mathrm{P}_4 - \sin\epsilon\mathrm{P}_1 \right) + \cos\phi \left(\cos\epsilon\mathrm{P}_2 + \sin\epsilon\mathrm{P}_3 \right) \right] , \tag{A.39}$$

$$\frac{d\phi}{dl} = \frac{k}{2\mathcal{E}_0\epsilon_0} \sec 2\epsilon \left[\cos\phi \left(\sin\epsilon\mathrm{P}_1 + \cos\epsilon\mathrm{P}_4 \right) - \sin\phi \left(\cos\epsilon\mathrm{P}_2 - \sin\epsilon\mathrm{P}_3 \right) \right] , \tag{A.40}$$

$$\frac{d\epsilon}{dl} = -\frac{k}{2\mathcal{E}_0\epsilon_0} \left[\sin\phi \left(\cos\epsilon\mathrm{P}_1 + \sin\epsilon\mathrm{P}_4 \right) + \cos\phi \left(\sin\epsilon\mathrm{P}_2 - \cos\epsilon\mathrm{P}_3 \right) \right] , \tag{A.41}$$

$$\frac{d\varphi}{dl} = \frac{k}{2\mathcal{E}_0\epsilon_0} \sec 2\epsilon \left[\cos\phi \left(\cos\epsilon\mathrm{P}_1 + \sin\epsilon\mathrm{P}_4 \right) + \sin\phi \left(\cos\epsilon\mathrm{P}_3 - \sin\epsilon\mathrm{P}_2 \right) \right] . \tag{A.42}$$

A.7   Calculating Observables

When modelling an atomic system using the density-matrix formalism, or in fact quantum mechanics in general, not every quantity is able to be directly measured in an experiment. However, quantities that *are* measurable are known as observables, and will have an associated operator, $\mathcal{O}$. Calculating the expectation value of an observable operator $\mathcal{O}$ is done via [225]

$$\langle \mathscr{O} \rangle = \sum_j P_j \langle \psi_j | \mathscr{O} | \psi_j \rangle$$

$$= \sum_j P_j \sum_k \langle \psi_j | \mathscr{O} | k \rangle \langle k | \psi_j \rangle$$

$$= \sum_k \sum_j P_j \langle k | \psi_j \rangle \langle \psi_j | \mathscr{O} | k \rangle$$

$$= \sum_k \langle k | \sum_j P_j | \psi_j \rangle \langle \psi_j | \mathscr{O} | k \rangle$$

$$= \sum_k \langle k | \rho \mathscr{O} | k \rangle$$

$$= \sum_k (\rho \mathscr{O})_{k,k}$$

$$= \mathrm{tr} (\rho \mathscr{O}) \,, \tag{A.43}$$

where we have made use of the fact that the population is normalised:

$$\sum_j P_j | \psi_j \rangle \langle \psi_j | = \sum_j P_j \langle \psi_j | \psi_j \rangle = 1 \,. \tag{A.44}$$

In Sec. A.6 it was shown that the change in light-field parameters, due to the propagation of the light through the atomic vapour, was dependent upon the polarisation components P_1, P_2, P_3 and P_4. In order to relate these polarisation components — and hence the change in light-field parameters — to calculated density-matrix elements, the expectation value of the polarisation operator must be calculated:

$$\mathbf{P} = n_v \langle \mathbf{d} \rangle \,, \tag{A.45}$$

where n_v is the atomic number density, and $\mathbf{d}$ is the dipole operator. Note that the density matrix in this case is that of the laboratory frame, which is related to the density matrix in the rotating frame via Eq. (4.24). Given that the polarisation of the medium is a *vector* quantity, each component of the polarisation operator must be calculated by tracing over the dipole operator for each Cartesian basis vector:

$$\mathbf{P} = n_v \mathrm{tr} (\rho d_x) \, \hat{\mathbf{x}} + n_v \mathrm{tr} (\rho d_y) \, \hat{\mathbf{y}} + n_v \mathrm{tr} (\rho d_z) \, \hat{\mathbf{z}} \,. \tag{A.46}$$

Given that density-matrix modelling is most easily performed in the spherical basis, the unitary transformation given by Eq. (A.8) must be used in order to transform the spherical components of the dipole operator into their corresponding Cartesian components. In the case of the $F = 1 \rightarrow F' = 0$ transition, the expectation values of the Cartesian components of the polarisation operator are given by

$$P_x = \frac{n_v \langle \xi 1 \| d \| \xi' 0' \rangle}{\sqrt{2}} e^{i\omega t} \left(\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'} \right) + \frac{n_v \langle \xi 1 \| d \| \xi' 0' \rangle}{\sqrt{2}} e^{-i\omega t} \left(\widetilde{\rho}_{0',-1} - \widetilde{\rho}_{0',1} \right), \qquad \text{(A.47)}$$

$$P_y = \frac{n_v \langle \xi 1 \| d \| \xi' 0' \rangle}{\sqrt{2}} i e^{i\omega t} \left(-\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'} \right) + \frac{n_v \langle \xi 1 \| d \| \xi' 0' \rangle}{\sqrt{2}} i e^{-i\omega t} \left(\widetilde{\rho}_{0',-1} + \widetilde{\rho}_{0',1} \right), \qquad \text{(A.48)}$$

$$P_z = -n_v \langle \xi 1 \| d \| \xi' 0' \rangle e^{i\omega t} \widetilde{\rho}_{0,0'} - n_v \langle \xi 1 \| d \| \xi' 0' \rangle e^{-i\omega t} \widetilde{\rho}_{0',0}. \qquad \text{(A.49)}$$

Significant simplification can made to these equations by noticing that, for each of the three components, the two terms are *complex conjugates* of each other — i.e. each term takes the form $P_i = (a + ib) + (a - ib)$. This is due to the fact that off-diagonal density-matrix elements are related to each other via $\rho_{m,n} = \rho_{n,m}^*$. This sum can therefore be simplified via $(a + ib) + (a - ib) = 2\,\text{Re}\,(a \pm ib) = 2a$. Performing this simplification therefore yields the following expression for the polarisation of the medium [242,283]:

$$\mathbf{P} = n_v\,\text{Re}\left(\sqrt{2} \langle \xi 1 \| d \| \xi' 0' \rangle e^{-i\omega t} \left[\left(\widetilde{\rho}_{0',-1} - \widetilde{\rho}_{0',1} \right) \hat{\boldsymbol{x}} + i \left(\widetilde{\rho}_{0',-1} + \widetilde{\rho}_{0',1} \right) \hat{\boldsymbol{y}} - \sqrt{2} \widetilde{\rho}_{0',0} \hat{\boldsymbol{z}} \right] \right). \qquad \text{(A.50)}$$

Through complex expansion of the density-matrix elements (i.e. $\widetilde{\rho}_{i,j} = \text{Re}\,\widetilde{\rho}_{i,j} + i\,\text{Im}\,\widetilde{\rho}_{i,j}$) and subsequent comparison of Eq. (A.35) with Eq. (A.50), it is possible to obtain the polarisation components, P_1, P_2, P_3 and P_4[3]. These components are given explicitly by

$$P_1 = \sqrt{2} \langle \xi 1 \| d \| \xi' 0' \rangle n_v\,\text{Re}\left(\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'} \right), \qquad \text{(A.51)}$$

$$P_2 = \sqrt{2} \langle \xi 1 \| d \| \xi' 0' \rangle n_v\,\text{Im}\left(\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'} \right), \qquad \text{(A.52)}$$

$$P_3 = \sqrt{2} \langle \xi 1 \| d \| \xi' 0' \rangle n_v\,\text{Im}\left(\widetilde{\rho}_{-1,0'} + \widetilde{\rho}_{1,0'} \right), \qquad \text{(A.53)}$$

$$P_4 = -\sqrt{2} \langle \xi 1 \| d \| \xi' 0' \rangle n_v\,\text{Re}\left(\widetilde{\rho}_{-1,0'} + \widetilde{\rho}_{1,0'} \right). \qquad \text{(A.54)}$$

These polarisation components, combined with Eqs. (A.39) through (A.42), are used to calculate the observables of the density matrix. Assuming that the incident light is linearly polarised along $\hat{\boldsymbol{x}}$ and propagating along $\hat{\boldsymbol{z}}$, and that the initial values of the rotation angle ϕ and ellipticity ϵ are zero, then the differential equations for the light-field parameters simplify to

---

[3]Note that the assumed form given by Eq. (A.35) contains no z-component, whereas Eq. (A.50) contains a z-component which is proportional to $\widetilde{\rho}_{0',0}$. Physically, one would expect that there should be no z-component of the polarisation since the light is a transverse plane wave and is propagating along $\hat{\boldsymbol{z}}$. After solving the Liouville equation, it turns out that $\widetilde{\rho}_{0',0} = 0$ and hence everything is consistent.

$$\frac{1}{\mathcal{E}_0}\frac{\mathrm{d}\mathcal{E}_0}{\mathrm{d}z} = \frac{\omega}{2\mathcal{E}_0\epsilon_0 c}P_2 , \tag{A.55}$$

$$\frac{\mathrm{d}\phi}{\mathrm{d}z} = \frac{\omega}{2\mathcal{E}_0\epsilon_0 c}P_4 , \tag{A.56}$$

$$\frac{\mathrm{d}\epsilon}{\mathrm{d}z} = \frac{\omega}{2\mathcal{E}_0\epsilon_0 c}P_3 , \tag{A.57}$$

$$\frac{\mathrm{d}\varphi}{\mathrm{d}z} = \frac{\omega}{2\mathcal{E}_0\epsilon_0 c}P_1 . \tag{A.58}$$

By substituting the expressions for the polarisation components into the differential equations for the light-field parameters, the expressions for the observables in terms of density-matrix elements are obtained [242,283]:

$$\frac{1}{\mathcal{E}_0}\frac{\mathrm{d}\mathcal{E}_0}{\mathrm{d}z} = \frac{n_\mathrm{v}\omega\langle\xi 1\|d\|\xi'0'\rangle}{\sqrt{2}\mathcal{E}_0\epsilon_0 c}\,\mathrm{Im}\left(\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'}\right) , \tag{A.59}$$

$$\frac{\mathrm{d}\phi}{\mathrm{d}z} = -\frac{n_\mathrm{v}\omega\langle\xi 1\|d\|\xi'0'\rangle}{\sqrt{2}\mathcal{E}_0\epsilon_0 c}\,\mathrm{Re}\left(\widetilde{\rho}_{-1,0'} + \widetilde{\rho}_{1,0'}\right) , \tag{A.60}$$

$$\frac{\mathrm{d}\epsilon}{\mathrm{d}z} = \frac{n_\mathrm{v}\omega\langle\xi 1\|d\|\xi'0'\rangle}{\sqrt{2}\mathcal{E}_0\epsilon_0 c}\,\mathrm{Im}\left(\widetilde{\rho}_{-1,0'} + \widetilde{\rho}_{1,0'}\right) , \tag{A.61}$$

$$\frac{\mathrm{d}\varphi}{\mathrm{d}z} = \frac{n_\mathrm{v}\omega\langle\xi 1\|d\|\xi'0'\rangle}{\sqrt{2}\mathcal{E}_0\epsilon_0 c}\,\mathrm{Re}\left(\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'}\right) . \tag{A.62}$$

Although these equations are quantitatively correct, they are expressed in terms of quantities which are either difficult or inconvenient to measure experimentally. Fortunately, these equations simplify rather nicely by substituting the expression for the Rabi frequency, Ω_R:

$$\Omega_\mathrm{R} \equiv \frac{d\cdot\mathcal{E}_0}{\hbar} , \tag{A.63}$$

and the spontaneous-decay rate given by Eq. (A.26). By making these substitutions, Eqs. (A.59) – (A.62) simplify to

$$\frac{1}{\mathcal{E}_0}\frac{\mathrm{d}\mathcal{E}_0}{\mathrm{d}z} = \left(\frac{3n_\mathrm{v}\Gamma_\mathrm{d}\lambda^2}{4\sqrt{2}\pi\Omega_\mathrm{R}}\frac{2J'+1}{2J+1}\right)\mathrm{Im}\left(\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'}\right) , \tag{A.64}$$

$$\frac{\mathrm{d}\phi}{\mathrm{d}z} = -\left(\frac{3n_\mathrm{v}\Gamma_\mathrm{d}\lambda^2}{4\sqrt{2}\pi\Omega_\mathrm{R}}\frac{2J'+1}{2J+1}\right)\mathrm{Re}\left(\widetilde{\rho}_{-1,0'} + \widetilde{\rho}_{1,0'}\right) , \tag{A.65}$$

$$\frac{\mathrm{d}\epsilon}{\mathrm{d}z} = \left(\frac{3n_\mathrm{v}\Gamma_\mathrm{d}\lambda^2}{4\sqrt{2}\pi\Omega_\mathrm{R}}\frac{2J'+1}{2J+1}\right)\mathrm{Im}\left(\widetilde{\rho}_{-1,0'} + \widetilde{\rho}_{1,0'}\right) , \tag{A.66}$$

$$\frac{\mathrm{d}\varphi}{\mathrm{d}z} = \left(\frac{3n_\mathrm{v}\Gamma_\mathrm{d}\lambda^2}{4\sqrt{2}\pi\Omega_\mathrm{R}}\frac{2J'+1}{2J+1}\right)\mathrm{Re}\left(\widetilde{\rho}_{-1,0'} - \widetilde{\rho}_{1,0'}\right) , \tag{A.67}$$

where λ is the wavelength of the light.

Saturated Absorption Spectroscopy

Saturated absorption spectroscopy (SAS) is a technique which, at least in principle, enables the natural linewidth of a transition to be resolved at temperatures for which the linewidth is usually dominated by Doppler broadening. In practice however, the measured linewidth is typically limited by: residual Doppler shifts from suboptimal beam overlap, power broadening, time-of-flight broadening, and collisional broadening. Irrespective of these limitations, SAS resonances can be as much as 100 times narrower than their Doppler-broadened counterparts. This is *immensely* useful for stabilising the optical frequency of a laser, as it enables the laser frequency to be locked *far more* tightly to the atomic resonance than would otherwise be the case.

B.1 Transmission Spectrum

In order to achieve sub-Doppler resolution of a transition, a counter-propagating pump-probe configuration is employed. In this setup — presented in Fig. B.1 — a reference cell is optically pumped using linearly polarised light. The intensity of the 'pump light' is large relative to the saturation intensity of the transition, such that the majority of the atoms have been excited from the ground state. At the output of the reference cell, the light is passed through a quarter-wave plate, which converts the light from linear polarisation to circular polarisation. A neutral density filter is then used to heavily attenuate the optical power of the pump light.

The pump light is then retro-reflected back along the initial propagation direction, again passing through both the neutral density filter and the quarter-wave plate. This light is now termed the 'probe beam'. When passing through the quarter-wave plate on the return path, the light is converted back to linear polarisation; except in this case the polarisation axis is *orthogonal* to that of the incident pump beam. The probe beam then passes back through the vapour cell, and is reflected into a photodetector through the use of a polarising beam splitter.

Although both the pump and probe beams are the same optical frequency — since they have both originated from the same laser source — they interact with different velocity classes of atoms due to the Doppler effect. If the incident beam is *red detuned* from the atomic transition, the pump will be absorbed by atoms *moving towards* the light at a particular velocity, while the probe beam will instead be absorbed by atoms *moving away* from the beam source at the same velocity. In the case of blue-detuned light, the same is true, except that the roles are reversed. However, in the case where the initial pump and probe beams are *exactly* on resonance according to the laboratory rest frame, they will both address the

Figure B.1: Key components of a saturated absorption spectroscopy measurement: HWP — half-wave plate, PBS — polarising beam splitter, NDF — neutral density filter, QWP — quarter-wave plate, PD — photodetector. The incident 'pump' beam (red) optically pumps the atomic vapour, while the retro-reflected 'probe' beam (blue) is used to detect the sub-Doppler SAS atomic resonances. All polarisation optical elements are used to separate the two spatially overlapped beams.

same velocity class of atoms — those with velocity vectors which are orthogonal to the propagation axis of the light (i.e. zero longitudinal velocity). Due to the pump beam having high optical intensity, many of the atoms in the beam path will either be in the excited state, or optically pumped to the other ground state. If both beams are tuned to resonance, then when the probe beam passes back along the same beam path, it will experience *less* absorption than usual, and hence there will be a 'spike' (local maximum) in the transmission spectrum of the probe beam. This spike corresponds to the hyperfine transition, and will exhibit a Lorentzian lineshape with a width that approaches the natural linewidth of the transition (cf. Sec. 3.2.1). An example of a partial SAS spectrum for the D_1 transition of rubidium is shown in Fig. B.2, where the spectrum has been fit using a function of the form

$$\mathcal{T}(\Delta\omega) = \exp\left(-\left[\sum_{i=1}^{4} a_i \mathcal{G}(\Delta\omega - \Delta\omega_i) - b_i \mathcal{L}(\Delta\omega - \Delta\omega_i) - \sum_{j=1}^{2} c_j \mathcal{L}(\Delta\omega - \Delta\omega_j)\right]\right), \quad \text{(B.1)}$$

where $\mathcal{T}(\Delta\omega)$ is the measured transmission as a function of optical detuning $\Delta\omega$, $\mathcal{L}(\Delta\omega - \Delta\omega_i)$ and $\mathcal{G}(\Delta\omega - \Delta\omega_i)$ describe the i^{th} Doppler-broadened transition, and a_i and b_i are the optical depths of the i^{th} transition. Note that the additional Lorentzian components $\mathcal{L}(\Delta\omega - \Delta\omega_j)$ are called 'crossover transitions', and are discussed in Sec. B.2.

B.2   Crossover Transitions

Upon inspection of Fig. B.2, it is clear that there are two unexpected Lorentzian features present inbetween hyperfine transitions which share the same ground state. These are known as 'crossover transitions' [221,224,346], and arise due to fact that the Doppler-broadened hyperfine transitions, which share the same ground state, are *not completely* resolved at room temperature. In this situation, atoms with a velocity of $-v$ will see the pump beam *blue*shifted to the transition frequency ω_1, and simultaneously see the probe beam

Figure B.2: Partial rubidium D_1 spectrum measured using saturated absorption spectroscopy. Transitions from $F = 2$ ground states correspond to the ^{87}Rb isotope, while transitions from $F = 3$ ground states correspond to the ^{85}Rb isotope. Measured data is shown in grey, while a fit using Eq. (B.1) is shown in blue. Discrepancies between the measured spectrum and the fit are associated with nonlinearities in the PZT, which results in a non-uniform sweep of the laser frequency.

*red*shifted to ω_2. This is shown diagrammatically in Fig. B.3, and is described mathematically using a system of two coupled equations:

$$\omega_1 = \omega_c \left(1 + \frac{v}{c}\right),$$
$$\omega_2 = \omega_c \left(1 - \frac{v}{c}\right),$$

where ω_c is the crossover frequency and c is the speed of light. Solving these simultaneous equations reveals that the crossover frequency is *exactly* halfway between the two transition frequencies:

$$\omega_c = \frac{\omega_1 + \omega_2}{2}. \tag{B.2}$$

Vapour cell

Figure B.3: Diagram showing the generation of crossover transitions in saturated absorption spectroscopy. This occurs in atoms which have two partially overlapping Doppler-broadened transitions that share the ground state. In this situation, atoms with a velocity of $-v$ will see the pump beam (blue) *blue*shifted to the transition frequency ω_1, and simultaneously see the probe beam (red) *red*shifted to ω_2. The labels '1' and '2' indicate that the atoms will observe the light of optical frequency ω_c Doppler shifted to ω_1 and ω_2, respectively. Note that the beams have been drawn 'separately' in this diagram, though in practice the beam depicted here will be retro-reflected on top of itself with maximal cross-sectional overlap.

Acousto-Optic Modulators

An acousto-optic modulator (AOM) is a device which is used to frequency shift and/or amplitude modulate light which is passed through it. An AOM is typically comprised of a birefringent crystal, with a PZT on one end and an acoustic absorber on the other. By modulating the PZT, a carrier sound wave (typically radio-frequency) is generated within the crystal, which can cause phonon-photon scattering (Brillouin scattering) of light which passes through the crystal. Since Brillouin scattering is an inelastic process, the photon may either *gain* (anti-Stokes process) or *lose* (Stokes process) energy, resulting in a frequency shift according to the Planck-Einstein relation. The frequency shift of the photon depends on the carrier frequency of the sound wave, and whether the frequency is blue- or red-shifted depends upon the launch angle of the light.

When a photon Brillouin scatters off a phonon, it is diffracted at an angle θ, which depends upon the wavelength of the light, λ, and the wavelength of the sound wave, Λ, via the relation

$$2\Lambda \sin \theta = m\lambda \,, \tag{C.1}$$

where $m \in \mathbb{Z}$ is the diffraction order. Typically only the first-order diffraction is used, as higher-order diffractions are associated with larger frequency shifts and *significantly* reduced conversion efficiency.

In all variations of the experiment, the pump beam is amplitude modulated using a high-bandwidth AOM. The AOM of choice here uses a TeO2 crystal with a $1.5 \times 2\,\mathrm{mm}$ aperture. It features a 9-MHz bandwidth and a $\sim$90% first-order diffraction efficiency. A graph of the measured first-order diffraction efficiency, as a function of rf driving power, is presented in Fig. C.1.

The AOM is driven by an 80 MHz carrier-frequency sound wave, applied to the crystal via a PZT. This carrier wave is modulated with a low-duty-cycle square wave, which has a modulation depth of 100%. This enables the pump beam to be rapidly switched on/off, with a typical modulation frequency between 30 kHz and 1 MHz used in the experiments detailed within this thesis. In addition to modulation of the pump beam, the AOMs are also used to stabilise the optical powers of both the pump and probe beams as described in Sec. 5.1.2.

Figure C.1: First-order diffraction efficiency of the acousto-optic modulators at 795 nm, measured as a function of rf driving power. Cubic spline interpolation has been included to guide the eye.

Laser Frequency-Stability Measurements

Measuring the frequency stability of an infrared laser is not *directly* possible as optical frequencies are $\mathcal{O}\left(100\,\mathrm{THz}\right)$, which is *many* orders of magnitude larger than the bandwidth of electronic devices. It is therefore necessary to, in some way, reduce this frequency to a region where electronics have sufficient bandwidth to measure it. This is typically achieved by mixing two lasers together in order to create a beat note via the trigonometric identity $2\sin\theta\sin\phi = \cos\left(\theta-\phi\right) - \cos\left(\theta+\phi\right)$. Provided that the two lasers are sufficiently close together in optical frequency, the beat note — given by the $\cos\left(\theta-\phi\right)$ term — will lie within the bandwidth of electronic devices[1]. Measurement of the beat-note frequency permits an indirect measurement of the laser's frequency stability.

Prior to the invention of the frequency comb, measuring the frequency stability of a laser involved the arduous task of frequency chaining [347]. As the name suggests, this technique involved chaining the laser of interest with numerous intermediary oscillators in order to systematically lower the beat-note frequency to the point at which it can be mixed with a caesium atomic clock. This was both a difficult and expensive task, which often inhibited this technique from being utilised in many laboratories.

Fortunately however, the invention of the optical frequency comb [348] has made measuring optical frequencies *significantly* easier [348–351]. An optical frequency comb is a laser whose spectrum contains hundreds or thousands of equally spaced frequency components, which is essentially equivalent to having hundreds or thousands of individual lasers in a single beam. This is incredibly useful when measuring the frequency stability of a cw laser — provided the laser frequency is either close to or contained within the comb envelope — as mixing a laser with the comb essentially *guarantees* a resulting low-frequency beat note *without* the need to chain numerous intermediary oscillators together. In practice there will be multiple beat notes generated when mixing a cw laser with a frequency comb; however, these may be rejected through the use of appropriate electronic filtering.

An experimental diagram depicting a frequency-stability measurement — performed by mixing a cw laser with an optical frequency comb — is presented in Fig. D.1. A fibre splitter is used to mix a cw laser with the output of a frequency comb, creating a slew of mixing products with the numerous comb modes. This mixed light is then shone onto a diffraction grating, which separates it into its various frequency components along the horizontal axis. In order to reduce the shot noise of the measurement, the vast majority of the light is discarded at this point — only a small subset of light with optical frequency close to that of the cw laser is retained and launched onto a photodetector. A combination of Butterworth and

---

[1]The $\cos\left(\theta+\phi\right)$ term, on the other hand, will occur at *twice* the optical frequency and can therefore be neglected.
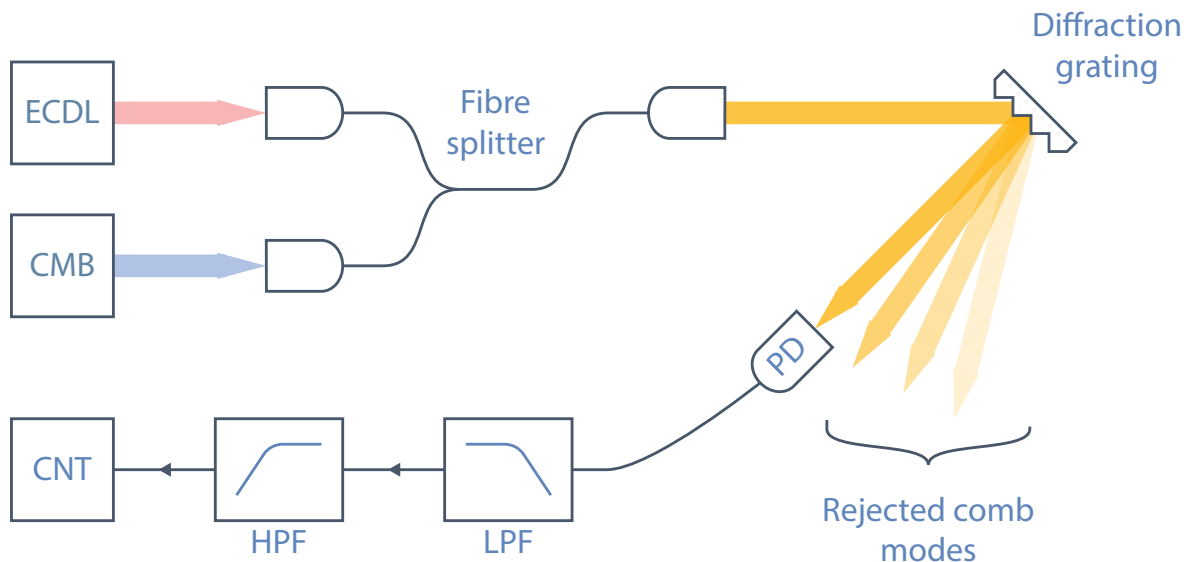
Figure D.1: Simplified frequency-stabilisation measurement setup, showing the external-cavity diode laser (ECDL), optical frequency comb (CMB), photodetector (PD), low-pass filter (LPF), high-pass filter (HPF), and frequency counter (CNT). The ECDL and frequency comb are mixed using a fibre splitter and subsequently shone onto a diffraction grating, separating the light into its various frequency components along the horizontal axis. In order to reduce shot noise, only a small subset of the comb modes near the ECDL frequency are launched onto the photodetector. A combination of low-pass and high-pass filtering is used to reject all but one beat note. Counting the beat-note frequency allows the frequency-stability of the ECDL to be inferred.

high-pass filtering is performed on the photodetector output, effectively creating a bandpass filter which enables just *one* of the mixing products (beat notes) to be selected. A frequency counter is then used to measure the frequency of the beat note as a function of time, and an Allan deviation (cf. Sec. H.2) is subsequently calculated. Given that the beat note is a mixing product of two lasers — the cw laser and the frequency comb — drifts in the frequency of *either* of these two lasers are *indistinguishable*. However, provided that the frequency stability of the optical frequency comb *greatly exceeds* that of the cw laser, any drifts in the beat-note frequency will be dominated by the cw laser and hence the effect of the frequency comb can be safely neglected.

A measurement of the frequency stability of the external-cavity diode laser used in this thesis — performed using an optical frequency comb and calculated using Eq. (H.17) — is presented in Fig. D.2. When the laser is unlocked, its fractional frequency stability is about 1 ppb for 1 s integration times, degrading as approximately $\tau^{0.3}$ to 5 ppb over minute timescales. In contrast, when the laser is locked using saturated absorption spectroscopy as described in Sec. 5.1.1, the fractional frequency stability is about 7×10^{-12} for 1 s integrations, dropping as low as 3×10^{-12} at 10 s, then degrading as approximately $\tau^{0.65}$ to 2×10^{-10} over hour timescales. In terms of absolute values when locked to the $F = 2 \to F' = 1$ hyperfine transition of the ^{87}Rb D$_1$ line, the laser frequency stability when locked is about 3 kHz at 1 s, dropping to 1 kHz at 10 s, and degrading to 70 kHz at 10^4 s. Such a level of absolute

Figure D.2: Allan deviation of the optical frequency of the external-cavity diode laser used in this thesis, measured using an optical frequency comb via the experimental setup presented in Fig. D.1. On the left axis is the fractional frequency stability, while the right axis shows the absolute frequency stability. During these measurements, the ECDL was either unlocked (red), or locked (blue) to the $F = 2 \to F' = 1$ hyperfine transition of the ^{87}Rb D$_1$ line using saturated absorption spectroscopy (cf. Sec. 5.1.1 and App. B), giving rise to an optical frequency of 377.11 THz. Black dashed lines correspond to integration-time dependencies of $\tau^{0.3}$ (unlocked) and $\tau^{0.65}$ (locked). Shown in grey is the frequency stability of the optical frequency comb, measured using a cryogenic sapphire oscillator.

frequency stability was deemed to be sufficient for the precision measurements performed in this thesis (cf. Sec. 8.2.8).

