



THE DEVELOPMENT OF AN FPGA-BASED NMR SPECTROMETER FOR CONDENSED MATTER PHYSICS AND MEASUREMENTS OF THE STATIC AND DYNAMIC PROPERTIES OF THE HALDANE CHAIN COMPOUND $\text{NII}_2(3,5\text{-LUT})_4$

By

JAKE HEAD

A thesis submitted to
the University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

Condensed Matter Physics Group
School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham
December 2023

**UNIVERSITY OF
BIRMINGHAM**

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

© Copyright by JAKE HEAD, 2023

All Rights Reserved

ABSTRACT

Nuclear magnetic resonance (NMR) techniques can be used to probe the local magnetic environment of a compound. Over time, advancements in computing have led to the digitisation of NMR spectrometers in order to provide more stable and efficient measurement instruments. This thesis presents a fully digitised NMR spectrometer based on a field programmable gate array (FPGA) for use in condensed matter NMR experiments. The transmitter can generate three, phase-configurable, RF pulses of frequencies up to 800 MHz, whilst the receiver can process signals up to 400 MHz. The full NMR set-up was tested by successfully measuring the spin echo signal from ^{59}Co powder at 213 MHz. Quantum spin models have been a subject of interest in condensed matter physics since their introduction in the early 20th century. The range of interactions within these models has led to the observation of a wide variety of exotic quantum phenomena. In particular, the Spin-1 chain hosts an exotic gapped state where there is a finite energy gap between the ground and excited states. This thesis presents magnetisation, heat capacity and ^{13}C NMR measurements on the Spin-1 chain sample $\text{NiL}_2(3,5\text{-lut})_4$ to characterise the predicted attractive Tomonaga-Luttinger liquid (TLL) state. The magnetisation measurements investigated scaling behaviour around the low-field quantum critical point at the critical field $H_c = 4.4$ T but demonstrated poor agreement with the theoretical predictions. The heat capacity measurements found no evidence for the existence of the TLL phase down to a temperature of 2 K. The ^{13}C relaxation time constant T_1 was found to follow a power law dependence on temperature with an exponent of approximately $\alpha = -1.25$, outside the range of an attractive TLL.

ACKNOWLEDGMENTS

I would like to thank my supervisor Dr. Mingee Chung from the Condensed Matter Physics group at the University of Birmingham for his advice and guidance over the last four years. Secondly, I would like to thank Jon Perrins for his help in constructing the PPMS compatible NMR probe and Prof. Ted Forgan and Dr. Chris Muirhead for helpful discussions. Additionally, I would like to thank our collaborators: Prof. Paul Goddard from the Superconductivity and Magnetism group at the University of Warwick for providing the sample of $\text{NiI}_2(3,5\text{-lut})_4$ and heat capacity measurements, Dr. Dominik Kubicki for performing the magic angle spinning NMR measurements and Dr. Toni Shiroka from ETH Zürich and Jonas Phillippe from Paul Scherrer Institute for measurements of the temperature dependent NMR spectra. Finally, I would like to thank my family, friends and my girlfriend, Tabatha, for their support throughout this period.

LIST OF PUBLICATIONS

1. J. Head, P. Manuel, F. Orlandi, M. Jeong, M. R. Lees, R. Li and C. Greaves, "Structural, Magnetic, Magnetocaloric, and Magnetostrictive Properties of $\text{Pb}_{1-x}\text{Sr}_x\text{MnBO}_4$ ($x = 0, 0.5$, and 1.0)" *Chem. Mater.*, vol. 32, no. 23, pp. 10184-10199, 2020.
<https://doi.org/10.1021/acs.chemmater.0c03701>
2. K. Sun, M. A. Abd el-Moez, L. Shen, M. Jeong, J. Head, and M. M. Attallah, "Laser powder bed fusion of the Ni-Mn-Sn Heusler alloy for magnetic refrigeration applications" *Additive Manufacturing*, vol. 69, p. 103536, 2023.
<https://doi.org/10.1016/j.addma.2023.103536>

Contents

1	Introduction	1
2	Theory	6
2.1	Digital Signal Processing	6
2.1.1	Direct Digital Synthesis	7
2.1.2	Phase Shift Keying	12
2.1.3	Sampling Rates	16
2.1.4	Digital Filtering	20
2.1.5	Field Programmable Gate Arrays	27
2.2	Nuclear Magnetic Resonance	32
2.2.1	Introduction to nuclear magnetic resonance	33
2.2.2	NMR pulse sequences	41
2.2.3	The NMR spectrometer	45
2.2.4	NMR electronics	48
2.2.5	NMR in condensed matter	52
2.3	Magnetism	56
2.3.1	Phase transitions and spin models	56
2.3.2	The Haldane chain	62
3	FPGA-based NMR spectrometer	69
3.1	System overview	69

3.2	SDR14 spectrometer	71
3.2.1	Transmitter	77
3.2.2	Pulse programmer	88
3.2.3	Receiver	92
3.3	Python control software	100
3.4	NMR probe	107
3.5	^{59}Co NMR	112
3.5.1	Quarter wave transformer	112
3.5.2	Resonance circuit	114
3.5.3	^{59}Co spin echo	117
3.6	Future work and conclusion	119
3.6.1	Verilog	119
3.6.2	Python control software	120
3.6.3	NMR Probe	120
3.6.4	NMR measurements	121
3.6.5	Conclusion	122
4	Spin Chain	123
4.1	Sample information and methods	123
4.1.1	$\text{NiL}_2(3,5\text{-lut})_4$ sample information	124
4.1.2	Experimental	126
4.1.3	Simulation	134
4.2	Results and Discussion	139
4.2.1	Magnetisation	139
4.2.2	Heat capacity	146
4.2.3	^{13}C NMR	156
4.3	Conclusion and Future Work	164

List of Figures

2.1	A diagram of a DDS system indicating the main components. Arrows represent signals flowing between components. LUT stands for ‘Look-up table’ which is described in the main text.	7
2.2	A sine wave sampled using a 3-bit LUT according to equation (2.1). Each sample is labelled with the binary representation of its memory address in the LUT.	9
2.3	A diagram of a DDS process. A 4-bit LUT is represented by the circle, where the inner ring denotes the memory address and the outer ring indicates the sine wave sample stored at that address. The inner arrows represent a tuning word after the labelled number of clock cycles. Each colour denotes a different output wave.	11
2.4	A diagram showing the available phase states in a QPSK modulation scheme calculated from equation (2.5). Each state is labelled with its binary representation.	14
2.5	a) A generic bandlimited signal. b) The spectrum of a bandlimited signal where f_b is the bandlimit and f_s is the sampling frequency. In this case the sampling frequency satisfies the Nyquist condition. c) The spectrum of an under-sampled bandlimited signal where the shaded regions indicate the presence of aliasing	17

LIST OF FIGURES

2.6 Demonstration of ‘folding back’ of aliased signals. The ‘original’ signal indicates a frequency component of the signal before sampling. The signal is reflected twice across multiples of $f_s/2$ until its frequency is below $f_s/2$. The ‘aliased’ signal displays the frequency the ‘original’ signal will appear at after sampling.	18
2.7 a) An example signal. b) The signal in a) downsampled by $M = 2$. Every other sample has been kept. The samples between have been removed.	20
2.8 The frequency response of a 100 Hz LPF for a filter with a varying number of taps.	23
2.9 The Kaiser window function with varying values of β . $\beta = 0$ is a rectangular window.	24
2.10 Passband and stopband ripples in the frequency response of a 100 Hz LPF for a filter with 29 taps.	26
2.11 The coefficients for a 17-tap 100 Hz halfband filter.	27
2.12 The architecture of a generic FPGA chip.	28
2.13 a) An unpipelined section of digital logic that performs operations A_1 , A_2 , A_3 on a set of inputs b , c , d and e , producing the output y . b) A pipelined version of a), with clocked registers indicated by the blue squares, placed on the output of all operations.	32
2.14 a) The energy levels of an $I = 1/2$ spin in a magnetic field. Splitting is observed for a non-zero applied field. b) The spectrum of an $I = 1/2$ spin in a magnetic field	34
2.15 A generic Bloch sphere for a two-level system. $ \psi\rangle$ indicates a general state found on the surface of the sphere with polar angle θ and azimuthal angle ϕ	36

LIST OF FIGURES

2.16 a) The magnetisation vector at thermal equilibrium. b) The magnetisation vector rotates into the x - y plane due to the application of an RF pulse, B_1 . c) The precession of the magnetisation vector in the x - y plane around the applied field B_0 d) The NMR signal generated by the transverse components of the precessing magnetisation vector.	38
2.17 a) The average magnetisation vector immediately after being rotated in the x - y plane by an RF pulse. All spins are aligned along the x -axis. b) The average magnetisation vector after experiencing dephasing. The blue arrows indicate the individual spin magnetisation vectors ‘fanning out’, due to precession at a range of Larmor frequencies.	40
2.18 The spin echo pulse sequence. It consists of a $\pi/2$ pulse around the x -axis followed by a π pulse around the x -axis, with a spacing of τ . This causes the appearance of a spin echo at time τ after the π pulse.	41
2.19 a) The magnetisation vector at thermal equilibrium. b) The magnetisation vector rotates into the x - y plane due to the application of an RF pulse, B_1 . c) The precession of the magnetisation vector in the x - y plane around the applied field B_0 . The individual spin vectors ‘fan-out’ as in figure 2.17b. d) The π pulse inverts the magnetisation in the x - y plane producing the spin echo. The reversal of the ‘fanning-out’ occurs here. The ‘fanning-out’ of the spins due to dephasing in c) and d) is not shown for clarity.	42
2.20 The CPMG pulse sequence consisting of a $\pi/2$ pulse followed by a sequence of π pulses generating successive spin echoes. Each echo follows a T_2^* decay envelope and the series of echoes follow a T_2 decay envelope.	43
2.21 The saturation recovery sequence used to measure T_1 . It consists of a pair of $\pi/2$ pulses around the x -axis separated by a time period τ . The signal amplitude follows an exponential recovery curve with time constant T_1	43

LIST OF FIGURES

2.22 a) The magnetisation vector at thermal equilibrium. b) The magnetisation vector rotates into the x-y plane due to the application of an RF pulse, B_1 . c) The magnetisation vector after the second $\pi/2$ pulse and $\tau \ll T_1$. d) The magnetisation vector after the second $\pi/2$ pulse and $\tau \gg T_1$. In c) and d) the green arrow represents the magnetisation before the second $\pi/2$ pulse. In d) the magnetisation vector has fully relaxed back to the z -direction.	45
2.23 The receiver path within the NMR spectrometer. The local oscillator provides a pair of IQ signals that are mixed with the NMR signal. A low pass filter is used to remove the sum frequency component of the mixed signal, leaving only the difference frequency. The resulting signals form the transverse components of the magnetisation.	46
2.24 A diagram of the components of a typical NMR transcoupler. The red arrows denote the path of the high-power RF pulse. The blue arrow indicates the path of the low-power NMR signal.	48
2.25 The I - V curve for a crossed diode. Current can only flow when the magnitude of the voltage exceeds the threshold voltage, V_t	49
2.26 A diagram of a quarter wave transformer connected to an open load. The blue line and red line are the transmitted and reflected signals. The black dashed line is the signal measured at the source.	50
2.27 The circuit diagram of the LC circuit used to detect the NMR signal.	51
2.28 a) The energy level diagram for a $I = 3/2$ spin in an external magnetic field with a quadrupole interaction as described in the text. b) The NMR spectrum for a $I = 3/2$ spin. There are three resonance peaks located at the frequencies given by equation (2.43).	52
2.29 An example of AFM ordering. The red and blue arrows indicate the anti-aligned sublattices.	57

LIST OF FIGURES

2.30 A phase diagram showing a classical and quantum phase transition. The main text describes the case when the tuning parameter r is an applied magnetic field. In this case r is a reduced tuning parameter where $r = 0$ at the critical value. Reproduced from [113].	61
2.31 The phase diagram for an isotropic Haldane chain in a magnetic field. H_c and H_s are the critical and saturation fields that denote the fields at which the system transitions between gapped and TLL phase and the TLL and saturated ferromagnet phase. Δ is the energy gap between the singlet ground state and triplet state in the gapped phase. This gap is closed by tuning the applied field to H_c . The grey hatched region indicates where the measurements presented in this thesis were taken.	64
2.32 The energy level diagram for the isotropic $S = 1$ Haldane chain in a magnetic field. For $H = 0$, the excited state is a $S = 1$ magnon. For $H > 0$, the Zeeman effect causes an overlap between the $S_z = -1$ branch of the magnon and the $S = 0$ singlet at the critical field H_c	65
2.33 The energy level diagram for the anisotropic $S = 1$ Haldane chain in a magnetic field with $D > 0$. For $H = 0$, there is a zero-field splitting in the $S = 1$ state. For $H > 0$ the behavior is the same as the isotropic case shown in figure 2.32.	66
2.34 A schematic diagram showing the energy gap, Δ , as a function of the single-ion anisotropy interaction strength D . The gap is closed at the critical value D_c . Below D_c the gapped phase is the same one as in figure 2.31. Above D_c the gap re-opens and is known as the ‘large-D’ phase.	67
3.1 A diagram of the full NMR setup. The red and blue arrows indicate data flow for the transmitter and receiver signals. Black arrows represent control signals.	70
3.2 An image of the SDR-14 hardware used in the spectrometer.	71

LIST OF FIGURES

3.3	A diagram of the internal FPGA logic. Each box represents a Verilog module that implements the labelled functionality. The red, blue and green arrows represent data flow for the transmitter, receiver and pulse programmer signals respectively. The reconstruction filter was an external inline component connected the DAC SMA output. The NMR environment block refers to every component shown in figure 3.1 after the SDR14 spectrometer	73
3.4	A diagram showing the layout of the FPGA user registers. A description of the parameters can be found in table 3.1	76
3.5	A 50 MHz sine wave generated using an 8-core DDS Python simulation. The grey line has to been superimposed in order to guide the eye.	79
3.6	a) The spectrum of the RF signal generated by the SDR14 at varying frequencies from 100 MHz - 800 MHz. The red and grey dotted lines represent the dominant and secondary frequency components expected from the ADC response. The solid black line indicates the sampling frequency. b) The spectrum of the RF signal generated by the SDR14 at varying frequencies from 100 MHz - 800 MHz, with a reconstruction LPF used to filter the high frequency components, as described in the text.	81
3.7	The output voltage of the SDR14 as function of frequency	83
3.8	An overview of the internal structure of the QPSK module. Each QPSK instance is labelled with its corresponding phase shift.	85
3.9	A 10 MHz RF signal output from the TX QPSK instance as a function of time. Each line is labelled with the phase shift and the corresponding binary encoding stored in the FPGA user registers.	86
3.10	A pair of 10 MHz LO signal outputs from the RX1 and RX2 QPSK instances labelled with their corresponding RX phase. a) The in-phase (I) signal. b) The quadrature (Q) signal.	87

3.11 A diagram showing the output of the pulse programmer module. P1, P2 and P3 are the TTL pulse lengths. G1 and G2 are gap lengths between pulses. The trigger signal is used to begin data acquisition at the start of the NMR signal.	88
3.12 A set of TTL pulse sequences generated by the pulse programmer module with the following parameters a) P1=20 μ s. b) P1=200 μ s, G1=200 μ s, P2=100 μ s. c) P1=2 ms, G1=2 ms, P2=1 ms, G2=2 ms, P3=2 ms.	90
3.13 a) A 213 MHz RF output after being gated by the TTL pulse sequence. b) A zoom-in of the start of P1. c) The spectrum of the signal during P1. There is a single component at 213 MHz as expected.	91
3.14 A schematic diagram of the output spectrum from the mixer module. The signal contained components at the sum and difference frequency as expected with an additional leakage frequency component at the LO frequency.	93
3.15 The initial state of a 5-tap 4-parallel FIR filter. The filters act on a set of input samples, $X[n]$, and produce output samples $Y[n]$. The coloured lines denoted the samples being used for each separate filter operation.	95
3.16 The next state of the 4-parallel FIR filter shown in figure 3.15. The inputs have been shifted by 4 samples.	96
3.17 A python simulation of a 29-tap, 4-parallel 200 MHz FIR LPF applied to an input signal with frequency components of 150 MHz and 250 MHz.	98
3.18 The frequency response of the 200 MHz FIR LPF implemented on the SDR14.	99
3.19 The main panel of the Python control GUI. The axes display a live-plot of acquired data during an experiment. The experimental command list is the box in the right panel. The environmental conditions are shown in the top right box, updating once a second. The buttons on the lower right side are used to modify the experiment list and start the experiment.	101

LIST OF FIGURES

3.20 The sample panel of the Python control GUI. This form allows the user to record a sample name, mass and shape.	102
3.21 The sequence tab of the Python control GUI. This form allows for entry of parameters for an NMR command. These correspond to the parameters in table 3.1.	102
3.22 An overview of the code structure used for the Python GUI. The red and blue entities are executed on the NMR and PPMS worker threads respectively. Black entities are executed on the main thread.	104
3.23 An image of the NMR probe designed for use in the QD PPMS. The labels indicate the three main sections of the probe.	108
3.24 An image of the probe head from a) the left side and b) the right side. The labelled components are discussed in the text.	110
3.25 An image of the sample chamber. The region where the sample and coil are mounted is highlighted in grey.	111
3.26 a) The 210 MHz quarter wave transformer used to obtain the ^{59}Co powder NMR signals. b) The VNA measurement of S_{11} phase shift of the quarter wave transformer shown in a).	113
3.27 The circuit diagram for the resonance circuit used for the ^{59}Co powder NMR measurements.	114
3.28 The S_{11} measurement of the tuned resonance circuit for the ^{59}Co powder NMR measurements.	115
3.29 a) The raw NMR signal obtained for the ^{59}Co powder sample. The red and blue lines indicate the M_x and M_y component respectively. Inset) A zoom-in of the dotted region. b) The magnitude of the spin echo signal obtained from the raw data.	117

LIST OF FIGURES

3.30	The ^{59}Co spectrum at measured using the FPGA-based NMR spectrometer. The spectrum was calculated from the raw data shown in figure 3.29 by performing the processing steps described in the main text.	118
4.1	a) A diagram showing the structure of $\text{NiI}_2(3,5\text{-lut})_4$. The chains run along the c direction. b) The chain viewed along the c direction. Pairs of lutidine ligands overlap when viewed from this angle. The dashed lines indicate the direction of the J and J_{\perp} intra/inter chain interactions.	124
4.2	A diagram to show the SQUID measurement setup. The sample is represented by the red square and is moved along the dotted line by the linear motor. The coil winding direction is denoted by the + and - symbols. The blue curve shows the SQUID response generated using the DC scan mode. The red curve shows the response from the SQUID-VSM operating mode. The arrows denote the range of movement during each mode of operation.	127
4.3	A schematic diagram of the heat capacity puck provided by QD. Reproduced from [165].	129
4.4	A schematic diagram of the heater, sample, platform and thermal bath entities in the two-tau model. K_s and K_p denote the sample-platform and platform-thermal bath respectively. $P(t)$ represents the pulse of heat applied to the platform by the heater.	131
4.5	A diagram showing a MAS NMR setup. The sample is contained within the sample holder (grey square) and is suspended at the ‘magic angle’ from the static field.	132
4.6	a) An example powder pattern spectrum caused by line broadening. b) The same spectrum measured using MAS NMR. Spinning sidebands are introduced at multiples of the rotational frequency f_{rot}	133

LIST OF FIGURES

4.7 a) The Hamiltonian file specifying the ‘spin’ model in ALPS. b) The lattice file specifying a square lattice with interactions along the edges and main diagonals in ALPS.	138
4.8 a) Susceptibility of the $\text{NiI}_2(3,5\text{-lut})_4$ powder sample at $\mu_0H = 0.1$ T as a function of temperature. The solid line denotes a Curie-Weiss fit. b) Inverse susceptibility of the $\text{NiI}_2(3,5\text{-lut})_4$ powder sample at $\mu_0H = 0.1$ T as a function of temperature.	140
4.9 The fitting parameters obtained by fitting the low-field susceptibility data to a Curie-Weiss law. The x -axis represents the start temperature of the fitting region. The end of the region was always 300 K.	141
4.10 The low temperature region of the low field susceptibility data showing the fit to equation (4.15). Inset) A zoom in of the low temperature region, a small upturn in the data can be seen at the lowest temperatures.	144
4.11 a) The theoretical magnetisation calculated using equation (4.18) for $\text{NiI}_2(3,5\text{-lut})_4$. b) Demonstration of magnetisation scaling by plotting the rescaled variables described in the text.	145
4.12 a) In-field magnetisation measurements of $\text{NiI}_2(3,5\text{-lut})_4$ from 3.8 T to 5.0 T. b) The data from a) plotted in terms of the scaling variables discussed in the text. The green dashed line is theoretical curve from equation (4.18). . .	146
4.13 The nuclear contribution to the specific heat for $\text{NiI}_2(3,5\text{-lut})_4$ at 14 T. The dashed lines indicate the contribution from each of the constituent nuclei and the black line shows the total contribution using equation (4.22).	148
4.14 An example calculation of the TLL heat capacity of $\text{NiI}_2(3,5\text{-lut})_4$ at 6 T, 9 T and 14 T using equation (2.67).	149

LIST OF FIGURES

4.15 a) The total low temperature heat capacity for $\text{NiI}_2(3,5-\text{lut})_4$ at 9 T. The black line indicates the total heat capacity and the red and blue dashed lines indicate the TLL and nuclear component respectively. b) The total calculated low temperature heat capacity for $\text{NiI}_2(3,5-\text{lut})_4$ at 6 T, 9 T and 14 T.	150
4.16 a) The zero-field heat capacity of $\text{NiI}_2(3,5-\text{lut})_4$ as a function of temperature. The circles denote the measured data, the blue and green dotted lines indicate the Debye-Einstein fits using parameter set 1 and set 2 respectively. inset) Zoomed into section highlighting the low temperature zero-field heat capacity. b) Percentage difference between the Debye-Einstein fits, QMC method and the experimental data.	152
4.17 The zero-field magnetic heat capacity of $\text{NiI}_2(3,5-\text{lut})_4$ as a function of temperature. The red circles represent the magnetic heat capacity calculated by subtracting the lattice heat capacity modelled using parameter set 2 and the blue circles indicate the result of a QMC simulation.	153
4.18 a) The in-field magnetic heat capacity of $\text{NiI}_2(3,5-\text{lut})_4$ from 0 T - 9 T. b) The in-field magnetic heat capacity divided by the temperature, to accentuate the low temperature behaviour	155
4.19 The diagonal components of the dipole tensor, A_{dip} , plotted as a function of r_{max} . The components converge for $r_{max} \approx 150 \text{ \AA}$	157
4.20 a) A diagram of a ^{28}Ni site in $\text{NiI}_2(3,5-\text{lut})_4$ viewed in the a - b plane. ^{13}C sites are numbered and their colour denotes equivalent sites in each ligand. b) The diagonal components of A_{dip} for each ^{13}C site. The site number corresponds to the numbering in a). The lines guide the eye.	158
4.21 The MAS ^{13}C NMR spectra of $\text{NiI}_2(3,5-\text{lut})_4$ at 300 K. The peaks labelled with a number indicate ‘true’ peaks. These labels are consistent with the labels in figure 4.22. The peaks labelled with a \star indicate spinning sidebands.	159

4.22 a) ^{13}C NMR spectra as a function of temperature measured using the spin echo sequence measured in a magnetic field of 7.075 T. Dotted lines are to guide the eye. b) Colour map generated from the data in a) highlighting the trajectory of each peak as a function of temperature. The dotted lines are to guide the eye to show the movement of the peak as a function of temperature. c) A comparison of the 300 K MAS NMR spectrum with the 220 K spin-echo signal. The peaks are numbered according to figure 4.21. The dotted lines indicated equivalent peaks between each spectrum. inset) The molecular structure of 3-5 lutidine, labelled with the ^{13}C sites found in the spectra. d) The NMR shift of the shifting and non-shifting peaks as a function of temperature overlaid with magnetic susceptibility measured at 7 T. Error bars have been omitted as they are too small to represent.	160
4.23 a) ^{13}C T_1 spectra obtained from MAS NMR at 300 K using the saturation recovery sequence, the arrows indicate the peaks used to measure T_1 . b) Magnetisation recovery curves calculated for a peak located at the indicated NMR shift. Error bars have been omitted as they are too small to represent.	162
4.24 The temperature dependence of ^{13}C T_1^{-1} . The dotted lines indicate power laws with varying exponents. The green region indicates the expected power law for an attractive TLL.	164

List of Tables

2.1	The relationship between the two-bit phase encoding and amplitude of the quadrature reference signals	14
2.2	Definitions for the critical exponents for a magnetic system. H is the magnetic field, t is the reduced temperature defined in the text and d is the dimensionality of the system. Table taken from [113].	59
3.1	A description of the parameters stored in the FPGA user registers.	75
3.2	Conversion table between the 5-bit TX phase and RX phase stored in the user registers and their 2-bit encoding forms.	84
4.1	Parameter values for $\text{NiI}_2(3,5\text{-lut})_4$. All values taken from [48].	125
4.2	A table showing the values of the coefficients A_i and B_i calculated using a Padé Approximation. Values taken from [183].	142
4.3	Parameters used to calculate the nuclear contribution to the heat capacity for $\text{NiI}_2(3,5\text{-lut})_4$ where NA is the natural abundance.	147
4.4	Fitting parameters used to characterise the lattice contribution to the the heat capacity using the Debye-Einstein model.	151
4.5	^{13}C T_1 values for the set of highlighted peaks in figure 4.23.	163

Chapter One

Introduction

Nuclear magnetic resonance (NMR) in solids and liquids was independently discovered in 1946 by Felix Bloch [1], [2] and Edward Mills Purcell [3]. Since then, applications for NMR techniques exist across many scientific disciplines including physics and chemistry. Most notably it underpins magnetic resonance imaging techniques used for a wide range of medical imaging applications. NMR is an invaluable tool in the field of condensed matter physics. NMR measurements provide an additional set of measurements to compare with the results of magnetization-derived data and neutron scattering studies. Additionally, they can provide direct information regarding the spin fluctuations of a system.