Spin Relaxation via Diffusion Through a Buffer Gas in a Vapour Cell

Optical atomic magnetometers determine the local magnetic-field strength by measuring the Larmor frequency of a spin-polarised atomic vapour and, as such, require long spin-relaxation times — typically the transverse spin-relaxation time T_2 — in order to achieve high sensitivity (cf. Sec. 6.4.1). One of the most dominant relaxation mechanisms in a vapour cell is due to collisions between the polarised atoms and the cell walls (cf. Sec. 5.2.2.4). For this reason, vapour cells are often either filled with an inert buffer gas (to restrict the mean free path), or coated with an anti-relaxation coating such as paraffin. In the case of a vapour cell containing buffer gas, the diffusion equation for the polarisation of an atomic vapour is given by [297,298]

$$\frac{\partial \mathrm{P}}{\partial t} = \mathcal{D}\nabla^2 \mathrm{P}, \tag{E.1}$$

where $\mathcal{D} = \mathcal{D}_0 P_\mathrm{A}/P_\mathrm{v}$ is the diffusion coefficient at a gas pressure P_v, $\mathcal{D}_0$ is the species-specific diffusion coefficient at atmospheric pressure P_A, and ∇^2 is the Laplace operator.

As would be expected, the temporal dependence of the atomic polarisation is dependent — among other things — upon the geometry of the cell. Two of the most common geometries, spherical and cylindrical, are considered below.

E.1 Spherical Cells

Due to both the radial and spherical symmetry of a spherical cell, we assume separable solutions of the form $\mathrm{P}(r,t) = \mathcal{R}(r)\mathcal{T}(t)$. Such a symmetry suggests the use of spherical polar coordinates, in which the Laplace operator, acting on a function f, is given by

$$\nabla^2 f = \frac{1}{r^2}\frac{\partial}{\partial r}\left(r^2\frac{\partial f}{\partial r}\right) + \frac{1}{r^2\sin\theta}\frac{\partial}{\partial\theta}\left(\sin\theta\frac{\partial f}{\partial\theta}\right) + \frac{1}{r^2\sin^2\theta}\frac{\partial^2 f}{\partial\phi^2}. \tag{E.2}$$

Inserting the separable-solution form of $\mathrm{P}(r,t)$ into Eq. (E.2), and discarding terms which are dependent upon θ and ϕ citing symmetry arguments, yields

$$\frac{1}{\mathcal{D}\mathcal{T}(t)}\frac{\mathrm{d}\mathcal{T}(t)}{\mathrm{d}t} = \frac{2}{r\mathcal{R}(r)}\frac{\mathrm{d}R(r)}{\mathrm{d}r} + \frac{1}{\mathcal{R}(r)}\frac{\mathrm{d}^2 R(r)}{\mathrm{d}r^2}. \tag{E.3}$$

Given that there are no shared variables between the two sides of Eq. (E.3), both sides must be equal to the *same* constant. If we define this constant to be $-\lambda^2$, this yields two equations:

$$\frac{d\mathcal{T}(t)}{dt} + \lambda^2 \mathcal{D}\mathcal{T}(t) = 0, \tag{E.4}$$

$$r^2 \frac{d^2 R(r)}{dr^2} + 2r \frac{dR(r)}{dr} + r^2 \lambda^2 R(r) = 0. \tag{E.5}$$

Solving Eqs. (E.4) and (E.5) yields the following general solutions:

$$\mathcal{R}(r) = c_1 \frac{\sin(\lambda r)}{\lambda r} + c_2 \frac{\cos(\lambda r)}{\lambda r}, \tag{E.6}$$

$$\mathcal{T}(t) = c_3 \exp\left(-\lambda^2 \mathcal{D}t\right), \tag{E.7}$$

where c_i are constant coefficients. It should be noted that, although $c_2 \cos(\lambda r)/\lambda r$ is part of the general solution to Eq. (E.5), it is *not* physical as $\cos(\lambda r)/\lambda r \to \infty$ as $r \to 0$. It is therefore necessary that $c_2 = 0$.

The term λ can be solved for using the boundary condition that arises from wall collisions completely depolarising the atoms. Explicitly, this boundary condition is given by $\mathcal{R}(R_{\text{cell}}) = 0$ and implies that $c_1 \text{sinc}(\lambda R_{\text{cell}}) = 0$; which, for a non-trivial solution, requires that $\lambda R_{\text{cell}} = n\pi$ for $n \in \mathbb{Z}$. This gives rise to an infinite set of modes with

$$\lambda_n = \frac{n\pi}{R_{\text{cell}}}. \tag{E.8}$$

The general solution to the diffusion equation is then given by a sum over all diffusion modes:

$$P(r,t) = \sum_{n=1}^{\infty} P_n \text{sinc}(\lambda_n r) \exp\left(-\lambda_n^2 \mathcal{D}t\right), \tag{E.9}$$

where P_n is the amplitude of the n^{th} mode.

Note that the polarisation has an exponential decay, with a time constant of the n^{th} mode given by $\tau_n = 1/\lambda_n^2 \mathcal{D}$. Explicitly, this is given by

$$\tau_n = \frac{1}{\mathcal{D}} \left(\frac{R_{\text{cell}}}{n\pi}\right)^2. \tag{E.10}$$

The lowest order (slowest) diffusion mode, $n = 1$, yields the time constant [298,299]

$$\tau_1 = \frac{1}{\mathcal{D}} \left(\frac{R_{\text{cell}}^2}{\pi^2}\right). \tag{E.11}$$

E.2   Cylindrical Cells

Calculating the diffusion of polarisation in a cylindrical cell is a bit more complicated than for a spherical cell, owing to the reduction in symmetry. However, there is still *cylindrical* symmetry, which suggests the use of cylindrical polar coordinates. In this coordinate system, the Laplace operator is given by

$$\nabla^2 f = \frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial f}{\partial r}\right) + \frac{1}{r^2}\frac{\partial^2 f}{\partial \phi^2} + \frac{\partial^2 f}{\partial z^2}.$$ (E.12)

Let us assume that the solution is separable, i.e. that $P(r,z,t) = \mathcal{R}(r)\,\mathcal{Z}(z)\,\mathcal{T}(t)$. Inserting this solution form into the Laplace operator, and ignoring the ϕ terms due to radial symmetry, yields

$$\frac{1}{\mathcal{D}\mathcal{T}(t)}\frac{d\mathcal{T}(t)}{dt} = \frac{1}{r\mathcal{R}(r)}\frac{d\mathcal{R}(r)}{dr} + \frac{1}{\mathcal{R}(r)}\frac{d^2\mathcal{R}(r)}{dr^2} + \frac{1}{\mathcal{Z}(z)}\frac{d^2\mathcal{Z}(z)}{dz^2}.$$ (E.13)

Given that there are no shared variables between the two sides of Eq. (E.13), *both* sides must be equal to the *same* constant. If we define this constant to be $-\lambda^2$, this yields two equations:

$$\frac{d\mathcal{T}(t)}{dt} + \lambda^2\mathcal{D}\mathcal{T}(t) = 0,$$ (E.14)

$$\frac{1}{r\mathcal{R}(r)}\frac{d\mathcal{R}(r)}{dr} + \frac{1}{\mathcal{R}(r)}\frac{d^2\mathcal{R}(r)}{dr^2} + \frac{1}{\mathcal{Z}(z)}\frac{d^2\mathcal{Z}(z)}{dz^2} = -\lambda^2.$$ (E.15)

As was the case previously, Eq. (E.15) is an equality with no shared variables and hence they must be equal to *another* constant. Defining this constant as μ^2 and subsequently rearranging yields two new differential equations:

$$r^2\frac{d^2\mathcal{R}(r)}{dr^2} + r\frac{d\mathcal{R}(r)}{dr} + r^2\left(\lambda^2 - \mu^2\right)\mathcal{R}(r) = 0,$$ (E.16)

$$\frac{d^2\mathcal{Z}(z)}{dz^2} + \mu^2\mathcal{Z}(z) = 0.$$ (E.17)

Solving Eqs. (E.14), (E.16) and (E.17) yields the following general solutions:

$$\mathcal{R}(r) = c_1 J_0\left(\sqrt{\lambda^2 + \mu^2}\,r\right) + c_2 Y_0\left(\sqrt{\lambda^2 + \mu^2}\,r\right),$$ (E.18)

$$\mathcal{Z}(z) = c_3\cos(\mu z) + c_4\sin(\mu z),$$ (E.19)

$$\mathcal{T}(t) = c_5\exp\left(-\lambda^2\mathcal{D}t\right),$$ (E.20)

where c_i are constant coefficients, and J_0 and Y_0 are the zeroth-order Bessel functions of the first and second kind, respectively. It should be noted that, although $c_2 Y_0\left(\sqrt{\lambda^2 + \mu^2}\,r\right)$ is part of the general solution to Eq. (E.16), it is *not* physical as $Y_0(r) \to -\infty$ as $r \to 0$. It is therefore necessary that $c_2 = 0$.

In terms of longitudinal boundary conditions, there are two key points to consider: the two edges of the cell at $z = 0$ and $z = L_{\text{cell}}$. The two boundary conditions at these points are given by $\mathcal{Z}(0) = \mathcal{Z}(L_{\text{cell}}) = 0$. Using the former condition results in $c_3 = 0$, and hence the solution for $\mathcal{Z}(z)$ reduces to $\mathcal{Z}(z) = c_4\sin(\mu z)$. The latter boundary condition yields $c_4\sin(\mu L_{\text{cell}}) = 0$ which, for a non-trivial solution, requires that

$$\mu_m = \frac{m\pi}{L_{\text{cell}}}, \tag{E.21}$$

where $m \in \mathbb{Z}$. Now that the solution for μ_m has been calculated, it is possible to determine λ. The radial boundary condition $\mathcal{R}(R_{\text{cell}}) = 0$ implies that $c_1 J_0\left(\sqrt{\lambda^2 + \mu^2}R_{\text{cell}}\right) = 0$ which, for a non-trivial solution, requires that

$$\lambda_{m,n} = \sqrt{\left(\frac{^{(n)}J_0^{-1}(0)}{R_{\text{cell}}}\right)^2 + \left(\frac{m\pi}{L_{\text{cell}}}\right)^2}, \tag{E.22}$$

where $^{(n)}J_0^{-1}(0)$ is the n^{th} solution to $J_0(r) = 0$, and $n \in \mathbb{Z}^* \setminus \{0\}$. The general solution to the diffusion equation is a sum over all diffusion modes:

$$P(r,z,t) = \sum_{m=1}^{\infty}\sum_{n=1}^{\infty} P_{m,n} J_0\left(\sqrt{\lambda_{m,n}^2 + \mu_n^2}\, r\right) \sin(\mu_n z) \exp\left(-\lambda_{m,n}^2 \mathcal{D}t\right), \tag{E.23}$$

where $P_{m,n}$ is the amplitude of the $(m,n)^{\text{th}}$ mode.

Note that the polarisation undergoes exponential decay, with a time constant of the $(m,n)^{\text{th}}$ mode given by $\tau_{m,n} = 1/\lambda_{m,n}^2 \mathcal{D}$. Explicitly, this is given by

$$\tau_{m,n} = \frac{1}{\mathcal{D}}\left[\left(\frac{R_{\text{cell}}}{^{(n)}J_0^{-1}(0)}\right)^2 + \left(\frac{L_{\text{cell}}}{m\pi}\right)^2\right]. \tag{E.24}$$

The lowest-order (slowest) diffusion mode, $m = n = 1$, yields the time constant [298,300]

$$\tau_{1,1} \approx \frac{1}{\mathcal{D}}\left(\frac{R_{\text{cell}}^2}{(2.40483)^2} + \frac{L_{\text{cell}}^2}{\pi^2}\right). \tag{E.25}$$

Characterisation of Anti-Relaxation-Coated Vapour Cells

One method for studying ground-state relaxation in anti-relaxation-coated vapour cells is by using an alternate pump-probe technique — first introduced in Ref. [297] — known as 'relaxation in the dark'. This essentially constitutes a free-induction decay measurement, in which the atomic vapour is optically pumped for a finite period of time, and then observed with a weak probe beam as it undergoes free temporal evolution. Depending on the optical configuration — i.e. the polarisation and propagation directions of both the pump and probe beams relative to the magnetic field — the observed probe signal will undergo (bi-)exponential decay.

In the magnetometry measurements performed in this thesis (cf. Secs. 3.6, 3.7, and 6.1), the observed relaxation is a *single* exponential characterised by a time constant T_2 — the transverse spin-relaxation time. As given by Eq. (6.1) and repeated here, the FID signal in the case of a static magnetic field is given by

$$\phi(t) = \phi_0 \sin(2\Omega_L t + \varphi) \exp\left(-\frac{t}{T_2}\right), \tag{F.1}$$

where ϕ_0 is the maximum optical-rotation amplitude at $t = 0$, Ω_L is the Larmor frequency, and φ is an arbitrary phase offset.

An example of an FID measurement performed in this thesis is presented in Fig. 6.1. By fitting the signal upon extinction of the pump beam, using Eq. (6.1), a transverse spin-relaxation time of $T_2 = 43.4\,\text{ms}$ was obtained, corresponding to a linewidth (FWHM) of $\Gamma/2\pi = 2/(2\pi T_2) = 7.33\,\text{Hz}$. This is *many* orders of magnitude smaller than the predicted linewidth one would obtain using an uncoated cell of the same dimensions, $\Gamma/2\pi \approx 3\,\text{kHz}$ (cf. Sec. 5.2.2.4).

The transverse spin-relaxation time of the atomic vapour is subject to a multitude of relaxation mechanisms within the vapour cell, including: spin-exchange collisions between alkali atoms [150,240], electron-randomisation (spin-destruction) collisions with the cell walls [150], exchange of alkali atoms between the volume of the cell and its stem (known as the 'reservoir effect') [352,353], power broadening [150,199,201,238], and gradient broadening [234].

The data presented in Fig. 6.1 was recorded for a probe power of $P_{\text{probe}} = 5\,\mu\text{W}$, a pump power of $P_{\text{pump}} = 0\,\mu\text{W}$ for $t \geq 0$, and at a magnetic-field strength of $B_{\text{dc}} = 2.188\,\mu\text{T}$. The measured transverse spin-relaxation time under these conditions is approximately the same as that which is measured in the limit as $P_{\text{probe}} \to 0\,\mu\text{W}$ and $B_{\text{dc}} \to 0\,\mu\text{T}$, which in-

dicates that power broadening and gradient broadening were not contributing significantly to the observed resonance linewidth in this measurement. Combined with the fact that spin-exchange collisions contribute negligibly at the temperature at which all measurements within this thesis are taken (cf. Sec. 5.2.2.1), the resonance width of $\Gamma_0/2\pi = 2/(2\pi T_2) \approx$ 7 Hz corresponds to the *intrinsic linewidth* of the cell and is limited by wall collisions — i.e. the quality of the anti-relaxation coating.

Fourier Analysis

This chapter provides the mathematical foundation required to understand Fourier analysis, beginning with an introduction to the Fourier series, and then a derivation of the Fourier transform. The discrete Fourier transform — useful when numerically computing the Fourier transform of experimental data — is then introduced, which heralds a discussion about Nyquist's sampling theorem, spectral leakage, and window functions.

G.1 The Fourier Series

The Fourier series is a way of representing any continuous-time function $f(t)$ — which is periodic with a period T — as a weighted linear combination of positive and negative harmonics. A function is said to be periodic if, for any $n \in \mathbb{Z}$, the following condition is satisfied:

$$f(t + nT) = f(t) .$$
(G.1)

If the above condition is satisfied, or at least approximately satisfied[1], then the function $f(t)$ can be calculated via a Fourier series:

$$f(t) = \frac{a_0}{2} + \sum_{n=1}^{\infty} a_n \cos(n\omega t) + b_n \sin(n\omega t) ,$$
(G.2)

where $\omega = 2\pi/T$ is the fundamental angular frequency of the signal, and a_0, a_n and b_n are weighting coefficients. Equation (G.2) can be written in complex-exponential form using trigonometric identities, resulting in an elegant formulation:

$$
\begin{aligned}
f(t) &= \frac{a_0}{2} + \sum_{n=1}^{\infty} a_n \cos(n\omega t) + b_n \sin(n\omega t) \\
&= \sum_{n=0}^{\infty} a_n \cos(n\omega t) + b_n \sin(n\omega t) \\
&= \sum_{n=0}^{\infty} \left(\frac{a_n}{2} - i\frac{b_n}{2} \right) e^{in\omega t} + \left(\frac{a_n}{2} + i\frac{b_n}{2} \right) e^{-in\omega t} \\
&= \sum_{n=-\infty}^{\infty} c_n e^{in\omega t} ,
\end{aligned}
$$
(G.3)

---

[1] In practice using real data of finite length, even if the signal is non-periodic (e.g. white noise), the Fourier series (or Fourier transform — cf. App. G.2) can be calculated under the assumption that the signal *is* periodic and that the measured data constitutes *one period* of the signal.

where $c_n \in \mathbb{C}$ is a complex coefficient. This formula allows any arbitrary function to be decomposed into the sum of simple, oscillatory functions. The coefficients c_n are called Fourier coefficients, and are used to *uniquely* describe the function $f(t)$. To calculate the coefficients we must find a general expression for the k^{th} coefficient, c_k. This is achieved by multiplying both sides of Eq. (G.3) by $\exp(-ik\omega t)$, which yields

$$
e^{-ik\omega t} f(t) = e^{-ik\omega t} \sum_{n=-\infty}^{\infty} c_n e^{in\omega t}
$$
$$
= \ldots + e^{-ik\omega t} c_{k-1} e^{i(k-1)\omega t} + e^{-ik\omega t} c_k e^{ik\omega t} + e^{-ik\omega t} c_{k+1} e^{i(k+1)\omega t} + \ldots .
$$

This summation can then be re-arranged to yield an expression for the coefficient c_k:

$$
c_k = e^{-ik\omega t} f(t) - \sum_{\substack{n=-\infty \\ n\neq k}}^{\infty} c_n e^{i(n-k)\omega t} . \tag{G.4}
$$

Further simplification is possible by integrating both sides over a period T:

$$
\int_{-T/2}^{T/2} c_k \mathrm{d}t = \int_{-T/2}^{T/2} e^{-ik\omega t} f(t)\, \mathrm{d}t - \int_{-T/2}^{T/2} \sum_{\substack{n=-\infty \\ n\neq k}}^{\infty} c_n e^{i(n-k)\omega t} \mathrm{d}t
$$
$$
= \int_{-T/2}^{T/2} e^{-ik\omega t} f(t)\, \mathrm{d}t - \sum_{\substack{n=-\infty \\ n\neq k}}^{\infty} c_n \int_{-T/2}^{T/2} e^{i(n-k)\omega t} \mathrm{d}t
$$
$$
= \int_{-T/2}^{T/2} e^{-ik\omega t} f(t)\, \mathrm{d}t - \sum_{\substack{n=-\infty \\ n\neq k}}^{\infty} c_n \left[\frac{1}{i(n-k)\omega t} \left(e^{i(n-k)\omega T/2} - e^{-i(n-k)\omega T/2} \right) \right]
$$
$$
= \int_{-T/2}^{T/2} e^{-ik\omega t} f(t)\, \mathrm{d}t - \sum_{\substack{n=-\infty \\ n\neq k}}^{\infty} c_n \left[\frac{2}{(n-k)\omega} \sin\left(\frac{1}{2}(n-k)\omega T \right) \right]
$$
$$
= \int_{-T/2}^{T/2} e^{-ik\omega t} f(t)\, \mathrm{d}t - \sum_{\substack{n=-\infty \\ n\neq k}}^{\infty} c_n \left[\frac{2}{(n-k)\omega} \sin\left(\pi(n-k) \right) \right]
$$
$$
= \int_{-T/2}^{T/2} e^{-ik\omega t} f(t)\, \mathrm{d}t ,
$$

where the above has been possible by noting that $\sin(a\pi) = 0$ for all $a \in \mathbb{Z}$. Solving this integral equation yields the solution

$$
c_k = \frac{1}{T} \int_{-T/2}^{T/2} f(t)\, e^{-ik\omega t} \mathrm{d}t . \tag{G.5}
$$

Equation (G.5), along with Eq. (G.3), is all that is required in order to calculate the Fourier series of any arbitrary continuous-time function $f(t)$. Notice however that Eq. (G.3) is an *infinite* sum, which in most cases is impossible to evaluate completely. Fortunately, given

that the Fourier series is generally a *convergent* sum, the higher-order terms (i.e. components with large values of n) make minimal contribution to the reproduction of the signal and hence the summation can be truncated once the signal has been approximated with sufficient accuracy[2].

G.2 The Fourier Transform

As T gets larger, $n\omega$ and $(n+1)\omega$ get closer together — infinitely so in the limit $T \to \infty$. As the linear frequency of the n^{th} harmonic is given by $\nu_n = n/T$, a differential change in frequency is thus given by the change in harmonic number via $\Delta\nu_n = \Delta n/T$. However, $n \in \mathbb{Z}$ (i.e. n is an integer) and hence $\Delta n = 1$. This enables Eq. (G.3) to be written as a Riemann sum:

$$
\begin{aligned}
f(t) &= \sum_{n=-\infty}^{\infty} c_n e^{2\pi i \nu_n t} \Delta n \\
&= \sum_{n=-\infty}^{\infty} c_n e^{2\pi i \nu_n t} T \Delta\nu_n \\
&= \sum_{n=-\infty}^{\infty} c_n T e^{2\pi i \nu_n t} \Delta\nu_n \\
&= \sum_{n=-\infty}^{\infty} F(\nu_n) e^{2\pi i \nu_n t} \Delta\nu_n .
\end{aligned}
$$

As $T \to \infty$, $\Delta\nu_n \to 0$ and hence the Riemann sum becomes an integral:

$$
f(t) = \int_{-\infty}^{\infty} F(\nu_n) e^{2\pi i \nu_n t} \mathrm{d}\nu_n . \tag{G.6}
$$

By specifying ν, n is also being specified indirectly and hence the superfluous n label may be dispensed with. This yields the following:

$$
f(t) = \int_{-\infty}^{\infty} F(\nu) e^{2\pi i \nu t} \mathrm{d}\nu . \tag{G.7}
$$

Equation (G.7) is known as the *inverse* Fourier transform[3], and is used to construct a continuous-time function $f(t)$ from its frequency spectrum $F(\nu)$. Although this is useful, it is more commonly the case that one wishes to decompose a continuous-time function into its frequency spectrum, in which case the Fourier transform is performed:

$$
F(\nu) = \int_{-\infty}^{\infty} f(t) e^{-2\pi i \nu t} \mathrm{d}t . \tag{G.8}
$$

---

[2]In practice, when using experimentally measured or computer-generated data, calculation of the Fourier transform only requires a finite number of terms to be evaluated. This is discussed in further detail in Sec. G.3.

[3]The Fourier transform is typically denoted by $\mathscr{F}\{f(t)\}$, while the inverse Fourier transform is denoted by $\mathscr{F}^{-1}\{F(\nu)\}$.

G.3 The Discrete Fourier Transform (DFT)

In practice, a signal $f(t)$ will *not* be a continuously sampled function, but will instead be sampled N times over a particular measurement period. This means that spectral information can *only* be obtained at N points. Given a sampling interval $dt = \Delta t$, the function is sampled at times $t_k = k\Delta t$, in which case the Fourier transform given by Eq. (G.8) is approximated by a *finite* sum:

$$
\begin{aligned}
F[n] &\approx \sum_{k=0}^{N-1} f[k]\, e^{-2\pi i \nu_n t_k} \Delta t \\
&= \Delta t \sum_{k=0}^{N-1} f[k]\, e^{-2\pi i k \frac{n}{N}} ,
\end{aligned}
$$

where $0 \le n \le N-1$ and $0 \le k \le N-1$. This gives rise to the definition of the discrete and inverse discrete Fourier transforms, respectively [354]:

$$
F[n] = \sum_{k=0}^{N-1} f[k]\, e^{-2\pi i k \frac{n}{N}} , \tag{G.9}
$$

$$
f[k] = \frac{1}{N} \sum_{n=0}^{N-1} F[n]\, e^{2\pi i k \frac{n}{N}} . \tag{G.10}
$$

G.3.1 Discrete Frequencies and Nyquist's Sampling Theorem

When computing the discrete Fourier transform[4] using Eq. (G.9), the first element ($n = 0$) corresponds to the dc component of the signal. The second element ($n = 1$) corresponds to the smallest (fundamental) frequency above zero, $\nu = (N\Delta t)^{-1}$, the third element ($n = 2$) corresponds to the second harmonic, $\nu = 2(N\Delta t)^{-1}$, and so on. Increasing values of n therefore correspond to higher frequencies; that is, *up until* the value $n = N/2 - 1$ is reached. Elements above this, i.e. in the range $N/2 \le n \le N-1$, are aliased (cf. App. H.4) to *negative frequencies* [354]. However, these negative-frequency components appear in *reverse order*; i.e. the final element ($n = N - 1$) corresponds to the frequency $\nu = -(N\Delta t)^{-1}$, which *should* appear *just prior* to the first element. Therefore, the output of the discrete Fourier transform must be *reorganised* by swapping the order of its first and second halves. Upon doing so, spectral information is obtained at the discrete frequencies [354]

$$
\nu_{n'} = \frac{n' - \frac{N}{2}}{N\Delta t} , \tag{G.11}
$$

---

[4]In practice, the discrete Fourier transform is typically calculated using the Cooley-Tukey fast Fourier transform (FFT) algorithm [355]. The benefit of the Cooley-Tukey algorithm is that the computational time of the discrete Fourier transform is decreased from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log N)$, which is *dramatically* beneficial when working with large datasets. A caveat is that the algorithm often requires the dataset be a power of 2 in size (i.e. $N = 2^m$ for $m \in \mathbb{Z}^* \setminus \{0\}$), though there do exist more general algorithms that allow for datasets of arbitrary length [355]. Datasets which are not a power of 2 can either be truncated to a power of 2 (by discarding points at either the beginning or the end of the set), or padded with zeros.

where $0 \leq n' \leq N - 1$ are now the *reordered* elements. The highest frequency component that can be resolved in a discrete Fourier transform occurs when $n' = 0$ [354]. Due to the periodicity of the discrete Fourier transform, the element $n' = 0$ corresponds to *both* positive *and* negative frequencies [354]. This maximum frequency component is known as the Nyquist frequency, and is given explicitly by [354]

$$|\nu_{\mathrm{N}}| = \frac{1}{2\Delta t} = \frac{\nu_s}{2} , \tag{G.12}$$

where $\nu_s = 1/\Delta t$ is the sampling frequency.

The fact that the highest resolvable frequency component is given by *half* the sampling frequency gives rise to the Nyquist sampling theorem, which states that a signal $f(t)$ is *perfectly* reconstructed provided that the sampling frequency is *at least* twice the highest frequency component in the signal's spectrum.

G.4   Spectral Leakage and Windowing

Spectral leakage is simply the introduction of unwanted frequency components into the spectrum of a signal, regardless of whether this occurs at the point of measurement or in post-processing [356]. Not only does this result in the apparent loss of signal amplitude/power, but small signals can be lost in the leakage caused by *other* signals. This is of course *extremely* problematic, and techniques should be employed to reduce the amount of spectral leakage for these reasons.

Embedded in the derivation of the Fourier transform is the assumption that the transformed data is both periodic (i.e. $f(t + nT) = f(t)$, for $n \in \mathbb{Z}$ and some period T) and of infinite extent. In practice however, measured data, $\hat{f}(t)$, will only be a 'window' of the original signal, $f(t)$. Mathematically, this means that the signal has been multiplied by a window function, $w(t)$:

$$\hat{f}(t) = f(t) w(t) . \tag{G.13}$$

The result of multiplication by a window function is that the Fourier transform of the measured signal, $\mathscr{F}\left\{\hat{f}(t)\right\}$, is given by the convolution theorem:

$$\mathscr{F}\left\{\hat{f}(t)\right\} = \mathscr{F}\left\{f(t)\right\} * \mathscr{F}\left\{w(t)\right\} . \tag{G.14}$$

In the typical case, the window function is simply a rectangular function, as a finite amount of data has been measured with a hard cut-off at the beginning and end points. That is, the window function takes the form $w(t) = u(t - t_0 + \Delta t) - u(t - t_0 - \Delta t)$, where u is the Heaviside function centred at t_0, and the resulting rectangular window has a width of $2\Delta t$. The Fourier transform of a rectangular function is a well-known result, and is in general proportional to the sinc function. If one chooses the rectangular window to exist over $t \in [9, 11]$ s for example, then the Fourier transform is given explicitly by

Figure G.1: Example of the spectral leakage that occurs when calculating the Fourier transform of a rectangular window.

$$\mathscr{F}\{w(t)\} = \sqrt{\frac{2}{\pi}}\Delta t \exp(i\omega t_0) \operatorname{sinc}(\omega \Delta t) \, . \tag{G.15}$$

A plot showing the rectangular window, as well as the normalised magnitude of its corresponding Fourier transform given by Eq. (G.15), is shown in Fig. G.1.

It is evident that in general, a rectangular window causes spectral leakage by spreading the signal energy over other frequencies. Note that if the sample window were to extend from $t = -\infty$ to $t = \infty$, then $w(t) = 1$, and hence $\mathscr{F}\{w(t)\} = \delta(\omega)$. The Fourier transform as given by the convolution theorem will therefore return the *original* signal, and hence there will be *no* spectral leakage. However, due to finite measurement times, the only practical case where there is no spectral leakage is when a periodic signal (e.g. a sine wave) is sampled over an *integer* number of cycles. In this case, the sampled frequencies of the rectangular window's frequency response are all zero, and hence the spectral leakage that would usually be introduced from using a rectangular window is effectively mitigated. However, if a non-integer number of cycles are sampled, the frequency response of the window function is now sampled at points where it is non-zero, which results in spectral leakage. An example of this phenomenon is presented in Fig. G.2.