On the simplest level, NMR effects are caused by the interactions of nuclear spins in a compound with an externally applied magnetic field. As a result, the nuclei act as sensitive probes of their environment [4]. This allows for the behaviour of a compound to be examined under a wide range of experimental parameters such as temperature, applied magnetic field strength [5] and pressure [6]. NMR measurements also probe magnetic phase changes within a compound, highlighting the usefulness of the technique when probing quantum phase transitions and quantum critical behaviour [7].

NMR experiments are performed using an NMR spectrometer. Classically, spectrom-

eters were assembled from a set of analogue components [8]. Due to the large number of components and complicated signal routing this often leads to a system that is costly to implement and maintain with many potential sources of failure. These issues have been addressed with the digitisation of NMR spectrometers. Over time, analogue components were replaced with digital chips designed to perform a specific function [9]–[11]. For example, direct digital synthesis (DDS) chips were used in RF signal generation [12]. Eventually, paired with advancements in computing, this led to the development of a fully digitized single-component NMR spectrometer [13]. This implementation was based on the use of a field programmable gate array (FPGA) chip.

An FPGA is a general purpose chip providing functionality for basic logical and arithmetic operations. FPGA software is developed and then uploaded to the chip by the user. In addition to user created software, manufacturers provide computationally efficient pre-made functions for common tasks such as multiplication called IP Cores and these allow for relatively fast software development. An FPGA based approach provides many benefits over the classical analogue NMR spectrometer. Firstly, as it is a software-based implementation, changing the functionality of the spectrometer only requires the software to be modified in contrast with having to buy additional analogue components. Secondly, updates to the software can be developed in parallel with the operation of the spectrometer. This results in minimal downtime when the version of the software is changed. The new software can be written on a separate PC that has access to the developer tools and then uploaded to the FPGA at a convenient date. Finally, the portability of a software approach is desirable as the software will run on several commercially available FPGAs, with minor adjustments to the code. Essentially, once the initial development time of the software concludes, future spectrometers can easily be set-up.

The study of magnetic interactions between the atoms within a compound has long been the focus of condensed matter physics. The nucleons and electrons within the atom

possess an intrinsic angular momentum called ‘spin’. The spin causes these particles to have an intrinsic magnetic moment that creates a magnetic field around the particle [14]. The development of quantum mechanics in the early 20th century led to the rise of quantum spin models to describe the microscopic behaviour of an ensemble of spins. The most famous of these models is the Ising model where each spin can take one of two states. It is named after Ernst Ising who proved that the one-dimensional Ising model contained no phase transitions in 1924 [15]. Further study of the Ising model in higher dimensions provided a rich landscape for theoretical physics with the ferromagnetic, zero-field Ising model on a 2D square lattice being solved exactly by Onsager in 1942, proving the presence of a phase transition from an ordered to disordered state at a finite temperature [16].

Heisenberg introduced the Heisenberg spin model in which the spins take the form of three-component vectors [17]. The Heisenberg model has been extended with additional terms capturing the next-nearest neighbour, anisotropic and external field interactions [18]. Additionally, the Heisenberg model has been applied to a wide variety of lattices and dimensions such as, the 1D chain [18], quasi-1D spin ladders [19], [20], triangular [21], [22] and Kagome [23] lattices. The combination of terms, lattice shape and dimensionality yield a variety of exotic magnetic states and phenomena such as the quantum spin liquid state in frustrated systems [23].

The development of the theoretical models has resulted in the synthesis of ‘toy’ compounds that are used to experimentally explore the behaviour of the models. These compounds include: the spin ladder compounds $(C_5H_{12}N)_2CuBr_4$ (BPCB) [24] and $(C_7H_{10}N)_2CuBr_4$ (DIMPY) [25], the Kagome lattice compounds $Rb_2Ni_3S_4$ [26] and $NaFe_3(SeO)_2(OH)_6$ [27], the triangular lattice compound $LiCrTe_2$ [28] and the spin-1 (or $S = 1$) chain compound $[Ni(en)_2(\mu-NO_2)](ClO_4)$ (NENP, en=ethylenediamine) [29].

The $S = 1$ chain system is of particular interest after the prediction of the presence

of a gapped excitation spectrum in integer spin chains by Duncan Haldane in 1983 [30]. A range of compounds have been produced to explore the phase diagram of the $S = 1$ chain, yielding a confirmation of the existence of a gapped, ‘Haldane’ state as well as a range of other exotic states [31]. These phases include the long range ordered phases in CsNiCl_3 and LiVSi_2O_6 [32]–[34], dimerized phases in NMOHP and NDOAP [35], [36], large- D phases in $\text{NiCl}_2\text{-4SC}(\text{NH}_2)_2$ (DTN) [37] and the Haldane phase in $\text{Ni}(\text{CH}_3)_4\text{N}(\text{NO}_2)_3$ (TMNIN), $\text{Ni}(\text{C}_3\text{H}_{10}\text{N}_2)_2\text{N}_3\text{ClO}_4$ (NINAZ) and $\text{Ni}(\text{C}_3\text{H}_{10}\text{N}_2)_2\text{NO}_2\text{ClO}_4$ (NINO) [38], [39].

When the $S = 1$ chain system is placed within a magnetic field, the gapped phase breaks down and the system transitions to a Tomonaga-Luttinger liquid (TLL) phase. The TLL phase is a low-energy description of applicable to gapless 1D quantum systems [40], [41]. The properties of TLL states have been measured across a range of 1D systems including: carbon nanotubes [42], [43], quantum wires [44], [45] and edge states in the fractional quantum hall effect [46], [47]. Consequently, the TLL state is said to be a ‘universal’ low energy description as it describes a number of systems regardless of their exact physical realisations.

This thesis begins with a theoretical overview of some common digital signal processing methods and an introduction to FPGAs. Then a summary of the physical background of NMR and the functionality of the NMR spectrometer will be discussed. Next, an overview of magnetism, quantum phase transitions and the Haldane spin chain will be presented. Chapter three will detail the design, implementation and testing of the FPGA-based NMR spectrometer, control software and NMR environment produced for the University of Birmingham Condensed Matter group. This chapter will conclude with further improvements to the hardware and software within the experimental setup. Chapter four will present static and dynamic measurements on the isotropic $S = 1$ Haldane chain compound $\text{NiI}_2(3,5\text{-lut})_4$ [48]. These measurements probe various regions of the Haldane chain phase diagram including the gapped phase, Tomonaga-Luttinger liquid phase and the quantum critical region. The aim of this chapter is to characterise the Tomonaga-Luttinger liquid state by calculat-

ing the Luttinger parameters and demonstrate scaling behaviour of the magnetisation in the quantum critical region. Finally, chapter five will provide a summary of the results obtained in this thesis.

Chapter Two

Theory

This chapter will introduce the theoretical background required for the following chapters in this thesis. Firstly, an overview of digital signal processing methods and architecture of FPGAs will be discussed. Secondly, there will be an introduction to nuclear magnetic resonance including both theoretical and experimental considerations. Finally, a summary of magnetism and quantum spin chains will be presented.

2.1 Digital Signal Processing

There are two types of signals: analogue and digital. An analogue signal is a continuous electrical signal where the voltage can take any value. A digital signal is a discretised electrical signal that consists of a string of ones and zeroes representing a high and low voltage state, typically 0V and 5V. The adoption of digital signals was driven by the rise of modern electronics and technologies such as the computer during the 20th century [49].

This also resulted in the creation of the field of digital signal processing (DSP) in this era [49]. Digital signal processing is the process of manipulating a digitised signal in order to convert the original signal into a more desirable signal, for example removing

noise or encoding information into the signal. This often involves multiple processes such as signal synthesis, filtering and resampling. DSP techniques are generally responsible for many cutting-edge technologies used today such as image and speech recognition [50], [51].

This section will provide an introduction to some of the DSP techniques and digital components used within the FPGA-based NMR spectrometer designed and implemented in chapter 3. Additionally, it will outline some of the topics typically encountered within the design of FPGA-based systems.

2.1.1 Direct Digital Synthesis

Direct Digital Synthesis (DDS) refers to a technique to synthesise a digital signal at a desired frequency, which can be converted to an analogue signal via the use of an digital to analogue converter (DAC) [52]. Typically, DDS generates a pair of quadrature signals at a specified frequency [52], [53]. Quadrature means a pair of signals separated by a 90° phase shift and these are often represented by a sine and a cosine signal.

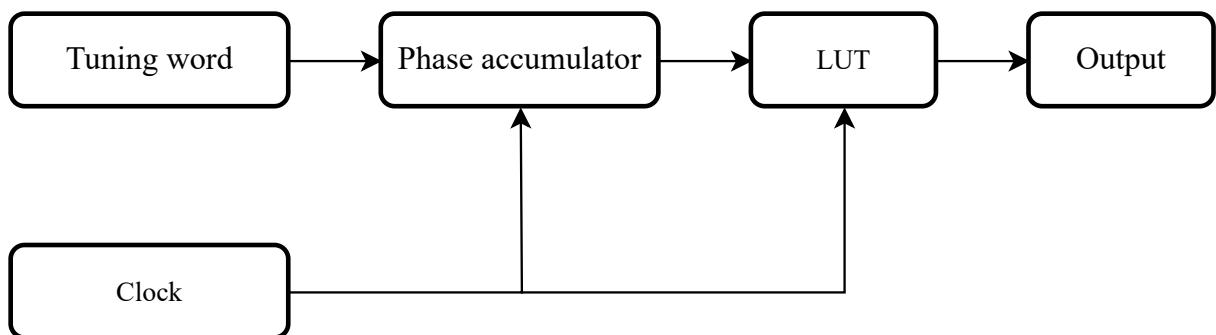


Figure 2.1: A diagram of a DDS system indicating the main components. Arrows represent signals flowing between components. LUT stands for ‘Look-up table’ which is described in the main text.

Figure 2.1 shows an overview of a typical DDS system with four key components.

The clock signal controls the operation of the DDS procedure. The frequency of this signal is determined by the hardware that generates the clock signal. A digital clock signal is typically generated using a piezoelectric crystal oscillator that generates a fixed frequency sine wave [54]. This wave is then converted to a square wave to produce the clock signal. At the lowest level, the clock signal is a 1-bit wide variable. This bit flips between 0 and 1 periodically. All internal DDS operations are executed on the rising edge (the 0 to 1 transition) of the clock signal. The frequency of the clock signal limits the maximum frequency of the operations and consequently the maximum frequency of the synthesized signal. Nyquist's theorem states that a wave must be sampled twice a cycle to be perfectly digitised [55]. Therefore DDS can only synthesize signals with a maximum frequency equal to half of the clock signal frequency.

The look up table (LUT) is a table stored within memory. Each row within the table contains two fields, a memory address and a sample of the digitised signal that the DDS module outputs. In the case of a sine wave output, each sample will be from of a digitised sine wave [52]. The function of the LUT is to convert an input memory address into the corresponding sine wave sample output. A sine LUT can be generated by sampling a full period of a sine wave into equally spaced samples. The number of samples (and the number of entries in the LUT) is specified by the length in bits of the allocated memory. For example, a sine wave sampled using a 16-bit LUT will contain 2^{16} equally spaced samples over a full period of the wave. The decimal value of a sine wave sample for a N -bit LUT can be obtained using the following formula

$$s_i = \sin\left(\frac{2\pi i}{2^N}\right) \quad i = 0, 1, \dots, (2^N - 1) \quad (2.1)$$

where s_i is the value of the sine wave at memory address i .

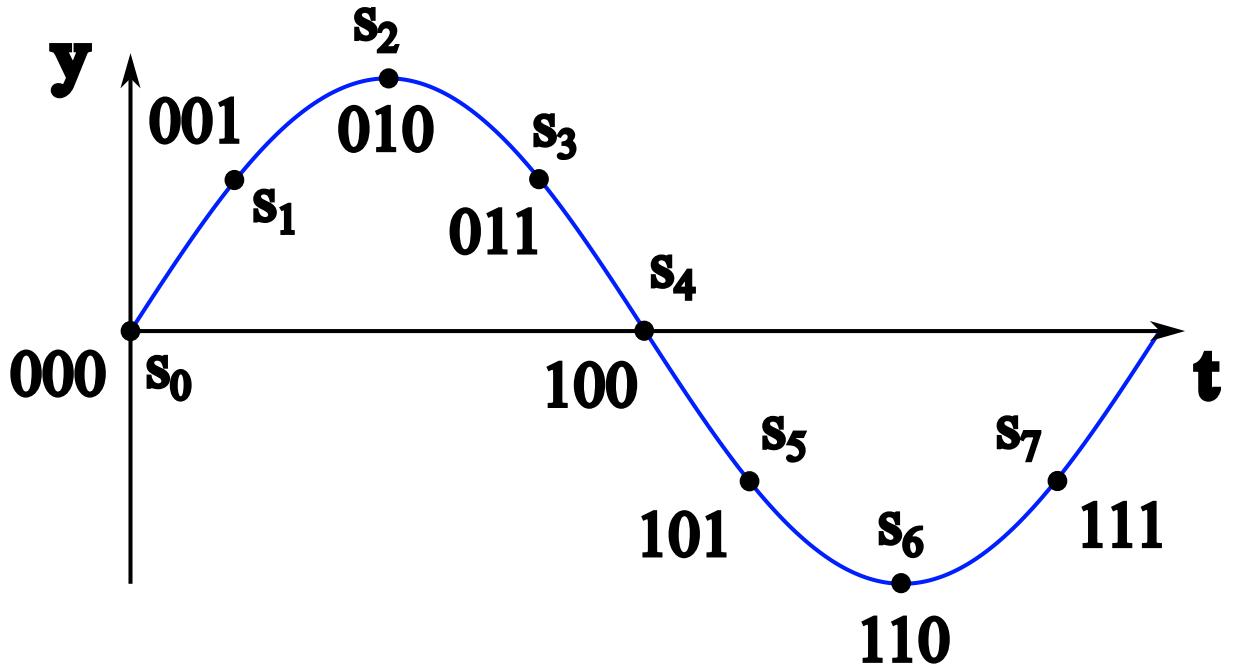


Figure 2.2: A sine wave sampled using a 3-bit LUT according to equation (2.1). Each sample is labelled with the binary representation of its memory address in the LUT.

Figure 2.2 demonstrates a sampled sine wave using a 3-bit LUT. The y-values of each sample were calculated using equation (2.1).

The sine function will output values between -1 and 1 so they need to be re-scaled in order to be represented in binary. This is done by multiplying the decimal values by a factor of $2^{(N_p-1)}$ where N_p is the width in bits of each sample. The scaled samples are then rounded to the nearest integer and converted to binary. The samples are stored in a two's complement binary format where the most significant bit is negative introducing a factor of -1 [56]. In general, N and N_p can be different. N refers to the number of bits allocated to the memory address whereas N_p is the number of bits allocated to a sample within an address. Increasing N reduces the space on the x -axis between samples in figure 2.2. N_p defines how precise the binary representation of the decimal value of a sample is. Increasing N_p , allows for a more precise decimal to binary conversion at the cost of increased memory usage [56].

The tuning word is an N -bit wide pointer. It stores the memory address of a specific sample of the sine wave in the LUT. The tuning word is the input to the phase accumulator. The phase accumulator adds a constant value to the tuning word every clock cycle, incrementing the tuning word over time. The constant is called the phase increment. The phase increment can be related to the frequency of the synthesized wave using [53]

$$\Delta\theta = \frac{f_{syn}}{f_{clk}} \times 2^N \quad (2.2)$$

where f_{syn} and f_{clk} are the frequencies of the synthesized signal and clock signal respectively. Furthermore, the frequency resolution of DDS is given by [53]

$$\Delta f_{syn} = \frac{f_{clk}}{2^N} \quad (2.3)$$

At the end of each clock cycle, the content of memory address in the LUT being referenced by the incremented tuning word pointer is outputted. By iteratively incrementing the tuning word, the pointer steps through the LUT and outputs a different sine wave sample on each positive clock edge. Connecting these samples results in the synthesis of sine wave. If the phase increment is small, the increment is smaller and therefore takes longer for the pointer to traverse the LUT, producing a lower frequency sine wave. If the phase increment is large, the pointer jumps through the LUT faster, generating a higher frequency signal. When the tuning word reaches the end of the LUT, binary overflow occurs and the pointers value wraps around to the start. A phase offset can be added to the synthesized signal by setting the initial value of the tuning word pointer to a non-zero value which begins the iteration from the middle of the LUT.

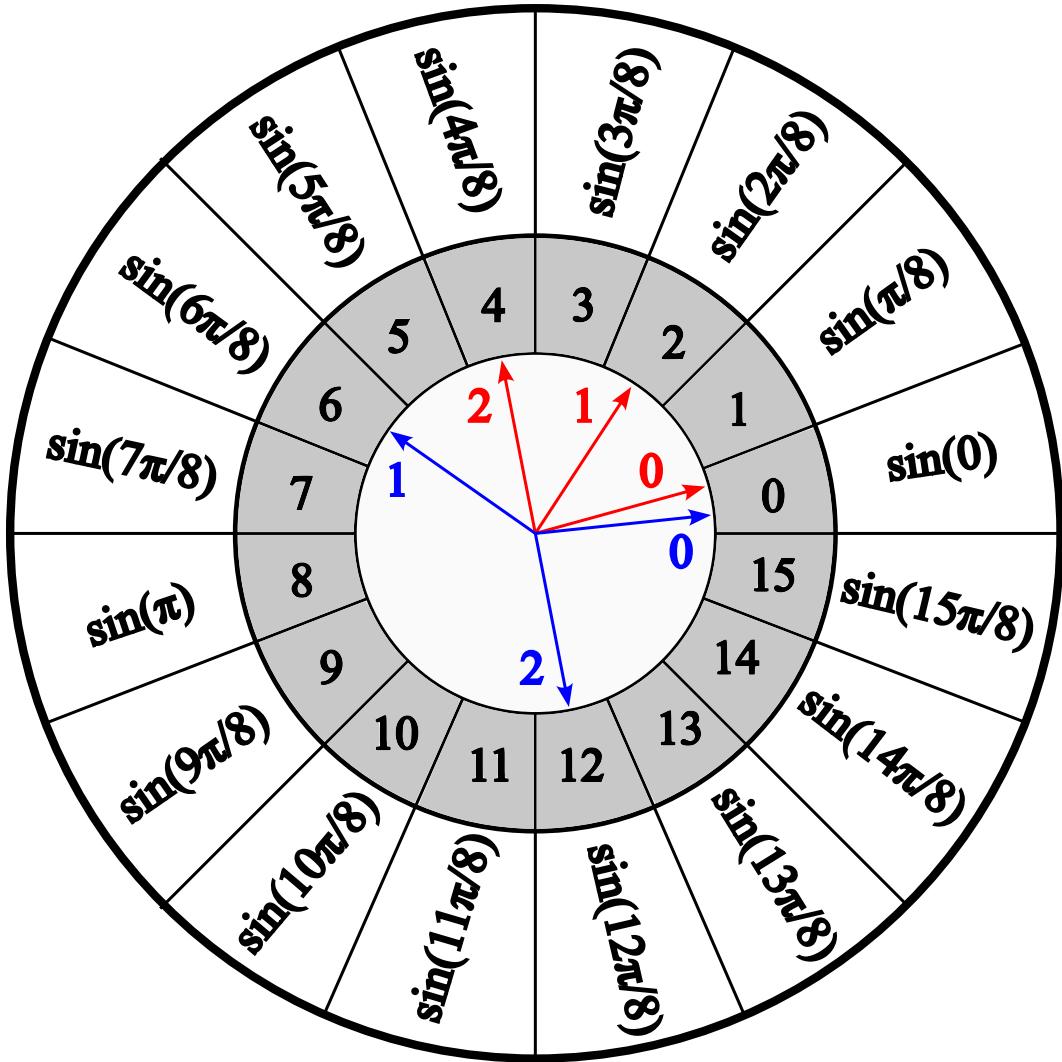


Figure 2.3: A diagram of a DDS process. A 4-bit LUT is represented by the circle, where the inner ring denotes the memory address and the outer ring indicates the sine wave sample stored at that address. The inner arrows represent a tuning word after the labelled number of clock cycles. Each colour denotes a different output wave.