Since it is evident that spectral leakage can arise from the way in which data is sampled — specifically from the discontinuities between the initial and final points in the dataset — logic dictates that spectral leakage can be reduced by weighting the data points in the *middle* of the set more than those on the *edges* of the set. This technique is called 'windowing', and the weighting functions are known as 'window functions'.

Figure G.2: Example of the spectral leakage that occurs when calculating the Fourier transform of non-periodic data. (a) 16 cycles of a 0.125 Hz sine wave with unity amplitude, sampled at a rate of 1 Hz and therefore resulting in a Nyquist frequency of 0.5 Hz. The corresponding Fourier transform shows a Dirac delta function of unity amplitude, at the frequency of the sine wave. (b) 16.5 cycles of the same sine wave. Due to the discontinuity between the start and end points of the time series, the Fourier transform exhibits significant spectral leakage, and the amplitude of the ∼0.125 Hz component is no longer accurate.

There are numerous window functions, each of which has intrinsic advantages and disadvantages. Perhaps the most commonly used window function is the Hanning window, which is mathematically given by

$$w(n) = \frac{1}{2} - \frac{1}{2}\cos\left(\frac{2\pi n}{N-1}\right),$$ (G.16)

where n is the sample number, and N is the total number of samples in the dataset. A graph of the Hanning window — calculated for the data shown in Fig. G.2(b) — is shown in Fig. G.3(a). Note that when using window functions, if the integral of the function is *not* unity, then the signal will either *gain* or *lose* energy. This is characterised by the normalised integral of the window function, which is known as the 'coherent gain':

$$G_{\text{coh}} = \frac{1}{N-1}\int_0^{N-1} w(n)\,\mathrm{d}n.$$ (G.17)

If $G_{\text{coh}} \neq 1$, the resulting Fourier transform *must* be scaled by the *reciprocal* of the coherent gain in order to conserve energy. In the case of the Hanning window, the coherent gain is given explicitly by

$$G_{\text{coh}} = \frac{1}{N-1}\int_0^{N-1} \frac{1}{2} - \frac{1}{2}\cos\left(\frac{2\pi n}{N-1}\right)\mathrm{d}n = \frac{1}{2}.$$ (G.18)

An example of the reduction in spectral leakage that occurs when using window functions is presented in Fig. G.3(b). This data corresponds to that shown in Fig. G.2(b); however, prior to calculating the Fourier transform, the time-domain data is windowed using the Hanning window shown in Fig. G.3(a). The amplitude of the frequency component at 0.125 Hz after reducing spectral leakage is about 0.85, which is closer to true value of unity than the value of around 0.65 calculated *without* using window functions.

Figure G.3: Example of the effectiveness of windowing techniques in reducing the spectral leakage that occurs when calculating the Fourier transform of non-periodic data. (a) Hanning window calculated using Eq. (G.16) for the data presented in Fig. G.2(b). (b) The time series and Fourier transform of the data presented in Fig. G.2(b) after applying the Hanning Window. Note that the spectral leakage is significantly reduced when compared to Fig. G.2(b), and the amplitude is now much closer to the true value.

Characterising Noise

In a signal processing context, noise is defined as random signals which carry no *useful* information. In most cases, random noise masks signals with small amplitude, and hence the presence of noise introduces an upper limit to the performance of a device. It is therefore useful to characterise both the amount of noise, as well as its origin.

There are two different domains in which to characterise the noise of a device: the time domain and the frequency domain. Perhaps the two most commonly used noise characterisation methods are the amplitude/power spectral density, and the Allan deviation/variance. These two methods are frequency-domain and time-domain methods, respectively. The power spectral density of a time series describes the power distribution amongst the frequency components which comprise the signal. On the other hand, the Allan variance gives the variance of measurements spaced between two points in time. These two methods are discussed in detail below.

H.1   The Power Spectral Density

Given a signal $f(t)$, we determine the average spectrum of the signal by carrying out a Fourier decomposition. The amplitude of the component with linear frequency ν is given by the Fourier transform of $f(t)$, $\mathscr{F}\{f(t)\} = F(\nu)$. The average energy per unit area of components with frequencies ν to $\nu + d\nu$ is $\langle |F(\nu)|^2 \rangle d\nu$, and so $\langle |F(\nu)|^2 \rangle$ is the *energy spectral density* of the signal. This has units of energy per unit area, per unit frequency.

If $f(t)$ is a stationary function (that is, its statistical averages are *invariant* as a function of time), then we must instead consider the *power* spectral density. Rather than integrate over $(-\infty, \infty)$, we instead integrate over an observation window of time width T by finding the truncated Fourier transform:

$$F_T(\nu) = \int_{-T/2}^{T/2} f(\nu) e^{-2\pi i \nu t} dt\,, \tag{H.1}$$

and then calculate the energy spectral density, $\langle |F_T(\nu)|^2 \rangle$. The power of a signal is simply the energy *per unit time*, and therefore the power spectral density is the energy spectral density per unit time. If we now extend the time window to infinity by taking the limit as $T \to \infty$, then we obtain the definition of the power spectral density [357]:

$$S(\nu) = \lim_{T \to \infty} \frac{1}{T} \langle |F_T(\nu)|^2 \rangle\,, \tag{H.2}$$

where the brackets $\langle\ \rangle$ denote the ensemble average over many realisations of $f(t)$.

H.1.1 Estimating the Power Spectral Density

In practice it is *impossible* to measure a signal for infinite time and hence the power spectral density can never *truly* be measured. However, it can be approximated to a good degree through the use of the discrete Fourier transform and a combination of averaging and windowing techniques. A few of the more frequently encountered power spectral density estimation methods are discussed below.

H.1.1.1 The Periodogram Method

The simplest approximation to the power spectral density is obtained by calculating the periodogram. The periodogram of a finitely sampled signal $f[k]$ is given by

$$P[n] = \frac{2|\mathscr{F}\{f[k]\}|^2}{N^2}, \tag{H.3}$$

where N is the number of points in the signal $f[k]$, and the factor of 2 accounts for power lost to the negative-frequency components when calculating a single-sided Fourier transform[1]. The power spectral density is then obtained from the periodogram by dividing by the frequency resolution (i.e. the bin width). Doing so normalises the power spectrum to the spectrum that would be measured by a 1 Hz-wide square filter, and hence the magnitude at each frequency component corresponds to what it would be if it were measured through a 1 Hz filter centred at that frequency. Mathematically, the power spectral density is calculated from the periodogram via

$$S_{\text{PSD}}[n] = \frac{P[n]}{\Delta \nu}, \tag{H.4}$$

where $\Delta \nu$ is the bin width. Although the periodogram technique is able to be used for calculating the power spectral density of a signal, it has a few drawbacks: it exhibits a large variance and suffers from spectral leakage (cf. Sec. G.4). Improvements can be made to the periodogram method in order to address these issues, and these techniques will be discussed in Secs. H.1.1.2 and H.1.1.3.

H.1.1.2 Bartlett's Method

Rather than calculating a *single* periodogram of the *entire* dataset, a more accurate approximation to the power spectrum is obtained by splitting the dataset into K non-overlapping sets of length M. For each segment, the periodogram is computed as per Eq. (H.3), replacing N with the length of each segment, M, instead. The result for each of the K sets is then averaged. Mathematically, power spectral density when calculated in this way is given by

$$S_{\text{PSD}}[n] = \frac{2}{\Delta \nu K M^2} \sum_{i=1}^{K} |\mathscr{F}\{f_i[k]\}|^2, \tag{H.5}$$

---

[1]Note that the factor of 2 should *not* be applied to the dc or Nyquist-frequency components as they are already present in the single-sided discrete Fourier transform

Figure H.1: Illustration of Welch's method. The time trace is split into *K* overlapping segments of length *M*. Each segment is multiplied by a window function to reduce spectral leakage and subsequently Fourier transformed. Averaging the Fourier transforms of all the segments results in a single spectrum with reduced variance. The power spectral density is then computed from the averaged Fourier transform.

where $f_i[k]$ is the i^{th} segment of the data, M is the number of points per segment, and K is the number of segments. Equation (H.5) is known as Bartlett's method [356,358,359], and results in a power spectral density which has *significantly* lower variance than the periodogram method described in Sec. H.1.1.1. In fact, the variance of the power spectral density is reduced by a factor of K, but this comes at the cost of reduced resolution since each Fourier transform is performed on fewer points. Additionally, the power spectral density does not extend down to as low Fourier frequency, since the 'measurement time' has been effectively reduced by a factor of K also.

H.1.1.3 Welch's Method

An even better approximation to the power spectrum is achieved by instead splitting the dataset into *K overlapping* sets of length *M*. These overlapping sets are then windowed (i.e the individual sets have a window function applied to them in the time domain) to reduce spectral leakage [360], and the periodograms are calculated and averaged as per Bartlett's method. This process — depicted in Fig. H.1 — is known as Welch's method [356,361], and is mathematically described by

$$S_{\text{PSD}}[k] = \frac{2}{\Delta\nu K M^2 N_w} \sum_{i=1}^{K} |\mathscr{F}\{f_i[k] \cdot w[k]\}|^2, \tag{H.6}$$

where $f_i[k] \cdot w[k]$ is the i^{th} segment of the data with an applied window function, and N_w is a normalisation factor to account for the loss of power as a result of applying the window function, which is given by

$$N_w = \frac{1}{M} \sum_{k=1}^{M} w[k]^2. \tag{H.7}$$

Welch's method amounts to averaging modified periodograms, and converges to the *true* power spectral density of the signal $f(t)$ as both the number of averages and the measurement period T approach infinity [362]. In principle it is possible to use *any* degree of overlap (provided that it is less than 100%, of course) when calculating the power spectral density using Welch's method. However, the most commonly implemented degree of overlap is 50%. In this case, the variance is reduced by an additional factor of approximately $18/11 \approx 64\%$ compared to Bartlett's method [361]. By overlapping the segments, the resulting power spectral density is safeguarded against the loss of information that arises due to the tapering of the window function applied to each segment.

H.1.2   Parseval's Theorem

Parseval's theorem provides a means of verifying that the power spectral density algorithm (or any algorithm based on the Fourier transform for that matter) is quantitatively correct. The theorem states that the integral of the square of the function $f(t)$ is equal to the integral of the square of its Fourier transform $\mathscr{F}\{f(t)\}$. This is essentially a statement that the Fourier transform does not violate conservation of energy — i.e. that power is conserved and no information is lost during the transform. Mathematically, Parseval's theorem is given by [363]

$$\int_{-\infty}^{\infty} |f(t)|^2 \mathrm{d}t = \int_{-\infty}^{\infty} |\mathscr{F}\{f(t)\}|^2 \mathrm{d}\nu. \tag{H.8}$$

In terms of the discrete Fourier transform and power spectral density, Eq. (H.8) becomes

$$\sum_{k=0}^{N-1} |f[k]|^2 = \frac{1}{N} \sum_{n=0}^{N-1} |F[n]|^2 = N\Delta\nu \sum_{n=0}^{N-1} S_{\text{PSD}}[n]. \tag{H.9}$$

Note that the discrete versions of Parseval's theorem, as given by Eq. (H.9), may in practice require additional pre-factors depending on how the discrete Fourier transform and/or power spectral density were calculated — e.g. contingent upon how their amplitude was scaled and whether or not window functions were used. In any case, both Bartlett's method and Welch's method of approximating the power spectral density — as described in Secs. H.1.1.2 and H.1.1.3 — have been shown to satisfy Parseval's theorem.

H.1.3   Root Mean Square

Provided that the amplitude and/or power spectral density is appropriately scaled (i.e. that it satisfies Parseval's theorem), it is possible to calculate the root mean square over a particular frequency range via

$$\sigma_{\text{rms}} = \sqrt{\int_{\nu_{\text{L}}}^{\nu_{\text{H}}} S_{\text{PSD}}\left(\nu\right) \mathrm{d}\nu} = \sqrt{\int_{\nu_{\text{L}}}^{\nu_{\text{H}}} \left[A_{\text{ASD}}\left(\nu\right)\right]^2 \mathrm{d}\nu}, \tag{H.10}$$

where ν_{L} is the lower frequency bound of interest, and ν_{H} is the upper frequency bound. If using the discrete Fourier transform to calculate the power spectral density, the integral can be numerically calculated using the trapezoidal technique, for example.

H.2   Allan Variance

As discussed in Sec. 6.2.2, the Allan variance is a time-domain measurement which quantifies the variance in the value of a signal as a function of integration time. This is useful when characterising the long-term performance of a measurement device — which in the context of this thesis, is the magnetometer. This section provides a mathematical derivation of the Allan variance, an estimation of the error associated with it, and discusses how it relates to the power spectral density.

It is a well-known result that the formula for calculating the variance of a random variable x is $\sigma^2 = \overline{x^2} - \overline{x}^2$. When sampling from a population of size N, the variance can be calculated explicitly via the Naïve algorithm [364]:

$$\sigma^2 = \frac{1}{N} \left\{ \sum_{i=1}^{N} x_i^2 - \frac{1}{N} \left[\sum_{n=1}^{N} x_i \right]^2 \right\}, \tag{H.11}$$

where x_i is the i^{th} sample of the random variable x. Using Bessel's correction to correct for biasing [365], the variance of a sub-sample of n observations ($n < N$) is then given by

$$\sigma^2 = \frac{1}{n-1} \left\{ \sum_{i=1}^{N} x_i^2 - \frac{1}{n} \left[\sum_{i=1}^{n} x_i \right]^2 \right\}. \tag{H.12}$$

If we consider x to be a time-dependent measured quantity, i.e. $x = x\left(t\right)$, then Eq. (H.12) can be used to calculate the variance of $x\left(t\right)$ over the observation period. If we measure M samples of $x\left(t\right)$, where T is the time between each sample and τ is the integration time of each observation, then we define an 'M-sample variance' [366]:

$$\sigma_{\text{M}}^2\left(M, T, \tau\right) = \frac{1}{M-1} \left\{ \sum_{i=0}^{M-1} \left[\frac{x\left(iT + \tau\right) - x\left(iT\right)}{\tau} \right]^2 - \frac{1}{M} \left[\sum_{i=0}^{M-1} \frac{x\left(iT + \tau\right) - x\left(iT\right)}{\tau} \right]^2 \right\}. \tag{H.13}$$

The M-sample variance is a powerful method for calculating the variance, as it enables

measurement dead time to be included by letting the time between measurements, T, to be different than the integration time, τ. The Allan variance is a modification of Eq. (H.13), and was first introduced as a measure of frequency stability in order to characterise the performance of frequency standards such as atomic clocks and oscillators. The Allan variance is defined as

$$\sigma_A^2(\tau) = \langle \sigma_M^2(2, \tau, \tau) \rangle,$$ (H.14)

where $\langle \dots \rangle$ denotes the expectation operator. The Allan variance is simply an M-sample variance, except with zero dead time (i.e. $T = \tau$), and only the difference between successive measurements is considered (i.e. $M = 2$). Through the substitution of Eq. (H.14) into Eq. (H.13), it can be shown that the Allan variance may be conveniently expressed as

$$\sigma_A^2(\tau) = \frac{1}{2} \langle (\overline{x}_{n+1} - \overline{x}_n)^2 \rangle,$$ (H.15)

where $\overline{x}_n$ is the n^{th} average measurement of the random variable x over the integration time τ. Given that the expectation operator is simply a weighted average, if we assume that each measurement is equally probable, then we can re-write Eq. (H.15) as

$$\sigma_A^2(\tau) = \frac{1}{2(M-1)} \sum_{n=1}^{M-1} (\overline{x}_{n+1} - \overline{x}_n)^2,$$ (H.16)

where again, M is the number of samples. Typically it is far more convenient to talk about standard deviation rather than variance, as in many cases, the standard deviation is relatable to the rms value. The Allan deviation is simply calculated by taking the square root of the Allan variance, and is generally plotted as a function of integration time, τ, where each point describes the standard deviation of observations spaced by the integration time. Explicitly, the Allan deviation can be calculated via

$$\sigma_A(\tau) = \sqrt{\frac{1}{2(M-1)} \sum_{n=1}^{M-1} (\overline{x}_{n+1} - \overline{x}_n)^2}.$$ (H.17)

Using this definition of the Allan deviation is both restrictive and cumbersome, as it means that the measurement must be performed *multiple times* using *different* integration times. There are multiple issues with this. Firstly, it would require that the measurement be performed many times with different integration times, which may be significantly time consuming. Secondly, it may not be physically possible to change the integration time using the measurement devices. Finally, this would require that the measured quantity have a time-independent reproducibility. In practice, one would choose an integration time that has a sufficient Nyquist frequency, and perform a single measurement for a given period of time. In this case, the Allan deviation can be *approximated* by splitting the data into bins to 'simulate' different integration times. Of course, the physical integration time of the measurement would place the lower limit on the integration time displayed in the Allan deviation.

H.2.1   Relation to the Power Spectral Density

Given that the Allan deviation/variance and the amplitude/power spectral density both quantify the noise/stability of a system, one would expect that the two quantities should be relatable to each other. The general M-sample variance, given by Eq. (H.13), is related to the power spectral density via [367]

$$\langle \sigma_{\mathrm{M}}^2 (M, T, \tau) \rangle = \frac{M}{M-1} \int_0^\infty \frac{\sin^2 (\pi \nu \tau)}{(\pi \nu \tau)^2} \left[1 - \frac{\sin^2 (\pi \nu T M)}{M^2 \sin^2 (\pi \nu T)} \right] S_{\mathrm{PSD}} (\nu) \, \mathrm{d}\nu, \qquad \text{(H.18)}$$

where, as per Eq. (H.13), M is the number of observations, T is the time between each sample, and τ is the integration time of each observation. Using Eqs. (H.14) and (H.18), the Allan variance $\sigma_{\mathrm{A}}^2 (\tau)$ is related to the power spectral density $S_{\mathrm{PSD}} (\nu)$ via the following integral [367,368]:

$$\sigma_{\mathrm{A}}^2 (\tau) = \int_0^\infty 2 \left[\frac{\sin^4 (\pi \nu \tau)}{(\pi \nu \tau)^2} \right] S_{\mathrm{PSD}} (\nu) \, \mathrm{d}\nu. \qquad \text{(H.19)}$$

Given that $\tau \in \mathbb{R}^*$, and assuming that $S_{\mathrm{PSD}} (\nu)$ is constant (i.e. white noise), then Eq. (H.19) simplifies to

$$\sigma_{\mathrm{A}}^2 (\tau) = \frac{S_{\mathrm{PSD}} (\nu)}{2\tau}. \qquad \text{(H.20)}$$

Equation (H.20) is a well-known result. It states that the Allan variance of white noise is inversely proportional to the integration time τ; or equivalently, the Allan deviation is inversely proportional to $\sqrt{\tau}$. This is *exactly* the reason why atomic clocks — which are tied to nature via an atomic transition — are more precise over long integration times. Given that an atomic transition frequency is a fixed value, long integration times result in higher precision via Eq. (H.20) since the mean of white noise is zero.

H.2.2   Estimating Error Bars

When calculating an Allan deviation, there will necessarily be less data points at larger integration times, in turn resulting in higher uncertainty for these points. It is therefore useful to calculate and plot error bars on each data point. In order to do this, consider the variance of a linear combination of random variables x_i:

$$\mathrm{var} \left(\sum_{i=1}^n a_i x_i \right) = \left\langle \left(\sum_{i=1}^n a_i x_i - \sum_{i=1}^n a_i \mu_i \right)^2 \right\rangle, \qquad \text{(H.21)}$$

where μ_i is the mean of x_i. Note that both terms have the same summation bounds, allowing this to be simplified by reducing to a single summation. Additionally, there is a common variable a_i which can be factored out. After performing these two operations, the variance simplifies to

$$\text{var}\left(\sum_{i=1}^{n} a_i x_i\right) = \left\langle \left(\sum_{i=1}^{n} a_i \left[x_i - \mu_i\right]\right)^2 \right\rangle. \tag{H.22}$$

Note that there is now a square of a summation. This can be re-written as the product of two summations using different 'dummy' indices:

$$\text{var}\left(\sum_{i=1}^{n} a_i x_i\right) = \left\langle \left(\sum_{i=1}^{n} a_i \left[x_i - \mu_i\right]\right)\left(\sum_{j=1}^{n} a_j \left[x_j - \mu_j\right]\right) \right\rangle. \tag{H.23}$$

The expectation operator is a linear operator, so Eq. (H.23) can be further simplified via factorisation:

$$\left\langle \left(\sum_{i=1}^{n} a_i \left[x_i - \mu_i\right]\right)\left(\sum_{j=1}^{n} a_j \left[x_j - \mu_j\right]\right) \right\rangle = \sum_{i=1}^{n}\sum_{j=1}^{n} a_i a_j \langle (x_i - \mu_i)(x_j - \mu_j)\rangle. \tag{H.24}$$

This may look like a difficult expression to evaluate, but there are some significant simplifications that can be made. Note that when evaluating the summations, when $i = j$ the term $\langle \dots \rangle$ is simply $\text{var}(x_i)$. On the other hand, when $i \neq j$, then $\langle \dots \rangle$ is simply $\text{cov}(x_i, x_j)$. Mathematically, this is described by

$$\langle (x_i - \mu_i)(x_j - \mu_j)\rangle = \begin{cases} \text{var}(x_i) & \text{for } i = j, \\ \text{cov}(x_i, x_j) & \text{for } i \neq j. \end{cases} \tag{H.25}$$

If we now consider the case that x_i and x_j are uncorrelated (as is typically the case when measuring noise, for example), then $\text{cov}(x_i, x_j) = 0$. This enables Eq. (H.24) to be greatly simplified to

$$\text{var}\left(\sum_{i=1}^{n} a_i x_i\right) = \sum_{i=1}^{n} a_i^2 \text{var}(x_i). \tag{H.26}$$

Equation (H.26) is a powerful equation, as it describes how the variance of a linear combination of random variables is calculated. A common example of when this may prove useful, is the case when the average of many measurements are performed. This is just a special case of Eq. (H.26) in which $a_i = n^{-1}$. Explicitly, the variance of a mean of measurements is given by

$$\text{var}\left(\sum_{i=1}^{n} \frac{x_i}{n}\right) = \frac{1}{n^2}\sum_{i=1}^{n} \text{var}(x_i). \tag{H.27}$$

In typical measurements, the variance will be the same in each measurement, and hence $\text{var}(x_i) = \text{var}(x)$. The summation will therefore simplify to $n\text{var}(x)$, enabling Eq. (H.27) to be written as

$$\text{var}\left(\sum_{i=1}^{n} \frac{x_i}{n}\right) = \frac{1}{n}\text{var}(x). \tag{H.28}$$

Figure H.2: Example of an Allan deviation of pseudo-randomly generated white noise with standard deviation $\sigma = 3$, calculated using Eq. (H.17). The blue markers correspond to the calculated Allan deviation, while the error bars indicate the uncertainty. Note that white noise averages down as $1/\sqrt{2\tau}$. Discrepancies from this scaling are attributed to the fact that this noise was generated using pseudo-random numbers.

This is a profound result. It implies that when performing an average of uncorrelated measurements, the variance in the measurements will decrease linearly as the number of measurements, n, increases. Typically the standard deviation will be discussed rather than the variance, in which case the standard deviation will decrease by $\sqrt{n}$. This is a well-known result that, when measuring white noise, the standard deviation will decrease as $\sqrt{n}$, and converge to zero as $n \to \infty$.

In many precision measurements, we are often concerned with measuring and quantifying noise. In optics experiments, quite often it is the case that measurements are limited by photon shot noise — which is white and uncorrelated. If we choose to quantify this noise using an Allan deviation, then it is useful to also calculate error bars. This is possible through the use of Eq. (H.28). For each point on the Allan deviation plot, $\sigma_A(\tau)$, we calculate error bars via $\sigma_A(\tau)/\sqrt{n}$, where n is the number of samples that were used to calculate that data point. Data points at longer integration times will therefore have larger error bars due to the reduction in sample points used in their calculation. An example of an Allan deviation calculated using generated white noise is shown in Fig. H.2.

Figure H.3: Illustration of the amplitude distribution of Gaussian white noise with mean μ and standard deviation σ. The probability of the noise falling within two standard deviations of the mean, i.e. $\mu \pm 2\sigma$, is given by erf $\left(\sqrt{2}\right) \approx 95.5\%$. Provided that $\mu = 0$, the rms of the noise amplitude corresponds to one standard deviation.

H.3   Types of Noise

As alluded to in Sec. 6.2.1, the frequency response of the noise produced by a source provides a fingerprint which can be used to both identify and characterise it. This is true *regardless* of whether the amplitude/power spectral density, or the Allan deviation/variance, is used to analyse the data — except in the latter case, the dependence will instead be on integration time instead of Fourier frequency.

The characteristic response of any noise source is typically described by a power law — that is, either the frequency or integration-time dependence is loosely described by ν^α or τ^α, where ν is the linear Fourier frequency, τ is the integration time, and $\alpha \in \mathbb{R}$ is some real parameter which typically takes values in the range $-2 \lesssim \alpha \lesssim 2$. Note that for a particular type of noise, the value of α will differ depending on whether it is analysed using an amplitude/power spectral density, or an Allan deviation/variance.

As an example, consider the case of Gaussian white noise depicted in Fig. H.3. This type of noise has an amplitude distribution which is Gaussian, characterised by a mean μ and standard deviation σ. In terms of the power spectral density, white noise is defined as $\alpha = 0$, as it has an equal power (or amplitude) in each frequency component (i.e. ν^0). However, in terms of Allan variance, white noise scales with integration time as τ^{-1}. A comparison between some common noise types, and their dependence on Fourier frequency and integration time in the power spectral density and Allan variance respectively, is shown in Fig. H.4.
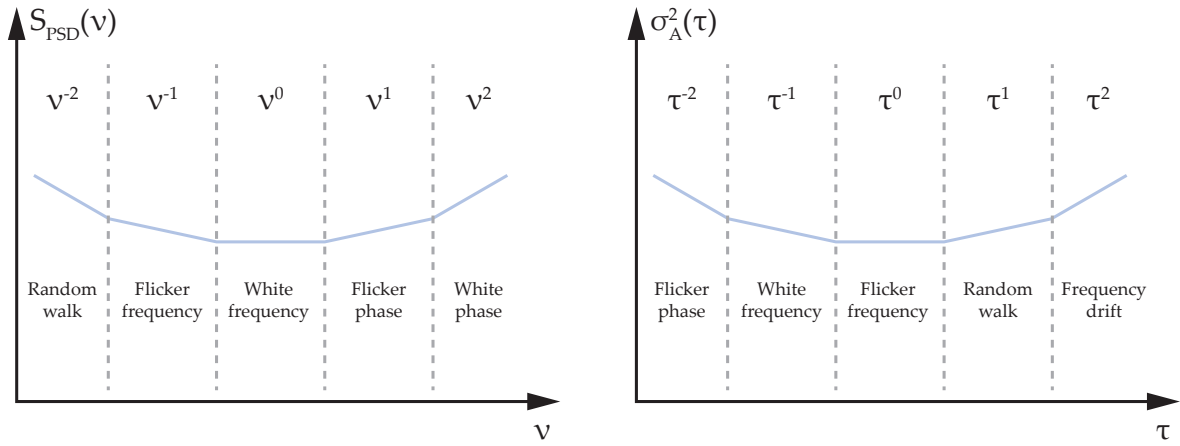
Figure H.4: Comparison of various types of noise, and their dependence on both Fourier frequency and integration time in the power spectral density (left) and Allan variance (right).

H.4 Aliasing

Aliasing is an effect which results in two (or more) different signals appearing *indistinguishable* when sampled at a finite rate (cf. Fig. H.5). These indistinguishable signals are known as *aliases* and, when sampling in the time domain, this phenomenon is called *temporal aliasing*.

Given a signal frequency ν and a sample rate of ν_s, aliases occur when the condition $\nu_a = \nu + n\nu_s$ is satisfied, where $n \in \mathbb{Z} \setminus \{0\}$. This is especially problematic when measuring *broadband* signals (e.g. white noise), as there are a *multitude* of frequency components which can *all* give rise to aliasing. For this reason, it is important to *always* low-pass filter the signal of interest in order to reduce aliasing. Generally, it is recommended to filter at a frequency *below* the Nyquist frequency of the measurement. If this is not performed, one can potentially *overestimate* the amount of noise in the system.

H.5 Cramér-Rao Lower Bound

As briefly discussed in Sec. 6.2.3, the Cramér-Rao lower bound is extremely useful for estimating the lower bound on the variance (and therefore standard deviation) of an unbiased estimator of a deterministic parameter [315,316,318,356]. In this section the CRLB will be derived, in general, for both the single-parameter and multivariate cases. Special cases which are common in the context of this thesis will also considered, with explicit expressions derived where possible.