Figure 2.3 shows a schematic diagram of the internal operation of a DDS process. The ring represents a 4-bit LUT. The inner grey area represents the memory address and the outer area represents the value of sine sample at that address (both using their decimal values for demonstration purposes). The inner arrows represent the value of the tuning word labelled by the elapsed number of clock cycles. A zero indicates the initial position of

the tuning word. After each clock cycle, the arrows step clockwise around the LUT by an increment equal to the phase increment. The blue arrow represents a high frequency output as it jumps between more memory address on each clock cycle than the red arrow.

Due to the symmetric nature of the sine function, in practice the memory requirements for the LUT can be reduced. Firstly, only the first quadrant of the sine wave needs to be digitised. Then samples in the remaining quadrants can be obtained by a combination of reversing the direction of the tuning word pointer or negating the value from the first quadrant [57]. Additionally, as a cosine function is just a phase-shifted sine function, a cosine signal can be synthesised by using a tuning word offset by a quarter of the sine wave LUT, at no extra memory cost [57].

2.1.2 Phase Shift Keying

Phase-shift keying (PSK) is a family of schemes that modulate a signal by varying a number of reference signals at specified times. PSK belongs to a class of phase modulation techniques which can be used to encode information within the phase of a signal, as a phase analogue of amplitude modulation and frequency modulation. This section will introduce the most general PSK technique M -ary phase-shift keying and a specific case of M -ary phase-shift keying, quadrature phase shift keying (QPSK).

M -ary phase-shift keying

PSK methods can be generalised to M -ary phase-shift keying (MPSK). M is the number of phase states used for the phase encoding and $\log_2(M)$ bits can be encoded per operation. The available phase states are maximally separated over the full range from 0 to 2π in order

to minimize encoding errors. The encoded signal can be written as [58]

$$s_{MPSK}(t) = A \cos \left(2\pi f_r t + \frac{(2i-1)}{M} \pi \right), \quad i = 1, 2, \dots, M \quad (2.4)$$

Decoding the signal requires a M -ary receiver that is synchronized to the transmitter in order to calculate the phase changes within the signal with respect to the same reference phase. Increasing M allows for an higher data rate, as more bits can be transmitted per operation, at the cost of increased power and complexity in the modulation and demodulation systems [58], [59].

Quadrature phase-shift keying

Quadrature phase-shift keying (QPSK) is a the specific case of $M = 4$, M -ary PSK. Therefore, four phase states are available for each bit of the encoded signal and the states are represent using a two-bit wide value. A QPSK encoded signal can be described by

$$s_{QPSK}(t) = A \cos \left(2\pi f_r t + \frac{(2i-1)}{4} \pi \right), \quad i = 1, 2, 3, 4 \quad (2.5)$$

The encoded signal uses four phase states separated by $\pi/2$ radians with a phase shift of: $\pi/4$, $3\pi/4$, $5\pi/4$ and $7\pi/4$ respectively [58].

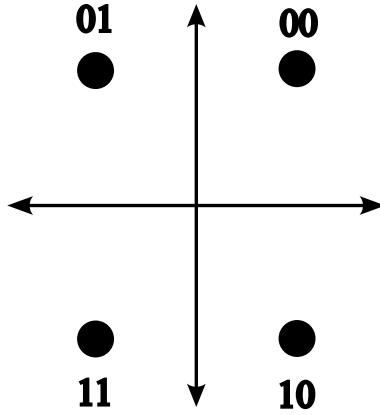


Figure 2.4: A diagram showing the available phase states in a QPSK modulation scheme calculated from equation (2.5). Each state is labelled with its binary representation.

Figure 2.4 shows the relationship between the four available phase states. Each state is labelled with its corresponding two-bit encoding. Neighbouring states are separated by $\pi/2$ radians.

These phase states are chosen as they can be generated by multiplying a pair of sine and cosine reference signals by the amplitude factor from table 2.1 and summing them. For example, to encode a $3\pi/4$ shift, the sine signal is inverted and the cosine signal is non-inverted and they are summed. This technique only relies on the relative phase shift of each of the phase states, not the absolute phase. Normally the $\pi/4$ state is designated the ‘00’ state [58], [60].

Table 2.1: The relationship between the two-bit phase encoding and amplitude of the quadrature reference signals

Phase shift	Two-bit encoding	Sin amplitude	Cos amplitude
$\pi/4$	00	Non-inverted	Non-inverted
$3\pi/4$	01	Inverted	Non-inverted
$5\pi/4$	10	Inverted	Inverted
$7\pi/4$	11	Non-inverted	Inverted

The sharp transitions in the output signal during a phase change introduce large distortions in the spectrum of the output signal. This can be removed by filtering the reference signals in order to smooth the transition between different phases [58].

2.1.3 Sampling Rates

The sampling rate or sampling frequency of a digital signal is a measure of how often the value of a digital signal is recorded. It is measured in Hz or samples per second (Sa/s). The sampling period (the time between samples) T_s is related to the sampling frequency, f_s by its inverse [61]

$$T_s = \frac{1}{f_s} \quad (2.6)$$

Nyquist's sampling theorem states that to fully reconstruct a bandlimited signal, the sampling frequency must be at least twice the highest frequency present in the signal

$$f_s \geq 2f_b \quad (2.7)$$

where f_b is the bandlimit. In practice, the sample rate is a property set by the system hardware.

A bandlimited signal contains no frequencies larger than the bandlimit (i.e. $X(f) = 0$ for $f > f_b$) [62]. Mathematically, a sampled bandlimited signal, $s(t)$, can be expressed as a multiplication between the original signal, $x(t)$, and a train of delta functions spaced by T as

$$s(t) = x(t) \sum_{n=-\infty}^{\infty} \delta(t - nT) \quad (2.8)$$

The spectrum of the signal, $S(f)$, can be found by Fourier transforming and using the convolution theorem (a multiplication in the time domain is equal to a convolution in the frequency domain).

$$S(f) = FT[x(t) \sum_{n=-\infty}^{\infty} \delta(t - nT)] = X(f) \star FT[\sum_{n=-\infty}^{\infty} \delta(t - nT)] \quad (2.9)$$

where the \star denotes the convolution. The Fourier transform of an equally spaced time-domain pulse train with spacing T can be shown to be a pulse train separated by f_s in frequency domain

$$FT[\sum_{n=-\infty}^{\infty} \delta(t - nT)] = \frac{1}{T} \sum_{n=-\infty}^{\infty} \delta(f - nf_s) \quad (2.10)$$

Performing the convolution on equation (2.9) yields the spectrum

$$S(f) = \int_{-\infty}^{\infty} X(f') \frac{1}{T} \sum_{n=-\infty}^{\infty} \delta(f' - nf_s) df' = \frac{1}{T} \sum_{n=-\infty}^{\infty} X(f - nf_s) \quad (2.11)$$

Therefore the spectrum of a bandlimited signal contains copies of the spectrum of the original signal centred at integer multiples of f_s . A bandlimited signal and its spectrum are shown in figure 2.5a and 2.5b respectively.

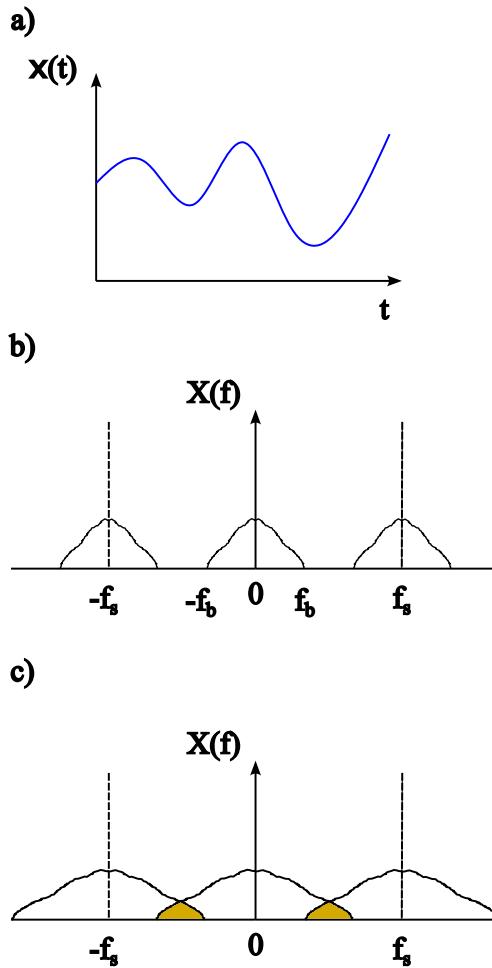


Figure 2.5: a) A generic bandlimited signal. b) The spectrum of a bandlimited signal where f_b is the bandlimit and f_s is the sampling frequency. In this case the sampling frequency satisfies the Nyquist condition. c) The spectrum of an under-sampled bandlimited signal where the shaded regions indicate the presence of aliasing

If the sampling frequency does not fulfil equation (2.7), the signal will contain distortions known as aliasing and the signal is said to be ‘under-sampled’. Aliasing is the appearance of frequencies within the spectrum of the sampled signal that were not present in the original signal [61], [63]. The aliasing is caused by regions of overlap between the copies of the spectrum of the sampled signal from equation (2.11). Figure 2.5c shows the spectrum of an under-sampled signal where overlap between copies of the spectrum cause aliasing. The frequency that the aliased signal will appear at can be found by ‘folding-back’ frequencies larger than $f_s/2$ until they fall into the range $0 < f < f_s/2$. Graphically, this is demonstrated in figure 2.6 by reflecting the signal across neighbouring lines of integer multiples of $f_s/2$ until it is in the correct range.

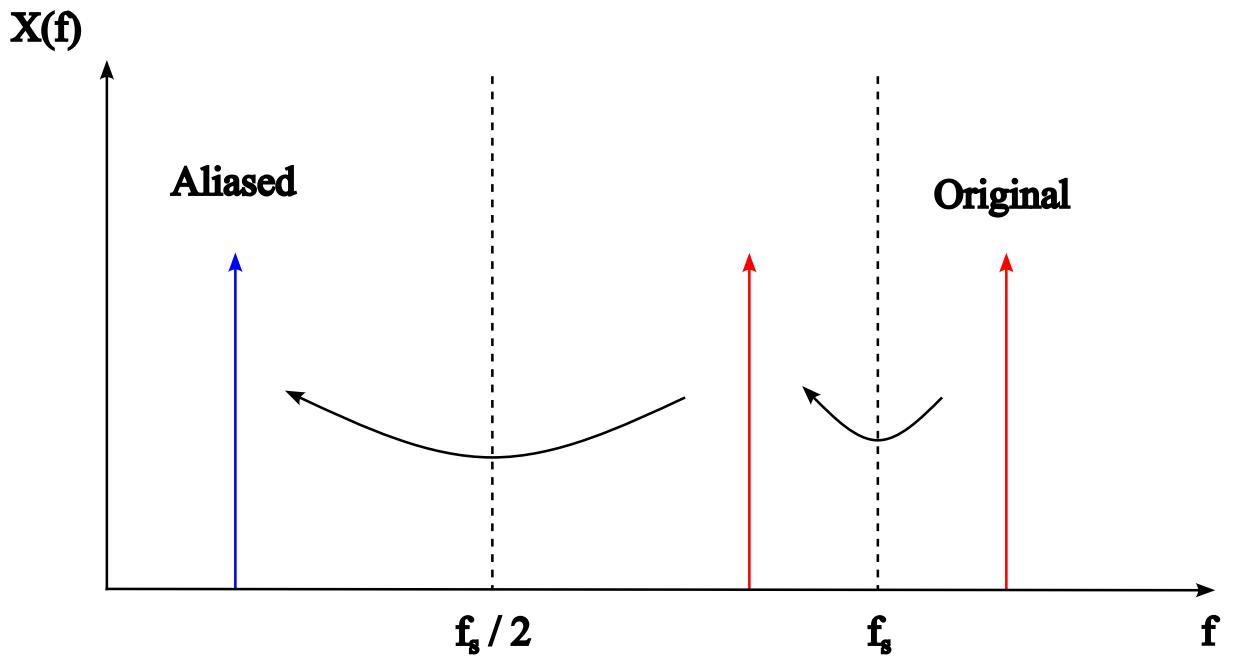


Figure 2.6: Demonstration of ‘folding back’ of aliased signals. The ‘original’ signal indicates a frequency component of the signal before sampling. The signal is reflected twice across multiples of $f_s/2$ until its frequency is below $f_s/2$. The ‘aliased’ signal displays the frequency the ‘original’ signal will appear at after sampling.

The sample rate of a signal can be modified by resampling the signal. Resampling

comes in two forms: interpolation and decimation, corresponding to the increase or decrease of the sample rate [64]. Resampling is typically used when interfacing between systems of differing sample rates. Interpolation will not be discussed any further here as it was not required for the FPGA-based spectrometer. Decimation is used when a signal is being sampled at a much larger frequency than required. If the useful information in a signal is contained in the low frequency region of the signal, the sample rate can be lowered without affecting the quality of the signal. Lowering the sample rate is beneficial in terms of memory usage to store the signal, and computational resources required to perform any further processing of the signal [65]. Both techniques require the use of a digital filter to prevent aliasing in the new signal, introduced by the manipulation of the samples in the resampling process.

Decimation consists of two steps:

1. Application of an anti-aliasing filter.
2. Downsampling.

The decimation factor, M , is defined as the ratio of input sampling frequency, $f_{s,i}$, to the output sampling frequency, $f_{s,o}$,

$$M = \frac{f_{s,i}}{f_{s,o}} \quad (2.12)$$

The signal must be filtered to remove any artefacts that can be introduced during downsampling [65]. The need for this becomes apparent with an example. Consider a signal consisting of 50 Hz and 75 Hz components sampled at a rate of 200 Hz, where the 50Hz component is the signal of interest. In this signal, both components lie within the range $0 < f < f_{s,i}/2$ so no aliasing occurs. After the application of a $M = 2$ decimation, the new sample rate becomes 100 Hz. The 50 Hz component lies within the new range $0 < f \leq f_{s,o}/2$ and still appears at the correct frequency. However, the 75 Hz component lies outside the new range and causes aliasing in the output signal. Therefore, the application of a anti-aliasing low

pass filter removes the 75 Hz frequency component that will cause aliasing in the output signal but will leave the 50 Hz component unchanged.

Downsampling is the process of removing samples from the input signal. For a factor M decimation, every M^{th} sample is kept and the $(M - 1)$ samples between are deleted. This process is shown in figure 2.7.

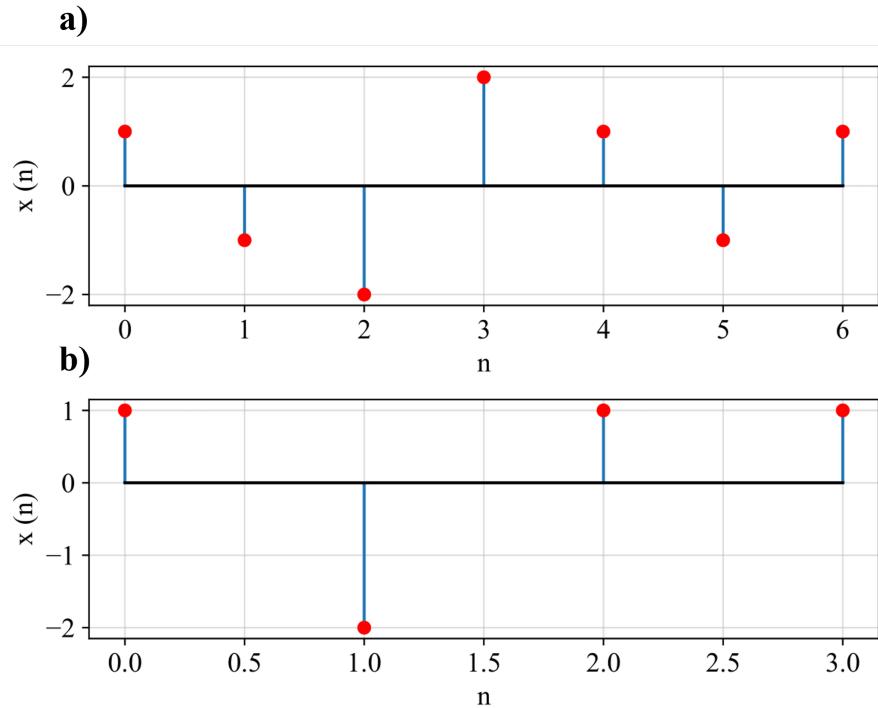


Figure 2.7: a) An example signal. b) The signal in a) downsampled by $M = 2$. Every other sample has been kept. The samples between have been removed.

2.1.4 Digital Filtering

Digital filtering is a process used to modify the frequency components of an input signal in order to produce a more desirable signal. These filters include low pass filters (LPFs), which attenuate signals above cut-off frequency and high pass filters (HPFs) which attenuate signals below a cut-off frequency [66]. LPFs and HPFs can be combined into bandpass or

bandstop filters that define a region of frequencies that will be passed or stopped based upon the characteristics of the filter.

Digital filters are an example of a linear time-invariant (LTI) system. This means that if a set of input signals pass through the system, the output signal can be written as a sum of the individually filtered signals [61]. Additionally, the output signal will be equivalent if the input signal is preceded by a time delay as the filter properties do not vary with time so the system is time-invariant [67]. LTI systems can have a linear phase delay, so the temporal order of samples within the input and output signals are preserved as all frequency components will be delayed by the same duration. LTI systems are commonly used for DSP as their behaviour can be fully characterised by their impulse response function that describes the behaviour of a system to an impulse input [66]. The output, $y(n)$ of a LTI digital filter with an impulse response, $h(n)$, operating on the input signal, $x(n)$, is given by their convolution [61]

$$y(n) = h(n) \star x(n) \quad (2.13)$$

In the frequency domain, the convolution theorem yields

$$Y(f) = H(f)X(f) \quad (2.14)$$

The impulse response of a filter consists of a set of coefficients known as taps. Due to linearity, the output of the filter can be written as the following linear superposition:

$$y(n) = \sum_{k=1}^N a_k y(n-k) + \sum_{k=0}^M b_k x(m-k) \quad (2.15)$$

Equation (2.15) states that the filter output is the sum of the N previous output samples and M previous input samples, weighted by the filter coefficients $\{a_k, b_k\}$. Note that there is no a_0 term in the sum as this would correspond to the current output sample, which has not been calculated yet. The ‘order’ of the filter is defined by the maximum of N and M . This equation is equivalent to the Laplace transform of the impulse response $h(n)$ given as

[66]

$$H(z) = \frac{\sum_{k=0}^M b_k z^{-k}}{1 + \sum_{k=1}^N a_k z^{-k}} \quad (2.16)$$

where the variable z is a complex frequency variable. $H(z)$ is obtained by performing a Z-transform, a Laplace transform for discrete functions, on $h(n)$. The Laplace transform is the generalisation of the Fourier transform of a complex frequency variable [68]. There are two main classes of digital filter indicated by equation (2.16) known as infinite impulse response (IIR) and finite impulse response (FIR).

IIR filters have $a_k \neq 0$ for all k [64]. These filters have a recursive output [64], where each sample is dependent on both the previous input samples and previous output samples. IIR filters tend to use fewer coefficients [61] to obtain a specified frequency response which implies smaller memory requirements and faster calculation times than the equivalent FIR implementation. Generally, IIR filters are more complex to design as they are not guaranteed to be stable, meaning that the filter output signal may grow or oscillate indefinitely, producing an unusable filter.

FIR filters are the class of filters that have $a_k = 0$ for all k so the denominator of equation (2.16) is 1 [66]. The number of taps of the FIR filter is equal to its order. The output of an FIR filter is only dependent on the previous input samples and is guaranteed to be stable regardless of the chosen coefficients. FIR filters require more coefficients than the equivalent IIR filter due to the lack of recursion increasing the memory storage and computation time of the filtering [67]. Additionally, FIR filters with symmetric coefficients about the centre tap have a linear phase change so no dephasing of the input signal occurs.

The filter coefficients can be generated using a variety of methods, the simplest of which is called the windowing method [69]. The method consists of generating the ideal impulse response of the filter and multiplying it with a window function. The window function ensures the impulse response is finite in time. The coefficients are then obtained by

sampling the truncated impulse response at regular intervals, generating an approximation of the ideal frequency response. Consequently, increasing the number of coefficients provides a sharper cutoff in the filter response due to the impulse response being sampled more often. A demonstration of the variation of a low pass FIR filter response with cutoff frequency of 100 Hz with the number of taps is shown in figure 2.8.

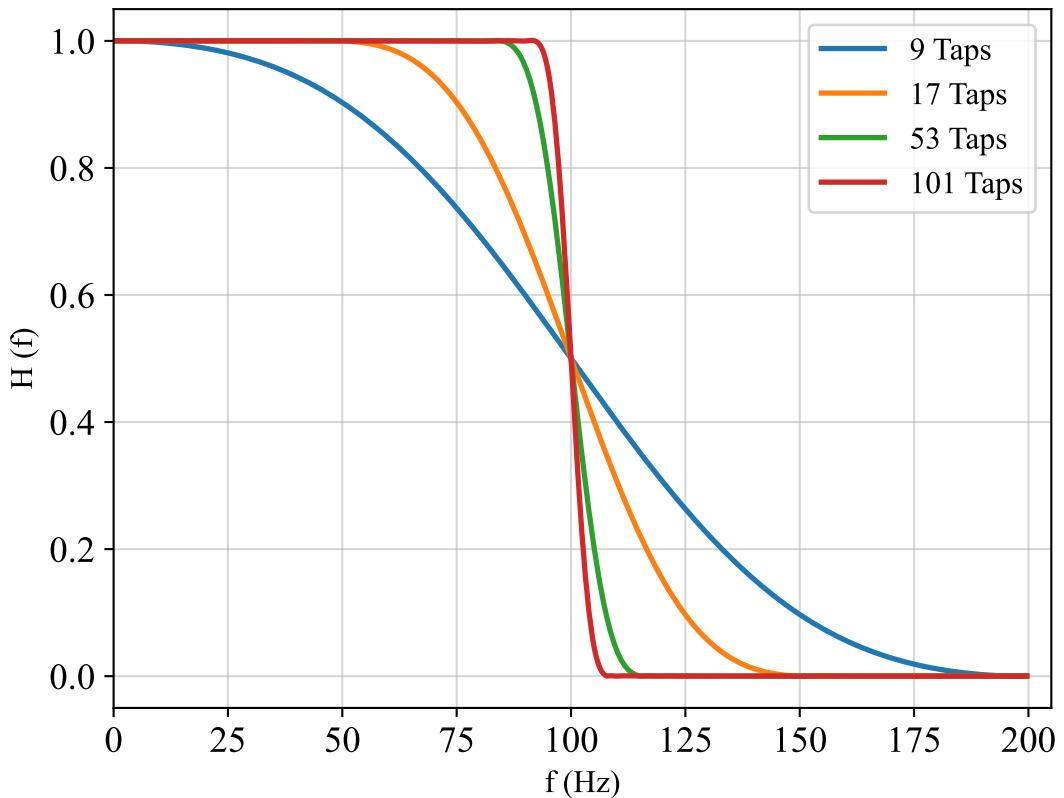


Figure 2.8: The frequency response of a 100 Hz LPF for a filter with a varying number of taps.