H.5.1 Single-Parameter Case

In the derivation of the CRLB, we will consider signals $s[n;\theta]$ which contain a Gaussian white noise component $w[n]$:

$$x[n] = s[n;\theta] + w[n] \,, \tag{H.29}$$

Figure H.5: Example of the aliasing that can occur when undersampling a signal. The black dots correspond to a 70 Hz sampling of a 60 Hz sine wave (blue trace). This results in alias frequencies given by $\nu_a = 60 + 70n$ for all $n \in \mathbb{Z} \setminus \{0\}$. The red trace is a fit to the black markers using an alias frequency of $\nu_a = 10$ Hz, which is the $n = -1$ alias frequency.

where $n \in [0, N-1]$ is the n^{th} measurement of N total measurements, and θ is parameter upon which the signal $s[n; \theta]$ depends [315,318,356]. The CRLB states that the variance of any *unbiased* estimator $\hat{\theta}$ of θ is bounded by the reciprocal of the Fisher information $\mathcal{I}(\theta)$ [315,318,356]:

$$\text{var}\left(\hat{\theta}\right) \geq \frac{1}{\mathcal{I}(\theta)}, \tag{H.30}$$

where the Fisher information is defined by [315,318,356]

$$\mathcal{I}(\theta) = \mathbb{E}\left[\left(\frac{\partial \ln p(x;\theta)}{\partial \theta}\right)^2\right] = -\mathbb{E}\left[\frac{\partial^2 \ln p(x;\theta)}{\partial \theta^2}\right], \tag{H.31}$$

and where $p(x; \theta)$ is the *likelihood function*, and $\mathbb{E}$ denotes the expectation value. Assuming that the dataset $s[n]$ comes from a Gaussian process, $p(x; \theta)$ is given by [315,356]

$$p(x;\theta) = \frac{1}{(2\pi\sigma^2)^{N/2}} \exp\left\{-\frac{1}{2\sigma^2}\sum_{n=0}^{N-1}(x[n] - s[n;\theta])^2\right\}, \tag{H.32}$$

where σ^2 is the variance of the Gaussian white noise. By explicitly computing the partial derivatives, and taking the expected value with respect to $p(x; \theta)$, the CRLB for the single-parameter case is obtained[2]:

---

[2]Note that the formulae derived for the single-parameter case is *not* restricted to the case where the signal

$$\text{var}\left(\hat{\theta}\right) \geq \frac{\sigma^2}{\sum\limits_{n=0}^{N-1}\left(\frac{\partial s[n;\theta]}{\partial \theta}\right)^2} \, . \tag{H.33}$$

H.5.2 Multivariate Case

Now consider multivariate signals $s\left[n;\boldsymbol{\theta}\right]$ which contain a Gaussian white noise component $w\left[n\right]$:

$$x\left[n\right] = s\left[n;\boldsymbol{\theta}\right] + w\left[n\right]\,, \tag{H.34}$$

where everything is defined as per the single-parameter case, except that $\boldsymbol{\theta} = \left[\theta_1, \theta_2, \ldots, \theta_p\right]^T$ is a set of p parameters [315,356]. The CRLB of the i^{th} parameter, θ_i, is given by the $(i,i)^{\text{th}}$ element of the inverse of the Fisher information matrix [315,356]:

$$\text{var}\left(\theta_i\right) = \left[\boldsymbol{\mathcal{I}}^{-1}\left(\boldsymbol{\theta}\right)\right]_{i,i}\,, \tag{H.35}$$

where $\boldsymbol{\mathcal{I}}$ is the $p \times p$ Fisher information matrix, with each element defined by [315,356]

$$\left[\boldsymbol{\mathcal{I}}\left(\boldsymbol{\theta}\right)\right]_{i,j} = -\mathbb{E}\left\{\frac{\partial^2 \ln p\left(\boldsymbol{x};\boldsymbol{\theta}\right)}{\partial\theta_i\partial\theta_j}\right\}\,, \tag{H.36}$$

and where $i, j \in [1, p]$, and $p\left(\boldsymbol{x};\boldsymbol{\theta}\right)$ is the likelihood function, defined by Eq. (H.32), but where $\theta \to \boldsymbol{\theta}$.

H.5.3 Special Cases

Although the Cramér-Rao lower bound has been derived in general for both the single-parameter and multivariate cases, we now consider a few special cases which are relevant in the context of this thesis.

H.5.3.1 Fixed Offset with White Noise

For a fixed offset in the presence of Gaussian white noise, the signal $s\left[n\right] = a$, where a is the offset. Therefore, Eq. (H.29) becomes

$$x\left[n\right] = a + w\left[n\right]\,. \tag{H.37}$$

By evaluating Eq. (H.33), one obtains

$$\text{var}\left(\hat{a}\right) \geq \frac{\sigma^2}{N}\,, \tag{H.38}$$

which is the familiar result of the variance of uncorrelated white Gaussian noise as derived in App. H.2.2 — specifically, Eq. (H.28).

---

form only has one variable, but rather, they apply to the case where only one parameter of the signal is unknown.

H.5.3.2 Sinusoid Amplitude with White Noise

For a sinusoidal signal, we have $s[n] = a\cos(2\pi\nu\Delta tn + \varphi)$, where a is the amplitude, ν is the linear frequency, Δt is the sample time, and φ is an arbitrary phase offset. In this case, we have

$$x[n] = a\cos(2\pi\nu\Delta tn + \varphi) + w[n]. \tag{H.39}$$

Upon calculating Eq. (H.36) for this signal, significant simplifications can be made using the following identities [369]:

$$\frac{1}{N^{i+1}}\sum_{n=0}^{N-1} n^i \sin(4\pi\nu\Delta tn + 2\varphi) \approx 0, \tag{H.40}$$

$$\frac{1}{N^{i+1}}\sum_{n=0}^{N-1} n^i \cos(4\pi\nu\Delta tn + 2\varphi) \approx 0, \tag{H.41}$$

for $i \in [0,2]$. By making these simplifications, and subsequently calculating the inverse of the Fisher information matrix, we obtain the following CRLBs [315]:

$$\text{var}(\hat{a}) \geq \frac{2\sigma^2}{N}, \tag{H.42}$$

$$\text{var}(\hat{\nu}) \geq \frac{24\sigma^2}{(2\pi a\Delta t)^2 N(N^2-1)}, \tag{H.43}$$

$$\text{var}(\hat{\varphi}) \geq \frac{4\sigma^2(2N-1)}{a^2 N(N+1)}. \tag{H.44}$$

H.5.3.3 Exponentially Damped Sinusoid with White Noise

An exponentially damped sinusoid takes the form $s[n] = a\cos(2\pi\nu\Delta tn + \varphi)\exp(-t/T_2)$, where a is the amplitude, ν is the linear frequency, Δt is the sample time, φ is an arbitrary phase offset, and T_2 is a decay constant. In this case, we have

$$x[n] = a\cos(2\pi\nu\Delta tn + \phi)\exp\left(-\frac{t}{T_2}\right) + w[n]. \tag{H.45}$$

The derivation of the CRLBs for an exponentially damped sinusoid is similar to that of a non-damped sinusoid. Using Eqs. (H.40) and (H.41), as well as the following identities:

$$\sum_{n=0}^{N-1} n\exp(-2\beta n) = -\frac{\mathrm{d}}{\mathrm{d}\beta}\left(\frac{1}{2}\sum_{n=0}^{N-1}\exp(-2\beta n)\right), \tag{H.46}$$

$$\sum_{n=0}^{N-1} n^2\exp(-2\beta n) = \frac{\mathrm{d}^2}{\mathrm{d}\beta^2}\left(\frac{1}{4}\sum_{n=0}^{N-1}\exp(-2\beta n)\right), \tag{H.47}$$

it is possible to show that [136,139]

$$\text{var}\left(\hat{v}\right) \gtrsim \frac{12\sigma^2}{\left(2\pi\right)^2 a^2 v_{\text{BW}} T^3} \mathcal{C}\left(N\right),\tag{H.48}$$

where $\mathcal{C}\left(N\right)$ is a factor given by [136,139]

$$\mathcal{C}\left(N\right) = \left\{\frac{N^3}{12} \frac{\left(1-\exp\left(-\frac{t}{T_2}\right)^2\right)^3 \left(1-\exp\left(-\frac{t}{T_2}\right)^{2N}\right)}{\exp\left(-\frac{t}{T_2}\right)^2 \left(1-\exp\left(-\frac{t}{T_2}\right)^{2N}\right)^2 - N^2 \exp\left(-\frac{t}{T_2}\right)^{2N} \left(1-\exp\left(-\frac{t}{T_2}\right)^2\right)^2}\right\}.\tag{H.49}$$

Note that the result given by Eq. (H.48) is equivalent to Eq. (6.6) up to the pre-factor $\mathcal{C}\left(N\right)$ given by Eq. (H.49). This factor depends on two parameters: the number N of data points, and the decay time T_2. For decreasing decay time, $\mathcal{C}\left(N\right)$ increases, yielding an increased uncertainty in the frequency estimate. However, for large T_2 we have

$$\lim_{T_2 \to \infty} \mathcal{C}\left(N\right) = \frac{N^2}{N^2 - 1} \approx 1,\tag{H.50}$$

in which case, Eq. (H.48) simplfies to Eq. (6.6) as expected.

H.5.3.4 Sinusoid on a Linear Gradient with White Noise

Quite often one may encounter a situation in which a sinusoidal signal is present on a linear polynomial background, in the presence of Gaussian white noise. A specific example in which this situation arises is due to sinusoidal magnetic-field modulation when using the instantaneous-phase-retrieval technique (cf. Ch. 7 and Ch. 11). When using this technique, the signal takes the form $s\left[n\right] = 2\pi v_{\text{c}}\Delta tn + \left(\delta v / v_{\text{m}}\right)\sin\left(2\pi v_{\text{m}}\Delta tn + \varphi_{\text{m}}\right) + \varphi_{\text{c}}$, where v_{c} is the linear 'carrier' frequency, Δt is the sample time, δv is the frequency deviation, v_{m} is the linear modulation frequency, and φ_{m} and φ_{c} are the arbitrary phase offsets of the modulation and carrier signals, respectively. In this case, we have

$$x\left[n\right] = 2\pi v_{\text{c}}\Delta tn + \left(\frac{\delta v}{v_{\text{m}}}\right)\sin\left(2\pi v_{\text{m}}\Delta tn + \varphi_{\text{m}}\right) + \varphi_{\text{c}} + w\left[n\right].\tag{H.51}$$

Evaluating the Fisher information matrix using Eq. (H.36), and performing simplifications given by Eqs. (H.40) and (H.41), yields a Fisher information matrix with elements given by

$$\mathcal{I}\left(\boldsymbol{\theta}\right) = \begin{pmatrix} \mathcal{I}_{11} & 0 & 0 & 0 & \mathcal{I}_{15} \\ 0 & \mathcal{I}_{22} & \mathcal{I}_{23} & 0 & 0 \\ 0 & \mathcal{I}_{32} & \mathcal{I}_{33} & \mathcal{I}_{34} & 0 \\ 0 & 0 & \mathcal{I}_{43} & \mathcal{I}_{44} & 0 \\ \mathcal{I}_{51} & 0 & 0 & 0 & \mathcal{I}_{55} \end{pmatrix}.\tag{H.52}$$

In this case, the variance in the frequency estimators, var$\left(\hat{v}_{\text{c}}\right)$ and var$\left(\hat{v}_{\text{m}}\right)$, are given by

$$\text{var}\left(\hat{v}_{\text{c}}\right) \geq \mathcal{I}_{11}^{-1} = \frac{\mathcal{I}_{55}}{\mathcal{I}_{11}\mathcal{I}_{55} - \mathcal{I}_{15}\mathcal{I}_{51}},\tag{H.53}$$

$$\text{var}\left(\hat{\nu}_{\text{m}}\right) \geq \mathcal{I}_{33}^{-1} = \frac{\mathcal{I}_{22}\mathcal{I}_{44}}{\mathcal{I}_{22}\left(\mathcal{I}_{33}\mathcal{I}_{44} - \mathcal{I}_{34}\mathcal{I}_{43}\right) - \mathcal{I}_{23}\mathcal{I}_{32}\mathcal{I}_{44}}, \tag{H.54}$$

and where the Fisher information matrix elements are explicitly given by

$$\mathcal{I}_{11} = \frac{2\pi^2\Delta t^2 N\left(2N-1\right)\left(N-1\right)}{3\sigma^2}, \tag{H.55}$$

$$\mathcal{I}_{15} = \frac{N\left(N-1\right)\pi\Delta t}{\sigma^2} = \mathcal{I}_{51}, \tag{H.56}$$

$$\mathcal{I}_{22} = \frac{N}{2\nu_{\text{m}}^2\sigma^2}, \tag{H.57}$$

$$\mathcal{I}_{23} = -\frac{N\delta\nu}{2\nu_{\text{m}}^3\sigma^2} = \mathcal{I}_{32}, \tag{H.58}$$

$$\mathcal{I}_{33} = \frac{N\delta\nu^2\left[2\pi^2\nu_{\text{m}}^2\Delta t^2\left(2N-1\right)\left(N-1\right)+3\right]}{6\nu_{\text{m}}^4\sigma^2}, \tag{H.59}$$

$$\mathcal{I}_{34} = \frac{N\left(N-1\right)\pi\delta\nu^2\Delta t}{2\nu_{\text{m}}^2\sigma^2} = \mathcal{I}_{43}, \tag{H.60}$$

$$\mathcal{I}_{44} = \frac{N\delta\nu^2}{2\nu_{\text{m}}^2\sigma^2}, \tag{H.61}$$

$$\mathcal{I}_{55} = \frac{N}{\sigma^2}. \tag{H.62}$$

Upon inserting the above Fisher information matrix elements into Eqs. (H.53) and (H.54), and simplifying using the fact that $N^3 \gg N$ as per Sec. 6.2.3, the following CRLBs for Eq. (H.51) are derived:

$$\text{var}\left(\hat{\nu}_{\text{c}}\right) \gtrsim \frac{6\sigma^2}{\left(2\pi\right)^2\nu_{\text{BW}}T^3}, \tag{H.63}$$

$$\text{var}\left(\hat{\nu}_{\text{m}}\right) \gtrsim \frac{12\nu_{\text{m}}^2\sigma^2}{\left(2\pi\right)^2\delta\nu^2\nu_{\text{BW}}T^3}. \tag{H.64}$$

where $\nu_{\text{BW}} = \left(2\Delta t\right)^{-1}$ is the Nyquist frequency (or measurement bandwidth). It's interesting to note that Eq. (H.64) is *exactly* the same as Eq. (6.6). This implies that the presence of a linear polynomial background *does not* affect the precision to which the frequency of a sinusoid can be determined.

Numerical Computation of the Hilbert Transform

Although the definition of the Hilbert transform given by Eq. (7.8) is rather simple, this integral becomes *exceedingly* difficult to compute for anything other than basic functions. This often necessitates the computation of the Hilbert transform *numerically*. Furthermore, when working with experimentally measured data, one must compute the Hilbert transform numerically in any case. This appendix introduces two methods — one time-domain technique, and one frequency-domain technique — with which to calculate the Hilbert transform numerically.

I.1 Time-Domain Techniques

Assume that the signal $x(t)$ has been sampled N times, in intervals of Δt, to give a sequence $x_k = x(k\Delta t)$ for $k = 1, 2, \ldots, N$. The corresponding Hilbert transform $\mathcal{H}_k$ of the sequence x_k, assuming linear variation over the sampling interval Δt to $N\Delta t$, is given by [370]

$$
\begin{aligned}
\mathcal{H}_k &= \frac{1}{\pi} \int_{\Delta t}^{N\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau \\
&= \frac{1}{\pi} \left[\int_{\Delta t}^{2\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau + \ldots + \int_{(k-2)\Delta t}^{(k-1)\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau + \int_{(k-1)\Delta t}^{k\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau \right. \\
&\quad \left. + \int_{k\Delta t}^{(k+1)\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau + \int_{(k+1)\Delta t}^{(k+2)\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau + \ldots + \int_{(N-1)\Delta t}^{N\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau \right] \\
&= \frac{1}{\pi} \left[\sum_{i=1}^{k-2} \int_{i\Delta t}^{(i+1)\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau + \int_{(k-1)\Delta t}^{k\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau \right. \\
&\quad \left. + \int_{k\Delta t}^{(k+1)\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau + \sum_{j=k+2}^{N} \int_{(j-1)\Delta t}^{j\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau \right],
\end{aligned}
\tag{I.1}
$$

where the Cauchy principal value integral is implied, though the notation has been dropped to improve readability. Through substitution and the use of summation notation, Eq. (I.1) can be simplified to

$$
\mathcal{H}_k = \frac{1}{\pi} \left[\sum_{i=1}^{k-2} I_i^{(f)} + I_k^{(c)} + \sum_{j=k+2}^{N} I_j^{(r)} \right],
\tag{I.2}
$$

where the functions $I_i^{(f)}$, $I_k^{(c)}$, and $I_j^{(r)}$ are given by [370,371]

$$I_i^{(f)} \equiv \int_{i\Delta t}^{(i+1)\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau, \tag{I.3}$$

$$I_k^{(c)} \equiv \int_{(k-1)\Delta t}^{k\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau + \int_{k\Delta t}^{(k+1)\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau, \tag{I.4}$$

$$I_j^{(r)} \equiv \int_{(j-1)\Delta t}^{j\Delta t} \frac{x(\tau)}{k\Delta t - \tau} d\tau. \tag{I.5}$$

By assuming that the signal $x(t)$ varies linearly over the sampling interval[1], Eqs. (I.3) to (I.5) can be written explicitly via slight modification [370] to the expressions found in Ref. [371]:

$$I_i^{(f)} = x_i \ln\left(\frac{k-i}{k-i-1}\right) + (x_{i+1} - x_i)\left[-1 + (k-i)\ln\left(\frac{k-i}{k-i-1}\right)\right], \tag{I.6}$$

$$I_k^{(c)} = x_{k-1} - x_{k+1}, \tag{I.7}$$

$$I_j^{(r)} = x_j \ln\left(\frac{j-k-1}{j-k}\right) + (x_{j-1} - x_j)\left[1 + (j-k)\ln\left(\frac{j-k-1}{j-k}\right)\right]. \tag{I.8}$$

The Hilbert transform of a sinusoid, calculated using Eq. (I.2), is presented in Fig. I.1.

I.2 Frequency-Domain Techniques

In addition to the time-domain technique described in Sec. I.1, it is also possible to use frequency-domain techniques to calculate the Hilbert transform. As discussed in multiple previous sections; by definition, the Hilbert transform is a convolution of the signal $x(t)$ and the function $(\pi t)^{-1}$. Using the convolution theorem, Eq. (7.10), yields the following expression for the Fourier transform of the Hilbert transform:

$$\begin{aligned} \mathscr{F}\{\mathcal{H}\{x(t)\}\} &= \mathscr{F}\left\{\frac{1}{\pi t}\right\} \mathscr{F}\{x(t)\} \\ &= -i\,\mathrm{sgn}(\omega)\,\mathscr{F}\{x(t)\}. \end{aligned} \tag{I.9}$$

It immediately follows that the Hilbert transform of the signal $x(t)$ can be calculated via

$$\begin{aligned} \mathcal{H}\{x(t)\} &= \mathscr{F}^{-1}\{\mathscr{F}\{\mathcal{H}\{x(t)\}\}\} \\ &= \mathscr{F}^{-1}\{-i\,\mathrm{sgn}(\omega)\,\mathscr{F}\{x(t)\}\}. \end{aligned} \tag{I.10}$$

In terms of computing the Hilbert transform analytically, Eq. (I.10) is computed by simply using the fast Fourier transform and inverse fast Fourier transform algorithms. In practice, this corresponds to discarding negative frequency components, which is achieved through the use of a window function [372]:

---

[1]Due to the Taylor expansion of most continuous functions being predominantly linear in a small region about the expansion point, this is a valid assumption provided that the sampling rate is sufficiently high relative to the rate of change of the function.

Figure I.1: Left: Input sinusoidal signal with a frequency of 1 Hz (blue), and the corresponding Hilbert transform, calculated using Eq. (I.2), using 10 Hz (purple) and 1 kHz (red) sample frequencies. Right: Absolute error between the analytic Hilbert transform, and the numerically computed Hilbert transform using Eq. (I.2). The purple trace corresponds to a sample frequency of 10 Hz, while the red trace corresponds to a sample frequency of 1 kHz. Note that there is significant error in the time-domain computation of the Hilbert transform, especially at the beginning and end points of the dataset.

$$\mathcal{H}[n] = \mathscr{F}^{-1}\left\{w[k] \cdot \mathscr{F}\left\{x[k]\right\}\right\}, \tag{I.11}$$

where the window function $w[k]$ is given by [372]

$$w[k] = \begin{cases} 2 & \text{for } k = 1, 2, \ldots, \frac{N}{2} - 1, \\ 1 & \text{for } k = 0 \text{ and } k = \frac{N}{2}, \\ 0 & \text{for } k = \frac{N}{2} + 1, \ldots, N - 1. \end{cases} \tag{I.12}$$

That is, all negative-frequency components (cf. App. G.3.1) are set to zero.

The Hilbert transform of a sinusoid, calculated using Eq. (I.10), is presented in Fig. I.2. Evident in Fig. I.2 is a Hilbert transform which is faithfully reproduced, in contrast to that produced by the time-domain technique shown in Fig. I.1. All Hilbert transforms calculated in this thesis have been done so via the frequency-domain technique unless stated otherwise.

Figure I.2: Left: Input sinusoidal signal with a frequency of 1 Hz (blue), and the corresponding Hilbert transform, calculated using Eq. (I.11), using a 1 kHz sample frequency (red). Right: absolute error between the analytic Hilbert transform, and the numerically computed Hilbert transform using Eq. (I.11). Note that the magnitude of the error is consistent with numerical noise, rather than any discrepancy between the analytic solution and the numerical result.

Magneto-Optical Resonance Shift due to the Nonlinear Zeeman Effect

As discussed in Sec. 3.3.2, the energy shift of the Zeeman sublevels, due to the nonlinear Zeeman effect, is given by

$$\Delta E^{(\mathrm{BR})} = -\frac{\Delta E_{\mathrm{hfs}}}{2(2I+1)} + g_I \mu_{\mathrm{B}} m_F B \pm \frac{\Delta E_{\mathrm{hfs}}}{2} \sqrt{1 + \frac{4 m_F x}{2I+1} + x^2}, \tag{J.1}$$

where $\Delta E_{\mathrm{hfs}} = A_{\mathrm{hfs}} (I + 1/2)$ is the hyperfine splitting, A_{hfs} is the magnetic dipole constant, g_I is the nuclear g-factor, μ_{B} is the Bohr magneton, B is the magnetic-field strength, I is the nuclear quantum number, and x is a dimensionless 'field-strength parameter' defined by

$$x \equiv \frac{(g_J - g_I) \mu_{\mathrm{B}} B}{\Delta E_{\mathrm{hfs}}}, \tag{J.2}$$

where g_F is the ground-state Landé g-factor. By dropping the first two terms of Eq. (J.1), substituting Eq. (J.2) and performing a Taylor-series expansion, we obtain

$$\Delta E^{(\mathrm{NL})} \approx \pm \frac{\Delta E_{\mathrm{hfs}}}{2} \pm \frac{(g_J - g_I) \mu_{\mathrm{B}} m_F B}{2I+1} \pm \frac{\left[\left(4 m_F{}^2 - 4I^2 - 4I - 1\right) \mu_{\mathrm{B}} (g_I - g_J)^2 \right] B^2}{4 (2I+1)^2 \Delta E_{\mathrm{hfs}}} + \mathcal{O}\left(B^3\right). \tag{J.3}$$

Now, consider only the term that is quadratic in the field strength. By substituting the nuclear quantum number for rubidium, $I = 3/2$, into this term, we obtain

$$\Delta E^{(\mathrm{NL})} \approx \pm \frac{\left(4 m_F{}^2 - 16\right) \mu_{\mathrm{B}}{}^2 (I - g_J)^2 B^2}{64 \Delta E_{\mathrm{hfs}}}. \tag{J.4}$$

For ^{87}Rb, the nuclear g-factor is given by $g_I = -0.0009951414$ [229]; and for the $F = 2$ ground state, the fine-structure Landé g-factor is $g_J = 2.00233113$ [229]. Substituting these into Eq. (J.4) yields

$$\Delta E^{(\mathrm{NL})} \approx \pm \frac{0.00627081 \left(4 m_F{}^2 - 16\right) \mu_{\mathrm{B}}{}^2 B^2}{\Delta E_{\mathrm{hfs}}}. \tag{J.5}$$

In the context of nonlinear magneto-optical rotation with linearly polarised light, coherences are generated between ground-state Zeeman sublevels with $|\Delta m_F| = 2$. That is, coherences form between $|m_F = 2\rangle$ and $|m_F = 0\rangle$, $|m_F = 1\rangle$ and $|m_F = -1\rangle$, and $|m_F = 0\rangle$ and $|m_F = -2\rangle$. In order to calculate the shift in the magneto-optical resonances due to the nonlinear Zeeman effect, we calculate the difference between $\Delta E^{(\mathrm{NL})}$ for the $|m_F = \pm 2\rangle$ and $|m_F = 0\rangle$ states. Doing so yields the following:

$$\delta E^{(\text{NL})} \approx \frac{0.00627081 \left(4 \left(\pm 2\right)^2 - 16\right) \mu_{\text{B}}{}^2 B^2}{\Delta E_{\text{hfs}}} - \frac{0.00627081 \left(-16\right) \mu_{\text{B}}{}^2 B^2}{\Delta E_{\text{hfs}}}$$
$$\approx \pm \frac{1.00333 \left(\mu_{\text{B}} B\right)^2}{\Delta E_{\text{hfs}}} \, . \tag{J.6}$$

Dropping the factor of order unity, taking the absolute value, and dividing both sides by the reduced Planck constant, yields the following equation for the resonance shift due to the nonlinear Zeeman effect:

$$\delta^{\text{NLZ}} \approx \frac{\left(\mu_{\text{B}} B\right)^2}{\hbar \Delta E_{\text{hfs}}} \, . \tag{J.7}$$

It should be noted that the resonance shift δ^{NLZ} for the coherence between $|m_F = 1\rangle$ and $|m_F = -1\rangle$ is zero, and hence the central resonance is *unaffected* by the quadratic Zeeman effect.

Phase-Sensitive Detection

Measurements detailed within this thesis are most commonly performed using phase-sensitive detection via a lock-in amplifier. Phase-sensitive detection provides a means of measuring small ac signals that are obscured by noise that is thousands of times larger than the signal. This is made possible by performing a homodyne detection in which the input signal to the lock-in amplifier is mixed with a reference signal of the same frequency as the signal of interest.

A lock-in amplifier takes an input signal, V_{sig}, amplifies it by a gain factor G, and then multiplies it by the reference signal, V_{ref}, using a mixer. This results in an output signal that is the product of two sine waves:

$$V_{\text{out}} = G V_{\text{sig}} V_{\text{ref}} \sin\left(\omega_{\text{sig}} t + \varphi_{\text{sig}}\right) \sin\left(\omega_{\text{ref}} t + \varphi_{\text{ref}}\right) . \tag{K.1}$$

Using a product-to-sum trigonemetric identity, Eq. (K.1) can be written as the sum of two cosine waves: one oscillating at the difference frequency, $\omega_{\text{sig}} - \omega_{\text{ref}}$, and the other at the sum frequency, $\omega_{\text{sig}} + \omega_{\text{ref}}$:

$$\begin{aligned}
V_{\text{out}} &= \frac{G}{2} V_{\text{sig}} V_{\text{ref}} \cos\left(\left[\omega_{\text{sig}} - \omega_{\text{ref}}\right] t + \varphi_{\text{sig}} - \varphi_{\text{ref}}\right) \\
&+ \frac{G}{2} V_{\text{sig}} V_{\text{ref}} \cos\left(\left[\omega_{\text{sig}} + \omega_{\text{ref}}\right] t + \varphi_{\text{sig}} + \varphi_{\text{ref}}\right) .
\end{aligned} \tag{K.2}$$

From here, the signal is passed through a Butterworth filter with a cut-off frequency set sufficiently low in order to remove the sum-frequency component, leaving only the difference-frequency component. However, if the reference signal is the *same* frequency as the signal of interest, i.e. $\omega_{\text{sig}} = \omega_{\text{ref}}$, then the difference-frequency component is simply a dc signal whose amplitude is proportional to the signal amplitude:

$$V_{\text{out}} = \frac{G}{2} V_{\text{sig}} V_{\text{ref}} \cos\left(\varphi_{\text{sig}} - \varphi_{\text{ref}}\right) . \tag{K.3}$$

Note that Eq. (K.3) is proportional to $\cos\left(\varphi_{\text{sig}} - \varphi_{\text{ref}}\right)$. This means that, if the phase of the reference signal is set such that $\varphi_{\text{sig}} - \varphi_{\text{ref}} = \pm\pi/2$, then the output of the lock-in amplifier will be zero. This unwanted side-effect can be eliminated by introducing a *second* phase-sensitive detector, and sending a $\pi/2$ phase-shifted copy of the signal to it. This results in two outputs: one of which is proportional to $\cos\left(\varphi_{\text{sig}} - \varphi_{\text{ref}}\right)$, and the other is proportional to $\sin\left(\varphi_{\text{sig}} - \varphi_{\text{ref}}\right)$.

After multiplying out the pre-factors, the two lock-in amplifier outputs are known as the "in-phase" (X) and "quadrature" (Y) components respectively, and are given by

$$X = GV_{\text{sig}} \cos\left(\varphi_{\text{sig}} - \varphi_{\text{ref}}\right) , \tag{K.4}$$

$$Y = GV_{\text{sig}} \sin\left(\varphi_{\text{sig}} - \varphi_{\text{ref}}\right) . \tag{K.5}$$

The magnitude, R, of the signal can then be calculated via $R = \sqrt{X^2 + Y^2}$, while the phase can be calculated via $\phi_{\text{sig}} - \phi_{\text{ref}} = \arctan\left(Y/X\right)$.

Monte Carlo Methods

Monte Carlo methods are a class of computational algorithms which enable analytically difficult (or impossible) problems to be numerically solved with relative ease. By assuming that the problem is deterministic, the law of large numbers is exploited in order to estimate the expected value of a random variable by computing the sample mean of a large set of independent samples of the variable.