The simplest window function is a rectangular function performing a sharp cutoff however other types of windowing functions, such as the Kaiser window, are also used [69].

Mathematically, an equally spaced sampled Kaiser window function has the form

$$w(n) = \frac{I_0\left(\beta \sqrt{1 - \left(\frac{2n}{N} - 1\right)^2}\right)}{I_0(\beta)}, \quad 0 \leq n \leq N \quad (2.17)$$

where $N + 1$ is the total number of samples in the window and $I_0(x)$ is the zeroth-order modified Bessel function of the first kind [70]. The parameter β controls the central lobe width and Kaiser windows with varying β and a window length of 100 samples are shown in figure 2.9. FIR filter design becomes a compromise between the number of taps, windowing function, sharpness of the cutoff and magnitude of the ripples in the filter response [66], [71].

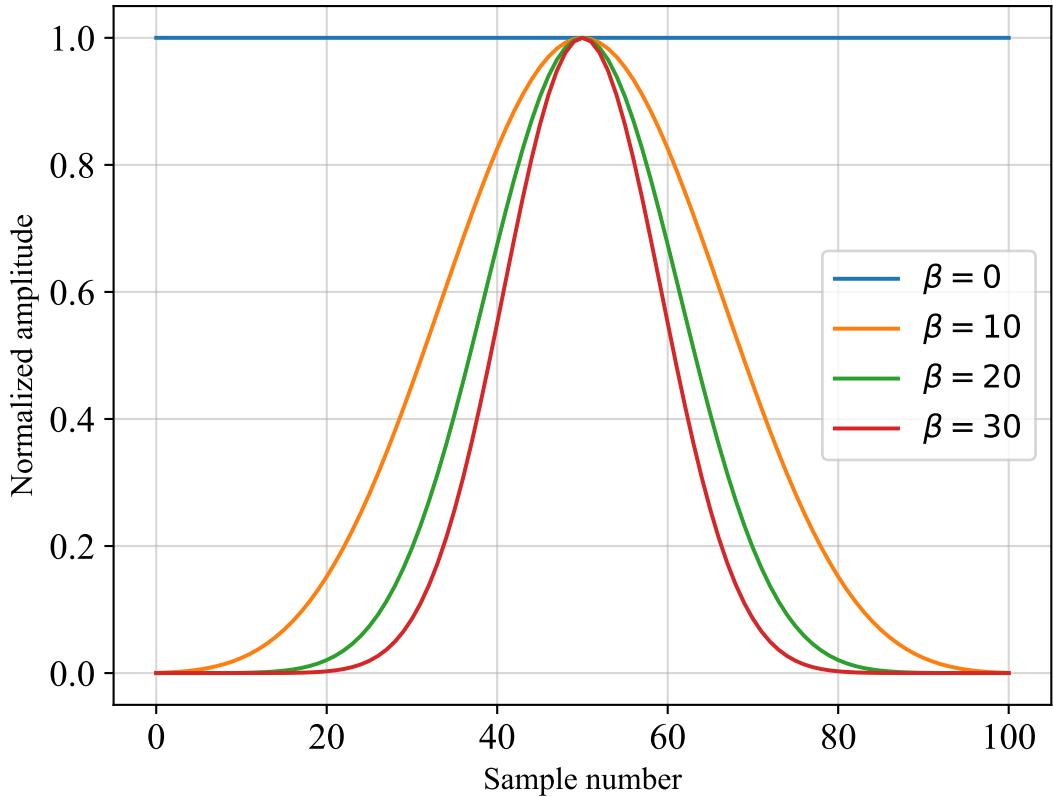


Figure 2.9: The Kaiser window function with varying values of β . $\beta = 0$ is a rectangular window.

The cut-off performed by the window function introduces distortions into the passband

and stopband of the filter response, called ripples, shown in figure 2.10. The magnitude of these ripples is generally a specified parameter when designing an FIR filter. The presence of the ripples is dependent on the type of windowing function used to generate the frequency response [69]. For example, a Chebyshev filter based on the Chebyshev polynomials, have a steep filter response at the cost of large passband and stopband ripples [66].

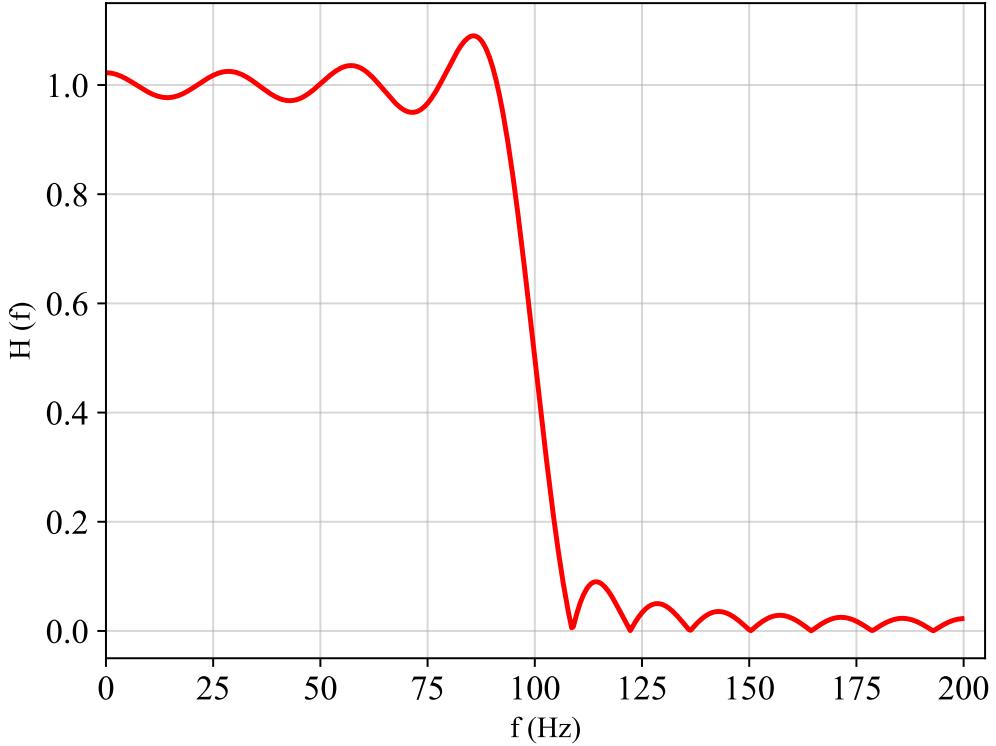


Figure 2.10: Passband and stopband ripples in the frequency response of a 100 Hz LPF for a filter with 29 taps.

A half-band FIR filter is a special case of FIR filter where the cutoff frequency, f_c , satisfies the following condition:

$$f_c = \frac{f_s}{4} \quad (2.18)$$

where f_s is the sampling rate of the signal. If equation (2.18) is satisfied, there are an odd number of coefficients of the filter which are symmetric about the centre tap and, excluding the centre tap, every other coefficient is equal to 0 [72]. Figure 2.11 shows the coefficients of a 17-tap half-band FIR filter calculated using the windowing method. Both of these properties lead to increase in computational efficiency of the filter [71], [72]. These are typically used as anti-aliasing filters when decimating a signal by a factor of two.

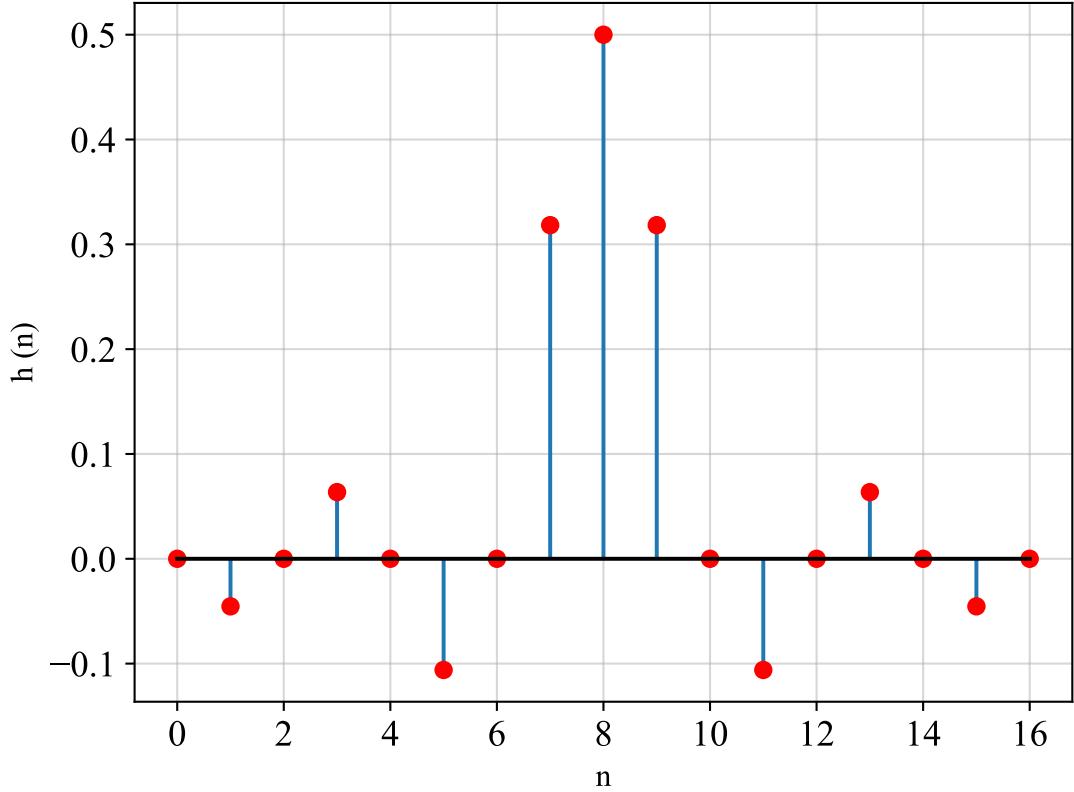


Figure 2.11: The coefficients for a 17-tap 100 Hz halfband filter.

2.1.5 Field Programmable Gate Arrays

A Field Programmable Gate Arrays (FPGA) is a chip that can be programmed by the user in order to implement custom digital logic [73]–[75]. The simplest FPGA architecture is shown in figure 2.12. It consists of an array of configurable logic blocks (CLBs), connected by a network of configurable interconnects and routing paths known as the fabric. Finally, there is a set of input/output (I/O) pads that allow for signals to be transferred to and from the internal FPGA logic to the external environment [76].

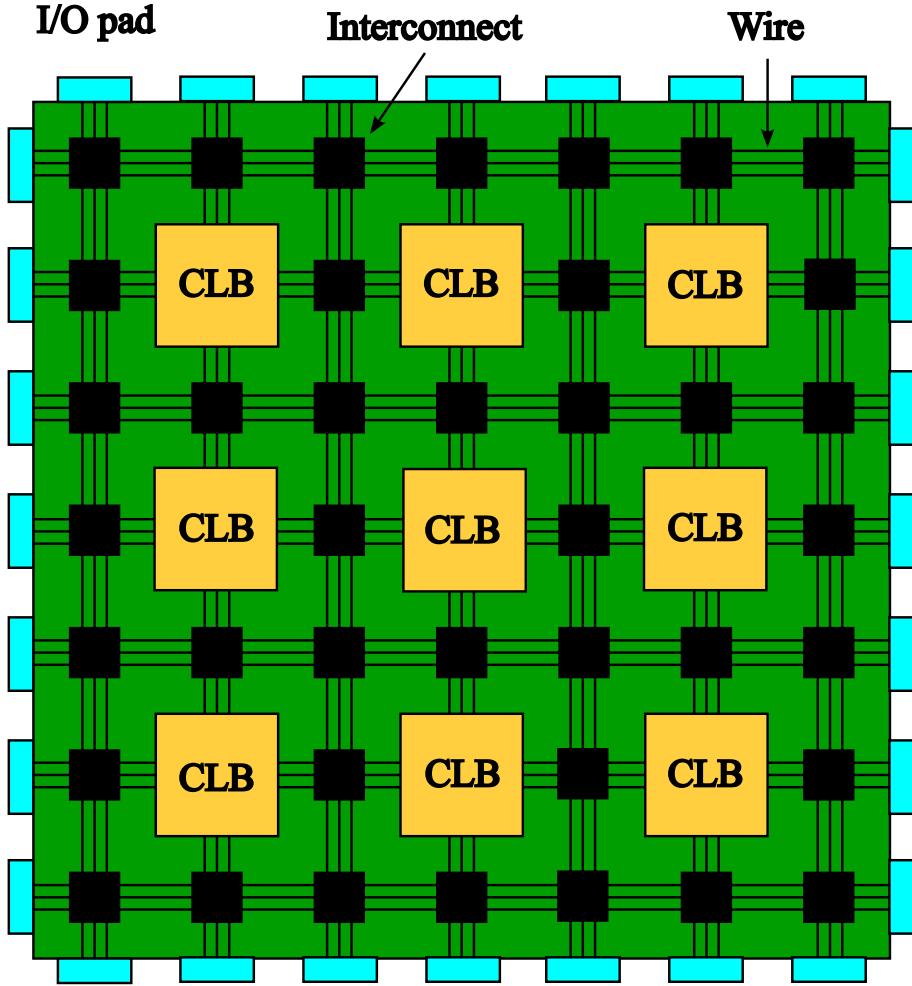


Figure 2.12: The architecture of a generic FPGA chip.

Modern FPGAs can contain specialised components known as FPGA primitives [76]. These include memory blocks (BRAM) and efficient DSP blocks. The range of primitives on a device are specific to the manufacturer and model of the FPGA. CLBs are combined to create user-defined digital logic blocks which are connected via the routing paths. This architecture allows for efficient parallelisation of digital logic as logical blocks act as black-boxes that can be duplicated across the device [77]. The logic is specified by the designer in a configuration file using a hardware description language (HDL) such as Verilog or VHDL [78], [79]. Additionally, a constraint file is provided to define physical characteristics of the design such as the system clock frequency, which cannot be specified in HDL. Developer

tools are then used to generate an FPGA bitstream file from the human-readable HDL files. The bitstream file provides a description of the configuration of CLBs, signal routing and initial values for any internal registers. The format of the bitstream file is dependent on the vendor of the chip. This file is then uploaded to the FPGA to configure the device.

There are a number of steps required to convert a HDL file into the bitstream file, namely synthesis, map and place and route (PAR) [76]. Synthesis converts the logic defined by the HDL files into a set of FPGA primitives that can realise the specified design. This step performs logic utilisation which counts the total number of each primitive used in the design. Finally, it specifies the connections between the components. The synthesis process outputs a ‘netlist’ file, which contains a list of all components and their connections within the design.

Map takes the netlist and converts it to a device-specific design [76]. This means it translates the components of the netlist into components available on a specified model of FPGA. Errors will occur here if the total number of a certain component on the netlist is larger than the amount available on the device.

PAR specifies the physical placement of each component in the netlist into the CLBs and other primitives on the FPGA board. After placement, the placed CLBs are routed using the programmable routing network and interconnects. The PAR procedure is designed to produce a design that is optimised according to the constraints of a design [76]. For example, a constraint on the clock period may cause all logic blocks requiring a clock signal to be located close to the clock source, to minimise the amount of time taken for the clock signal to reach the components. The PAR procedure then checks the routed design against all specified constraints, if a constraint is not met the design may be unreliable. If successful, the PAR outputs a final netlist which is converted into the bitstream file.

For large designs, the PAR process can take many hours due to the placement of

many components, so testing and validating the design before execution of PAR is essential. This is done by creating a HDL ‘testbench’. A large design is typically split across multiple HDL files known as modules. A testbench is normally created for each module and is used to test the module in isolation using a simulator. A set of inputs into the module under test are defined and the output signals (or any internal signals) of the module can be monitored and compared to the expected result. This process is similar to the unit testing process in software development.

One important concept when creating FPGA designs is timing closure [80]. Timing closure is when all timing constraints are satisfied after PAR. Within the FPGA, there are two main causes of delay for a signal: wire delay and gate delay. Wire delay is the time taken for the electrical signal to propagate around the chip, whereas gate delay is the inherent reaction time of a component. This delay leads to a phenomenon known as ‘clock-skew’ [80], [81] which is where the clock signal that drives all on-chip operations is delayed by the route it takes around the chip. A common cause of timing problems in large designs, is performing too many mathematical operations on a single clock cycle, which causes the maximum operating frequency of the design to decrease.

Pipelining is a common solution to this class of timing problems [82], [83]. Pipelining consists of adding extra registers to the logic to break up complex operations across multiple clock cycles. Figure 2.13 shows an example of a unpipelined and pipelined version of a piece of digital logic. In figure 2.13a the logic is unpipelined. Due to this, later operations in the chain cannot be executed until earlier operations have finished. For example, operation A_2 cannot be executed until A_1 has finished as it relies on the output from A_1 . This causes timing issues because operations A_2 and A_3 are reliant on A_1 so inputs d and e get delayed. This prevents any new inputs being inserted into the logic until A_2 and A_3 have been executed. If all operations take one clock cycle, the output y will only change every three cycles.

Figure 2.13b shows the pipelined version of this logic using a series of clocked registers. On every clock cycle, each input is transferred to the neighbouring column of registers. For example, on the first clock cycle inputs d and e are shifted along one register and the output of operation A_1 is deposited into its output register. Finally, a new set of inputs are inserted into the four leftmost registers. This implementation results in the output y changing every clock cycle because every set of inputs are synchronised between the columns as they traverse the logic so no input is delayed. Therefore, the operating frequency is three times higher than the unpipelined logic.

Pipelining adds latency to the logic. Latency is the delay between the first input and the first valid output of the logic [81]. For a pipeline containing L registers, there will be a latency of L clock cycles. This is caused as the pipeline registers are initially empty. After L clock cycles, there will be a new output on every clock cycle. In figure 2.13b the logic has a latency of three clock cycles as the longest path through the logic contains three registers.

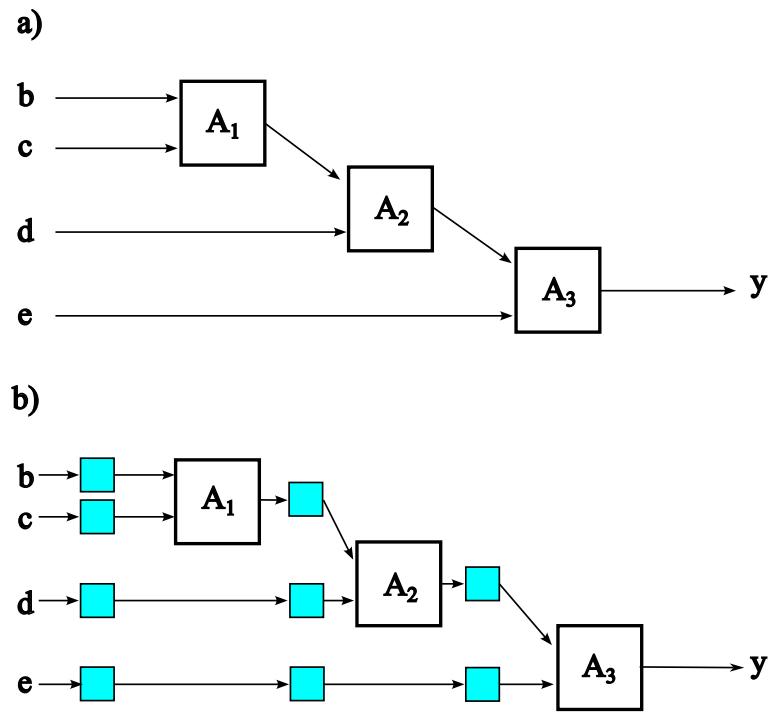


Figure 2.13: a) An unpipelined section of digital logic that performs operations A_1 , A_2 , A_3 on a set of inputs b , c , d and e , producing the output y . b) A pipelined version of a), with clocked registers indicated by the blue squares, placed on the output of all operations.

2.2 Nuclear Magnetic Resonance

This section will introduce the theory behind nuclear magnetic resonance (NMR). It begins with an explanation of the interaction between a magnetic moment and an applied magnetic field. Next, it will introduce the origins of the NMR signal and a set of common NMR experiments. Then, there will be an overview of the functionality of the NMR spectrometer and electronics used in an NMR experiment. Finally, the use of NMR in condensed matter physics will be discussed.

2.2.1 Introduction to nuclear magnetic resonance

Nuclear magnetic resonance (NMR) occurs due to the interactions of nuclear spins with an externally applied magnetic field. The Hamiltonian for a nuclear spin in a magnetic field is known as the Zeeman Hamiltonian and is of the form

$$\mathcal{H}_{zee} = -\boldsymbol{\mu} \cdot \mathbf{B} \quad (2.19)$$

where $\boldsymbol{\mu}$ is the magnetic moment of the nuclear spin and \mathbf{B} is the applied magnetic field [84].

The most common NMR schemes consist of a unidirectional applied magnetic field. Consequently, the z-direction is defined along the direction of the magnetic field and $\mathbf{B} = (0, 0, B_0)$ where B_0 is the magnitude of the applied field. The magnetic moment of a nuclear spin is related to its spin operator, \mathbf{I} , by

$$\boldsymbol{\mu} = \gamma_n \mathbf{I} \quad (2.20)$$

where γ_n is the gyromagnetic ratio of the species of nuclear spin [84]. The gyromagnetic ratio is a constant for a given species of nucleus. Another consequence of equation (2.20) is that $I = 0$ spin nuclei have no magnetic moment and are not NMR-active.

Substitution into equation (2.19) results in

$$\mathcal{H}_Z = \hbar \gamma_n B_0 I_z = \hbar \omega_0 I_z \quad (2.21)$$

where $\omega_0 = \gamma_n B_0$ is called the Larmor frequency. The application of a magnetic field to a single $I = 1/2$ nucleus yields a pair of energy levels with energies

$$E_{\pm} = \pm \frac{1}{2} \hbar \omega_0 \quad (2.22)$$

These levels represent parallel and anti-parallel alignment of the nuclear spin with the magnetic field. The energy splitting between the levels is given by

$$\Delta E = E_+ - E_- = \hbar \omega_0 \quad (2.23)$$

Figure 2.14a shows an energy level diagram of a $I = 1/2$ nucleus in a magnetic field. Therefore, the $I = 1/2$ system has a single resonance line located at the Larmor frequency. The spectrum of this system is shown in 2.14b. In the general case of a spin I nucleus in a magnetic field, the spectrum will contain $2I$ resonance lines [85]. These lines originate from quadrupolar coupling if $I > 1/2$.

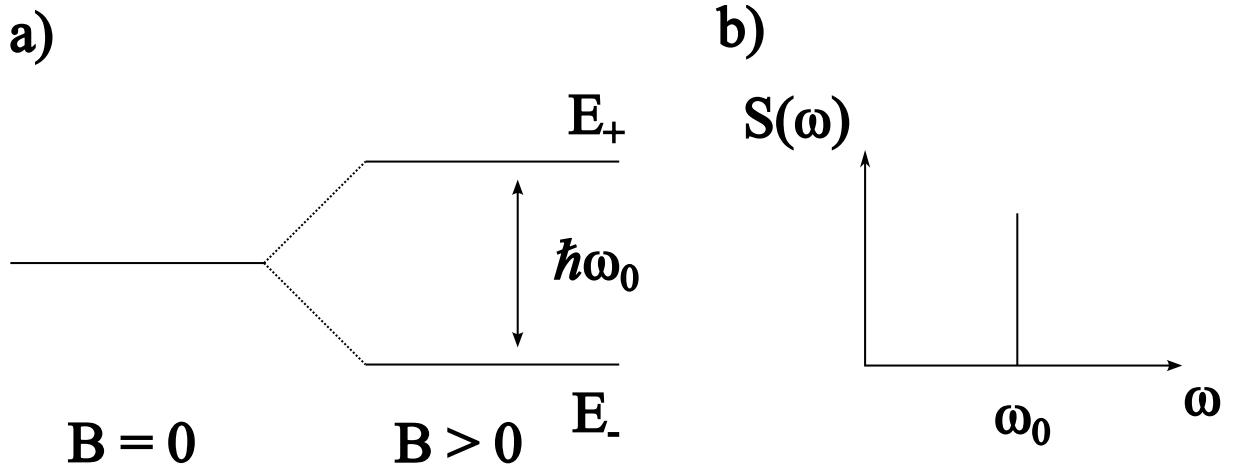


Figure 2.14: a) The energy levels of an $I = 1/2$ spin in a magnetic field. Splitting is observed for a non-zero applied field. b) The spectrum of an $I = 1/2$ spin in a magnetic field

The energy gap causes the spin to transition between energy levels from the ground state to the excited state if it absorbs a photon with energy equal to the gap, E_{\pm} . Alternately, the system can ‘relax’ to the lower energy state by emitting a photon with energy E_{\pm} . This mechanism allows the system to be driven by the application of a radiofrequency (RF) pulse with a frequency equal to the Larmor frequency. Therefore, the application of the RF pulses forms a key component of the NMR experiment.

In reality, NMR is performed on systems with $N \sim 10^{23}$ spins so a density matrix approach is typically used to mathematically describe the system [85]. This approach allows for ideas from statistical mechanics to be easily incorporated into the description. At

equilibrium, the density matrix of a system is given by

$$\rho_{eq} = \frac{1}{Z} e^{-\beta \mathcal{H}} \quad (2.24)$$

where Z is the partition function and $\beta = \frac{1}{k_B T}$, k_B is Boltzmann's constant and T is the temperature [85].

The dynamical evolution of the density matrix is described by the quantum master equation [86]

$$\frac{\partial}{\partial t} \rho(t) = -\frac{i}{\hbar} [\mathcal{H}, \rho(t)] \quad (2.25)$$

where \mathcal{H} is the interaction Hamiltonian of the system. This Hamiltonian is constructed from terms that represent different interaction pathways within the system, of which the Zeeman interaction from equation (2.19) is an example. Additional terms include the quadrupole interaction for nuclei with $I > 1/2$ and the hyperfine interaction between nuclear and electron spins. These terms will be discussed in a section 2.2.5. Furthermore, equation (2.25) can be expanded with additional terms known as relaxation superoperators in order to model relaxation within the system [87]. The time evolution of experimental observables, including the magnetisation can be calculated from the trace of ρ_{eq} using

$$\langle M_\alpha(t) \rangle = \gamma_n \hbar \operatorname{Tr}[\rho(t) I_\alpha] \quad (2.26)$$

where $\alpha = x, y, z$ represents the component of magnetisation or spin vector in the α direction and Tr is the trace operation. [86]

The time evolution of the magnetisation vector of the ensemble of nuclear spins throughout an NMR experiment is visualised using the vector model. This vector describes the direction of the average magnetisation of the system. It is defined by $\mathbf{M} = (\langle M_x \rangle, \langle M_y \rangle, \langle M_z \rangle)$ where the component in each direction is calculated using equation (2.26). Plotting the vector on a set of Cartesian axes produces a geometrical representation of the state of the system known as the Bloch sphere [88].

The Bloch sphere is more generally used to represent any two-level quantum system, such as a qubit in quantum computing. Figure 2.15 shows a diagram of a Bloch sphere. The poles of the sphere are labelled with the eigenstates of the two-level system which form an orthonormal basis $\{|0\rangle, |1\rangle\}$. All possible pure quantum states of the system are found on the surface of the unit sphere. A generic state of the system, $|\psi\rangle$ can then be described as

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle \quad (2.27)$$

where θ and ϕ are the polar and azimuthal angle of the vector respectively [88].

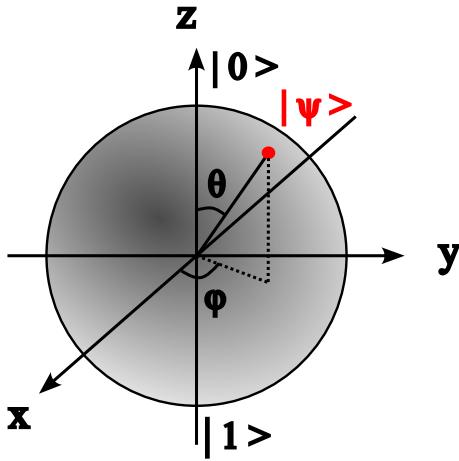


Figure 2.15: A generic Bloch sphere for a two-level system. $|\psi\rangle$ indicates a general state found on the surface of the sphere with polar angle θ and azimuthal angle ϕ .

In thermal equilibrium, the nuclear spins tend to align in the direction of the magnetic field, creating a new magnetisation in the z -direction (for a magnetic field applied along the z direction). The vector becomes $\mathbf{M} = (0, 0, M_0)$, where M_0 is the magnetisation at thermal equilibrium. This is calculated from ρ_{eq} at the system temperature. If \mathbf{M} is tilted away from the z -direction, it begins to precess around the static field, B_0 at the Larmor frequency. The NMR signal is formed by the measuring the oscillating transverse components of \mathbf{M} through magnetic induction in a nearby coil.

The application of an RF pulse along the x -direction on the system can be considered

by applying a perturbation to the Hamiltonian [85].

$$\mathcal{H}_{pulse} = \hbar\omega_0 I_z + \hbar\gamma_n I_x B_1 \cos(\omega_{TX} t) \quad (2.28)$$

where B_1 , ω_{TX} are the pulse amplitude and frequency respectively. Since $B_0 \gg B_1$, time-dependent perturbation theory can be used.

Switching reference frame from the lab frame to a rotating frame that rotates at frequency ω_{TX} removes the time dependence from equation (2.28). In the rotating frame, the Larmor frequency is transformed to [14]

$$\Omega = \omega_0 - \omega_{TX}. \quad (2.29)$$

The Hamiltonian in the rotating frame can then be found by substitution of Ω into equation (2.28)

$$\mathcal{H}_{rot,x} = \hbar\Omega I_z + \hbar\gamma_n B_1 I_x. \quad (2.30)$$

For an on-resonance RF pulse, $\omega_{TX} = \omega_0$ and the first term in equation (2.30) vanishes. Therefore, an equivalent equation for the application of an RF pulse along the y -axis can be obtained by substituting I_x for I_y . It can be shown by calculating the time evolution of ρ under \mathcal{H}_{rot} , that for an on-resonance pulse around the x -axis, the magnetisation vector rotates around the x -axis. The flip angle, β , between the z -axis and \mathbf{M} after the pulse is given by

$$\beta = \gamma_n B_1 t_p \quad (2.31)$$

where t_p is the duration of the RF pulse [14].

The most common set of NMR pulses are the $\pi/2$ and π pulse. These describe the cases where the magnetisation vector is rotated by $\beta = \pi/2$ and $\beta = \pi$ respectively. The $\pi/2$ pulse is used because it generates the largest NMR signal as it rotates \mathbf{M} into the $x - y$ plane. Conversely, the π pulse inverts \mathbf{M} along the z -axis, producing no NMR signal and is typically used in more complex NMR experiments which will be discussed in section 2.2.2.

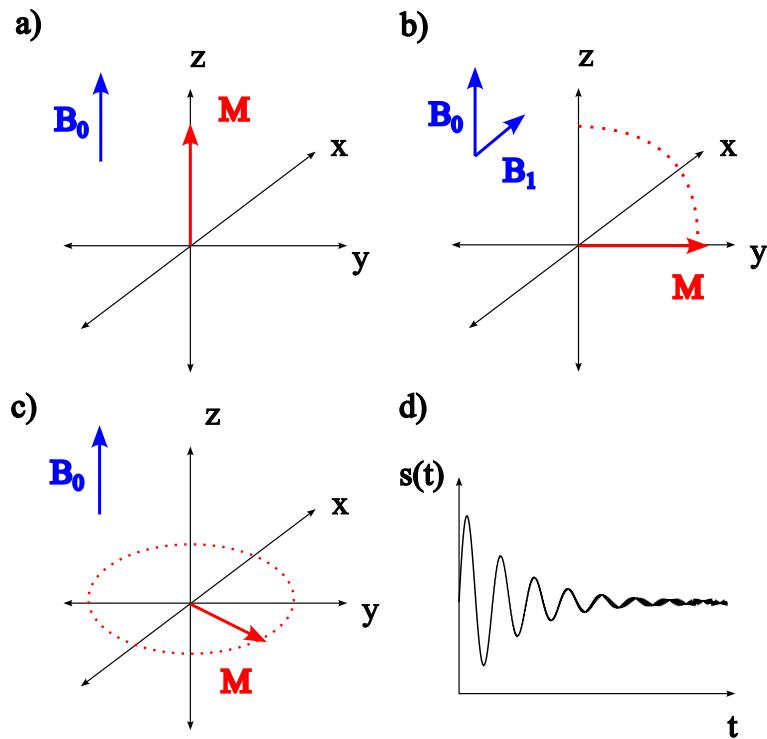


Figure 2.16: a) The magnetisation vector at thermal equilibrium. b) The magnetisation vector rotates into the x - y plane due to the application of an RF pulse, B_1 . c) The precession of the magnetisation vector in the x - y plane around the applied field B_0 d) The NMR signal generated by the transverse components of the precessing magnetisation vector.

Figure 2.16 demonstrates the most basic type of NMR experiment: the pulse-acquire sequence. The pulse-acquire sequence consists of the application of a single $\pi/2$ RF pulse followed by acquisition of the NMR signal. Figure 2.16a shows the system at thermal equilibrium. Figure 2.16b indicates the rotation of \mathbf{M} into $x - y$ plane due to the $\pi/2$ pulse. Figure 2.16c displays the precession of \mathbf{M} in the $x - y$ plane, generating the NMR signal. Figure 2.16d shows an example NMR signal obtained from a pulse-acquire experiment.

After the excitation by an RF pulse, the system will eventually return to its equilibrium state, through a process known as ‘relaxation’. This is shown in figure 2.16d by the decay of the NMR signal amplitude over time. Relaxation is governed by the Bloch equations

[89]

$$\frac{dM_x(t)}{dt} = \gamma_n (\mathbf{M}(t) \times \mathbf{B}(t))_x - \frac{M_x(t)}{T_2} \quad (2.32)$$

$$\frac{dM_y(t)}{dt} = \gamma_n (\mathbf{M}(t) \times \mathbf{B}(t))_y - \frac{M_y(t)}{T_2} \quad (2.33)$$

$$\frac{dM_z(t)}{dt} = \gamma_n (\mathbf{M}(t) \times \mathbf{B}(t))_z - \frac{M_z(t) - M_0}{T_1} \quad (2.34)$$

where \mathbf{M} is the usual magnetisation vector with components M_x , M_y and M_z in the x, y and z -directions respectively. M_0 is the thermal equilibrium magnetisation. T_1 and T_2 are the relaxation time constants for the longitudinal and transverse components of \mathbf{M} .

If an RF pulse rotates \mathbf{M} into the $x - y$ plane, it begins to relax back to the z -axis. This process is the longitudinal, or spin-lattice relaxation [90] characterized by T_1 in equation (2.34). The solution to this equation is

$$M_z(t) = M_0(1 - e^{-\frac{t}{T_1}}) \quad (2.35)$$

where t is the elapsed time from the RF pulse. Consequently, for large t , $M_z = M_0$ as previously described. Longitudinal relaxation can be driven by dipolar interactions within the sample. Dipolar interactions are interactions between nuclei in the sample. Hence, T_1 can be used as a probe for the local magnetic environment of the NMR-active nucleus [4].

The time constant T_2 characterizes how long the transverse, or spin-spin relaxation process takes. This refers to the decay of the M_x and M_y components of the NMR signal after the application of an RF pulse. T_2 relaxation processes can be decomposed into components. One component is the spatial magnetic field inhomogeneity [91] and is denoted by the time constant T_2^* . Nuclear spins within the sample are located at different points in space and experience a small variation in the static magnetic field. Then, by equation (2.21), each spin with precess at slightly different frequencies. This causes the individual spin magnetisation vectors to fan out over time and destroy the average magnetisation vector as shown in figure 2.17a and 2.17b. The second component is the interactions between spins in the sample. This mechanism allows for neighbouring spins to swap their spin states through a ‘spin diffusion’ process [92].

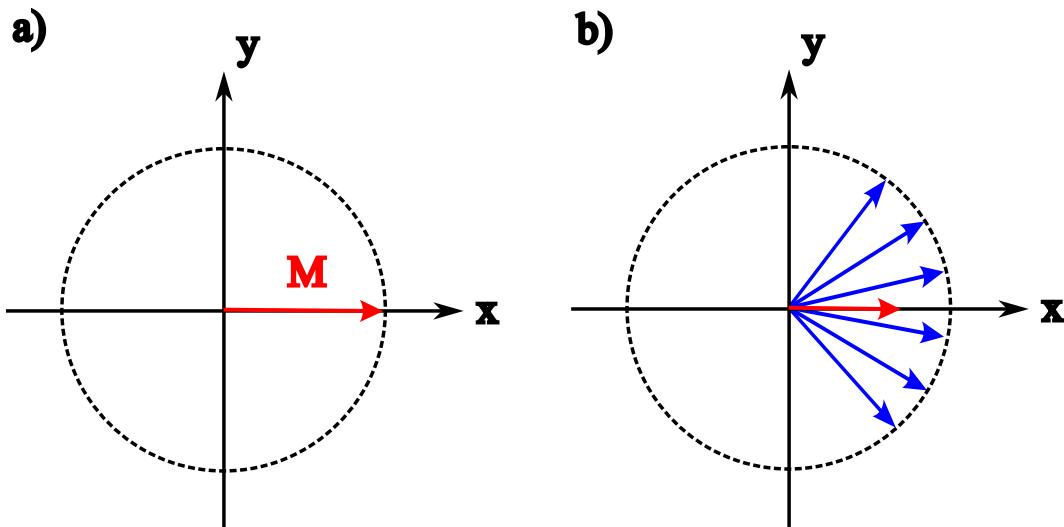


Figure 2.17: a) The average magnetisation vector immediately after being rotated in the x - y plane by an RF pulse. All spins are aligned along the x -axis. b) The average magnetisation vector after experiencing dephasing. The blue arrows indicate the individual spin magnetisation vectors ‘fanning out’, due to precession at a range of Larmor frequencies.

2.2.2 NMR pulse sequences

There are a wide range of NMR experiments that can be performed by varying the duration and number of RF pulses applied to the system. Each pulse sequence manipulates the magnetisation vector in a different way, allowing for various NMR quantities to be measured. This section will introduce the two main NMR sequences used in condensed matter experiments: the spin-echo and saturation recovery sequence.

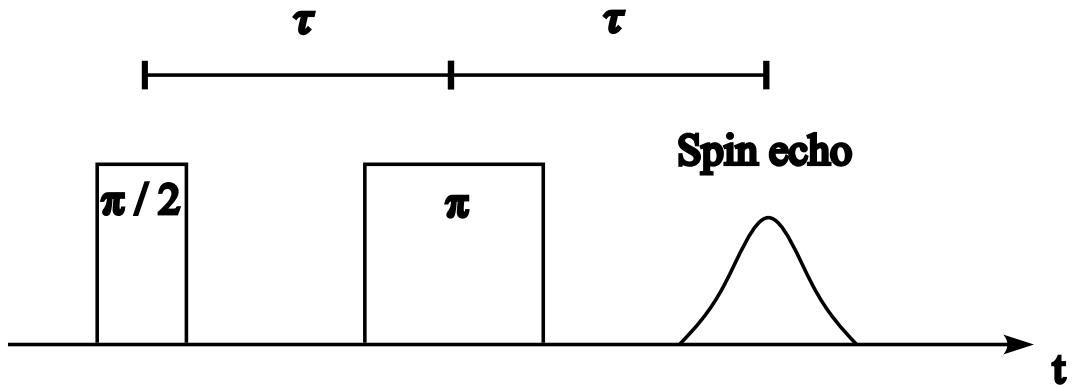


Figure 2.18: The spin echo pulse sequence. It consists of a $\pi/2$ pulse around the x -axis followed by a π pulse around the x -axis, with a spacing of τ . This causes the appearance of a spin echo at time τ after the π pulse.

The spin echo sequence was first used in 1950 by E. L. Hahn [93] to measure a spin echo signal. The spin echo sequence is shown in figure 2.18. A schematic diagram showing the magnetisation vector is shown in figure 2.19. The sequence consists of two RF pulses, a $\pi/2$ pulse followed by a π pulse, with a separation of τ . Initially, the magnetisation vector is in thermal equilibrium along the z -direction (figure 2.19a). The $\pi/2$ pulse rotates the magnetisation vector into the $x - y$ plane (figure 2.19b). Between the pulses, the dephasing of the individual spin vectors due to T_2^* occurs (figure 2.19c). The π pulse inverts the magnetisation vectors within the $x - y$ plane (2.19d). This causes an exact reversal of the component of dephasing caused by magnetic field inhomogeneity. The ‘fanned-out’ spins recombine to restore the magnetisation vector. Therefore, a spin echo signal is detected at

a duration τ after the π pulse. The spin echo decay can therefore be used to measure T_2^* and T_2 . For condensed matter NMR, samples typically have a small T_1 time constant, so measuring the spin echo signals is preferred over a pulse-acquire signal. This is because the rapid amplitude decay of the pulse-acquire signal causes the signal to be hidden within the ‘leakthrough’ signal from the high-power RF pulse. This is discussed further in section 2.2.4.

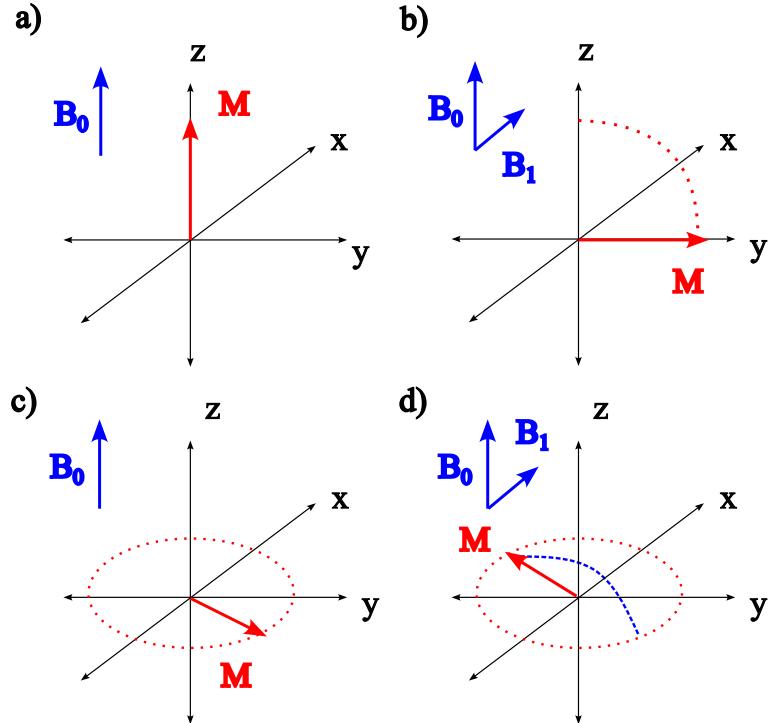


Figure 2.19: a) The magnetisation vector at thermal equilibrium. b) The magnetisation vector rotates into the x-y plane due to the application of an RF pulse, B_1 . c) The precession of the magnetisation vector in the x-y plane around the applied field B_0 . The individual spin vectors ‘fan-out’ as in figure 2.17b. d) The π pulse inverts the magnetisation in the x - y plane producing the spin echo. The reversal of the ‘fanning-out’ occurs here. The ‘fanning-out’ of the spins due to dephasing in c) and d) is not shown for clarity.

The spin echo sequence can be extended by the application of an arbitrary amount of π pulses in a sequence known as a Carr-Purcell-Meiboom (CPMG) sequence [94], [95]. This sequence is used for efficiently calculating T_2 as demonstrated in figure 2.20. Each pulse is

applied around the x -axis to repeatedly invert the magnetisation vector in the x - y plane. Each individual spin echo follows T_2^* decay whilst the decay envelope of the sequence of echoes follows the T_2 decay.

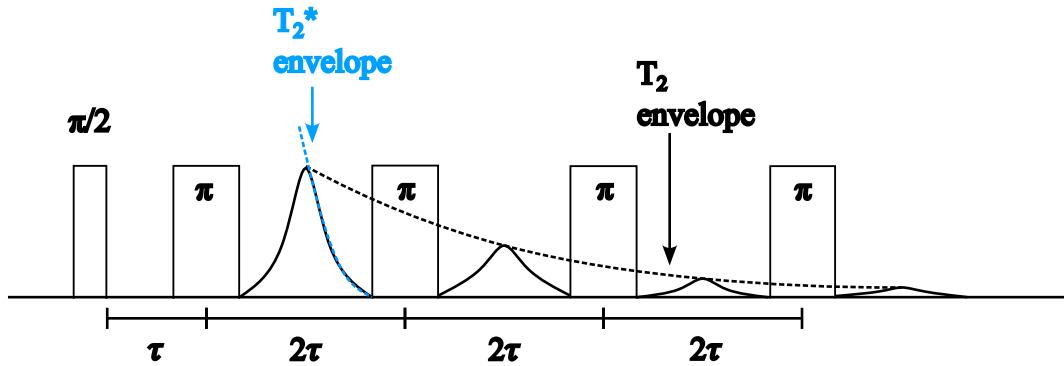


Figure 2.20: The CPMG pulse sequence consisting of a $\pi/2$ pulse followed by a sequence of π pulses generating successive spin echoes. Each echo follows a T_2^* decay envelope and the series of echoes follow a T_2 decay envelope.