This chapter explains in detail the Monte Carlo calculations that were performed in order to estimate the fraction of atoms that travelled between the pump and probe beams in the frequency-comb measurements described in Ch. 9 and Ref. [344], as well as their mean transit time. Additionally, Monte Carlo simulations of atom-wall collisions in cylindrical vapour cells are described, the results of which are used to estimate the wall-collision-induced resonance width in Sec. 5.2.2.4.

L.1 Generation of Random Initial Atomic Positions

Given that the observation of nonlinear magneto-optical rotation in the probe beam depends on the presence of ground-state coherence (cf. Ch. 4) — which is generated via optical pumping with linearly polarised light — the Monte Carlo modelling performed in this thesis typically begins by generating atoms within the pump volume. Since the pump beam is well approximated by a cylinder (i.e. it is a transverse electromagnetic (TEM)$_{00}$ mode), it is important to assign random atomic positions within a cylindrical volume.

Given N atoms, an initial position array $\vec{p}_0 = \{\vec{p}_1, \vec{p}_2, \ldots, \vec{p}_N\}$ is calculated by generating three random numbers for each atom: two which specify its value in the xy-plane (via a radius r and angle θ), and one which specifies its value along the z-axis. One might naively believe that by generating random positions using a radius $r \in [0, 1]$ and an angle $\theta \in [0, 2\pi)$, and subsequently transforming to Cartesian coordinates, the result would be a set of points which are *uniformly* distributed. However, while the points are uniform in (r, θ) coordinates, the Cartesian point distribution (x, y) will be *too dense* near the centre and *too sparse* near the outer edge — as shown in Fig. L.1(a). This is due to the fact that for a uniform distribution, the probability of generating a point in any finite region is proportional to the area of the region. Since the infinitesimal area element of a disk is given by $dA = 2\pi r dr$, which scales as $\mathcal{O}\left(r^2\right)$, the probability of generating a point inside the disk at some radius r should also be proportional to r^2 in order to be uniformly distributed. It is possible to adjust the probability by simply changing the distribution of the r variable. Rather than sampling r uniformly, one should sample as the square root of the uniform distribution [373], i.e. $\sqrt{r}$. The square-root distribution decreases the probability of points being generated near the

(a) Non-uniform   **(b)** Uniform

Figure L.1: Example of non-uniform (a) and uniform (b) sampling of points on a disk. The former was obtained by generating uniformly distributed positions using a radius $r \in [0, 1]$ and an angle $\theta \in [0, 2\pi)$, and subsequently transforming to Cartesian coordinates via $(r \cos \theta, r \sin \theta)$. The latter was obtained in a similar fashion, but with a transformation via $(\sqrt{r} \cos \theta, \sqrt{r} \sin \theta)$.

centre of the disk, while simultaneously increasing the probability of points being generated near the outer edge. An example of sampling points from the square-root distribution can be seen in Fig. L.1(b).

Denoting a random number between n_1 and n_2 by $\mathcal{R}(n_1, n_2)$, the initial position vector of the i^{th} atom in Cartesian components is calculated via

$$
\vec{p}_i = \left\{ -r\sqrt{\mathcal{R}(0,1)} \cos\left[2\pi\mathcal{R}(0,1)\right] - \frac{d}{2}, r\sqrt{\mathcal{R}(0,1)} \sin\left[2\pi\mathcal{R}(0,1)\right], \left(\frac{l}{2}\right)\mathcal{R}(-1,1) \right\},
\tag{L.1}
$$

where d is the separation between the centres of the pump and probe beams, r is the radius of the pump beam, and l is the beam length.

L.2 Assigning Random Initial Velocities

In a similar fashion to the way initial positions were allocated, each atom is also assigned an initial velocity vector $v_0 = \{v_1, v_2, \ldots, v_N\}$. However, in this case, each atom was assigned velocity components by random sampling of a three-dimensional Maxwell-Boltzmann distribution. Given that the atoms have no preferred direction, the Maxwell-Boltzmann distribution in each Cartesian component is simply a normal distribution with a mean of $\mu = 0$, and a standard deviation of $\sigma = \sqrt{k_B T / m}$, where k_B is the Boltzmann constant, T is the absolute vapour temperature, and m is the atomic mass. Mathematically, the velocity probability

Figure L.2: Probability density function for the one-dimensional velocity of ^{87}Rb atoms at $T = 42°C$, sampled for $N = 5 \times 10^5$ atoms and sorted into 100 bins (blue), with a fit using Eq. (L.2) shown in grey.

density function for each Cartesian component is given by (cf. Sec. 3.2.3)

$$\mathscr{P}(v_i) = \sqrt{\frac{m}{2\pi k_{\mathrm{B}}T}} \exp\left(\frac{-mv_i^2}{2k_{\mathrm{B}}T}\right), \qquad \text{(L.2)}$$

where the probability density function is normalised according to

$$\int_{-\infty}^{\infty} \mathscr{P}(v_i)\,\mathrm{d}v_i = 1. \qquad \text{(L.3)}$$

A graph showing the Maxwell-Boltzmann probability density function for the velocity of ^{87}Rb atoms in a single Cartesian component, i.e. Eq. (L.2), is shown in Fig. L.2.

L.3 Calculating Trajectories

In a vapour cell which does not contain buffer gas, the vapour pressure is relatively low and hence the dynamics of the atoms are well approximated by a ballistic trajectory. This enables the position of the i^{th} atom, at a time $t_n = n\Delta t$, to be calculated via

$$\vec{p}_i\,[t_n] = \vec{p}_i + n\Delta t \boldsymbol{v}_i, \qquad \text{(L.4)}$$

where n is the step number and Δt is the temporal resolution. At each time step, the new position of each atom is calculated using Eq. (L.4), and subsequently checked to see whether it is contained within the probe volume. This results in an array, for each atom, which shows the time steps in which the atom was (or wasn't) contained within the probe region. Two

Figure L.3: Probability density function for the mean transit time of ^{87}Rb atoms at $T = 42°$C from the pump beam to the probe beam when their centres are spaced $d = 10\,\text{mm}$ apart. The simulation was performed for $N = 5 \times 10^5$ atoms, with a temporal resolution of $\Delta t = 0.1\,\mu\text{s}$. The transit time is sorted into 100 bins (blue), and fitted using Eq. (L.5) (grey). A numerical peak-finding algorithm reveals that the expected transit time under these conditions is approximately 31 μs.

things can be done with this array. Firstly, the fraction of atoms which pass directly from the pump beam to the probe beam can be calculated; and secondly, the transit time of each atom can be calculated. Given that there is some ambiguity in regards to how the transit time should be calculated, the mean transit time has been chosen — i.e. the time step at which the atom is *halfway* between its entry and exit points of the probe volume.

L.4 Transit of Atoms from the Pump to the Probe

For the experiments discussed in Ch. 9, the atomic vapour had a temperature of $T = 42°$C. The pump and probe beams were spaced by $10 \pm 4\,\text{mm}$, and had $1/e^2$ diameters of $1.5\,\text{mm}$ and $3.5\,\text{mm}$, respectively. Using a temporal resolution of $\Delta t = 0.1\,\mu\text{s}$, and a simulation involving 5×10^5 atoms, the distribution of transit times from the pump beam to the probe beam is presented in Fig. L.3. Evident in Fig. L.3 is a probability density function for the atomic transit times, τ_{tr}, which is well approximated by a gamma distribution:

$$\mathscr{P}\left(\tau_{\text{tr}}\right) = \begin{cases} \frac{\gamma}{\beta\Gamma(\alpha)} \left(\frac{\tau_{\text{tr}}-\mu}{\beta}\right)^{\alpha\gamma-1} \exp\left[-\left(\frac{\tau_{\text{tr}}-\mu}{\beta}\right)^{\gamma}\right] & \text{for } \tau_{\text{tr}} > \mu\,, \\ 0 & \text{otherwise}\,, \end{cases} \tag{L.5}$$

where α and γ are shape parameters, β is a scale parameter, μ is a location parameter, and $\Gamma(\alpha)$ is the gamma function evaluated at α.

As discussed in Ch. 9 and Ref. [344], the narrow spectral contribution to the optical-rotation signal decays within $\mathcal{O}\left(20\,\mu s\right)$. Using numerical peak-finding methods to calculate the maxima of the generated Monte Carlo distributions yields an expected transit time of $33 \pm 14\,\mu s$, which is consistent with the observed decay time of the optical-rotation signal. Additionally, as discussed in Sec. L.3, it is possible to calculate the fraction of polarised atoms within the atomic vapour. Monte Carlo simulations would suggest that this fraction is $4.7 \pm 3.9\%$ under the conditions presented in Ch. 9, which is consistent with the measured value of $4.4 \pm 0.16\%$ using the novel technique introduced in Ref. [344].

L.5 Transit Relaxation

One of the major relaxation processes in light-atom interactions is transit relaxation, i.e. the *effective* relaxation experienced as a result of atoms escaping the probe beam (cf. Sec. 5.2.2.3). In the derivation of the Liouville equation found in Sec. 4.3, transit relaxation is approximated as taking the form of an exponential decay, characterised by a transit rate γ_t. A better approximation to the form of the transit rate can be achieved using Monte Carlo techniques.

In a similar fashion to the Monte Carlo simulation described in Sec. L.3, atoms are randomly generated within a cylindrical volume defined by the beam length and $1/e^2$ diameter of the probe beam. These atoms are then given a three-dimensional velocity in accordance with the Maxwell-Boltzmann distribution, Eq. (L.2), and allowed to propagate along ballistic trajectories out of the beam. The number of atoms within the probe beam is calculated as a function of time, yielding the temporal evolution shown in Fig. L.4. The $1/e$ time of the transit-time relaxation, as a function of beam $1/e^2$ diameter, is also presented.

L.6 Wall Collisions in Cylindrical Vapour Cells

The time between wall collisions for atoms in a vapour cell is readily modelled using Monte Carlo techniques. Given that wall collisions result in a redistribution of atomic velocity classes in accordance with the Maxwell-Boltzmann distribution [342–344], the simplest way to model the time between wall collisions is to randomly generate atoms on the boundary of the cylinder (i.e. on the walls of the cell), assign each atom a random three-dimensional velocity according to Eq. (L.2), and propagate the atoms along a ballistic trajectory[1] until they collide with another wall. The distance between the initial position and the intersection between the trajectory and the boundary of the cell can then be used to calculate the time between collisions, by dividing by the magnitude of the atoms' velocity.

Due to the radial and angular symmetry of a sphere, the average time between wall collisions is linearly proportional to the surface-area-to-volume ratio. Explicitly, the average time between wall collisions in a sphere is given by [299]

---

[1]This is only an approximation, as recent experiments have determined that there is a non-negligible amount of collisions with background gases in buffer-gas-free vapour cells [301]. In paraffin-coated cells, the background gas is mostly comprised of C3 (and higher) hydrocarbons, as well as hydrogen [374].

Figure L.4: Left: Transit relaxation (blue markers) of ^{87}Rb atoms at $T = 23°$C from a $3.5\,$mm ($1/e^2$) diameter probe beam, calculated using a Monte Carlo simulation described in Sec. L.5, with an exponential fit (red). The transit-relaxation time ($1/e$) obtained from the fit is $2\pi\gamma_t^{-1} = 9.44\,\mu$s. Right: Transit-relaxation time of ^{87}Rb atoms at $T = 23°$C in beams of various diameters (blue), calculated via cubic-spline interpolation of the temporal decay (left), with a linear fit shown in red. Residuals of the linear fit (top) indicate slight nonlinearity in the dependence of $2\pi\gamma_t^{-1}$ on beam size. Both simulations were performed for $N = 1 \times 10^6$ atoms, with a temporal resolution of $\Delta t = 0.1\,\mu$s.

$$T_{\text{wall}} \approx \frac{4V}{\bar{v}_{\text{th}} A}\,, \tag{L.6}$$

where V is the volume of the cell, A is the area of the cell, and $\bar{v}_{\text{th}}$ is the mean thermal velocity of the atoms. However, due to the reduced symmetry in a cylindrical cell, the mean time between wall collisions is nonlinearly dependent upon both the radius and the length of the cell. Examples of the mean time between wall collisions of ^{87}Rb atoms at room temperature in cylindrical cells of different surface-area-to-volume ratios is presented in Fig. L.5.

Figure L.5: Left: Time between wall collisions for ^{87}Rb atoms at $T = 23°C$, in a cylindrical cell with radius $R_{\text{cell}} = 20\,\text{mm}$ and length $L_{\text{cell}} = 40\,\text{mm}$. The vertical dashed line corresponds to the mean time between collisions, $T_{\text{wall}} = 106\,\mu\text{s}$. Right: Mean time between wall collisions, T_{wall}, for cylinders of various surface-area-to-volume ratios. For reference, the mean time between wall collisions in a spherical cell is indicated by the dashed grey line. These simulations were performed for $N = 1 \times 10^6$ atoms.

Visualisation of the Density Matrix

One method which has traditionally been used to visualise quantum beats is to draw the average angular momentum vector as a function of time. This is a relatively straightforward exercise when the quantum beat is a result of the linear Zeeman effect (i.e. Larmor precession); however, for more complicated quantum beats, such as Stark beats, one must trace the evolution of multiple vectors corresponding to each polarisation moment [284]. In this case it is more convenient to draw a surface in three dimensions representing the probability distribution of the angular momentum.

This appendix gives a brief description of rotations in quantum mechanics and how they're performed, and then applies quantum-mechanical rotation operators to the visualisation of atomic angular momentum states via angular momentum probability surfaces.

M.1   Euler Angles

Any arbitrary orientation of a rigid body can be described by three elementary rotations, i.e. rotations about the three orthogonal axes of a coordinate system [375]. These rotations may either be *extrinsic* or *intrinsic*. Extrinsic rotations are those which are performed with respect to the original coordinate system, which is assumed to remain motionless. Intrinsic rotations, on the other hand, are rotations about the axes of the rotating coordinate system which is 'bound' to the rigid body.

If we consider an intrinsic coordinate system, in which the three orthogonal axes remain motionless with respect to the reference frame of the rigid body, then we define the Euler angles as three elementary rotations [375]:

- A rotation by an angle $\alpha \in [0, 2\pi)$ about the z-axis, bringing the frame axes from the initial position S into the position S'.

- A rotation $\beta \in [0, \pi)$ about the y-axis of the frame S'. The resulting position of the frame of axes is then symbolised as S''.

- A rotation $\gamma \in [0, 2\pi)$ about the z-axis of the frame of axes S'', where the frame of axes S'' is dependent on the previous rotations α, β. The final position of the frame is denoted S'''.

Evident from the definition of the above Euler angles, α and β are identical to the spherical polar coordinates ϕ and θ, respectively.

M.2   Rotation Operators

Rotations in quantum mechanics are performed using a unitary rotation operator. From the definition of the Euler angles from Sec. M.1, it can be shown [375] that the rotation operator, $\mathscr{D}\left(\alpha, \beta, \gamma\right)$, takes the form

$$\mathscr{D}\left(\alpha, \beta, \gamma\right) = \exp\left(\frac{i\alpha}{\hbar}J_z\right) \exp\left(\frac{i\beta}{\hbar}J_y\right) \exp\left(\frac{i\gamma}{\hbar}J_z\right), \tag{M.1}$$

where J_y and J_z are the total angular momentum operators for the y and z Cartesian coordinates, respectively. Due to the fact that the total angular momentum operators can be expressed in the form $J_a = \left(\hbar/2\right)\sigma_a$, where σ_a is a Pauli spin matrix, it follows that the rotation operator can itself be expressed in matrix form. We thus construct a matrix where the elements are given by $\langle J\, m_J{}' |\, \mathscr{D}\left(\alpha, \beta, \gamma\right) | J\, m_J \rangle$. These matrix elements are denoted $\mathscr{D}^{(J)}_{m_J{}',m_J}\left(\alpha, \beta, \gamma\right)$, and are known as **Wigner D-functions**.

The Pauli spin matrix σ_z is a diagonal matrix, and hence we can write the Wigner D-functions in the form [375]

$$\mathscr{D}^{(J)}_{m_J{}',m_J}\left(\alpha, \beta, \gamma\right) = \exp\left(im_J{}'\alpha\right) d^{(J)}_{m_J{}',m_J}\left(\beta\right) \exp\left(im_J\gamma\right), \tag{M.2}$$

where $d^{(J)}_{m_J{}',m_J}\left(\beta\right)$ is **Wigner's (small) d-function**, and is defined by [375]

$$d^{(J)}_{m_J{}',m_J}\left(\beta\right) = \langle J\, m_J{}' |\, \exp\left(\frac{i\beta}{\hbar}J_y\right) | J\, m_J \rangle . \tag{M.3}$$

The construction of the quantum-mechanical rotation operator therefore comes down to calculating Wigner's (small) d-functions. These functions can be calculated explicitly from **Wigner's formula** [376]:

$$
\begin{aligned}
d^{(J)}_{m_J{}',m_J}\left(\beta\right) = \sum_{n=n_{\min}}^{n_{\max}} (-1)^n & \frac{\sqrt{\left(J+m_J{}'\right)!\left(J-m_J{}'\right)!\left(J+m_J\right)!\left(J-m_J\right)!}}{\left(J+m_J-n\right)!\,n!\left(J-m_J{}'-n\right)!\left(n+m_J{}'-m_J\right)!} \\
\times & \left(\cos\frac{\beta}{2}\right)^{2J-2n+m_J-m_J{}'}\left(-\sin\frac{\beta}{2}\right)^{2n+m_J{}'-m_J},
\end{aligned}
\tag{M.4}
$$

where the values $n_{\min}$ and $n_{\max}$ are given by

$$n_{\min} = \max\left(0, m_J - m_J{}'\right), \tag{M.5}$$

$$n_{\max} = \min\left(J - m_J{}', J + m_J\right), \tag{M.6}$$

and the sum is *only* evaluated for values of n such that the factorials are defined (i.e. no values of n such that there exist factorials of negative integers in Eq. M.4).

M.3   Angular Momentum Probability Surfaces (AMPS)

When studying atomic physics, it is often insightful to consider the angular momentum distribution of an atomic ensemble. Asphericity of such a distribution is associated with optical anisotropy of the medium, and has physical consequences in regards to measured optical signals. In the context of quantum beats, for example, symmetries in the angular momentum distribution with respect to the quantisation axis correspond to coherences in the density matrix [242,245,257,283,285,286].

In this section a method for visualising the density matrix is introduced — the angular momentum probability surface (AMPS) — which relies on calculating a mathematical surface that describes the angular momentum distribution of the atoms. Angular momentum probability surfaces of commonly encountered density-matrix states are computed, and their insight into underlying physical processes are discussed.

M.3.1   Definition and Calculation

In order to visualise the angular momentum state of atoms with total angular momentum J, we draw a three dimensional surface where the distance r from the origin is equal to the probability of finding the projection $m_J = J$ along the radial direction [242,283,284]. To calculate the radius, r, in a given direction, we rotate the density matrix ρ such that the quantisation axis lies along this direction, and then take the $\rho_{J,J}$ element. That is, we calculate the following matrix element:

$$r\left(\theta,\phi\right) = \langle m_J = J | \, \mathscr{D}^{-1}\left(\phi,\theta,0\right) \rho \mathscr{D}\left(\phi,\theta,0\right) | m_J = J \rangle\,, \tag{M.7}$$

where $\mathscr{D}\left(\phi,\theta,0\right)$ is the quantum-mechanical rotation operator defined by Eq. (M.1) and $\mathscr{D}^{-1}\left(\phi,\theta,0\right)$ is its inverse. The inverse rotation operator is calculated by performing the rotations through negative angles (about the same axes) but in opposite order. Since $\mathscr{D}\left(\phi,\theta,0\right)$ is a unitary matrix (i.e. $\mathscr{D}\mathscr{D}^{\dagger} = \mathscr{D}^{\dagger}\mathscr{D} = I$, where I is the identity matrix), the matrix elements of $\mathscr{D}^{-1}\left(\phi,\theta,0\right)$ are *identical* to those of the conjugate transpose of $\mathscr{D}\left(\phi,\theta,0\right)$ [278,377].

For a general density matrix with elements given by

$$\rho = \begin{pmatrix} \rho_{J,J} & \rho_{J,J-1} & \cdots & \rho_{J,-J} \\ \rho_{J-1,J} & \rho_{J-1,J-1} & \cdots & \rho_{J-1,-J} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{-J,J} & \rho_{-J,J-1} & \cdots & \rho_{-J,-J} \end{pmatrix}, \tag{M.8}$$

the quantum-mechanical rotation operator takes the form

$$\mathscr{D}\left(\phi,\theta,0\right) = \begin{pmatrix} \mathscr{D}_{J,J} & \mathscr{D}_{J,J-1} & \cdots & \mathscr{D}_{J,-J} \\ \mathscr{D}_{J-1,J} & \mathscr{D}_{J-1,J-1} & \cdots & \mathscr{D}_{J-1,-J} \\ \vdots & \vdots & \ddots & \vdots \\ \mathscr{D}_{-J,J} & \mathscr{D}_{-J,J-1} & \cdots & \mathscr{D}_{-J,-J} \end{pmatrix}. \tag{M.9}$$

For a state with total angular momentum J, the density matrix and quantum-mechanical rotation operator are $(2J + 1) \times (2J + 1)$ square matrices. Using Wigner's formula given by Eq. (M.4), we can evaluate the matrix elements of the rotation operator for the case of $J = 1$, and thus our rotation operators are given by

$$\mathscr{D}\left(\phi,\theta,0\right) = \begin{pmatrix} \frac{1}{2}\left(1+\cos\theta\right)e^{i\phi} & \frac{1}{\sqrt{2}}\sin\theta e^{i\phi} & \sin^2\left(\frac{\theta}{2}\right)e^{i\phi} \\ -\frac{1}{\sqrt{2}}\sin\theta & \cos\theta & \frac{1}{\sqrt{2}}\sin\theta \\ \sin^2\left(\frac{\theta}{2}\right)e^{-i\phi} & -\frac{1}{\sqrt{2}}\sin\theta e^{-i\phi} & \frac{1}{2}\left(1+\cos\theta\right)e^{-i\phi} \end{pmatrix}, \tag{M.10}$$

$$\mathscr{D}^{-1}\left(\phi,\theta,0\right) = \begin{pmatrix} \frac{1}{2}\left(1+\cos\theta\right)e^{-i\phi} & -\frac{1}{\sqrt{2}}\sin\theta & \sin^2\left(\frac{\theta}{2}\right)e^{i\phi} \\ \frac{1}{\sqrt{2}}\sin\theta e^{-i\phi} & \cos\theta & -\frac{1}{\sqrt{2}}\sin\theta e^{i\phi} \\ \sin^2\left(\frac{\theta}{2}\right)e^{-i\phi} & \frac{1}{\sqrt{2}}\sin\theta & \frac{1}{2}\left(1+\cos\theta\right)e^{i\phi} \end{pmatrix}. \tag{M.11}$$

In the case of $J = 2$, the rotation operators are 5×5 square matrices given by

$$\mathscr{D}\left(\phi,\theta,0\right) = \begin{pmatrix} \cos^4\left(\frac{\theta}{2}\right)e^{2i\phi} & \frac{1}{2}\left(1+\cos\theta\right)\sin\theta e^{2i\phi} & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\theta e^{2i\phi} & \sin^2\left(\frac{\theta}{2}\right)\sin\theta e^{2i\phi} & \sin^4\left(\frac{\theta}{2}\right)e^{2i\phi} \\ -\frac{1}{2}\left(1+\cos\theta\right)\sin\theta e^{i\phi} & \frac{1}{2}\left[\cos\theta+\cos\left(2\theta\right)\right]e^{i\phi} & \sqrt{\frac{3}{2}}\cos\theta\sin\theta e^{i\phi} & \frac{1}{2}\left[\cos\theta-\cos\left(2\theta\right)\right]e^{i\phi} & \sin^2\left(\frac{\theta}{2}\right)\sin\theta e^{i\phi} \\ \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\theta & -\sqrt{\frac{3}{2}}\cos\theta\sin\theta & \frac{1}{4}\left[1+3\cos\left(2\theta\right)\right] & \sqrt{\frac{3}{2}}\cos\theta\sin\theta & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\theta \\ -\sin^2\left(\frac{\theta}{2}\right)\sin\theta e^{-i\phi} & \frac{1}{2}\left[\cos\theta-\cos\left(2\theta\right)\right]e^{-i\phi} & -\sqrt{\frac{3}{2}}\cos\theta\sin\theta e^{-i\phi} & \frac{1}{2}\left[\cos\theta+\cos\left(2\theta\right)\right]e^{-i\phi} & \frac{1}{2}\left(1+\cos\theta\right)\sin\theta e^{-i\phi} \\ \sin^4\left(\frac{\theta}{2}\right)e^{-2i\phi} & -\sin^2\left(\frac{\theta}{2}\right)\sin\theta e^{-2i\phi} & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\theta e^{-2i\phi} & -\frac{1}{2}\left(1+\cos\theta\right)\sin\theta e^{-2i\phi} & \cos^4\left(\frac{\theta}{2}\right)e^{-2i\phi} \end{pmatrix}, \tag{M.12}$$

$$
\mathscr{D}^{-1}\left(\phi,\theta,0\right) = \begin{pmatrix}
\cos^4\left(\frac{\theta}{2}\right)e^{-2i\phi} & -\frac{1}{2}\left(1+\cos\theta\right)\sin\theta e^{-i\phi} & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\theta & -\sin^2\left(\frac{\theta}{2}\right)\sin\theta e^{i\phi} & \sin^4\left(\frac{\theta}{2}\right)e^{2i\phi} \\
\frac{1}{2}\left(1+\cos\theta\right)\sin\theta e^{-2i\phi} & \frac{1}{2}\left[\cos\theta+\cos\left(2\theta\right)\right]e^{-i\phi} & -\sqrt{\frac{3}{2}}\cos\theta\sin\theta & \frac{1}{2}\left[\cos\theta-\cos\left(2\theta\right)\right]e^{i\phi} & -\sin^2\left(\frac{\theta}{2}\right)\sin\theta e^{2i\phi} \\
\frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\theta e^{-2i\phi} & \sqrt{\frac{3}{2}}\cos\theta\sin\theta e^{-i\phi} & \frac{1}{4}\left[1+3\cos\left(2\theta\right)\right] & -\sqrt{\frac{3}{2}}\cos\theta\sin\theta e^{i\phi} & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\theta e^{2i\phi} \\
\sin^2\left(\frac{\theta}{2}\right)\sin\theta e^{-2i\phi} & \frac{1}{2}\left[\cos\theta-\cos\left(2\theta\right)\right]e^{-i\phi} & \sqrt{\frac{3}{2}}\cos\theta\sin\theta & \frac{1}{2}\left[\cos\theta+\cos\left(2\theta\right)\right]e^{i\phi} & -\frac{1}{2}\left(1+\cos\theta\right)\sin\theta e^{2i\phi} \\
\sin^4\left(\frac{\theta}{2}\right)e^{-2i\phi} & \sin^2\left(\frac{\theta}{2}\right)\sin\theta e^{-i\phi} & \frac{1}{2}\sqrt{\frac{3}{2}}\sin^2\theta & \frac{1}{2}\left(1+\cos\theta\right)\sin\theta e^{i\phi} & \cos^4\left(\frac{\theta}{2}\right)e^{2i\phi}
\end{pmatrix} . \quad \text{(M.13)}
$$

When calculating angular momentum probability surfaces for density matrices representing hyperfine states — where F and m_F are the relevant quantum numbers — all equations derived in this appendix are readily applicable upon substituting $J \to F$ and $m_J \to m_F$.

(a) Unpolarised state.　　　**(b)** Oriented state.　　　**(c)** Aligned state.

Figure M.1: Angular momentum probability surfaces for common $F = 1$ ground-state density matrices: (a) an unpolarised state, (b) an oriented state, and (c) an aligned state. These surfaces were calculated using Eqs. (M.15), (M.17), and (M.19), respectively.

M.3.2　Static Probability Surfaces

Perhaps the simplest example of a density-matrix state is that which corresponds to an atomic ensemble that has *not* been optically pumped, i.e. an unpolarised state. Such a state is observed when the atomic ensemble is in thermal equilibrium, in which case the atoms are *equally distributed* amongst the Zeeman ground states. The corresponding density matrix is therefore diagonal and, for an $F = 1$ ground state, is given explicitly by

$$\rho = \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix} . \tag{M.14}$$

Evaluating the angular momentum probability surface for an unpolarised state given by Eq. (M.14) — by inserting Eq. (M.14) into Eq. (M.7) — yields a distance from the origin to the surface given by

$$r\left(\theta, \phi\right) = \frac{1}{3} . \tag{M.15}$$

The resulting surface is therefore a sphere of radius $1/3$, as shown in Fig. M.1(a). Given that the angular momentum probability surface for an unpolarised state is spherically symmetric, the atomic ensemble is optically isotropic and hence the probability of finding an atom in the 'stretched state' is the same *regardless* of the choice of quantisation axis.

　　When optically pumping an atomic ensemble with circularly polarised light propagating along $\hat{z}$, it is possible to pump the atoms into a stretched state where $|m_F = F\rangle$. In this case, the density matrix is given by

$$\rho = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \tag{M.16}$$

Evaluating the AMPS in this case results in

$$r\left(\theta, \phi\right) = \cos^4\left(\frac{\theta}{2}\right), \tag{M.17}$$

which yields the surface shown in Fig. M.1(b). Evident in Fig. M.1(b) is an angular momentum distribution which is cylindrically symmetric about the z-axis. However, despite the *symmetry* about the z-axis, there is *asymmetry* with respect to *inversion* of the axis. The surface is therefore said to be *oriented* along $\hat{z}$, and such a state is therefore referred to as *atomic orientation*. Atomic orientation corresponds to a $\kappa = 1$ (dipole) polarisation moment [242,283]. With an appropriate choice of quantisation axis and light-propagation axis, an oriented atomic ensemble can give rise to an optical signal that is modulated at the Larmor frequency.

In contrast to optical pumping with circularly polarised light, the use of linearly polarised light generates an atomic distribution which is symmetric about the $|m_F = 0\rangle$ ground state, and coherences between ground states which satisfy $|\Delta m_F| = 2$. For light which is linearly polarised light along $\hat{x}$ and propagating along $\hat{z}$ the steady-state density matrix, calculated by solving the modified Liouville equation given by Eq. (4.8), is given by

$$\rho = \begin{pmatrix} \frac{1}{4} & 0 & \frac{1}{4} \\ 0 & \frac{1}{2} & 0 \\ \frac{1}{4} & 0 & \frac{1}{4} \end{pmatrix}. \tag{M.18}$$

Calculating the angular momentum distribution for the above density-matrix state yields a distance of the surface from the origin given by

$$r\left(\theta, \phi\right) = \frac{1}{16}\left[2\cos\left(2\phi\right)\sin^2\left(\theta\right) - \cos\left(2\theta\right) + 5\right]. \tag{M.19}$$

The resulting surface, shown in Fig. M.1(c), exhibits two-fold symmetry about the z-axis. In this case, the angular momentum distribution has a preferred *axis*, rather than a preferred *direction*. This distribution is referred to as *atomic alignment*, and corresponds to a $\kappa = 2$ (quadrupole) polarisation moment [242,283]. Atomic alignment can also give rise to modulation of an optical signal given appropriate experimental conditions; however, this will instead occur at *twice* the Larmor frequency.

M.3.3   Dynamic Probability Surfaces

In addition to visualising static density-matrix states, angular momentum probability surfaces are also a powerful tool for visualising states with temporal dependence. Of particular interest in the context of this thesis, is the temporal evolution of an aligned state in response

Figure M.2: Angular momentum probability surfaces of an aligned state undergoing Larmor precession counter-clockwise in the xy-plane due to a magnetic field pointing along $\hat{z}$, calculated using Eq. (M.21). The rotation of the angular momentum distribution occurs at the Larmor frequency Ω_L; however, due to the two-fold symmetry of an aligned state, the optical anisotropy modulates optical signals at a frequency $2\Omega_L$.

to an external magnetic field — Larmor precession. Solving the modified Liouville equation given by Eq. (4.8), for a magnetic-field Hamiltonian given by Eq. (4.21) and a density-matrix state initially aligned along the x-axis as described by Eq. (M.18), the time-dependent state is

$$
\rho\left(t\right) = \begin{pmatrix} \frac{1}{4} & 0 & \frac{1}{4}e^{-i2\Omega_L t} \\ 0 & \frac{1}{2} & 0 \\ \frac{1}{4}e^{i2\Omega_L t} & 0 & \frac{1}{4} \end{pmatrix}.
\tag{M.20}
$$

The time-dependent aligned state given by Eq. (M.20) has a corresponding dynamic angular momentum probability surface described by

$$
r\left(\theta,\phi,t\right) = \frac{1}{16}\left[2\cos\left(2\Omega_L t + 2\phi\right)\sin^2\left(\theta\right) - \cos\left(2\theta\right) + 5\right].
\tag{M.21}
$$

The dynamic angular momentum probability surface given by Eq. (M.21) is plotted in Fig. M.2. Evident in Fig. M.2 is an aligned state undergoing Larmor precession in the xy-plane, in response to a magnetic field directed along $\hat{z}$. Although the angular momentum distribution is seen to oscillate at the Larmor frequency Ω_L, the rotation of the optical anisotropy occurs at twice the Larmor frequency, $2\Omega_L$, due to two-fold symmetry about the quantisation axis.

The two-fold symmetry observed in the angular momentum distribution of an aligned state is associated with $|\Delta m_F| = 2$ coherences in the density matrix [242,245,257,283,285,286]. In general, a κ-fold symmetry with respect to the quantisation axis of the angular momentum probability surface corresponds to $|\Delta m_F| = \kappa$ coherences [242,283]. With an appropriate choice of quantisation axis and light-propagation axis, κ-fold symmetries give rise to modulation of the optical signals at $\kappa\Omega_L$ in the presence of an external magnetic field [242,245,283,285,286].

Although only Larmor precession has been explicitly discussed in this section, angular momentum probability surfaces are also useful for visualising more complicated quantum beats such as Stark beats [242,283,284] and alignment-to-orientation conversion [242,283,284].

References

[1] E. Boto, N. Holmes, J. Leggett, G. Roberts, V. Shah, S. S. Meyer, L. Duque Muñoz, K. J. Mullinger, T. M. Tierney, S. Bestmann, G. R. Barnes, R. Bowtell, and M. J. Brookes. Moving magnetoencephalography towards real-world applications with a wearable system. Nature, **555**, 657 (2018).

[2] T. R. Clem, G. J. Kekelis, J. D. Lathrop, D. J. Overway, and W. M. Wynn. Superconducting Magnetic Gradiometers for Mobile Applications with an Emphasis on Ordnance Detection. In *SQUID Sensors: Fundamentals, Fabrication and Applications*, pages 517–568 (Springer, 1996).

[3] J. E. McFee, Y. Das, and R. O. Ellingson. Locating and identifying compact ferrous objects. IEEE Trans. Geosci. Remote Sens., **28**, 182 (1990).

[4] A. Sheinker, L. Frumkis, B. Ginzburg, N. Salomonski, and B.-Z. Kaplan. Magnetic Anomaly Detection Using a Three-Axis Magnetometer. IEEE Trans. Magn., **45**, 160 (2009).

[5] E. Paperno, I. Sasada, and E. Leonovich. A new method for magnetic position and orientation tracking. IEEE Trans. Magn., **37**, 1938 (2001).

[6] H. P. Kalmus. A New Guiding and Tracking System. IRE Trans. Aerosp. Navig. Electron., **ANE-9**, 7 (1962).

[7] F. H. Raab, E. B. Blood, T. O. Steiner, and H. R. Jones. Magnetic Position and Orientation Tracking System. IEEE Trans. Aerosp. Electron. Syst., **AES-15**, 709 (1979).

[8] C. P. Du, M. Y. Xia, S. X. Huang, Z. H. Xu, X. Peng, and H. Guo. Detection of a Moving Magnetic Dipole Target Using Multiple Scalar Magnetometers. IEEE Geosci. Remote Sens. Lett., **14**, 1166 (2017).

[9] J. Zhou, J. Chen, and Z. Shan. Spatial Signature Analysis of Submarine Magnetic Anomaly at Low Altitude. IEEE Trans. Magn., **53**, 1 (2017).

[10] R. Wiegert. Magnetic anomaly guidance system for mine countermeasures using autonomous underwater vehicles. In *Proc. Oceans*, vol. 4, pages 2002–2010 (IEEE, 2003).

[11] R. Wiegert and J. Oeschger. Generalized magnetic gradient contraction based method for detection, localization and discrimination of underwater mines and unexploded ordnance. In *Proc. Oceans*, vol. 2, pages 1325–1332 (IEEE, 2005).

[12] R. Wiegert and J. Oeschger. Portable Magnetic Gradiometer for Real-Time Localization and Classification of Unexploded Ordnance. In *Proc. Oceans*, pages 1–6 (IEEE, 2006).

[13] R. Wiegert, J. Oeschger, and E. Tuovila. Demonstration of a Novel Man-Portable Magnetic STAR Technology for Real Time Localization of Unexploded Ordnance. In *Proc. Oceans*, pages 1–7 (IEEE, 2007).

[14] R. Wiegert, K. Lee, and J. Oeschger. Improved magnetic STAR methods for real-time, point-by-point localization of unexploded ordnance and buried mines. In *Proc. Oceans*, pages 1–7 (IEEE, 2008).

[15] H. H. Nelson and J. R. McDonald. Multisensor towed array detection system for UXO detection. IEEE Trans. Geosci. Remote Sens., **39**, 1139 (2001).

[16] S. D. Billings. Discrimination and classification of buried unexploded ordnance using magnetometry. IEEE Trans. Geosci. Remote Sens., **42**, 1241 (2004).

[17] G. Schultz, R. Mhaskar, M. Prouty, and J. Miller. Integration of micro-fabricated atomic magnetometers on military systems. In *Proc. SPIE*, vol. 9823, page 982318 (2016).

[18] T. R. Clem. Progress in magnetic sensor technology for sea mine detection. In *Proc. SPIE*, vol. 3079, pages 354–371 (1997).

[19] T. R. Clem, R. H. Koch, and G. A. Keefe. Superconducting magnetic sensors for mine detection and classification. In *Proc. SPIE*, vol. 2496, pages 374–383 (1995).

[20] L. Merlat and P. Naz. Magnetic localization and identification of vehicles. In *Proc. SPIE*, vol. 5090, pages 174–186 (2003).

[21] N. Wahlström and F. Gustafsson. Magnetometer Modeling and Validation for Tracking Metallic Targets. IEEE Trans. Signal Process., **62**, 545 (2013).

[22] J. Lan and Y. Shi. Vehicle detection and recognition based on a MEMS magnetic sensor. In *Proc. 4th IEEE Int. Conf. NEMS*, pages 404–408 (IEEE, 2009).

[23] N. Wahlström, J. Callmer, and F. Gustafsson. Single target tracking using vector magnetometers. In *Proc. Int. Conf. Acoust. Speech Signal Process.*, pages 4332–4335 (IEEE, 2011).

[24] C. Wang, X. Zhang, X. Qu, X. Pan, G. Fang, and L. Chen. A Modified Magnetic Gradient Contraction Based Method for Ferromagnetic Target Localization. Sensors, **16**, 2168 (2016).

[25] C. Wang, X. Qu, X. Zhang, W. Zhu, and G. Fang. A Fast Calibration Method for Magnetometer Array and the Application of Ferromagnetic Target Localization. IEEE Trans. Instrum. Meas., **66**, 1743 (2017).

[26] K.-M. Lee and M. Li. Magnetic Tensor Sensor for Gradient-Based Localization of Ferrous Object in Geomagnetic Field. IEEE Trans. Magn., **52**, 1 (2016).

[27] L. Fan, C. Kang, X. Zhang, and S. Wan. Real-Time Tracking Method for a Magnetic Target Using Total Geomagnetic Field Intensity. Pure Appl. Geophys., **173**, 2065 (2016).

[28] L.-M. Fan, Q. Zheng, X.-Y. Kang, X.-J. Zhang, and C. Kang. Baseline optimization for scalar magnetometer array and its application in magnetic target localization. Chin. Phys. B, **27**, 060703 (2018).

[29] L. Fan, X. Kang, Q. Zheng, X. Zhang, X. Liu, Z. Geng, and C. Kang. Tracking of moving magnetic target based on magnetic gradient system with total field magnetometers. Sens. Review, **38**, 501 (2018).

[30] H.-G. Meyer, K. Hartung, S. Linzen, M. Schneider, R. Stolz, W. Fried, and S. Hauspurg. Detection of buried magnetic objects by a SQUID gradiometer system. In *Proc. SPIE*, vol. 7303, page 73031X (2009).

[31] M. D. Prouty and M. Tchernychev. Real-time threat detection using magnetometer arrays. In *Proc. SPIE*, vol. 9825, page 98250A (2016).

[32] J. A. Baldoni and B. B. Yellen. Magnetic Tracking System: Monitoring Heart Valve Prostheses. IEEE Trans. Magn., **43**, 2430 (2007).

[33] T. R. Clem, D. J. Overway, J. W. Purpura, J. T. Bono, P. J. Carroll, R. H. Koch, J. R. Rozen, G. A. Keefe, S. Willen, and R. A. Mohling. Enhanced magnetic anomaly detection using a nitrogen-cooled superconducting gradiometer. In *Proc. SPIE*, vol. 4039, pages 70–84 (2000).

[34] T. R. Clem, C. P. Foley, and M. N. Keene. SQUIDs for Geophysical Survey and Magnetic Anomaly Detection. In *The SQUID Handbook: Applications of SQUIDs and SQUID Systems*, pages 481–543 (Wiley, 2006).

[35] M. Hirota, T. Furuse, K. Ebana, H. Kubo, K. Tsushima, T. Inaba, A. Shima, M. Fujinuma, and N. Tojyo. Magnetic detection of a surface ship by an airborne LTS SQUID MAD. IEEE Trans. Appl. Supercond., **11**, 884 (2001).

[36] K. P. Humphrey, T. J. Horton, and M. N. Keene. Detection of mobile targets from a moving platform using an actively shielded, adaptively balanced SQUID gradiometer. IEEE Trans. Appl. Supercond., **15**, 753 (2005).

[37] T. R. Clem, D. J. Overway, J. W. Purpura, J. T. Bono, R. H. Koch, J. R. Rozen, G. A. Keefe, S. Willen, and R. A. Mohling. High-T_c SQUID gradiometer for mobile magnetic anomaly detection. IEEE Trans. Appl. Supercond., **11**, 871 (2001).

[38] S. T. Keenan, K. R. Blay, and E. J. Romans. Mobile magnetic anomaly detection using a field-compensated high-T_c single layer SQUID gradiometer. Supercond. Sci. Technol., **24**, 085019 (2011).

[39] H. Pang, M. Pan, C. Wan, J. Chen, X. Zhu, and F. Luo. Integrated Compensation of Magnetometer Array Magnetic Distortion Field and Improvement of Magnetic Object Localization. IEEE Trans. Geosci. Remote Sens., **52**, 5670 (2013).

[40] M. Birsan. Recursive Bayesian Method for Magnetic Dipole Tracking With a Tensor Gradiometer. IEEE Trans. Magn., **47**, 409 (2010).

[41] W. Andrä and H. Nowak. *Magnetism in Medicine* (Wiley, 2007), 2nd ed.

[42] G. Bison, R. Wynands, and A. Weis. Dynamical mapping of the human cardiomagnetic field with a room-temperature, laser-optical sensor. Opt. Express, **11**, 904 (2003).

[43] E. Boto, S. S. Meyer, V. Shah, O. Alem, S. Knappe, P. Kruger, T. M. Fromhold, M. Lim, P. M. Glover, P. G. Morris, R. Bowtell, G. R. Barnes, and M. J. Brookes. A new generation of magnetoencephalography: Room temperature measurements using optically-pumped magnetometers. NeuroImage, **149**, 404 (2017).

[44] T. M. Tierney, N. Holmes, S. Mellor, J. D. López, G. Roberts, R. M. Hill, E. Boto, J. Leggett, V. Shah, M. J. Brookes, R. Bowtell, and G. R. Barnes. Optically pumped magnetometers: From quantum origins to multi-channel magnetoencephalography. NeuroImage, **199**, 598 (2019).

[45] R. Wyllie, M. Kauer, G. S. Smetana, R. T. Wakai, and T. G. Walker. Magnetocardiography with a modular spin-exchange relaxation-free atomic magnetometer array. Phys. Med. Biol., **57**, 2619 (2012).

[46] R. T. Wakai, J. M. Lengle, and A. C. Leuthold. Transmission of electric and magnetic foetal cardiac signals in a case of ectopia cordis: the dominant role of the *vernix caseosa*. Phys. Med. Biol., **45**, 1989 (2000).

[47] H. Xia, A. Ben-Amar Baranga, D. Hoffman, and M. V. Romalis. Magnetoencephalography with an atomic magnetometer. Appl. Phys. Lett., **89**, 211104 (2006).

[48] C. N. Johnson, P. D. D. Schwindt, and M. Weisend. Multi-sensor magnetoencephalography with atomic magnetometers. Phys. Med. Biol., **58**, 6065 (2013).

[49] J. Sheng, S. Wan, Y. Sun, R. Dou, Y. Guo, K. Wei, K. He, J. Qin, and J.-H. Gao. Magnetoencephalography with a Cs-based high-sensitivity compact atomic magnetometer. Rev. Sci. Instrum., **88**, 094304 (2017).

[50] V. K. Shah and R. T. Wakai. A compact, high performance atomic magnetometer for biomedical applications. Phys. Med. Biol., **58**, 8153 (2013).

[51] T. H. Sander, J. Preusser, R. Mhaskar, J. Kitching, L. Trahms, and S. Knappe. Magnetoencephalography with a chip-scale atomic magnetometer. Biomed. Opt. Express, **3**, 981 (2012).

[52] J. Osborne, J. Orton, O. Alem, and V. Shah. Fully integrated standalone zero field optically pumped magnetometer for biomagnetism. In *Proc. SPIE*, vol. 10548, page 105481G (2018).

[53] C. Johnson, P. D. D. Schwindt, and M. Weisend. Magnetoencephalography with a two-color pump-probe, fiber-coupled atomic magnetometer. Appl. Phys. Lett., **97**, 243703 (2010).

[54] J. Iivanainen, R. Zetter, M. Grön, K. Hakkarainen, and L. Parkkonen. On-scalp MEG system utilizing an actively shielded array of optically-pumped magnetometers. NeuroImage, **194**, 244 (2019).

[55] N. Holmes, J. Leggett, E. Boto, G. Roberts, R. M. Hill, T. M. Tierney, V. Shah, G. R. Barnes, M. J. Brookes, and R. Bowtell. A bi-planar coil system for nulling background magnetic fields in scalp mounted magnetoencephalography. NeuroImage, **181**, 760 (2018).

[56] A. Borna, T. R. Carter, J. D. Goldberg, A. P. Colombo, Y.-Y. Jau, C. Berry, J. McKay, J. Stephen, M. Weisend, and P. D. D. Schwindt. A 20-channel magnetoencephalography system based on optically pumped magnetometers. Phys. Med. Biol., **62**, 8909 (2017).

[57] S. Baillet. Magnetoencephalography for brain electrophysiology and imaging. Nat. Neurosci., **20**, 327 (2017).

[58] O. Alem, A. M. Benison, D. S. Barth, J. Kitching, and S. Knappe. Magnetoencephalography of Epilepsy with a Microfabricated Atomic Magnetrode. J. Neurosci., **34**, 14324 (2014).

[59] G. Zhang, S. Huang, F. Xu, Z. Hu, and Q. Lin. Multi-channel spin exchange relaxation free magnetometer towards two-dimensional vector magnetoencephalography. Opt. Express, **27**, 597 (2019).

[60] G. Zhang, S. Huang, and Q. Lin. Magnetoencephalography using a compact multi-channel atomic magnetometer with pump-probe configuration. AIP Adv., **8**, 125028 (2018).

[61] K. Kim, S. Begus, H. Xia, S.-K. Lee, V. Jazbinsek, Z. Trontelj, and M. V. Romalis. Multi-channel atomic magnetometer for magnetoencephalography: A configuration study. NeuroImage, **89**, 143 (2014).

[62] E. Labyt, M.-C. Corsi, W. Fourcault, A. P. Laloy, F. Bertrand, F. Lenouvel, G. Cauffet, M. Le Prado, F. Berger, and S. Morales. Magnetoencephalography With Optically Pumped ^{4}He Magnetometers at Ambient Temperature. IEEE Trans. Med. Imaging, **38**, 90 (2018).

[63] D. N. Barry, T. M. Tierney, N. Holmes, E. Boto, G. Roberts, J. Leggett, R. Bowtell, M. J. Brookes, G. R. Barnes, and E. A. Maguire. Imaging the human hippocampus with optically-pumped magnetoencephalography. NeuroImage, **203**, 116192 (2019).

[64] Y. J. Kim, I. Savukov, and S. Newman. Magnetocardiography with a 16-channel fiber-coupled single-cell Rb optically pumped magnetometer. Appl. Phys. Lett., **114**, 143702 (2019).

[65] K. Kamada, Y. Ito, and T. Kobayashi. Human MCG measurements with a high-sensitivity potassium atomic magnetometer. Physiol. Meas., **33**, 1063 (2012).

[66] G. Lembke, S. N. Erné, H. Nowak, B. Menhorn, A. Pasquarelli, and G. Bison. Optical multichannel room temperature magnetic field imaging system for clinical application. Biomed. Opt. Express, **5**, 876 (2014).

[67] G. Bison, R. Wynands, and A. Weis. A laser-pumped magnetometer for the mapping of human cardiomagnetic fields. Appl. Phys. B, **76**, 325 (2003).

[68] W. Zheng, S. Su, G. Zhang, X. Bi, and Q. Lin. Vector magnetocardiography measurement with a compact elliptically polarized laser-pumped magnetometer. Biomed. Opt. Express, **11**, 649 (2020).

[69] O. Alem, T. H. Sander, R. Mhaskar, J. LeBlanc, H. Eswaran, U. Steinhoff, Y. Okada, J. Kitching, L. Trahms, and S. Knappe. Fetal magnetocardiography measurements with an array of microfabricated optically pumped magnetometers. Phys. Med. Biol., **60**, 4797 (2015).

[70] R. T. Wakai, J. F. Strasburger, Z. Li, B. J. Deal, and N. L. Gotteiner. Magnetocardiographic Rhythm Patterns at Initiation and Termination of Fetal Supraventricular Tachycardia. Circulation, **107**, 307 (2003).

[71] L. Marmugi and F. Renzoni. Optical Magnetic Induction Tomography of the Heart. Sci. Rep., **6**, 23962 (2016).

[72] K. Jensen, R. Budvytyte, R. A. Thomas, T. Wang, A. M. Fuchs, M. V. Balabas, G. Vasilakis, L. D. Mosgaard, H. C. Stærkind, J. H. Müller, T. Heimburg, S.-P. Olesen, and E. S. Polzik. Non-invasive detection of animal nerve impulses with an atomic magnetometer operating near quantum limited sensitivity. Sci. Rep., **6**, 29638 (2016).

[73] J. F. Barry, M. J. Turner, J. M. Schloss, D. R. Glenn, Y. Song, M. D. Lukin, H. Park, and R. L. Walsworth. Optical magnetic detection of single-neuron action potentials using quantum defects in diamond. Proc. Natl. Acad. Sci. U.S.A., **113**, 14133 (2016).

[74] V. V. Yashchuk, J. Granwehr, D. F. Jackson Kimball, S. M. Rochester, A. H. Trabesinger, J. T. Urban, D. Budker, and A. Pines. Hyperpolarized Xenon Nuclear Spins Detected by Optical Atomic Magnetometry. Phys. Rev. Lett., **93**, 160801 (2004).

[75] D. F. Jackson Kimball, I. Lacey, J. Valdez, J. Swiatlowski, C. Rios, R. Peregrina-Ramirez, C. Montcrieffe, J. Kremer, J. Dudley, and C. Sanchez. A dual-isotope rubidium comagnetometer to search for anomalous long-range spin-mass (spin-gravity) couplings of the proton. Ann. Phys., **525**, 514 (2013).

[76] T. J. Kane, P. D. Hillman, C. A. Denman, M. Hart, R. Phillip Scott, M. E. Purucker, and S. J. Potashnik. Laser Remote Magnetometry Using Mesospheric Sodium. J. Geophys. Res., **123**, 6171 (2018).

[77] F. P. Bustos, D. B. Calia, D. Budker, M. Centrone, J. Hellemeier, P. Hickson, R. Holzlöhner, and S. Rochester. Polarization-driven spin precession of mesospheric sodium atoms. Opt. Lett., **43**, 5825 (2018).

[78] F. P. Bustos, D. B. Calia, D. Budker, M. Centrone, J. Hellemeier, P. Hickson, R. Holzlöhner, and S. Rochester. Remote sensing of geomagnetic fields and atomic collisions in the mesosphere. Nat. Commun., **9**, 3981 (2018).

[79] R. S. Grewal, G. Pati, R. Tripathi, A. W. Yu, M. Krainak, and M. Purucker. Magneto-optical resonances in fluorescence from sodium D_2 manifold. Opt. Express, **28**, 1114 (2020).

[80] C. J. Berglund, L. R. Hunter, D. Krause Jr., E. O. Prigge, M. S. Ronfeldt, and S. K. Lamoreaux. New Limits on Local Lorentz Invariance from Hg and Cs Magnetometers. Phys. Rev. Lett., **75**, 1879 (1995).

[81] M. Smiciklas, J. M. Brown, L. W. Cheuk, S. J. Smullin, and M. V. Romalis. New Test of Local Lorentz Invariance Using a ^{21}Ne-Rb-K Comagnetometer. Phys. Rev. Lett., **107**, 171604 (2011).

[82] S. Xu, V. V. Yashchuk, M. H. Donaldson, S. M. Rochester, D. Budker, and A. Pines. Magnetic resonance imaging with an optical atomic magnetometer. Proc. Natl. Acad. Sci. U.S.A., **103**, 12668 (2006).

[83] Y. Hu, G. Z. Iwata, M. Mohammadi, E. V. Silletta, A. Wickenbrock, J. W. Blanchard, D. Budker, and A. Jerschow. Battery Diagnostics with Sensitive Magnetometry. arXiv 1905.12507 (2019).

[84] K. Jensen, M. Zugenmaier, J. Arnbak, H. Stærkind, M. V. Balabas, and E. S. Polzik. Detection of low-conductivity objects using eddy current measurements with an optical magnetometer. Phys. Rev. Res., **1**, 033087 (2019).

[85] L. Marmugi, C. Deans, and F. Renzoni. Electromagnetic induction imaging with atomic magnetometers: Unlocking the low-conductivity regime. Appl. Phys. Lett., **115**, 083503 (2019).

[86] C. Deans, L. Marmugi, S. Hussain, and F. Renzoni. Electromagnetic induction imaging with a radio-frequency atomic magnetometer. Appl. Phys. Lett., **108**, 103503 (2016).

[87] C. Deans, L. Marmugi, and F. Renzoni. Sub-Sm^{-1} electromagnetic induction imaging with an unshielded atomic magnetometer. arXiv 2002.04943 (2020).

[88] C. Deans, L. D. Griffin, L. Marmugi, and F. Renzoni. Machine Learning Based Local-
     ization and Classification with Atomic Magnetometers. Phys. Rev. Lett., **120**, 033204
     (2018).

[89] L. Marmugi, L. Gori, S. Hussain, C. Deans, and F. Renzoni. Remote detection of rotat-
     ing machinery with a portable atomic magnetometer. Appl. Opt., **56**, 743 (2017).

[90] P. Bevington, R. Gartman, W. Chalupczak, C. Deans, L. Marmugi, and F. Renzoni.
     Non-destructive structural imaging of steelwork with atomic magnetometers. Appl.
     Phys. Lett., **113**, 063503 (2018).

[91] P. Bevington, R. Gartman, and W. Chalupczak. Imaging of material defects with a
     radio-frequency atomic magnetometer. Rev. Sci. Instrum., **90**, 013103 (2019).

[92] P. Bevington, R. Gartman, and W. Chalupczak. Enhanced material defect imaging
     with a radio-frequency atomic magnetometer. J. Appl. Phys., **125**, 094503 (2019).

[93] C. Deans, L. Marmugi, and F. Renzoni. Through-barrier electromagnetic imaging with
     an atomic magnetometer. Opt. Express, **25**, 17911 (2017).

[94] C. Deans, L. Marmugi, and F. Renzoni. Active underwater detection with an array of
     atomic magnetometers. Appl. Opt., **57**, 2346 (2018).

[95] M. Pospelov, S. Pustelny, M. P. Ledbetter, D. F. Jackson Kimball, W. Gawlik, and
     D. Budker. Detecting Domain Walls of Axionlike Models Using Terrestrial Experi-
     ments. Phys. Rev. Lett., **110**, 021803 (2013).

[96] S. Pustelny, D. F. Jackson Kimball, C. Pankow, M. P. Ledbetter, P. Wlodarczyk,
     P. Wcislo, M. Pospelov, J. R. Smith, J. Read, W. Gawlik, and D. Budker. The Global
     Network of Optical Magnetometers for Exotic physics (GNOME): A novel scheme to
     search for physics beyond the Standard Model. Ann. Phys., **525**, 659 (2013).

[97] Y. V. Stadnik. *Manifestations of Dark Matter and Variations of the Fundamental Constants
     in Atoms and Astrophysical Phenomena* (Springer, 2017).

[98] Y. V. Stadnik and V. V. Flambaum. Axion-induced effects in atoms, molecules, and
     nuclei: Parity nonconservation, anapole moments, electric dipole moments, and spin-
     gravity and spin-axion momentum couplings. Phys. Rev. D, **89**, 043522 (2014).

[99] C. Abel, N. J. Ayres, G. Ban, G. Bison, K. Bodek, V. Bondar, M. Daum, M. Fairbairn,
     V. V. Flambaum, P. Geltenbort, K. Green, W. Clark Griffith, M. van der Grinten, Z. D.
     Grujić, P. G. Harris, N. Hild, P. Iaydjiev, S. N. Ivanov, M. Kasprzak, Y. Kermaidic,
     K. Kirch, H.-C. Koch, S. Komposch, P. A. Koss, A. Kozela, J. Krempel, B. Lauss,
     T. Lefort, Y. Lemière, D. J. E. Marsh, P. Mohanmurthy, A. Mtchedlishvili, M. Musgrave,
     F. M. Piegsa, G. Pignol, M. Rawlik, D. Rebreyend, D. Ries, S. Roccia, D. Rozpędzik,
     P. Schmidt-Wellenburg, N. Severijns, D. Shiers, Y. V. Stadnik, A. Weis, E. Wursten,

J. Zejma, and G. Zsigmond. Search for Axionlike Dark Matter through Nuclear Spin Precession in Electric and Magnetic Fields. Phys. Rev. X, **7**, 041034 (2017).