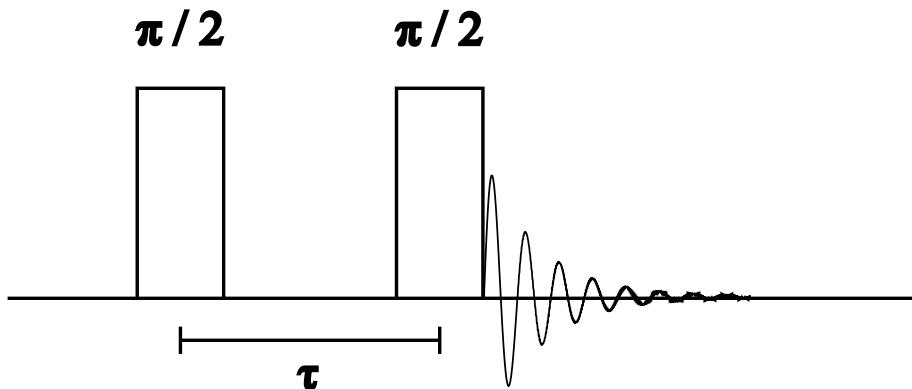


Figure 2.21: The saturation recovery sequence used to measure T_1 . It consists of a pair of $\pi/2$ pulses around the x -axis separated by a time period τ . The signal amplitude follows an exponential recovery curve with time constant T_1 .

The saturation recovery pulse sequence is a method to measure T_1 . The sequence is shown in figure 2.21 and consists of a pair of $\pi/2$ pulses separated by a duration τ [96]. Initially, the magnetisation vector is at thermal equilibrium and points along the z direction. (figure 2.22a) The first pulse rotates the magnetisation vector in the $x - y$ plane (figure

2.22b). Between pulses, the magnetisation begins to relax towards the z -direction. The second pulse then rotates the magnetisation again and the signal is measured. If $\tau \ll T_1$, the magnetisation still be in the $x - y$ plane so the second pulse will rotate it to the $-z$ -direction (figure 2.22c). The x, y components of the magnetisation are equal to zero when the magnetisation is parallel to $-z$ hence there will be no signal. If $\tau \gg T_1$, the magnetisation will have fully relaxed between the pulses and the signal will have maximal amplitude (figure 2.22d). Therefore, measuring the signal amplitude as a function of τ yields a recovery curve that follows equation (2.35). Furthermore, an additional π pulse can be used to generate a spin echo signal that follows the same decay as the original signal.

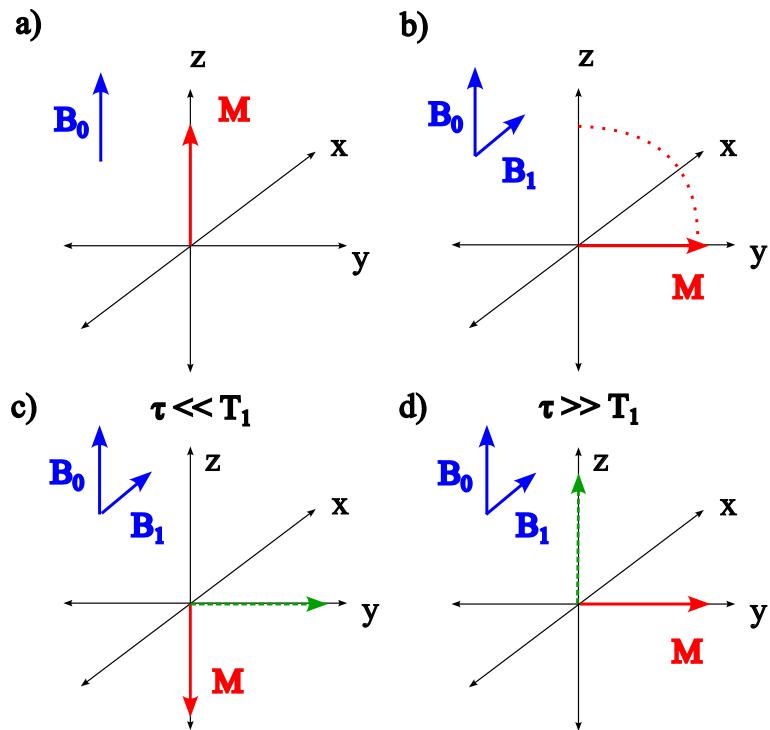


Figure 2.22: a) The magnetisation vector at thermal equilibrium. b) The magnetisation vector rotates into the x-y plane due to the application of an RF pulse, B_1 . c) The magnetisation vector after the second $\pi/2$ pulse and $\tau \ll T_1$. d) The magnetisation vector after the second $\pi/2$ pulse and $\tau \gg T_1$. In c) and d) the green arrow represents the magnetisation before the second $\pi/2$ pulse. In d) the magnetisation vector has fully relaxed back to the z -direction.

2.2.3 The NMR spectrometer

The NMR spectrometer has two main functions: to generate the RF pulses to manipulate M around the Bloch sphere to cause Larmor precession and to process the FID. The RF pulses are normally passed through a high-power amplifier in order to generate high power pulses. The amplifier output is controlled by a pulse programmer module within the spectrometer allowing for variable pulse duration. The amplified pulses are then sent to a resonance circuit tuned to the Larmor frequency on the NMR probe containing the sample. The circuit then

broadcasts the RF pulse to the sample and detects the induced NMR signal.

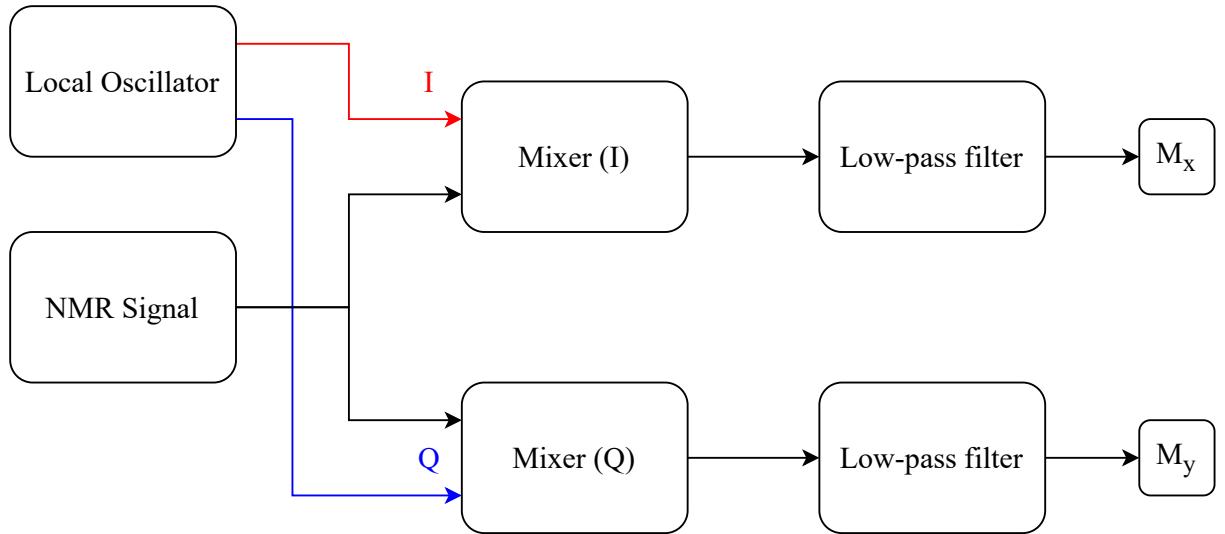


Figure 2.23: The receiver path within the NMR spectrometer. The local oscillator provides a pair of IQ signals that are mixed with the NMR signal. A low pass filter is used to remove the sum frequency component of the mixed signal, leaving only the difference frequency. The resulting signals form the transverse components of the magnetisation.

The spectrometer performs a set of operations on the acquired NMR signal shown in figure 2.23. The NMR signal originates from the transverse components of \mathbf{M} . This precession induces a voltage within the resonance circuit which is then amplified and transmitted to the spectrometer. The signal is then ‘mixed down’ to baseband using a quadrature mixer. The mixer is required in order to lower the frequency of the signal to a range that conventional electronic components can handle. This is required as the frequency range for condensed matter NMR experiments is of the order 100 MHz [13]. However, an NMR spectrum normally spans a narrow range of frequencies so it can be shifted down to baseband without losing information. The mixer multiplies the NMR signal with set of quadrature signals generated from the transmitter known as the local oscillator (LO) signals or IQ signals. I denotes the in-phase signal represented by $A \cos(\omega_{rx}t)$ and Q denotes the quadrature-phase

signal given by $A \sin(\omega_{rx}t)$. If the NMR signal with a Larmor frequency of ω_0 is modelled by $A \cos(\omega_0 t)$, the mixed signals are given by [14]

$$A \cos(\omega_0 t) \cos(\omega_{rx} t) = \frac{1}{2} A [\cos([\omega_0 + \omega_{rx}]t) + \cos([\omega_0 - \omega_{rx}]t)] \quad (2.36)$$

$$A \cos(\omega_0 t) \sin(\omega_{rx} t) = \frac{1}{2} A [\sin([\omega_0 + \omega_{rx}]t) - \sin([\omega_0 - \omega_{rx}]t)] \quad (2.37)$$

The pair of mixed signals contain two frequency components at the sum frequency, $\omega_0 + \omega_{rx}$, and difference frequency $\omega_0 - \omega_{rx}$. If the spectrometer frequency is matched to the Larmor frequency, $\omega_0 - \omega_{rx}$ is located at baseband. The sum frequency component is removed using a low pass filter. Additionally, when matched equation (2.36) and equation (2.37) are the x and y components of the magnetisation in the rotating frame. These components are then combined to form a complex signal, the FID, given by

$$S(t) = M_x(t) + iM_y(t) \quad (2.38)$$

The spectrum is obtained by applying a Fourier Transform to the FID, converting it to the frequency domain. The peaks in the spectrum represent the resonant frequencies of the system. The spectrum is a function of the offset frequency, Ω because the spectrometer measures the signal in the rotating frame. For example, in the simplest case of a single resonance line at the Larmor frequency and a well-matched receiver, the spectrum shows a zero frequency peak [14]. There are a number of additional processing stages that can be used to further process the FID and spectrum. Phase correction is a common technique used to correct artefacts in the spectrum [14]. Zeroth order phase correction is frequency independent and consists of multiplying the FID with a complex phase factor. This correction accounts for the arbitrary location from which the transverse magnetisation begins its precession in the $x - y$ plane. Additionally, frequency dependent phase factors can also be applied to correct distortions within the baseline of the spectrum [14].

2.2.4 NMR electronics

The RF pulse required to generate an NMR signal travels through a transcoupler [97]. The transcoupler acts as a splitter between the transmitted RF pulse source, the NMR probe containing the sample and the receiver path for the induced signal. A diagram of a typical NMR transcoupler is shown in figure 2.24. The transcoupler consists of two electrical components: crossed diodes and a quarter wave (or $\lambda/4$) transformer.

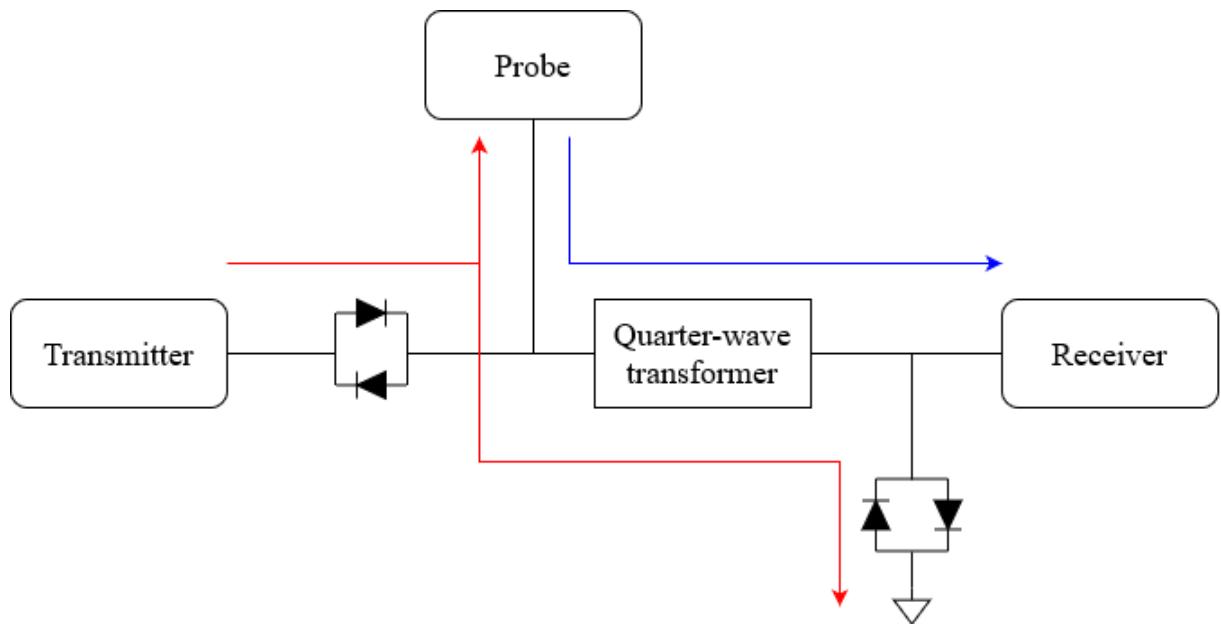


Figure 2.24: A diagram of the components of a typical NMR transcoupler. The red arrows denote the path of the high-power RF pulse. The blue arrow indicates the path of the low-power NMR signal.

The crossed diodes are formed by connecting two diodes in parallel. Current will only flow through a cross diode when the magnitude of the voltage exceeds the voltage threshold, V_t . Consequently, crossed diodes block low voltage signals and only activate for high voltage signals [98]. The I - V curve for a crossed diode is shown in figure 2.25. The first set of crossed diodes blocks the low-power NMR signal from reflecting into the transmitter. The second set of crossed diodes diverts the high-power RF signal away from the sensitive receiver electronics, to ground. Diodes with a short recovery time are typically used in order to support short RF pulses [97]. In reality, the receiver will not be perfectly shielded from the RF pulse which causes a leakthrough signal to be detected in the NMR output signal. The leakthrough signal will block the observation of any NMR signals until its amplitude has decayed away.

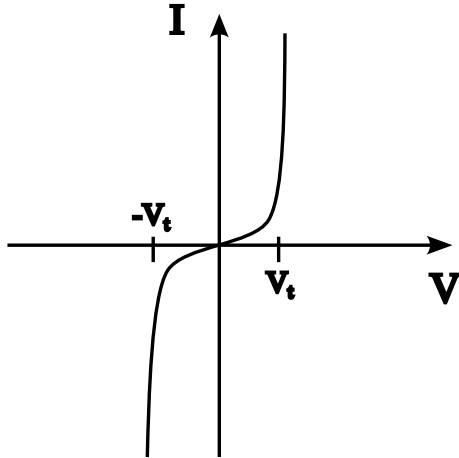


Figure 2.25: The I - V curve for a crossed diode. Current can only flow when the magnitude of the voltage exceeds the threshold voltage, V_t .

The quarter wave transformer is used to match a load impedance to the characteristic impedance to prevent reflections within the transmission line [99]. If a quarter wave transformer is connected to an open load, the transmitted signal is fully reflected back to the source and the load acts as a short circuit as the source detects the superposition of outgoing and reflected signals [99]. The case where the frequency of the signal causes the

outgoing and reflected signals to fully cancel out at the source (ignoring cable losses), the transformer is said to be a quarter wavelength long for that frequency. A diagram of this behaviour is shown in figure 2.26 Conversely, if the quarter wave transformer is connected to a shorted load, the transmitted and reflected signals are in phase and the source detects an open circuit. Therefore, the quarter wave transformer acts like an RF bandpass filter. Frequencies that correspond to the quarter wavelength can pass through, whereas other frequencies are attenuated. The simplest quarter wave transformers are constructed from a length of coaxial cable, where the length of the cable is related to the target frequency via the following equation

$$l = \frac{Vc}{4f} \quad (2.39)$$

where V is the velocity factor of the cable, c is the speed of light and f is the frequency of which a cable of length l is a quarter wave transformer [97]. The quarter wave transformer is used to filter out noisy signals located at frequencies not equal to the NMR frequency. Consequently, the quarter wave transformer must be tuned to the NMR frequency being used for a given NMR experiment.

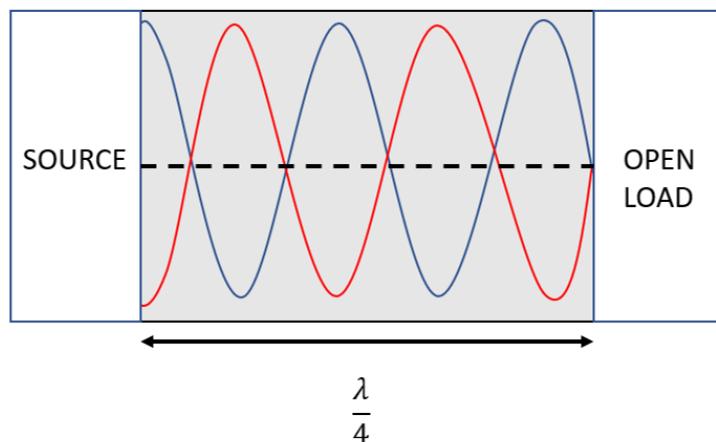


Figure 2.26: A diagram of a quarter wave transformer connected to an open load. The blue line and red line are the transmitted and reflected signals. The black dashed line is the signal measured at the source.

The NMR signal is detected using a resonance circuit located within the NMR probe. The simplest resonance circuit is the parallel LC circuit shown in figure It consists of coil with inductance L which is wrapped around the sample and simultaneously transmits the RF signal and detects the NMR signal and a pair of tuneable capacitors [97]. C_m is the matching capacitor used to impedance match the parallel components to the rest of the circuitry. A well-matched resonance circuit results in maximal signal amplitude within the system. The tuning capacitor, C_t , is used to adjust the resonant frequency of the LC circuit.

The resonant frequency of the circuit can be found by equating the reactance of each component. The reactance of the coil and tuning capacitor is given by $X_L = 2\pi fL$ and $X_{ct} = (2\pi fC_t)^{-1}$ respectively, where f is the frequency [100]. At the resonant frequency, $f = f_0$ and $X_L = X_{ct}$ yielding

$$f_0 = \frac{1}{\sqrt{2\pi LC_t}} \quad (2.40)$$

In real circuits, there is always a small resistance from the cable and coil windings present in the circuit. The resistance causes damping of the resonance and is quantified using the Q factor calculated as

$$Q = \frac{2\pi f L}{R} \quad (2.41)$$

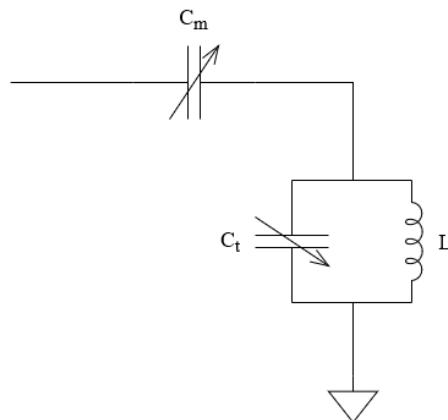


Figure 2.27: The circuit diagram of the LC circuit used to detect the NMR signal.

2.2.5 NMR in condensed matter

Solid state NMR experiments commonly probe nuclear spins with $I > 1/2$. These nuclei have a non-zero electric quadrupole moment that interacts with the local electronic environment generated by electrons within the compound. This causes a quadrupolar term to be added to the total Hamiltonian of the system of the form [85].

$$\mathcal{H}_Q = \frac{e^2 q Q}{4I(2I-1)} \left[(3I_z^2 - I^2) + \eta(I_x^2 - I_y^2) \right] \quad (2.42)$$

Q is the quadrupolar moment of the nucleus, η is the asymmetry parameter defined as $\eta = \frac{V_{xx} - V_{yy}}{V_{zz}}$ and $eq = V_{zz}$. $V_{\alpha\beta}$ are the components of the electric field gradient (EFG) tensor defined as $V_{\alpha\beta} = \frac{\partial^2 V(0)}{\partial x_\alpha \partial x_\beta}$ where $V(0)$ is the electrostatic potential generated by the electronic environment at the position of the nuclear spin. Equation (2.42) is only valid in the basis in which the EFG tensor is diagonal and V_{xx}, V_{yy}, V_{zz} are the principal components of the EFG. Conventionally, the z -direction is chosen such that $|V_{zz}| > |V_{xx}| > |V_{yy}|$.

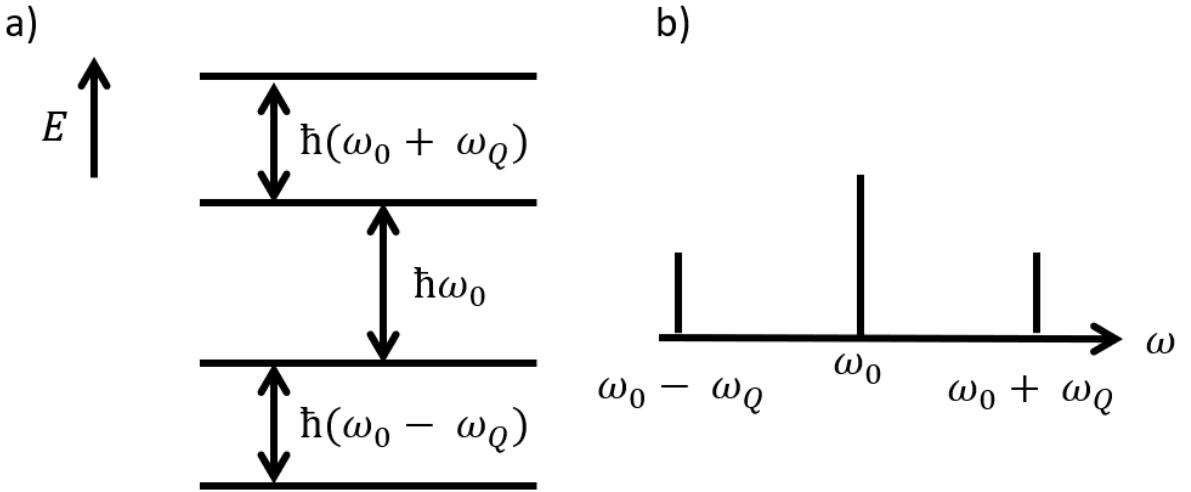


Figure 2.28: a) The energy level diagram for a $I = 3/2$ spin in an external magnetic field with a quadrupole interaction as described in the text. b) The NMR spectrum for a $I = 3/2$ spin. There are three resonance peaks located at the frequencies given by equation (2.43).

The presence of the quadrupolar interaction within the Hamiltonian complicates the

NMR spectrum. The energy levels of the system have an angular dependence on the angles between the applied field, B_0 , and the principal directions of the EFG tensor. The simplest case assumes an axially symmetric EFG with $\eta = 0$, $I = 3/2$ and B_0 aligned with a principle direction of the EFG [101]. This results in four energy levels with energy separation:

$$\Delta E_n = \hbar\omega_0 \pm n\hbar\omega_Q \quad (2.43)$$

where ω_Q is the quadrupole frequency and $n = -1, 0, 1$. The energy level diagram and spectrum is shown in figure 2.28. The relative intensities of each peak can be calculated from the time evolution of the density matrix under the Hamiltonian $\mathcal{H} = \mathcal{H}_{zee} + \mathcal{H}_Q$. When probing systems with multiple resonance peaks, many pulses are applied to the system. Each pulse targets a different transition. The set of spectra from each pulse are then combined to give the single NMR spectrum. The transition separated by ω_0 are called the ‘central’ transition and those which involve multiples of ω_Q are called ‘satellite’ transitions. Measurements of the frequency position of the resonance peaks, as a function of angle between the applied field and crystallographic directions in a compound, provide an experimental technique for calculating the elements and principle directions of the EFG [102], [103]. Furthermore, ω_Q is not dependent on the external magnetic field, so zero-field measurements are possible. This technique is known as Nuclear Quadrupole Resonance (NQR) and provides information on the internal structure of a system [104], [105].