[100] T. Wang, D. F. Jackson Kimball, A. O. Sushkov, D. Aybas, J. W. Blanchard, G. Centers, S. R. O'Kelley, A. Wickenbrock, J. Fang, and D. Budker. Application of spin-exchange relaxation-free magnetometry to the Cosmic Axion Spin Precession Experiment. Phys. Dark Universe, **19**, 27 (2018).

[101] E. Corsini, V. Acosta, N. Baddour, J. Higbie, B. Lester, P. Licht, B. Patton, M. Prouty, and D. Budker. Search for plant biomagnetism with a sensitive atomic magnetometer. J. Appl. Phys., **109**, 074701 (2011).

[102] A. Nikulin and T. S. de Smet. A UAV-based magnetic survey method to detect and identify orphaned oil and gas wells. Lead. Edge, **38**, 418 (2019).

[103] M. Hayakawa, K. Hattori, and K. Ohta. Monitoring of ULF (Ultra-Low-Frequency) Geomagnetic Variations Associated with Earthquakes. Sensors, **7**, 1108 (2007).

[104] A. C. Fraser-Smith, A. Bernardi, P. R. McGill, M. E. Ladd, R. A. Helliwell, and O. G. Villard Jr. Low-frequency magnetic field measurements near the epicenter of the M_s 7.1 Loma Prieta Earthquake. Geophys. Res. Lett., **17**, 1465 (1990).

[105] M. Hayakawa, R. Kawate, O. A. Molchanov, and K. Yumoto. Results of ultra-low-frequency magnetic field measurements during the Guam Earthquake of 8 August 1993. Geophys. Res. Lett., **23**, 241 (1996).

[106] S. J. Ingleby, I. C. Chalmers, T. E. Dyer, P. F. Griffin, and E. Riis. Resonant Very Low- and Ultra Low Frequency Digital Signal Reception Using a Portable Atomic Magnetometer. arXiv 2003.03267 (2020).

[107] V. Gerginov, F. C. S. da Silva, and D. Howe. Prospects for magnetic field communications and location using quantum sensors. Rev. Sci. Instrum., **88**, 125005 (2017).

[108] M. Janosek. Parallel Fluxgate Magnetometers. In *High Sensitivity Magnetometers*, pages 41–61 (Springer, 2017).

[109] F. Primdahl. The fluxgate magnetometer. J. Phys. E: Sci. Instrum., **12**, 241 (1979).

[110] J. L. M. J. van Bree, J. A. Poulis, and F. N. Hooge. Barkhausen noise in fluxgate magnetometers. Appl. Sci. Res., **29**, 59 (1974).

[111] H. H. J. van Dijk, J. A. Poulis, and F. N. Hooge. Noise in the measurements of magnetic susceptibilities. Appl. Sci. Res., **32**, 305 (1976).

[112] C. Hinnrichs, J. Stahl, K. Kuchenbrandt, and M. Schilling. Dependence of sensitivity and noise of fluxgate sensors on racetrack geometry. IEEE Trans. Magn., **37**, 1983 (2001).

[113] D. Scouten. Sensor noise in low-level flux-gate magnetometers. IEEE Trans. Magn., **8**, 223 (1972).

[114] M. Butta, M. Janosek, P. Ripka, L. Kraus, and R. El Kammouni. Influence of Magnetostriction of NiFe Electroplated Film on the Noise of Fluxgate. IEEE Trans. Magn., **50**, 1 (2014).

[115] F. London. *Superfluids* (John Wiley & Sons, 1954).

[116] B. S. Deaver Jr. and W. M. Fairbank. Experimental Evidence for Quantized Flux in Superconducting Cylinders. Phys. Rev. Lett., **7**, 43 (1961).

[117] R. Doll and M. Näbauer. Experimental Proof of Magnetic Flux Quantization in a Superconducting Ring. Phys. Rev. Lett., **7**, 51 (1961).

[118] P. W. Anderson and J. M. Rowell. Probable Observation of the Josephson Superconducting Tunneling Effect. Phys. Rev. Lett., **10**, 230 (1963).

[119] J. M. Rowell. Magnetic Field Dependence of the Josephson Tunnel Current. Phys. Rev. Lett., **11**, 200 (1963).

[120] R. C. Jaklevic, J. Lambe, A. H. Silver, and J. E. Mercereau. Quantum Interference Effects in Josephson Tunneling. Phys. Rev. Lett., **12**, 159 (1964).

[121] L. N. Cooper. Bound Electron Pairs in a Degenerate Fermi Gas. Phys. Rev., **104**, 1189 (1956).

[122] M. Schmelz. Development of a high sensitive receiver system for transient electromagnetics. Ph.D. thesis, University of Twente (2014).

[123] M. Podt. Wideband Low-Noise Integrated SQUID Systems. Ph.D. thesis, University of Twente (2003).

[124] M. J. van Duuren. Advanced Relaxation Oscillation SQUIDs. Ph.D. thesis, University of Twente (1998).

[125] J. W. Rohlf. *Modern Physics from α to Z^0* (Wiley, 1994).

[126] B. D. Josephson. Possible new effects in superconductive tunnelling. Phys. Lett., **1**, 251 (1962).

[127] J. Clarke and A. I. Braginski. *The SQUID Handbook*, vol. 1 (Wiley, 2004).

[128] M. Schmelz and R. Stolz. Superconducting Quantum Interference Device (SQUID) Magnetometers. In *High Sensitivity Magnetometers*, pages 279–311 (Springer, 2017).

[129] D. Drung. Advanced SQUID Read-Out Electronics. In *SQUID Sensors: Fundamentals, Fabrication and Applications*, pages 63–116 (Springer, 1996).

[130] R. L. Fagaly. Superconducting quantum interference device instruments and applications. Rev. Sci. Instrum., **77**, 101101 (2006).

[131] M. Mück. Increasing the dynamic range of a SQUID amplifier by negative feedback. Physica C, **368**, 141 (2002).

[132] C. D. Tesche and J. Clarke. dc SQUID: Noise and optimization. J. Low Temp. Phys., **29**, 301 (1977).

[133] M. Schmelz, R. Stolz, V. Zakosarenko, T. Schönau, S. Anders, L. Fritzsch, M. Mück, and H.-G. Meyer. Field-stable SQUID magnetometer with sub-fT $Hz^{-1/2}$ resolution based on sub-micrometer cross-type Josephson tunnel junctions. Supercond. Sci. Technol., **24**, 065009 (2011).

[134] T. H. Maiman. Stimulated Optical Radiation in Ruby. Nature, **187**, 493 (1960).

[135] C. Cohen-Tannoudji, J. Dupont-Roc, S. Haroche, and F. Laloë. Detection of the Static Magnetic Field Produced by the Oriented Nuclei of Optically Pumped ^{3}He Gas. Phys. Rev. Lett., **22**, 758 (1969).

[136] C. Gemmel, W. Heil, S. Karpuk, K. Lenz, C. Ludwig, Y. Sobolev, K. Tullney, M. Burghoff, W. Kilian, S. Knappe-Grüneberg, W. Müller, A. Schnabel, F. Seifert, L. Trahms, and S. Baeßler. Ultra-sensitive magnetometry based on free precession of nuclear spins. Eur. Phys. J. D, **57**, 303 (2010).

[137] S. Afach, G. Ban, G. Bison, K. Bodek, Z. Chowdhuri, Z. D. Grujić, L. Hayen, V. Hélaine, M. Kasprzak, K. Kirch, P. Knowles, H.-C. Koch, S. Komposch, A. Kozela, J. Krempel, B. Lauss, T. Lefort, Y. Lemière, A. Mtchedlishvili, O. Naviliat-Cuncic, F. M. Piegsa, P. N. Prashanth, G. Quéméner, M. Rawlik, D. Ries, S. Roccia, D. Rozpedzik, P. Schmidt-Wellenburg, N. Severjins, A. Weis, E. Wursten, G. Wyszynski, J. Zejma, and G. Zsigmond. Highly stable atomic vector magnetometer based on free spin precession. Opt. Express, **23**, 22108 (2015).

[138] G. Bison, V. Bondar, P. Schmidt-Wellenburg, A. Schnabel, and J. Voigt. Sensitive and stable vector magnetometer for operation in zero and finite fields. Opt. Express, **26**, 17350 (2018).

[139] Z. D. Grujić, P. A. Koss, G. Bison, and A. Weis. A sensitive and accurate atomic magnetometer based on free spin precession. Eur. Phys. J. D, **69**, 135 (2015).

[140] D. Hunter, S. Piccolomo, J. D. Pritchard, N. L. Brockie, T. E. Dyer, and E. Riis. Free-Induction-Decay Magnetometer Based on a Microfabricated Cs Vapor Cell. Phys. Rev. Appl., **10**, 014002 (2018).

[141] N. Behbood, F. Martin Ciurana, G. Colangelo, M. Napolitano, M. W. Mitchell, and R. J. Sewell. Real-time vector field tracking with a cold-atom magnetometer. Appl. Phys. Lett., **102**, 173504 (2013).

[142] M. Jasperse, M. J. Kewming, S. N. Fischer, P. Pakkiam, R. P. Anderson, and L. D. Turner. Continuous Faraday measurement of spin precession without light shifts. Phys. Rev. A, **96**, 063402 (2017).

[143] P.-X. Miao, W.-Q. Zheng, S.-Y. Yang, B. Wu, B. Cheng, J.-H. Tu, H.-L. Ke, W. Yang, J. Wang, J.-Z. Cui, and Q. Lin. Wide-range and self-locking atomic magnetometer based on free spin precession. J. Opt. Soc. Am. B, **36**, 819 (2019).

[144] D. Hunter, R. Jiménez-Martínez, J. Herbsommer, S. Ramaswamy, W. Li, and E. Riis. Waveform reconstruction with a Cs based free-induction-decay magnetometer. Opt. Express, **26**, 30523 (2018).

[145] V. G. Lucivero, W. Lee, M. E. Limes, E. L. Foley, T. W. Kornack, and M. V. Romalis. A femtotesla quantum-noise-limited pulsed gradiometer at finite fields. In *Quantum Inf. Meas.*, vol. V, page T3C.3 (2019).

[146] H. G. Dehmelt. Slow Spin Relaxation of Optically Polarized Sodium Atoms. Phys. Rev., **105**, 1487 (1957).

[147] H. G. Dehmelt. Modulation of a Light Beam by Precessing Absorbing Atoms. Phys. Rev., **105**, 1924 (1957).

[148] W. E. Bell and A. L. Bloom. Optical Detection of Magnetic Resonance in Alkali Metal Vapor. Phys. Rev., **107**, 1559 (1957).

[149] F. Bloch. Nuclear Induction. Phys. Rev., **70**, 460 (1946).

[150] D. Budker and D. F. Jackson Kimball. *Optical Magnetometry* (Cambridge University Press, 2013).

[151] S. Groeger, G. Bison, J.-L. Schenker, R. Wynands, and A. Weis. A high-sensitivity laser-pumped M_x magnetometer. Eur. Phys. J. D, **38**, 239 (2006).

[152] D. Arnold, S. Siegel, E. Grisanti, J. Wrachtrup, and I. Gerhardt. A rubidium M_x-magnetometer for measurements on solid state spins. Rev. Sci. Instrum., **88**, 023103 (2017).

[153] V. Schultze, R. IJsselsteijn, T. Scholtes, S. Woetzel, and H.-G. Meyer. Characteristics and performance of an intensity-modulated optically pumped magnetometer in comparison to the classical M_x magnetometer. Opt. Express, **20**, 14201 (2012).

[154] E. B. Aleksandrov, M. V. Balabas, A. K. Vershovskii, A. E. Ivanov, N. N. Yakobson, V. L. Velichanskii, and N. V. Senkov. Laser Pumping in the Scheme of an M_x-Magnetometer. Opt. Spectrosc., **78**, 292 (1995).

[155] N. Castagna, G. Bison, G. Di Domenico, A. Hofer, P. Knowles, C. Macchione, H. Saudan, and A. Weis. A large sample study of spin relaxation and magnetometric sensitivity of paraffin-coated Cs vapor cells. Appl. Phys. B, **96**, 763 (2009).

[156] P. D. D. Schwindt, B. Lindseth, S. Knappe, V. Shah, J. Kitching, and L.-A. Liew. Chip-scale atomic magnetoemter with improved sensitivity by use of the M_x technique. Appl. Phys. Lett., **90**, 081102 (2007).

[157] R. Jiménez-Martínez, W. Clark Griffith, Y.-J. Wang, S. Knappe, J. Kitching, K. Smith, and M. D. Prouty. Sensitivity Comparison of M_x and Frequency-Modulated Bell-Bloom Cs Magnetometers in a Microfabricated Cell. IEEE Trans. Instrum. Meas., **59**, 372 (2009).

[158] T. Scholtes, V. Schultze, R. IJsselsteijn, S. Woetzel, and H.-G. Meyer. Light-narrowed optically pumped M_x magnetometer with a miniaturized Cs cell. Phys. Rev. A, **84**, 043416 (2011).

[159] T. Scholtes, V. Schultze, R. IJsselsteijn, S. Woetzel, and H.-G. Meyer. Light-shift suppression in a miniaturized M_x optically pumped Cs magnetometer array with enhanced resonance signal using off-resonant laser pumping. Opt. Express, **20**, 29217 (2012).

[160] V. Schultze, B. Schillig, R. IJsselsteijn, T. Scholtes, S. Woetzel, and R. Stolz. An Optically Pumped Magnetometer Working in the Light-Shift Dispersed M_z Mode. Sensors, **17**, 561 (2017).

[161] W. Happer and H. Tang. Spin-Exchange Shift and Narrowing of Magnetic Resonance Lines in Optically Pumped Alkali Vapors. Phys. Rev. Lett., **31**, 273 (1973).

[162] W. Happer and A. C. Tam. Effect of rapid spin exchange on the magnetic-resonance spectrum of alkali vapors. Phys. Rev. A, **16**, 1877 (1977).

[163] S. Appelt, A. Ben-Amar Baranga, A. R. Young, and W. Happer. Light narrowing of rubidium magnetic-resonance lines in high-pressure optical-pumping cells. Phys. Rev. A, **59**, 2078 (1999).

[164] J. C. Allred, R. N. Lyman, T. W. Kornack, and M. V. Romalis. High-Sensitivity Atomic Magnetometer Unaffected by Spin-Exchange Relaxation. Phys. Rev. Lett., **89**, 130801 (2002).

[165] D. Budker, D. F. Jackson Kimball, and D. P. DeMille. *Atomic Physics: An Exploration through Problems and Solutions* (Oxford University Press, 2008), 2nd ed.

[166] I. M. Savukov. Spin Exchange Relaxation Free (SERF) Magnetometers. In *High Sensitivity Magnetometers*, pages 451–491 (Springer, 2017).

[167] M. P. Ledbetter, I. M. Savukov, V. M. Acosta, D. Budker, and M. V. Romalis. Spin-exchange-relaxation-free magnetometry with Cs vapor. Phys. Rev. A, **77**, 033408 (2008).

[168] H. B. Dang, A. C. Maloof, and M. V. Romalis. Ultrahigh sensitivity magnetic field and magnetization measurements with an atomic magnetometer. Appl. Phys. Lett., **97**, 151110 (2010).

[169] W. Clark Griffith, S. Knappe, and J. Kitching. Femtotesla atomic magnetometry in a microfabricated vapor cell. Opt. Express, **18**, 27167 (2010).

[170] D. Sheng, S. Li, N. Dural, and M. V. Romalis. Subfemtotesla Scalar Atomic Magnetometry Using Multipass Cells. Phys. Rev. Lett., **110**, 160802 (2013).

[171] V. Shah and M. V. Romalis. Spin-exchange relaxation-free magnetometry using elliptically polarized light. Phys. Rev. A, **80**, 013416 (2009).

[172] A. P. Colombo, T. R. Carter, A. Borna, Y.-Y. Jau, C. N. Johnson, A. L. Dagel, and P. D. D. Schwindt. Four-channel optically pumped atomic magnetometer for magnetoencephalography. Opt. Express, **24**, 15403 (2016).

[173] J. Fang, T. Wang, H. Zhang, Y. Li, and S. Zou. Optimizations of spin-exchange relaxation-free magnetometer based on potassium and rubidium hybrid optical pumping. Rev. Sci. Instrum., **85**, 123104 (2014).

[174] T. Karaulanov, I. Savukov, and Y. J. Kim. Spin-exchange relaxation-free magnetometer with nearly parallel pump and probe beams. Meas. Sci. Technol., **27**, 055002 (2016).

[175] H. J. Lee, J. H. Shim, H. S. Moon, and K. Kim. Flat-response spin-exchange relaxation free atomic magnetometer under negative feedback. Opt. Express, **22**, 19887 (2014).

[176] S. P. Krzyzewski, A. R. Perry, V. Gerginov, and S. Knappe. Characterization of noise sources in a microfabricated single-beam zero-field optically-pumped magnetometer. J. Appl. Phys., **126**, 044504 (2019).

[177] V. Shah, S. Knappe, P. D. D. Schwindt, and J. Kitching. Subpicotesla atomic magnetometry with a microfabricated vapour cell. Nat. Photonics, **1**, 649 (2007).

[178] S. J. Seltzer and M. V. Romalis. Unshielded three-axis vector operation of a spin-exchange-relaxation-free atomic magnetometer. Appl. Phys. Lett., **85**, 4804 (2004).

[179] K. Jensen, P. Kehayias, and D. Budker. Magnetometry with Nitrogen-Vacancy Centers in Diamond. In *High Sensitivity Magnetometers*, pages 553–576 (Springer, 2017).

[180] A. Wickenbrock, H. Zheng, L. Bougas, N. Leefer, S. Afach, A. Jarmola, V. M. Acosta, and D. Budker. Microwave-free magnetometry with nitrogen-vacancy centers in diamond. Appl. Phys. Lett., **109**, 053505 (2016).

[181] K. Jensen, N. Leefer, A. Jarmola, Y. Dumeige, V. M. Acosta, P. Kehayias, B. Patton, and D. Budker. Cavity-Enhanced Room-Temperature Magnetometry Using Absorption by Nitrogen-Vacancy Centers in Diamond. Phys. Rev. Lett., **112**, 160802 (2014).

[182] J. R. Maze, A. Gali, E. Togan, Y. Chu, A. Trifonov, E. Kaxiras, and M. D. Lukin. Properties of nitrogen-vacancy centers in diamond: the group theoretic approach. New J. Phys., **13**, 025025 (2011).

[183] M. W. Doherty, N. B. Manson, P. Delaney, and L. C. L. Hollenberg. The negatively charged nitrogen-vacancy centre in diamond: the electronic solution. New J. Phys., **13**, 025019 (2011).

[184] H. Clevenson, L. M. Pham, C. Teale, K. Johnson, D. Englund, and D. Braje. Robust high-dynamic-range vector magnetometry with nitrogen-vacancy centers in diamond. Appl. Phys. Lett., **112**, 252406 (2018).

[185] L. Rondin, J.-P. Tetienne, T. Hingant, J.-F. Roch, P. Maletinsky, and V. Jacques. Magnetometry with nitrogen-vacancy defects in diamond. Rep. Prog. Phys., **77**, 056503 (2014).

[186] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. L. Hollenberg. The nitrogen-vacancy colour centre in diamond. Phys. Rep., **528**, 1 (2013).

[187] V. M. Acosta, E. Bauch, M. P. Ledbetter, A. Waxman, L.-S. Bouchard, and D. Budker. Temperature Dependence of the Nitrogen-Vacancy Magnetic Resonance in Diamond. Phys. Rev. Lett., **104**, 070801 (2010).

[188] J. M. Schloss, J. F. Barry, M. J. Turner, and R. L. Walsworth. Simultaneous Broadband Vector Magnetometry Using Solid-State Spins. Phys. Rev. Appl., **10**, 034044 (2018).

[189] G. Chatzidrosos, A. Wickenbrock, L. Bougas, H. Zheng, O. Tretiak, Y. Yang, and D. Budker. Eddy-Current Imaging with Nitrogen-Vacancy Centers in Diamond. Phys. Rev. Appl., **11**, 014060 (2019).

[190] A. Hegyi and E. Yablonovitch. Molecular Imaging by Optically Detected Electron Spin Resonance of Nitrogen-Vacancies in Nanodiamonds. Nano Lett., **13**, 1173 (2013).

[191] S. K. Sarkar, A. Bumb, X. Wu, K. A. Sochacki, P. Kellman, M. W. Brechbiel, and K. C. Neuman. Wide-field *in vivo* background free imaging by selective magnetic modulation of nanodiamond fluorescence. Biomed. Opt. Express, **5**, 1190 (2014).

[192] D. A. Simpson, E. Morrisroe, J. M. McCoey, A. H. Lombard, D. C. Mendis, F. Treussart, L. T. Hall, S. Petrou, and L. C. L. Hollenberg. Non-Neurotoxic Nanodiamond Probes for Intraneuronal Temperature Mapping. ACS Nano, **11**, 12077 (2017).

[193] G. Kucsko, P. C. Maurer, N. Y. Yao, M. Kubo, H. J. Noh, P. K. Lo, H. Park, and M. D. Lukin. Nanometre-scale thermometry in a living cell. Nature, **500**, 54 (2013).

[194] J. R. Maze, P. L. Stanwix, J. S. Hodges, S. Hong, J. M. Taylor, P. Cappellaro, L. Jiang, M. V. Gurudev Dutt, E. Togan, A. S. Zibrov, A. Yacoby, R. L. Walsworth, and M. D. Lukin. Nanoscale magnetic sensing with an individual electronic spin in diamond. Nature, **455**, 644 (2008).

[195] K. Arai, C. Belthangady, H. Zhang, N. Bar-Gill, S. J. DeVience, P. Cappellaro, A. Ya-coby, and R. L. Walsworth. Fourier magnetic imaging with nanoscale resolution and compressed sensing speed-up using electronic spins in diamond. Nat. Nanotechnol., **10**, 859 (2015).

[196] G. Balasubramanian, I. Y. Chan, R. Kolesov, M. Al-Hmoud, J. Tisler, C. Shin, C. Kim, A. Wojcik, P. R. Hemmer, A. Krueger, T. Hanke, A. Leitenstorfer, R. Bratschitsch, F. Jelezko, and J. Wrachtrup. Nanoscale imaging magnetometry with diamond spins under ambient conditions. Nature, **455**, 648 (2008).

[197] H. C. Davis, P. Ramesh, A. Bhatnagar, A. Lee-Gosselin, J. F. Barry, D. R. Glenn, R. L. Walsworth, and M. G. Shapiro. Mapping the microscale origins of magnetic reso-nance image contrast with subcellular diamond magnetometry. Nat. Commun., **9**, 131 (2018).

[198] D. R. Glenn, K. Lee, H. Park, R. Weissleder, A. Yacoby, M. D. Lukin, H. Lee, R. L. Walsworth, and C. B. Connolly. Single-cell magnetic imaging using a quantum dia-mond microscope. Nat. Methods, **12**, 736 (2015).

[199] W. Gawlik and S. Pustelny. Nonlinear Magneto-Optical Rotation Magnetometers. In *High Sensitivity Magnetometers*, pages 425–450 (Springer, 2017).

[200] D. Budker, D. F. Jackson Kimball, S. M. Rochester, V. V. Yashchuk, and M. Zolotorev. Sensitive magnetometry based on nonlinear magneto-optical rotation. Phys. Rev. A, **62**, 043403 (2000).

[201] D. Budker and M. Romalis. Optical magnetometry. Nat. Phys., **3**, 227 (2007).

[202] S. Pustelny, A. Wojciechowski, M. Gring, M. Kotyrba, J. Zachorowski, and W. Gaw-lik. Magnetometry based on nonlinear magneto-optical rotation with amplitude-modulated light. J. Appl. Phys., **103**, 063108 (2008).

[203] D. Budker, W. Gawlik, D. F. Jackson Kimball, S. M. Rochester, V. V. Yashchuk, and A. Weis. Resonant nonlinear magneto-optical effects in atoms. Rev. Mod. Phys., **74**, 1153 (2002).

[204] V. Acosta, M. P. Ledbetter, S. M. Rochester, D. Budker, D. F. Jackson Kimball, D. C. Hovde, W. Gawlik, S. Pustelny, J. Zachorowski, and V. V. Yashchuk. Nonlinear magneto-optical rotation with frequency-modulated light in the geophysical field range. Phys. Rev. A, **73**, 053404 (2006).

[205] A. Wojciechowski, E. Corsini, J. Zachorowski, and W. Gawlik. Nonlinear Faraday rotation and detection of superposition states in cold atoms. Phys. Rev. A, **81**, 053420 (2010).

[206] N. Wilson, P. Light, A. Luiten, and C. Perrella. Ultrastable Optical Magnetometry. Phys. Rev. Appl., **11**, 044034 (2019).

[207] V. G. Lucivero, P. Anielski, W. Gawlik, and M. W. Mitchell. Shot-noise-limited magnetometer with sub-picotesla sensitivity at room temperature. Rev. Sci. Instrum., **85**, 113108 (2014).

[208] B. Patton, E. Zhivun, D. C. Hovde, and D. Budker. All-Optical Vector Atomic Magnetometer. Phys. Rev. Lett., **113**, 013001 (2014).

[209] C. Deans, L. Marmugi, and F. Renzoni. Sub-picotesla widely tuneable atomic magnetometer operating at room-temperature in unshielded environments. Rev. Sci. Instrum., **89**, 083111 (2018).

[210] Y. Cohen, K. Jadeja, S. Sula, M. Venturelli, C. Deans, L. Marmugi, and F. Renzoni. A cold atom radio-frequency magnetometer. Appl. Phys. Lett., **114**, 073505 (2019).

[211] W. Chalupczak, R. M. Godun, S. Pustelny, and W. Gawlik. Room temperature femtotesla radio-frequency atomic magnetometer. Appl. Phys. Lett., **100**, 242401 (2012).

[212] J. Rajroop. Radio-frequency atomic magnetometers: an analysis of interrogation regimes. Ph.D. thesis, University College London (2018).

[213] I. Savukov, T. Karaulanov, and M. G. Boshier. Ultra-sensitive high-density Rb-87 radio-frequency magnetometer. Appl. Phys. Lett., **104**, 023504 (2014).

[214] S.-K. Lee, K. L. Sauer, S. J. Seltzer, O. Alem, and M. V. Romalis. Subfemtotesla radio-frequency atomic magnetometer for detection of nuclear quadrupole resonance. Appl. Phys. Lett., **89**, 214106 (2006).

[215] D. A. Keder, D. W. Prescott, A. W. Conovaloff, and K. L. Sauer. An unshielded radio-frequency atomic magnetometer with sub-femtoTesla sensitivity. AIP Adv., **4**, 127159 (2014).

[216] M. P. Ledbetter, V. M. Acosta, S. M. Rochester, D. Budker, S. Pustelny, and V. V. Yashchuk. Detection of radio-frequency magnetic fields using nonlinear magneto-optical rotation. Phys. Rev. A, **75**, 023405 (2007).

[217] I. M. Savukov, S. J. Seltzer, M. V. Romalis, and K. L. Sauer. Tunable Atomic Magnetometer for Detection of Radio-Frequency Magnetic Fields. Phys. Rev. Lett., **95**, 063004 (2005).

[218] F. D. Colegrove and P. A. Franken. Optical Pumping of Helium in the $^3S_1$ Metastable State. Phys. Rev., **119**, 680 (1960).

[219] L. D. Schearer, F. D. Colegrove, and G. K. Walters. Optically Pumped Nuclear Magnetometer. Rev. Sci. Instrum., **34**, 1363 (1963).

[220] O. Nebel, E. E. Scherer, and K. Mezger. Evaluation of the ^{87}Rb decay constant by age comparison against the U–Pb system. Earth Planet. Sci. Lett., **301**, 1 (2011).

[221] C. J. Foot. *Atomic Physics* (Oxford University Press, 2005).

[222] A. Corney. *Atomic and Laser Spectroscopy* (Clarendon Press, 1978).

[223] I. I. Sobelman. *Atomic Spectra and Radiative Transitions* (Springer, 1992), 2nd ed.

[224] D. A. Steck. Quantum and Atom Optics. `http://atomoptics-nas.uoregon.edu/~dsteck/teaching/quantum-optics/`. (revision 0.9.10, 3 April 2014).

[225] G. W. F. Drake. *Springer Handbook of Atomic, Molecular and Optical Physics* (Springer, 2006).

[226] D. J. Griffiths. *Introduction to Quantum Mechanics* (Pearson Prentice Hall, 2005), 2nd ed.

[227] D. Suter. *The Physics of Laser-Atom Interactions* (Cambridge University Press, 1997).

[228] B. Shore. *The Theory of Coherent Atomic Excitation*, vol. 1 (Wiley, 1990).

[229] D. A. Steck. Rubidium 87 D Line Data. `http://steck.us/alkalidata`. (revision 2.1.5, 13 January 2015).

[230] D. R. Lide. *CRC Handbook of Chemistry and Physics* (CRC Press, Boca Raton, 2001), 82nd ed.

[231] M. S. Silberberg. *Principles of General Chemistry* (McGraw-Hill, 2007), 2nd ed.

[232] P. Siddons, C. S. Adams, C. Ge, and I. G. Hughes. Absolute absorption on rubidium D lines: comparison between theory and experiment. J. Phys. B, **41**, 155004 (2008).

[233] G. Rigolin. A simple derivation of the Schrödinger uncertainty relation. Eur. J. Phys., **36**, 065007 (2015).

[234] S. Pustelny, D. F. Jackson Kimball, S. M. Rochester, V. V. Yashchuk, and D. Budker. Influence of magnetic-field inhomogeneity on nonlinear magneto-optical resonances. Phys. Rev. A, **74**, 063406 (2006).

[235] D. Budker, D. F. Jackson Kimball, V. V. Yashchuk, and M. Zolotorev. Nonlinear magneto-optical rotation with frequency-modulated light. Phys. Rev. A, **65**, 055403 (2002).