Another common interaction in many condensed matter systems is the hyperfine interaction. The hyperfine interaction originates from coupling between nuclear spins and unpaired electron spins within the sample. The magnetic field generated by the unpaired electrons is called the ‘hyperfine field’. Typically, the hyperfine field is included as an additional term in equation (2.19) as

$$\mathcal{H}_{zee} = -\boldsymbol{\mu} \cdot (\mathbf{B} + \mathbf{B}_{hyp}) \quad (2.44)$$

where \mathbf{B}_{hyp} is the hyperfine field [4]. The full hyperfine interaction Hamiltonian between a

nuclear spin, \mathbf{I} and electron spin \mathbf{S} has the form [106].

$$\mathcal{H}_{hyp} = -\frac{\mu_0}{4\pi} \gamma_e \gamma_n \hbar^2 \sum_j \frac{\mathbf{I} \cdot \mathbf{S}_j}{r_i^3} + \frac{3(\mathbf{I} \cdot \mathbf{r}_j)(\mathbf{S}_j \cdot \mathbf{r}_j)}{r_j^5} \quad (2.45)$$

Where \mathbf{S}_j is the electron spin at site j and \mathbf{r}_j is the vector between \mathbf{I} and \mathbf{S}_j .

The hyperfine interaction causes a shift in the frequency of a resonance peak in the NMR spectrum, known as the Knight shift. The Knight shift is defined as the percent shift of the Larmor frequency of the nucleus in a compound from the Larmor frequency of the nucleus in free space. Consequently, the measurement of the Larmor frequency of a nucleus within a compound is a direct probe of the hyperfine interaction. K can be related to the magnetic susceptibility, χ of the system by [85].

$$K = \frac{A \chi}{\hbar \gamma_n \mu_B} \quad (2.46)$$

Equation (2.46) shows the linear dependence of K on χ allowing for the numerical value of the hyperfine coupling constant, A , to be measured from K - χ plots [107], [108]. A has units of magnetic field and the quantity $A/g\mu_B$ is quoted as the ‘hyperfine field’ [4]. If the hyperfine interaction is dependent on the orientation of the static magnetic field K , χ and A are tensors. The off-diagonal components of these tensors are equal to zero when the static field is aligned with a principal axis of the tensor. Additionally the temperature dependence of the Knight shift can be measured and used to indicate the existence of a temperature driven phase transition in a compound [108].

If a compound has no quadrupolar interaction, hyperfine interactions drive the NMR longitudinal relaxation rate, T_1 [4]. Fermi’s Golden Rule can be used to relate T_1 to the spin correlation functions within a quantum system as [107].

$$\frac{1}{T_1} = \frac{\gamma_n^2}{2} \int_{-\infty}^{\infty} \langle h_+(t)h_-(0) \rangle e^{i\omega_0 t} dt \quad (2.47)$$

where $\langle \dots \rangle$ denotes the time averaged value over the NMR timescale and h_+ and h_- are the transverse fluctuations of the hyperfine field. These fluctuations are in the two orthogonal

directions to the static field direction. Equation (2.47) can be expanded in terms of the components of the hyperfine coupling tensor

$$\begin{aligned} \frac{1}{T_1} = & \frac{\gamma_n^2}{2} \int_{-\infty}^{\infty} e^{i\omega_0 t} \left[(A_{xz}^2 + A_{yz}^2) \langle S_z(t) S_z(0) \rangle \right. \\ & + (A_{xx}^2 + A_{yx}^2) \langle S_x(t) S_x(0) \rangle \\ & \left. + (A_{yy}^2 + A_{xy}^2) \langle S_y(t) S_y(0) \rangle \right] dt \end{aligned} \quad (2.48)$$

where A_{ii} are the directional components of the A tensor and $i = \{x, y, z\}$ [4].

In general, A is non-diagonal which gives T_1 a complex angular dependence on the static field orientation as it depends on both the transverse and longitudinal components of the hyperfine field fluctuations. If A is diagonal, equation (2.48) simplifies and T_1 is solely dependent on the transverse fluctuations. Consequently, changing the orientation of the sample with respect to the static field allows for measurements of the longitudinal and transverse spin fluctuations [107]. This can be used when studying field induced magnetic phase transitions, as the amplitude of each component may dramatically vary between each magnetic phase [107].

2.3 Magnetism

2.3.1 Phase transitions and spin models

Almost all magnetic materials display a phase transition between an ordered and disordered magnetic state. The most common phase transitions are driven by a change of temperature. A phase transition causes a spontaneous symmetry breaking within the system. Above a critical temperature, T_c , the material is paramagnetic where the individual magnetic moments from each spin are randomly oriented forming a disordered state. Below T_c , the moments will magnetically order forming an ordered magnetic state. Ferromagnetic (FM) and antiferromagnetic (AFM) states are examples of ordered magnetic states. FM order causes neighbouring moments to align whereas AFM states have anti-aligned moments [109]. These phase transitions are a result of ‘competition’ between the entropy and internal energy of the moments. For $T > T_c$, the entropy contribution dominates the free energy of the system, resulting in a disordered state. For $T < T_c$, the internal energy dominates, preferring an ordered state [110].

Phase transitions are described by an order parameter. The order parameter is a quantity that is strictly zero in the disordered phase and holds a value in the ordered phase. In an AFM system, the order parameter is the staggered magnetisation [111]. The staggered magnetisation is the magnetisation between the two anti-aligned sublattices. This meets the criteria for a suitable order parameter as the staggered magnetisation is zero in the paramagnetic phase due to the random orientation of the spins. The evolution of the order parameter as a function of temperature across T_c allows the nature of a thermal phase transition to be examined. A discontinuous change of the order parameter across the transition is classified as a first-order phase transition [111]. A continuous change is known as a second-order transition [111].

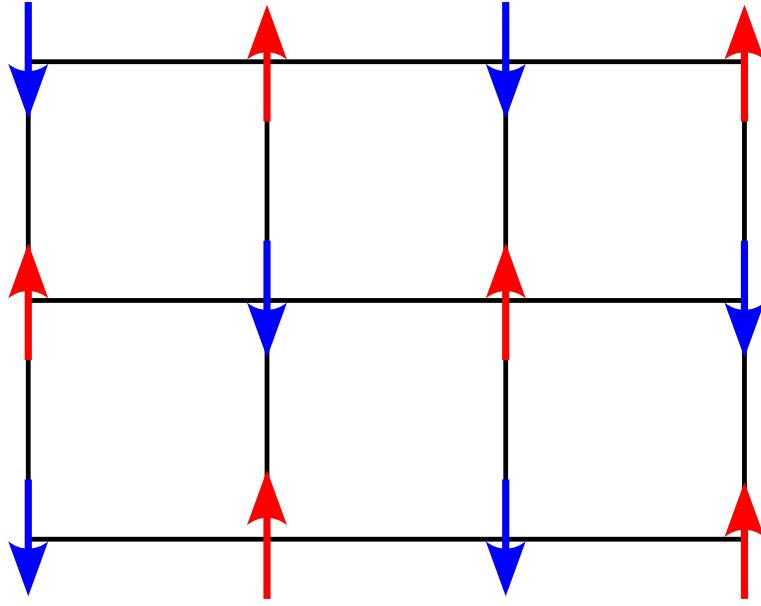


Figure 2.29: An example of AFM ordering. The red and blue arrows indicate the anti-aligned sublattices.

The nature of a phase transition is characterised by a set of critical exponents. These critical exponents describe how experimental quantities behave in the vicinity of T_c . For a thermal phase transition, these are defined in terms of the reduced temperature, t , as

$$t = \frac{T - T_c}{T_c} \quad (2.49)$$

Table 2.2 shows the most common set of critical exponents along with their associated observables for a magnetic system. Generally, an exponent will be different depending on if T_c is approached from above or below [112]. The critical exponents α , β , γ and δ are related to the thermodynamic observables of the system. The change in Gibbs free energy, G , of a magnetic system is

$$dG = -SdT - mdH \quad (2.50)$$

where S is the entropy, T is the temperature, m is the magnetisation (the order parameter) and H is an external magnetic field that acts on the system. The magnetisation, heat

capacity and susceptibility have the following relations with the Gibbs free energy

$$m = -\frac{\partial G}{\partial H} \quad (2.51)$$

$$C = T \left(\frac{\partial S}{\partial T} \right)_H \quad (2.52)$$

$$\chi = \left(\frac{\partial m}{\partial H} \right)_T \quad (2.53)$$

As these quantities are derived from the free energy, the thermodynamic exponents are not independent of each other. Consequently, these exponents satisfy a pair of scaling relations [113]

$$2 - \alpha = 2\beta + \gamma, \quad (2.54)$$

$$2 - \alpha = \beta(\delta + 1). \quad (2.55)$$

The exponents ν , η and z describe the behaviour of the correlations within the system. These correlations describe the fluctuations of the order parameter that occur over a characteristic length and time scale ξ and τ_c . Note that there still be fluctuations of the order parameter within the disordered phase however the thermally averaged value of the order parameter must remain zero. As the temperature approaches the critical temperature the correlation length and time diverge. At the critical temperature, the system is said to be ‘scale-invariant’ as the fluctuations occur at every length and time scale [113]. As a result, the system is completely characterised by the critical exponents in table 2.2. These exponents also follow a set of ‘hyper-scaling’ relations [113]

$$2 - \alpha = d\nu, \quad (2.56)$$

$$\gamma = (2 - \eta)\nu \quad (2.57)$$

where d is the dimensionality of the system.

Table 2.2: Definitions for the critical exponents for a magnetic system. H is the magnetic field, t is the reduced temperature defined in the text and d is the dimensionality of the system. Table taken from [113].

Quantity	Exponent	Relation	Conditions
Specific heat (C)	α	$C \propto t ^{-\alpha}$	$t \rightarrow 0$ and $H = 0$
Order parameter (m)	β	$m \propto (-t)^\beta$	$t \rightarrow 0^-$ and $H = 0$
Magnetic susceptibility (χ)	γ	$\chi \propto t ^{-\gamma}$	$t \rightarrow 0$ and $H = 0$
Critical isotherm	δ	$H \propto m ^\delta \text{sgn}(m)$	$H \rightarrow 0$ and $t = 0$
Correlation length (ξ)	ν	$\xi \propto t ^{-\nu}$	$t \rightarrow 0$ and $H = 0$
Correlation function (G)	η	$G(r) \propto r ^{-d+2-\eta}$	$t = 0$ and $H = 0$
Correlation time (τ_c)	z	$\tau_c \propto \xi^z$	$t \rightarrow 0$ and $H = 0$

The critical exponents of a system relate to a wider theoretical concept known as ‘Universality’. Universality means the physics of a system is not restricted to the physical realisation of the system [112]. This allows for systems which have very different physical realisations to be classified by their set of critical exponents. If two systems have the same set of critical exponents, they are said to belong to the same ‘universality class’. A common example is a van der Waals liquid-gas phase transition of a chemical system has the same set of critical exponents as the ferromagnetic-paramagnetic transition in the Ising model. These two systems have a different physical realisation however at the critical point, both are governed by the same physical theory. This is a consequence of the diverging correlation length of the fluctuations at the critical point. The microscopic details of the system are effectively averaged out and become unimportant [113]. This implies systems that belong to a specified universality class can all be described using the same mathematical ideas.

A quantum phase transition is a non-thermal phase transition driven by quantum fluctuations at $T = 0$ [114], [115]. They require a quantum mechanical description and cannot

be fully analysed using classical thermodynamics. Typically, quantum phase transitions are induced experimentally by the application of a magnetic field but can also be induced using other non-thermal parameters such as pressure [115].

Figure 2.30 shows a generic phase diagram for a system featuring a thermal and quantum phase transition. The quantum phase transition occurs at a quantum critical point (QCP). Changing the temperature of the system can induce a classical thermal phase transition between an ordered state and thermally disordered state. Likewise, adjusting the applied magnetic field can induce a transition between an ordered state and quantum disordered state. The field at which the quantum phase transition occurs is called the critical field, denoted by H_c .

The quantum critical region describes where quantum fluctuations dominate over thermal effects. This leads to a scaling behaviour of various observables [113]. The scaling relations for the heat capacity, C , the magnetic susceptibility, χ , and magnetisation M are given by [116]

$$C(T, h) = T^{d/z} \mathcal{C}(\tilde{h}) \quad (2.58)$$

$$M(T, h) = T^{\frac{(d+z)\nu-1}{z\nu}} \mathcal{M}(\tilde{h}) \quad (2.59)$$

$$\chi(T, h) = T^{\frac{(d+z)\nu-2}{z\nu}} \chi(\tilde{h}) \quad (2.60)$$

The RHS of these equations include a scaling function of the parameter \tilde{h} defined as $\tilde{h} = hT^{-1/z\nu}$. h is the reduced field defined in an analogous way to the reduced temperature from equation (2.49). The exponents are the same as defined in table 2.2. The effective dimension of a QCP is given by $d_{eff} = d + z$ where z is the dynamical exponent [117] defined in table 2.2. The combination of thermodynamical and the dynamical exponents within the scaling relations generate the exotic physics within this region.

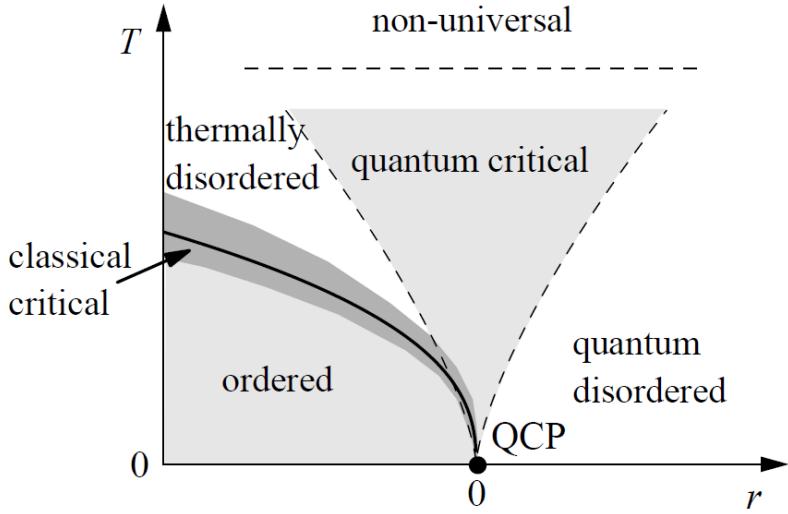


Figure 2.30: A phase diagram showing a classical and quantum phase transition. The main text describes the case when the tuning parameter r is an applied magnetic field. In this case r is a reduced tuning parameter where $r = 0$ at the critical value. Reproduced from [113].

The interactions in a magnetic system are mathematically represented by a spin model Hamiltonian. The most well-known spin model is the Heisenberg model with the Hamiltonian

$$\mathcal{H} = - \sum_{\langle i,j \rangle} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j \quad (2.61)$$

where J_{ij} is interaction strength between the spins on lattice site i and j , the sum runs over all nearest neighbour spins and \mathbf{S}_i is the spin vector on lattice site i . The sign of J_{ij} characterizes the type of magnetic ordering within the system. Conventionally, $J_{ij} > 0$ represents a ferromagnetic interaction and $J_{ij} < 0$ is an antiferromagnetic interaction [109]. This model reduces to the well-known Ising model if only the z -component spin interactions are considered.

The dimensionality, d , of the underlying lattice also affects the physics within the

system. For example, a spin chain is a $d = 1$ system whereas the square lattice is $d = 2$. Quasi-1D systems are an intermediate class of systems known as ‘spin ladders’ A spin ladder consists of a set of spin chains (‘legs’) that interact with neighbouring chains (‘rungs’). The behaviour of these systems are highly dependent on the relative strengths of the leg and rung interactions [118]. To study these phenomena, target materials have to be found that host low-dimensional magnetism. In reality, magnetic materials contain three spatial dimensions. These systems can host low-dimensional magnetism if the directional dependence of J is determined [119], for example, the system can be treated as one-dimensional if the magnitude of J is large in one direction.

2.3.2 The Haldane chain

The $S = 1$ Heisenberg AFM spin chain is more commonly known as the Haldane chain following Haldane’s conjecture that integer spin chains contained a gapped excitation spectrum [30]. The most general Hamiltonian for a Haldane chain system is

$$\mathcal{H}_{Hal} = J \sum_i [(\mathbf{S}_i \mathbf{S}_{i+1} + \beta (\mathbf{S}_i \mathbf{S}_{i+1})^2] + \sum_i [D(S_i^z)^2 + E[(S_i^x)^2 - (S_i^y)^2] - g\mu_B S_i^\alpha H^\alpha] \quad (2.62)$$

where \mathbf{S}_i is the spin at site i , the z axis runs parallel to the chain direction, $\alpha = x, y, z$ is the direction of the applied field H^α and J , β , D and E are the magnitudes of the Heisenberg, biquadratic, single-ion anisotropy and rhombic anisotropy interactions respectively [18]. The Heisenberg interaction strength between neighbouring spins, J , generally has units of energy. The other interaction strengths β , D and E are typically written in units of J so the dominant contributions of equation (2.62) can be easily identified. Synthesised Haldane compounds will not contain all of these terms, for example, the biquadratic term is often neglected due to its small interaction strength in 1D systems [120].

The third and fourth terms represent the anisotropic interactions in the system. Phys-

ically, the magnetic anisotropy describes how the magnetic properties of a compound are dependent on its crystallographic directions [121]. In this case $D < 0$ favours a spin orientation parallel to the z -axis whereas $D > 0$ results in a spin orientation perpendicular to the chain direction. The isotropic Haldane chain refers to the simplified case of no magnetic anisotropic interaction with $D = E = 0$. This model is only exactly solvable in the finite chain limit [122].

The Mermin-Wagner theorem states that there is no long range ordering in 1D spin chains at a finite temperature [123]. Under the application of a magnetic field, the Haldane chain system demonstrates a rich phase diagram. Figure 2.31 shows the phase diagram of an isotropic Haldane chain as a function of temperature and magnetic field. It contains three distinct regions: the eponymous gapped ‘Haldane’ phase, a Tomonaga-Luttinger spin liquid phase and a saturated ferromagnetic phase [18].

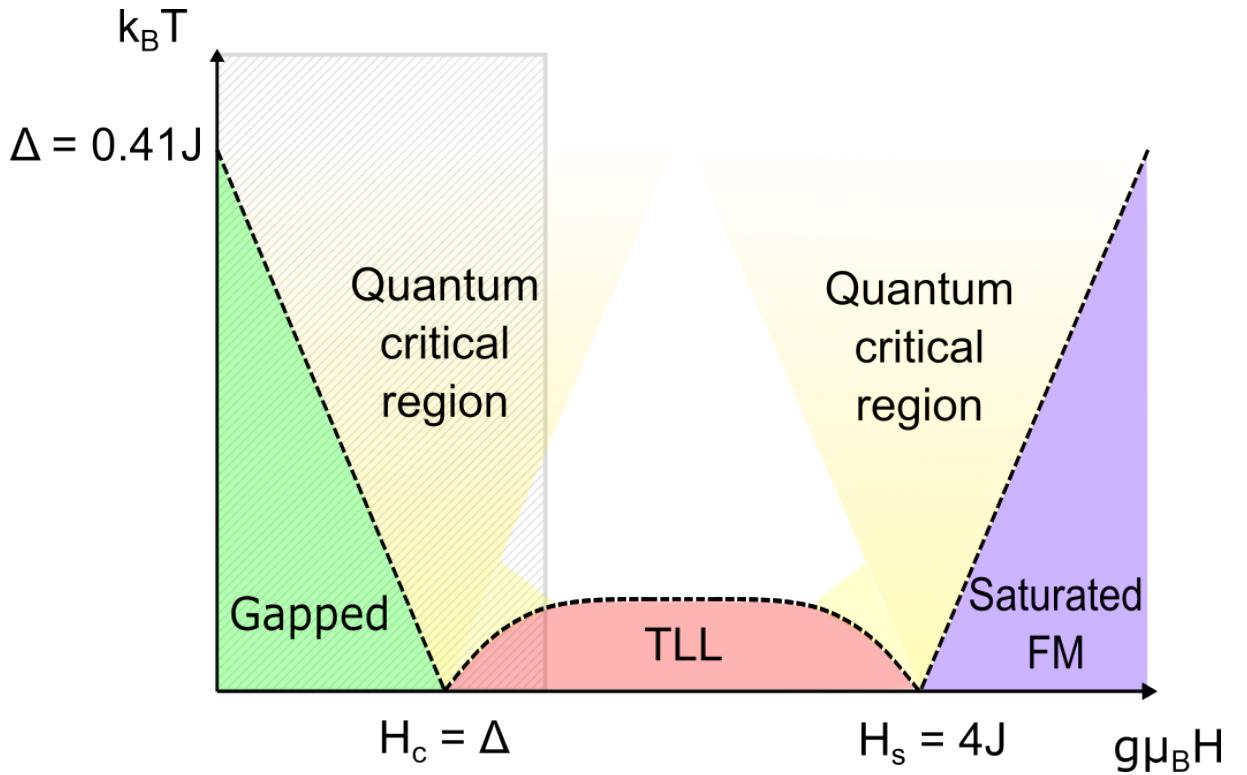


Figure 2.31: The phase diagram for an isotropic Haldane chain in a magnetic field. H_c and H_s are the critical and saturation fields that denote the fields at which the system transitions between gapped and TLL phase and the TLL and saturated ferromagnet phase. Δ is the energy gap between the singlet ground state and triplet state in the gapped phase. This gap is closed by tuning the applied field to H_c . The grey hatched region indicates where the measurements presented in this thesis were taken.

Figure 2.32 shows a schematic diagram of the energy levels in the isotropic Haldane chain. In the gapped phase, the ground state is a non-magnetic singlet state. The first excited state is a $S = 1$ triplet state, known as a magnon, is separated from the ground state by an energy gap of Δ . Quantum Monte-Carlo (QMC) and Density Matrix Renormalization Group (DMRG) numerical methods have been used to calculate the energy gap yielding $\Delta = 0.41 J$ [124], [125]. Upon the application of a magnetic field, the energy of $S_z = -1$ branch of the magnon begins to linearly decrease due to the Zeeman effect, suppressing the gap. At

the critical field, H_c , the energy of the $S_z = -1$ branch becomes equal to the ground state energy. The critical field is given by [126]

$$H_c = \frac{\Delta}{g\mu_B} \quad (2.63)$$

This causes a quantum phase transition from the gapped state to a Tomonaga-Luttinger liquid (TLL) state.