[236] D. Budker, V. Yashchuk, and M. Zolotorev. Nonlinear Magneto-optic Effects with Ultranarrow Widths. Phys. Rev. Lett., **81**, 5788 (1998).

[237] D. Budker, D. F. Jackson Kimball, S. M. Rochester, and V. V. Yashchuk. Nonlinear electro- and magneto-optical effects related to Bennett structures. Phys. Rev. A, **65**, 033401 (2002).

[238] W. Gawlik and S. Pustelny. *New Trends in Quantum Coherence*, chap. Nonlinear Faraday Effect and Its Applications, pages 47–84 (2009).

[239] M. V. Balabas, T. Karaulanov, M. P. Ledbetter, and D. Budker. Polarized Alkali-Metal Vapor with Minute-Long Transverse Spin-Relaxation Time. Phys. Rev. Lett., **105**, 070801 (2010).

[240] R. Jiménez-Martínez, W. Clark Griffith, S. Knappe, J. Kitching, and M. Prouty. High-bandwidth optical magnetometer. J. Opt. Soc. Am. B, **29**, 3398 (2012).

[241] R. J. Knize. Spin destruction in rubidium-rubidium and potassium-potassium collisions. Phys. Rev. A, **40**, 6219 (1989).

[242] M. Auzinsh, D. Budker, and S. M. Rochester. *Optically Polarized Atoms: Understanding Light-Atom Interactions* (Oxford University Press, 2010).

[243] A. B. Matsko, I. Novikova, G. R. Welch, D. Budker, D. F. Jackson Kimball, and S. M. Rochester. Vacuum squeezing in atomic media via self-rotation. Phys. Rev. A, **66**, 043815 (2002).

[244] F. Mandl. *Statistical Physics* (Wiley, 1988).

[245] S. Pustelny. Nonlinear magneto-optical effects. Ph.D. thesis, Jagiellonian University (2008).

[246] D. F. Jackson Kimball. Nonlinear Magneto-Optical Rotation with Frequency-Modulated Light. Ph.D. thesis, University of California (2005).

[247] V. Parthasarathi, S. Ganapathy, and C. A. McDowell. The evaluation and characterization of Voigt lineshape. J. Mol. Struct., **64**, 29 (1980).

[248] G. K. Wertheim, M. A. Butler, K. W. West, and D. N. E. Buchanan. Determination of the Gaussian and Lorentzian content of experimental line shapes. Rev. Sci. Instrum., **45**, 1369 (1974).

[249] J. Humlíček. Optimized computation of the Voigt and complex probability functions. J. Quant. Spectrosc. Radiat. Transf., **27**, 437 (1982).

[250] S. R. Drayson. Rapid computation of the Voigt profile. J. Quant. Spectrosc. Radiat. Transf., **16**, 611 (1976).

[251] S. M. Abrarov and B. M. Quine. Accurate Approximations for the Complex Error Function with Small Imaginary Argument. J. Math. Res., **7**, 44 (2015).

[252] J. J. Olivero and R. L. Longbothum. Empirical fits to the Voigt line width: A brief review. J. Quant. Spectrosc. Radiat. Transf., **17**, 233 (1977).

[253] G. Breit and I. I. Rabi. Measurement of Nuclear Spin. Phys. Rev., **38**, 2082 (1931).

[254] E. B. Alexandrov, M. Auzinsh, D. Budker, D. F. Jackson Kimball, S. M. Rochester, and V. V. Yashchuk. Dynamic effects in nonlinear magneto-optics of atoms and molecules. J. Opt. Soc. Am. B, **22**, 7 (2005).

[255] G. Bao, A. Wickenbrock, S. Rochester, W. Zhang, and D. Budker. Suppression of the Nonlinear Zeeman Effect and Heading Error in Earth-Field-Range Alkali-Vapor Magnetometers. Phys. Rev. Lett., **120**, 033202 (2018).

[256] K. Jensen, V. M. Acosta, J. M. Higbie, M. P. Ledbetter, S. M. Rochester, and D. Budker. Cancellation of nonlinear Zeeman shifts with light shifts. Phys. Rev. A, **79**, 023406 (2009).

[257] S. Pustelny, M. Koczwara, Ł. Cincio, and W. Gawlik. Tailoring quantum superpositions with linearly polarized amplitude-modulated light. Phys. Rev. A, **83**, 043832 (2011).

[258] A. Ben-Kish and M. V. Romalis. Dead-Zone-Free Atomic Magnetometry with Simultaneous Excitation of Orientation and Alignment Resonances. Phys. Rev. Lett., **105**, 193601 (2010).

[259] S. Pustelny, W. Gawlik, S. M. Rochester, D. F. Jackson Kimball, V. V. Yashchuk, and D. Budker. Nonlinear magneto-optical rotation with modulated light in tilted magnetic fields. Phys. Rev. A, **74**, 063420 (2006).

[260] G. Oelsner, V. Schultze, R. IJsselsteijn, F. Wittkämper, and R. Stolz. Sources of heading errors in optically pumped magnetometers operated in the Earth's magnetic field. Phys. Rev. A (2019).

[261] M. Faraday. On the magnetization of light and the illumination of magnetic lines of force. Philos. Trans. Royal Soc., **136**, 1 (1846).

[262] W. B. Garn, R. S. Caird, C. M. Fowler, and D. B. Thomson. Measurement of Faraday Rotation in Megagauss Fields over the Continuous Visible Spectrum. Rev. Sci. Instrum., **39**, 1313 (1968).

[263] D. L. Carr, N. L. R. Spong, I. G. Hughes, and C. S. Adams. Measuring the Faraday effect in olive oil using permanent magnets and Malus' law. Eur. J. Phys., **41**, 025301 (2020).

[264] A. Jain, J. Kumar, F. Zhou, L. Li, and S. Tripathy. A simple experiment for determining Verdet constants using alternating current magnetic fields. Am. J. Phys., **67**, 714 (1999).

[265] L. Sun, S. Jiang, J. D. Zuegel, and J. R. Marciante. Effective Verdet constant in a terbium-doped-core phosphate fiber. Opt. Lett., **34**, 1699 (2009).

[266] P. A. Williams, A. H. Rose, G. W. Day, T. E. Milner, and M. N. Deeter. Temperature dependence of the Verdet constant in several diamagnetic glasses. Appl. Opt., **30**, 1176 (1991).

[267] A. B. Villaverde, D. A. Donatti, and D. G. Bozinis. Terbium gallium garnet Verdet constant measurements with pulsed magnetic field. J. Phys. C, **11**, L495 (1978).

[268] L. Weller, K. S. Kleinbach, M. A. Zentile, S. Knappe, I. G. Hughes, and C. S. Adams. Optical isolator using an atomic vapor in the hyperfine Paschen–Back regime. Opt. Lett., **37**, 3405 (2012).

[269] G. Labeyrie, C. Miniatura, and R. Kaiser. Large Faraday rotation of resonant light in a cold atomic cloud. Phys. Rev. A, **64**, 033402 (2001).

[270] G. B. Thomas Jr. and R. L. Finney. *Calculus and Analytic Geometry* (Addison-Wesley, 1996), 9th ed.

[271] W. R. Bennett Jr. Hole Burning Effects in a He-Ne Optical Maser. Phys. Rev., **126**, 580 (1962).

[272] D. Budker, D. F. Jackson Kimball, S. M. Rochester, and V. V. Yashchuk. Nonlinear Magneto-optics and Reduced Group Velocity of Light in Atomic Vapor with Slow Ground State Relaxation. Phys. Rev. Lett., **83**, 1767 (1999).

[273] W. E. Bell and A. L. Bloom. Optically Driven Spin Precession. Phys. Rev. Lett., **6**, 280 (1961).

[274] D. Budker, D. F. Jackson Kimball, S. M. Rochester, and V. V. Yashchuk. Nonlinear Magneto-optical Rotation via Alignment-to-Orientation Conversion. Phys. Rev. Lett., **85**, 2088 (2000).

[275] K. Blum. *Density Matrix Theory and Applications* (Springer, 2012), 3rd ed.

[276] S. Stenholm. *Foundations of Laser Spectroscopy* (John Wiley & Sons, 1984).

[277] W. Demtröder. *Laser Spectroscopy: Basic Principles* (Springer, 2008), 4th ed.

[278] J. J. Sakurai. *Modern Quantum Mechanics* (Addison-Wesley, 1995).

[279] S. C. Rand. *Lectures on Light: Nonlinear and Quantum Optics using the Density Matrix* (Oxford University Press, 2010).

[280] R. W. Boyd. *Nonlinear Optics* (Elsevier, 2008), 3rd ed.

[281] M. H. Levitt. *Spin Dynamics: Basis of Nuclear Magnetic Resonance* (John Wiley & Sons, 2008), 2nd ed.

[282] R. Loudon. *The Quantum Theory of Light* (Oxford University Press, 2000), 3rd ed.

[283] S. M. Rochester. Modeling Nonlinear Magneto-optical Effects in Atomic Vapors. Ph.D. thesis, University of California (2010).

[284] S. M. Rochester and D. Budker. Atomic polarization visualized. Am. J. Phys., **69**, 450 (2001).

[285] V. V. Yashchuk, D. Budker, W. Gawlik, D. F. Jackson Kimball, Y. P. Malakyan, and S. M. Rochester. Selective Addressing of High-Rank Atomic Polarization Moments. Phys. Rev. Lett., **90**, 253001 (2003).

[286] S. Pustelny, D. F. Jackson Kimball, S. M. Rochester, V. V. Yashchuk, W. Gawlik, and D. Budker. Pump-probe nonlinear magneto-optical rotation with frequency-modulated light. Phys. Rev. A, **73**, 023817 (2006).

[287] K. Sycz, A. M. Wojciechowski, and W. Gawlik. Magneto-optical effects and rf magnetic field detection in cold rubidium atoms. J. Phys. Conf. Ser., **497**, 012006 (2014).

[288] K. Mouloudakis, M. Loulakis, and I. K. Kominis. Quantum trajectories in spin-exchange collisions reveal the nature of spin-noise correlations in multispecies alkali-metal vapors. Phys. Rev. Res., **1**, 033017 (2019).

[289] W. Happer. Optical Pumping. Rev. Mod. Phys., **44**, 169 (1972).

[290] S. J. Smullin, I. M. Savukov, G. Vasilakis, R. K. Ghosh, and M. V. Romalis. Low-noise high-density alkali-metal scalar magnetometer. Phys. Rev. A, **80**, 033420 (2009).

[291] T. G. Walker and W. Happer. Spin-exchange optical pumping of noble-gas nuclei. Rev. Mod. Phys., **69**, 629 (1997).

[292] H. N. de Freitas, M. Oria, and M. Chevrollier. Spectroscopy of cesium atoms adsorbing and desorbing at a dielectric surface. Appl. Phys. B, **75**, 703 (2002).

[293] J. Vanier, J.-F. Simard, and J.-S. Boulanger. Relaxation and frequency shifts in the ground state of Rb^{85}. Phys. Rev. A, **9**, 1031 (1974).

[294] F. A. Franz and C. Volk. Spin relaxation of rubidium atoms in sudden and quasimolecular collisions with light-noble-gas atoms. Phys. Rev. A, **14**, 1711 (1976).

[295] M. E. Wagshul and T. E. Chupp. Laser optical pumping of high-density Rb in polarized ^{3}He targets. Phys. Rev. A, **49**, 3854 (1994).

[296] R. J. McNeal. Disorientation Cross Sections in Optical Pumping. J. Chem. Phys., **37**, 2726 (1962).

[297] W. Franzen. Spin Relaxation of Optically Aligned Rubidium Vapor. Phys. Rev., **115**, 850 (1959).

[298] R. J. Knize, Z. Wu, and W. Happer. Optical Pumping and Spin Exchange in Gas Cells. Adv. At. Mol. Phys., **24**, 223 (1988).

[299] S. J. Seltzer. Developments in Alkali-Metal Atomic Magnetometry. Ph.D. thesis, Princeton University (2008).

[300] M. A. Rosenberry, J. P. Reyes, D. Tupa, and T. J. Gay. Radiation trapping in rubidium optical pumping at low buffer-gas pressures. Phys. Rev. A, **75**, 023401 (2007).

[301] N. Sekiguchi and A. Hatakeyama. Non-negligible collisions of alkali atoms with background gas in buffer-gas-free cells coated with paraffin. Appl. Phys. B, **122**, 81 (2016).

[302] C. G. Constable and S. C. Constable. Satellite Magnetic Field Measurements: Applications in Studying the Deep Earth. In *The State of the Planet: Frontiers and Challenges in Geophysics* (American Geophysical Union, 2004).

[303] T. J. Sumner, J. M. Pendlebury, and K. F. Smith. Convectional magnetic shielding. J. Phys. D, **20**, 1095 (1987).

[304] E. Baum and J. Bork. Systematic design of magnetic shields. J. Magn. Magn. Mater., **101**, 69 (1991).

[305] G. D. Cates, S. R. Schaefer, and W. Happer. Relaxation of spins due to field inhomogeneities in gaseous samples at low magnetic fields and low pressures. Phys. Rev. A, **37**, 2877 (1988).

[306] S. Jeon, G. Jang, H. Choi, and S. Park. Magnetic Navigation System With Gradient and Uniform Saddle Coils for the Wireless Manipulation of Micro-Robots in Human Blood Vessels. IEEE Trans. Magn., **46**, 1943 (2010).

[307] J. L. Kirschvink. Uniform magnetic fields and double-wrapped coil systems: Improved techniques for the design of bioelectromagnetic experiments. Bioelectromagnetics, **13**, 401 (1992).

[308] E. E. Callaghan and S. H. Maslen. The Magnetic Field of a Finite Solenoid. Tech. rep., National Aeronautics and Space Administration (1960).

[309] M. A. Green and J. M. de Olivares. Design of solenoid magnets using Legendre functions. In *Proc. 6th Int. Conf. Magn. Technol.* (1977).

[310] V. M. Acosta, M. Auzinsh, W. Gawlik, P. Grisins, J. M. Higbie, D. F. Jackson Kimball, L. Krzemien, M. P. Ledbetter, S. Pustelny, S. M. Rochester, V. V. Yashchuk, and D. Budker. Production and detection of atomic hexadecapole at Earth's magnetic field. Opt. Express, **16**, 11423 (2008).

[311] R. S. Grewal, S. Pustelny, A. Rybak, and M. Florkowski. Transient dynamics of a nonlinear magneto-optical rotation. Phys. Rev. A, **97**, 043832 (2018).

[312] G. Jin, Y. Xu, and Z. Wang. Transient evolution of optical magnetic resonance in rubidium vapor. Opt. Express, **27**, 7087 (2019).

[313] R. S. Grewal and S. Pustelny. Transient dynamics of nonlinear magneto-optical rotation in the presence of a transverse magnetic field. Phys. Rev. A, **101**, 033825 (2020).

[314] R. Li, C. Perrella, P. Light, and A. Luiten. A high-sensitivity, high-bandwidth atomic magnetometer for geophysical field measurements. In Preparation (2020).

[315] S. M. Kay. *Fundamentals of Statistical Signal Processing: Estimation Theory* (Prentice Hall, 1993).

[316] S. T. Smith. Statistical Resolution Limits and the Complexified Cramér-Rao Bound. IEEE Trans. Signal Process., **53**, 1597 (2005).

[317] C. R. Rao. Information and the accuracy attainable in the estimation of statistical parameters. Bull. Calcutta Math. Soc., **37**, 81 (1945).

[318] L. L. Scharf. *Statistical Signal Processing* (Addison-Wesley, 1991).

[319] A. A. Wood, L. M. Bennie, A. Duong, M. Jasperse, L. D. Turner, and R. P. Anderson. Magnetic tensor gradiometry using Ramsey interferometry of spinor condensates. Phys. Rev. A, **92**, 053604 (2015).

[320] D. Robbes. Highly sensitive magnetometers – a review. Sens. Actuators A, **129**, 86 (2006).

[321] M. W. Mitchell and S. P. Alvarez. Quantum limits to the energy resolution of magnetic field sensors. arXiv 1905.00618 (2019).

[322] E. M. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan. *Magnetic Resonance Imaging: Physical Principles and Sequence Design* (John Wiley & Sons, 1999).

[323] N. Otterstrom, R. C. Pooser, and B. J. Lawrie. Nonlinear optical magnetometry with accessible *in situ* optical squeezing. Opt. Lett., **39**, 6533 (2014).

[324] T. Horrom, R. Singh, J. P. Dowling, and E. E. Mikhailov. Quantum-enhanced magnetometer with low-frequency squeezing. Phys. Rev. A, **86**, 023803 (2012).

[325] F. Wolfgramm, A. Cerè, F. A. Beduini, A. Predojević, M. Koschorreck, and M. W. Mitchell. Squeezed-Light Optical Magnetometry. Phys. Rev. Lett., **105**, 053601 (2010).

[326] N. Wilson, R. Li, C. Perrella, P. Light, R. Anderson, and A. Luiten. A high-bandwidth atomic magnetometer. Proc. SPIE, **11200**, 112000M (2019).

[327] J. Belfi, G. Bevilacqua, V. Biancalana, S. Cartaleva, Y. Dancheva, K. Khanbekyan, and L. Moi. Dual channel self-oscillating optical magnetometer. J. Opt. Soc. Am. B, **26**, 910 (2009).

[328] J. M. Higbie, E. Corsini, and D. Budker. Robust, high-speed, all-optical atomic magnetometer. Rev. Sci. Instrum., **77**, 113106 (2006).

[329] P. D. D. Schwindt, L. Hollberg, and J. Kitching. Self-oscillating rubidium magnetometer using nonlinear magneto-optical rotation. Rev. Sci. Instrum., **76**, 126103 (2005).

[330] E. P. Corsini. The Saga of Light-Matter Interaction and Magneto-optical Effects: Applications to Atomic Magnetometry, Laser-cooled Atoms, Atomic Clocks, Geomagnetism, and Plant Bio-magnetism. Ph.D. thesis, University of California (2012).

[331] N. Wilson, C. Perrella, R. Anderson, A. Luiten, and P. Light. Wide-bandwidth atomic magnetometry via instantaneous-phase retrieval. Phys. Rev. Res., **2**, 013213 (2020).

[332] C. Perrella, N. Wilson, R. Li, P. Light, R. Anderson, and A. Luiten. Quantum sensing of rapidly varying magnetic fields. Proc. SPIE, **XXX**, XXX (2019).

[333] H. J. Korsch, A. Klumpp, and D. Witthaut. On two-dimensional Bessel functions. J. Phys. A, **39**, 14947 (2006).

[334] J. R. Carson. Notes on the Theory of Modulation. Proc. IRE, **10**, 57 (1922).

[335] R. P. Anderson, M. J. Kewming, and L. D. Turner. Continuously observing a dynamically decoupled spin-1 quantum gas. Phys. Rev. A, **97**, 013408 (2018).

[336] W. Chalupczak, A. Wojciechowski, S. Pustelny, and W. Gawlik. Competition between the tensor light shift and nonlinear Zeeman effect. Phys. Rev. A, **82**, 023417 (2010).

[337] H. Nyquist. Thermal Agitation of Electric Charge in Conductors. Phys. Rev., **32**, 110 (1928).

[338] P. J. Mohr, D. B. Newell, and B. N. Taylor. CODATA Recommended Values of the Fundamental Physical Constants: 2014. J. Phys. Chem. Ref. Data, **45**, 043102 (2016).

[339] S.-K. Lee and M. V. Romalis. Calculation of magnetic field noise from high-permeability magnetic shields and conducting objects with simple geometry. J. Appl. Phys., **103**, 084904 (2008).

[340] T. W. Kornack, S. J. Smullin, S.-K. Lee, and M. V. Romalis. A low-noise ferrite magnetic shield. Appl. Phys. Lett., **90**, 223501 (2007).

[341] E. Oelker, R. B. Hutson, C. J. Kennedy, L. Sonderhouse, T. Bothwell, A. Goban, D. Kedar, C. Sanner, J. M. Robinson, G. E. Marti, D. G. Matei, T. Legero, M. Giunta, R. Holzwarth, F. Riehle, U. Sterr, and J. Ye. Demonstration of 4.8×10^{-17} stability at 1 s for two independent optical clocks. Nat. Photonics, **13**, 714 (2019).

[342] C. H. Holbrow, A. P. Ghosh, D. Heinzen, X. Zhu, W. W. Quivers Jr., G. Shimkaveg, P. G. Pappas, J. E. Thomas, and M. S. Feld. Complete Doppler coverage in laser optical pumping by wall-induced velocity-changing collisions. Phys. Rev. A, **34**, 2477 (1986).

[343] N. Sekiguchi, A. Hatakeyama, K. Okuma, and H. Usui. Scattering of an alkali-metal atomic beam on anti-spin-relaxation coatings. Phys. Rev. A, **98**, 042709 (2018).

[344] N. Wilson, N. Bourbeau Hébert, C. Perrella, P. Light, J. Genest, S. Pustelny, and A. Luiten. Simultaneous Observation of Nonlinear Magneto-Optical Rotation in the Temporal and Spectral Domains with an Electro-Optic Frequency Comb. Phys. Rev. Appl., **10**, 034012 (2018).

[345] B. Patton, O. Neitzke, S. Rochester, E. Bahr, S. Guttikonda, D. F. Jackson Kimball, B. Coste, I. Novikova, and D. Budker. Transfer of atomic alignment in alkali systems. In *43rd Annu. Meeting APS Div. At. Mol. Opt. Phys.* (2012).

[346] W. Demtröder. *Laser Spectroscopy 2: Experimental Techniques* (Springer, 2015), 5th ed.

[347] A. N. Luiten. *Frequency Measurement and Control: Advanced Techniques and Future Trends* (Springer, 2003).

[348] T. Udem, R. Holzwarth, and T. W. Hänsch. Optical frequency metrology. Nature, **416**, 233 (2002).

[349] T. Udem, J. Reichert, R. Holzwarth, and T. W. Hänsch. Accurate measurement of large optical frequency differences with a mode-locked laser. Opt. Lett., **24**, 881 (1999).

[350] T. Udem, J. Reichert, R. Holzwarth, and T. W. Hänsch. Absolute Optical Frequency Measurement of the Cesium D_1 Line with a Mode-Locked Laser. Phys. Rev. Lett., **82**, 3568 (1999).

[351] J. N. Eckstein, A. I. Ferguson, and T. W. Hänsch. High-Resolution Two-Photon Spectroscopy with Picosecond Light Pulses. Phys. Rev. Lett., **40**, 847 (1978).

[352] M. A. Bouchiat and J. Brossel. Relaxation of Optically Pumped Rb Atoms on Paraffin-Coated Walls. Phys. Rev., **147**, 41 (1966).

[353] M. T. Graf, D. F. Jackson Kimball, S. M. Rochester, K. Kerner, C. Wong, D. Budker, E. B. Alexandrov, M. V. Balabas, and V. V. Yashchuk. Relaxation of atomic polarization in paraffin-coated cesium vapor cells. Phys. Rev. A, **72**, 023401 (2005).

[354] I. Amidror. *Mastering the Discrete Fourier Transform in One, Two or Several Dimensions* (Springer, 2013).

[355] J. W. Cooley and J. W. Tukey. An Algorithm for the Machine Calculation of Complex Fourier Series. Math. Comput., **19**, 297 (1965).

[356] P. Stoica and R. Moses. *Spectral Analysis of Signals* (Pearson Prentice Hall, 2005).

[357] B. E. A. Saleh and M. C. Teich. *Fundamentals of Photonics* (Wiley, 1991).

[358] M. S. Bartlett. Smoothing Periodograms from Time-Series with Continuous Spectra. Nature, **161**, 686 (1948).

[359] M. S. Bartlett. Periodogram Analysis and Continuous Spectra. Biometrika, **37**, 1 (1950).

[360] F. J. Harris. On the use of windows for harmonic analysis with the discrete Fourier transform. Proc. IEEE, **66**, 51 (1978).

[361] P. Welch. The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. IEEE Trans. Audio Electroacoust., **15**, 70 (1967).

[362] R. G. Brown and P. Y. C. Hwang. *Introduction to Random Signals and Applied Kalman Filtering* (Wiley, 2012).

[363] D. W. Kammler. *A First Course in Fourier Analysis* (Cambridge University Press, New York, 2008).

[364] T. F. Chan, G. H. Golub, and R. J. LeVeque. Algorithms for Computing the Sample Variance: Analysis and Recommendations. Am. Stat., **37**, 242 (1983).

[365] W. Mendenhall, R. J. Beaver, and B. M. Beaver. *Introduction to Probability & Statistics* (Cengage Learning, 2012), 14th ed.

[366] D. W. Allan. Statistics of atomic frequency standards. Proc. IEEE, **54**, 221 (1966).

[367] J. A. Barnes, A. R. Chi, L. S. Cutler, D. J. Healey, D. B. Leeson, T. E. McGunigal, J. A. Mullen, W. L. Smith, R. L. Sydnor, R. F. C. Vessot, and G. M. R. Winkler. Characterization of Frequency Stability. IEEE Trans. Instrum. Meas. (1971).

[368] S. R. Stein. *Precision Frequency Control*, vol. 2, chap. Frequency and Time – Their Measurement and Characterization, pages 191–416 (Academic Press, New York, 1985).

[369] P. Stoica, R. L. Moses, B. Friedlander, and T. Söderström. Maximum likelihood estimation of the parameters of multiple sinusoids from noisy measurements. IEEE Trans. Acoust. Speech Signal Process., **37**, 378 (1989).

[370] X. Wang. Numerical implementation of the Hilbert transform. Master's thesis, University of Saskatchewan (2006).

[371] J. A. Inaudi and J. M. Kelly. Linear Hysteretic Damping and the Hilbert Transform. J. Eng. Mech., **121**, 626 (1995).

[372] H. Olkkonen, P. Pesola, and J. T. Olkkonen. Computation of Hilbert Transform via Discrete Cosine Transform. J. Signal Inf. Process., **1**, 18 (2010).

[373] L. Devroye. *Non-Uniform Random Variate Generation* (Springer-Verlag, 1986).

[374] A. Hatakeyama, T. Kuroda, N. Sekiguchi, and K. Ishikawa. Analysis of background gas in an alkali-metal vapor cell coated with paraffin. Appl. Phys. B, **125**, 133 (2019).

[375] A. R. Edmonds. *Angular Momentum in Quantum Mechanics* (Princeton University Press, 1957).

[376] N. Tajima. Analytical formula for numerical evaluations of the Wigner rotation matrices at high spins. Phys. Rev. C, **91**, 014320 (2015).

[377] M. E. Rose. *Elementary Theory of Angular Momentum* (Dover Publications, 1995).

Index

M-sample variance, 241–243

M_x magnetometer, 15–17, 24

M_z magnetometer, 15–17, 24

μ-metal, 83, 84, 87, 90, 139, 140, 142, 143, 194, 195

Absolute sensitivity, 13, 24, 105–107, 111, 124, 125, 169

Absorption coefficient, 33–35

ac Stark effect, 36, 37, 42, 132

Acousto-optic modulator, 75, 76, 102, 150, 215, 216

Active stabilisation, 74–76, 91, 169, 194

Aliasing, 14, 96, 230, 247, 248

Alignment-to-orientation conversion, 52, 132, 276

Allan deviation, 14, 76, 102–104, 114, 135, 136, 146, 148, 218, 219, 237, 242, 243, 245, 246

Allan variance, 102, 103, 114, 237, 241–243, 246

Amplitude modulation, 52, 68–70, 74, 101, 102, 126, 132, 145, 147, 215

Amplitude response, xxxiii, 12, 96, 115, 117–119, 124, 133, 181, 193

Amplitude spectral density, 12, 83, 84, 102, 103, 105, 114, 123, 139, 141, 146, 151, 237, 241, 243, 246

Analytic representation, 121, 122, 126, 131

Angular momentum probability surface, 63, 64, 269, 271, 273–276

Anti-commutator, 32, 57

Anti-Helmholtz coil, 89

Anti-reflection coating, 76

Anti-relaxation coating, xxxiii, 13, 18, 35, 40, 78, 81–83, 96, 150, 155, 221, 225, 226, 265

Atmospheric pressure, 81, 221

Atomic alignment, 63, 64, 275, 276

Atomic ensemble, 56, 63, 64, 73, 108, 119, 274, 275

Atomic number, 27

Atomic orientation, 13, 16, 18, 22, 275

Atomic polarisation, 13, 17, 18, 22, 30, 35, 36, 46, 50–52, 61, 64, 65, 78, 79, 97, 118, 133, 150, 155, 181, 206, 208–210, 221, 222, 224, 269

Atomic population, 16, 17, 19, 33–35, 50, 56, 57, 61–63, 202, 211

Atomic transition, 30, 31, 33, 212

Atomic vapour, xxxiii, 17, 18, 27, 30, 31, 45, 63–65, 67, 73, 78, 79, 83, 87, 88, 95, 107, 110, 113, 116, 120, 147, 150, 155, 194, 212, 221, 225, 262, 264

Avogadro constant, 30

Balanced polarimeter, 73, 92, 95, 112, 144, 147

Bandpass filter, 73, 74, 218

Bandwidth, 5, 6, 9, 12–14, 17, 22, 24, 74, 95, 98, 101, 103, 105–107, 110, 112, 115–119, 124, 129, 130, 137, 138, 193, 215, 217

Barkhausen noise, 9

Barkhausen volume, 9

Bartlett's method, 239, 240

Beat note, 217, 218

Beer-Lambert law, 33

Bennett effect, 50

Bessel function, 119–121, 128, 129, 141, 142, 223

Bessel's correction, 241

Bias current, 10, 12

Biot-Savart law, 89, 90