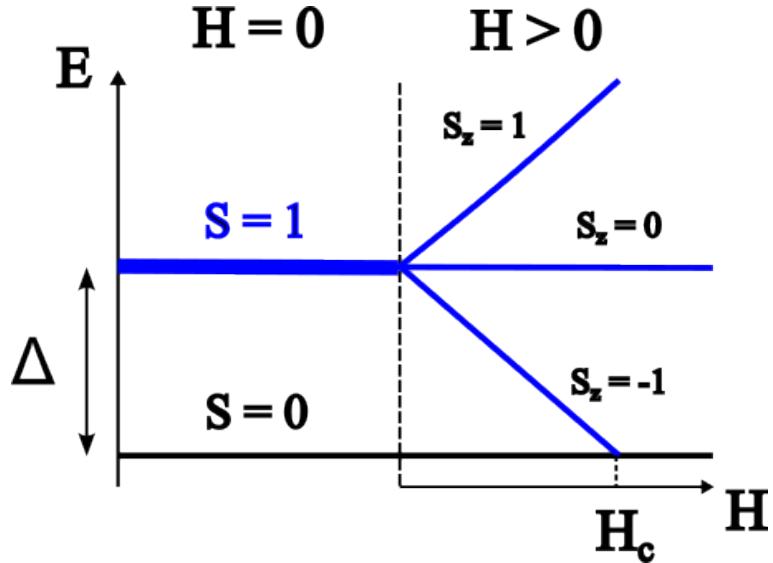


Figure 2.32: The energy level diagram for the isotropic $S = 1$ Haldane chain in a magnetic field. For $H = 0$, the excited state is a $S = 1$ magnon. For $H > 0$, the Zeeman effect causes an overlap between the $S_z = -1$ branch of the magnon and the $S = 0$ singlet at the critical field H_c .

In the case of an anisotropic $S = 1$ chain ($D \neq 0$), anisotropic interactions introduce an additional zero field splitting of the $S = 1$ triplet into a single and doublet [18]. Figure 2.33 shows the energy level diagram for an anisotropic $S = 1$ chain with $D > 0$. Additionally, the energy gap Δ , is now dependant on the field orientation with respect to the chain direction [18]. For the region $0 < D < 0.25J$, the following linear relations have been shown numerically [127]

$$\Delta_z = \Delta + 1.41D \quad (2.64)$$

$$\Delta_{xy} = \Delta - 0.57D \quad (2.65)$$

where Δ_z and Δ_{xy} are the energy gaps when the field is parallel and perpendicular to the chain direction respectively. In zero-field, the anisotropy interaction can be large enough to close the gap. Figure 2.34 shows the energy gap as a function of the single-ion anisotropy. The gap is closed for the critical value $D_c = 0.968J$ [128]. Above D_c , the energy gap re-opens and this signals a phase transition into the ‘large-D’ phase [18].

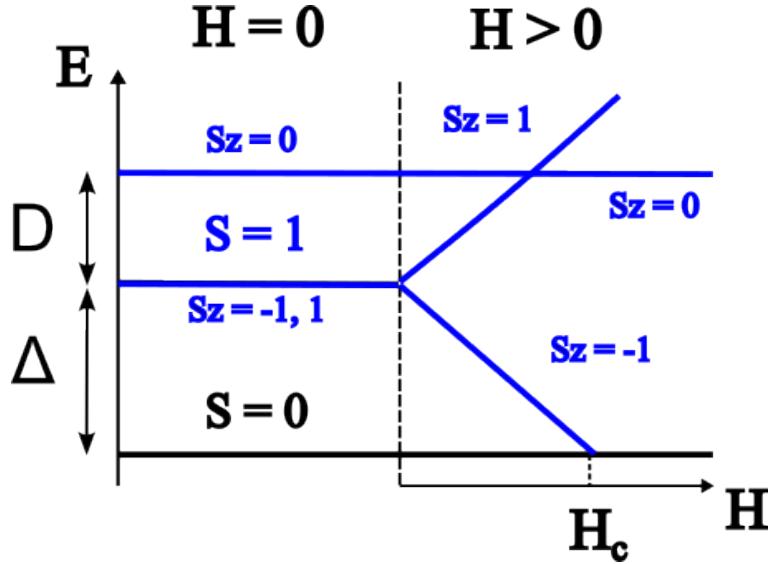


Figure 2.33: The energy level diagram for the anisotropic $S = 1$ Haldane chain in a magnetic field with $D > 0$. For $H = 0$, there is a zero-field splitting in the $S = 1$ state. For $H > 0$ the behavior is the same as the isotropic case shown in figure 2.32.

A TLL state is special case of spin liquid state found in systems with a gapless excitation spectrum. For $d > 1$, these systems are typically described using Landau’s Fermi liquid theory which describes the system in terms of interactions between quasiparticles [129]. The main assumption of this theory is that low energy states of the interacting system are adiabatically related to the system with no interactions causing a continuous change in the excitation spectrum. In $d = 1$, Fermi liquid breaks down due to the equivalence between collective and individual excitations caused by the lack of degrees of freedom [130]. The

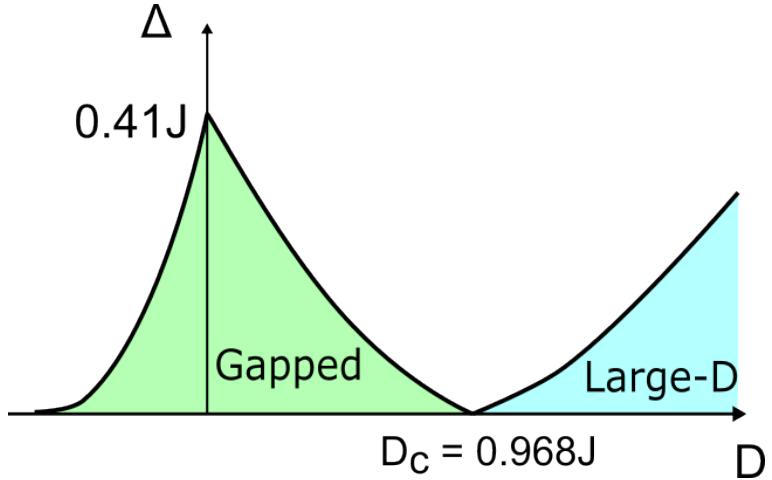


Figure 2.34: A schematic diagram showing the energy gap, Δ , as a function of the single-ion anisotropy interaction strength D . The gap is closed at the critical value D_c . Below D_c the gapped phase is the same one as in figure 2.31. Above D_c the gap re-opens and is known as the ‘large- D ’ phase.

excitation of a single spin in the systems perturbs all other spins in the system. Therefore, the TLL description of the system replaces Fermi liquid theory for $d = 1$ systems.

A TLL state is characterized by two parameters: the Fermi velocity of excitations, u , and the Luttinger parameter, K . The value of K determines whether the interactions are uninteractive ($K = 1$), attractive ($K > 1$) or repulsive ($K < 1$) [131]. The TLL in the Haldane chain is expected to show attractive behaviour [132]. All thermodynamic quantities and correlation functions in a TLL state can be expressed in terms of K and u [133]. The magnetisation and magnetic heat capacity of a TLL can be written as [131]

$$\frac{\partial m_z}{\partial(\mu_0 H)} = \frac{g\mu_B}{2\pi k_B} \left(\frac{K}{u} \right) \quad (2.66)$$

$$C_{mag} = N_A \frac{\pi k_B T}{6u} \quad (2.67)$$

where m_z is the reduced magnetisation in the field direction, g is the g-factor and C_{mag} is the magnetic heat capacity.

The correlation functions in a spin chain TLL are dominated by the transverse correlations [133]. These correlations are measured by the NMR relaxation rate, T_1 . The relaxation rate has the following scaling form

$$T_1^{-1}(T) = (k_B T)^\alpha F\left(\frac{\hbar\omega}{k_B T}\right) \quad (2.68)$$

where $\alpha = 1/2K - 1$ and F is a universal function [133]. In the limit of $\omega \rightarrow 0$, F loses its temperature dependence and T_1 approximately follows a power law

$$T_1^{-1}(T) \propto T^\alpha. \quad (2.69)$$

The TLL state persists up to a second critical field, known as the saturation field H_s given by

$$H_s = \frac{4J}{g\mu_B} \quad (2.70)$$

For fields $H > H_s$ the isotropic Haldane chain transition to saturated ferromagnetic or fully polarised state. This is caused by the magnetic field being large enough to force all spins to align in the field direction. The excitation spectrum becomes gapped in this regime and can be exactly calculated due to the magnetisation of this state being known [18].

The quantum phase transitions at H_c and H_s both form quantum critical points. Both quantum critical points are described by a $z = 2$ dynamical exponent and therefore show scaling relations of experimental observables [126]. The yellow areas in figure 2.31 indicate the regime where the physics of the quantum critical point is expected to govern the behaviour. The quantum critical behaviour persists up to a temperature of $T = J/k_B$, above which the system displays paramagnetic behaviour [110].

Chapter Three

FPGA-based NMR spectrometer

This section will present the FPGA-based NMR spectrometer designed and implemented during this project. It will cover the Verilog design and testing used to verify the operation of the FPGA source code. Then the Python control GUI and the bespoke NMR probe will be introduced. Finally, NMR measurements on a cobalt powder sample will be presented to demonstrate the functionality of the system.

3.1 System overview

Figure 3.1 shows a diagram of the NMR set-up. The FPGA-based spectrometer was built using an SP Devices SDR14 radioprocessor board containing a Xilinx FPGA. This controlled the NMR experiment and had the necessary functionality to perform NMR experiments as outlined in section 2.2.3. This was controlled using a Python GUI on a nearby host PC. A Tomco high-power RF amplifier was used to provide the high power RF signals required for NMR. Finally, a Quantum Design Physical Property Management System (PPMS) was used to provide a sample environment with controllable temperature and applied magnetic field. The control of the PPMS was integrated into the Python GUI allowing for control of

the sample environment during an NMR experiment. A bespoke NMR probe, containing a sample holder and the NMR circuitry was designed in accordance with the specifications of the PPMS.

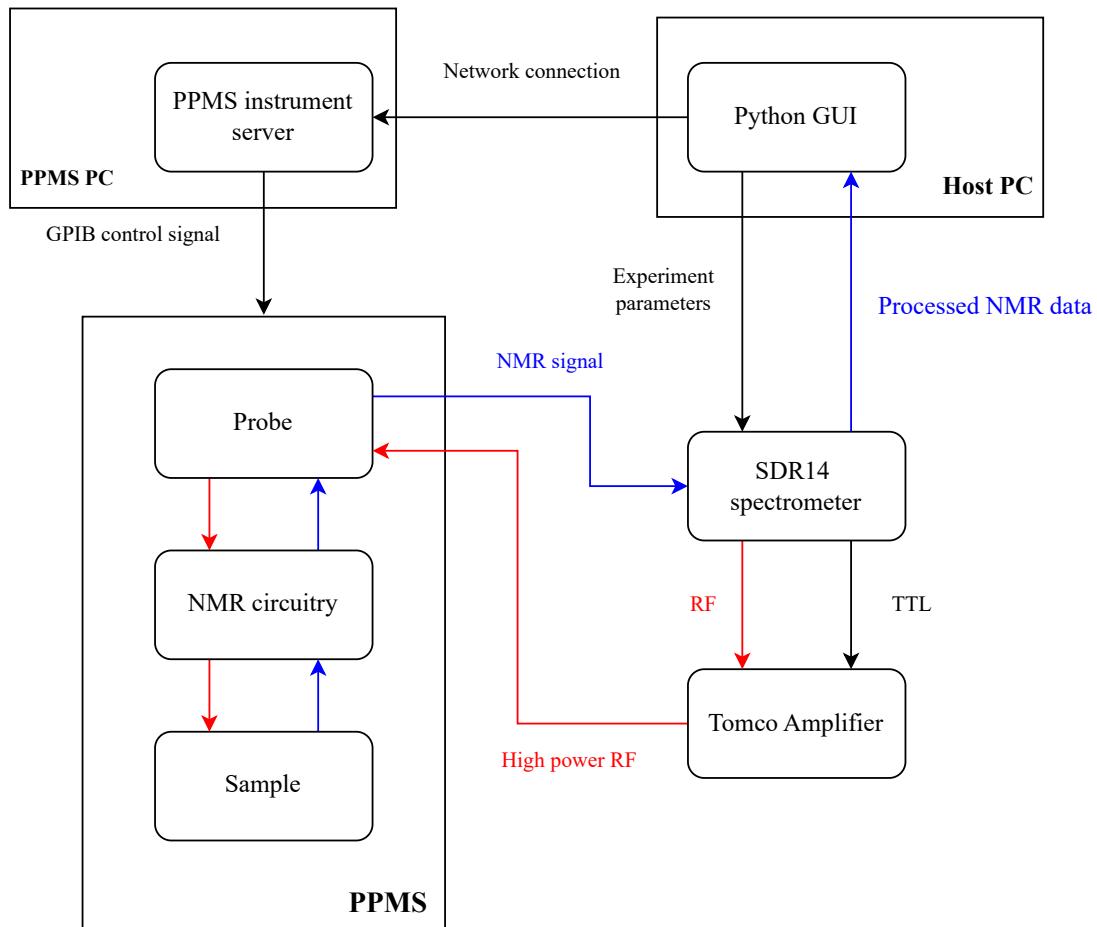


Figure 3.1: A diagram of the full NMR setup. The red and blue arrows indicate data flow for the transmitter and receiver signals. Black arrows represent control signals.

3.2 SDR14 spectrometer

The FPGA implementation of the spectrometer was based on the method found in Pikulski et al [13]. The FPGA was embedded on a SDR14 arbitrary waveform generator and digitiser radio processor board manufactured by Teledyne SP Devices [134]. An image of the device is shown in figure 3.2.



Figure 3.2: An image of the SDR-14 hardware used in the spectrometer.

The SDR-14 contained a Xilinx Virtex-6 LX240T FPGA [135] to control the operation of the device. The FPGA logic was controlled by a 200 MHz (5 ns time period) system clock. As seen in figure 3.2 the SDR14 contains a number of input and output channels that can be addressed by the FPGA logic. There are two analogue to digital converters (ADCs) with a sampling rate of 800 MHz. Additionally, there are two digital to analogue converters (DACs) with a sampling rate of 1.6 GHz. The ADCs and DACs both have a 14-bit resolution, which ensured an accurate conversion between the analogue and digital domains. The

sampling rates of these ports was sufficient for most condensed matter NMR experiments. Furthermore, the SDR-14 contains several general-purpose input/output (GPIO) lines which were used to output TTL control signals for external equipment. The SDR14 features a PCIe interface, which was used to connect the device to a Keysight M9005a PCIe chassis [136]. The PCIe chassis was connected to the host PC using a PCIe connector. The use of the chassis provided a portable device that could be transported and connected to any PC with a PCIe slot.

The main advantage of the SDR-14 board was that all the components on the board were integrated with the FPGA. The main challenge was to design and implement the FPGA logic that provided the functionality to run NMR experiments. The FPGA source code was programmed using Verilog [78]. The code was primarily developed using the Xilinx ISE design suite software package [137]. This contained a set of tools to create, design, visualize and simulate FPGA projects. Xilinx ISim simulator [138] was used to simulate designs which allowed for simple prototyping of Verilog modules. The software license also included the Xilinx COREGEN library [139] which provided efficient prefabricated Verilog modules for common data processing tasks known as IP Cores. These cores were used a for a range of functions across the spectrometer design to save development time.

The design of the spectrometer was split into three distinct entities in order to modularise the design for efficient testing purposes. These entities were the transmitter, receiver and pulse programmer. Figure 3.3 shows a block diagram of the FPGA spectrometer design and the connections between the host PC and NMR environment. The red, blue and green arrows indicate the data flow for the transmitter, receiver and pulse programmer entities within the spectrometer. The yellow blocks represent components that were located on the SDR14 board, but not within the internal FPGA logic. The transmitter and receiver sections were further divided into smaller modules for simplicity due to their complexity.

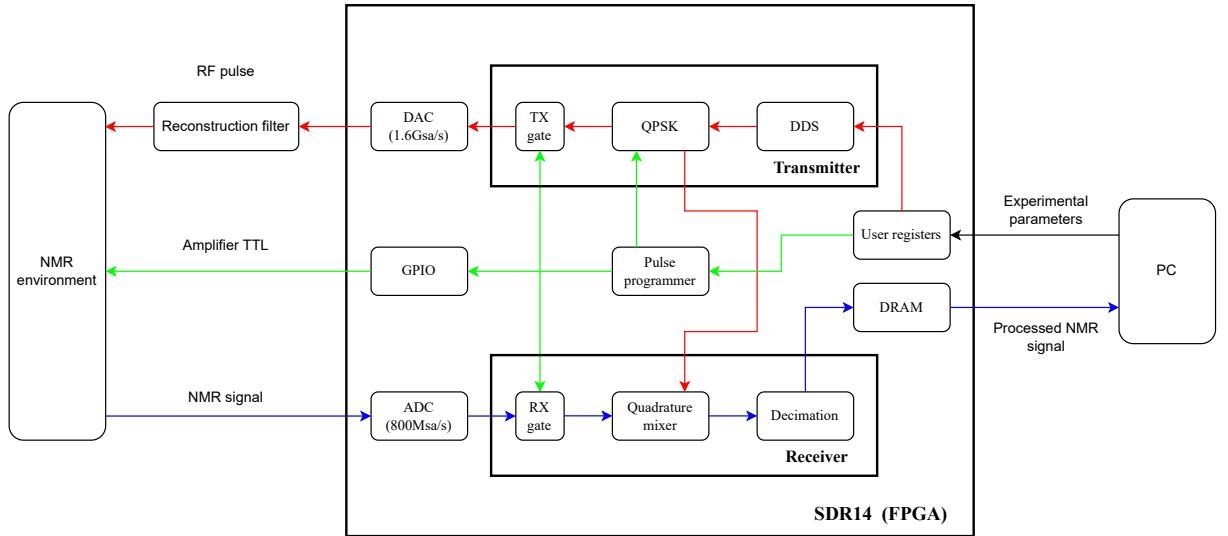


Figure 3.3: A diagram of the internal FPGA logic. Each box represents a Verilog module that implements the labelled functionality. The red, blue and green arrows represent data flow for the transmitter, receiver and pulse programmer signals respectively. The reconstruction filter was an external inline component connected the DAC SMA output. The NMR environment block refers to every component shown in figure 3.1 after the SDR14 spectrometer

The transmitter section was responsible for generating the RF output signals and the local oscillator (LO) signals for use in the receiver. The transmitter section was composed of a DDS module and a QPSK module. The DDS module generated a pair of quadrature signals. These signals were transformed into three signals: the RF output, and LO signals for the quadrature mixer in the receiver. The QPSK module was used to change the phase of the three signals before being sent to their respective destinations. The TX gate controls the output of the generated RF signal to the DAC.

The receiver section processed the incoming NMR signal and sent the processed signal to the distributed random access memory (DRAM) in order for data to be transmitted to the host PC. The receiver section consisted of a quadrature mixer and decimation chain. The combination of these components formed a digital down converter (DDC). The DDC

was used to lower the sample rate of the incoming NMR signal to a suitable rate for the host PC. The downsampled signal was then stored in the FPGA DRAM and transmitted to the PC.

The pulse programmer module controlled the number of RF pulses and their duration. Additionally, the pulse programmer was used to control the TX and RX gates. These are one-bit variables that, when asserted, indicate that the respective section is ready to send or receive data, ensuring for controlled data flow within the device. Furthermore, the pulse programmer generated the TTL signal that controlled the Tomco amplifier. This was required to amplify the low power RF pulse from the DDS to perform the NMR experiment. The flag ensured the amplifier was turned off when the spectrometer was receiving an NMR signal to reduce noise.

The experimental parameters, required for each module to perform their functionality were stored in the FPGA user registers. The user registers are a set of 16, 32-bit wide on-chip registers that can be addressed by the host PC. The full layout of the registers is shown in the register map in figure 3.4. A description of each parameter is provided in table 3.1

Table 3.1: A description of the parameters stored in the FPGA user registers.

Parameter	Width (bits)	Description
PC enable	1	An enable signal used to begin an NMR experiment from the host PC.
freq syn	32	The frequency of the RF output in Hz
P1 length	32	Duration of the first RF pulse in ns
P2 length	32	Duration of the second RF pulse in ns
P3 length	32	Duration of the third RF pulse in ns
G1 length	32	Duration of the gap between pulse 1 and pulse 2 in ns
G2 length	32	Duration of the gap between pulse 2 and pulse 3 in ns
Record length	32	Duration of the ADC enable pulse in ns
TX phase	15	The phase of the RF output signal. Contains three 5-bit phase encodings, one for each pulse.
RX phase	5	Phase of LO signals.

The specifications for the FPGA-based spectrometer functionality were as follows

- The transmitter must be able to generate RF signals up to 800 MHz, with a frequency resolution of the order 10 Hz.
- The phase of the RF signal must be individually configurable for each pulse.
- The phase of the LO signals used within the receiver must be configurable for each pulse.
- The generated RF pulse must be gated, to ensure the output is only non-zero within each TTL pulse, in order for more consistent performance from the Tomco amplifier.
- The pulse programmer must be able to generate up to three separate TTL pulses, of varying duration from 10 ns - 10 ms.

Bit	Register											
	0	1	2	3	4	5	6	7	8	9	10	11 - 15
0	PC_enable	TX phase	freq syn	P1 length	P2 length	P3 length	G1 length	G2 length	Record length	RX phase	UNUSED	UNUSED
1												
2												
3												
4												
5												
6												
7												
8												
9												
10												
11												
12												
13												
14												
15												
16												
17												
18												
19												
20												
21												
22												
23												
24												
25												
26												
27												
28												
29												
30												
31												

Figure 3.4: A diagram showing the layout of the FPGA user registers. A description of the parameters can be found in table 3.1

- The gap duration between each TTL pulse must be configurable over the same timescales.
- The pulse programmer must provide a trigger signal at the correct time for the data acquisition logic.
- The receiver must decimate the signal to a data rate of the order 10 MSa/s for the PC to handle.

The following subsections will discuss each of the modules in more detail. Details

about the implementation of each module will be discussed. Additionally, test data will demonstrate the performance of each module in the context of the above specifications.

3.2.1 Transmitter

The transmitter consisted of two modules, the DDS and QPSK modules. This section had two functions: the generation of the the RF and LO signals, and to apply the phase shift to these signals. This section used the ‘freq syn’ and ‘TX phase’ and ‘RX phase’ from the user registers as inputs. The RF signal was output to the onboard DAC and the LO signals were sent to the receiver section.

DDS

The DDS module was used to synthesise the RF output and LO signal for the receiver at a specified frequency. The implementation of the DDS module was built around the Xilinx DDS Compiler IP core [57]. This IP Core provided a customisable framework that implements the DDS method outlined in section 2.1.1. It included the addition of pipeline registers into the logic, so a ready signal was also implemented to indicate when the logic had overcome the initial latency.

The IP Core used a 32-bit wide tuning word for the DDS and generated output samples of width 16-bits. By equation 2.3, this gave a frequency resolution of less than 0.5 Hz, fulfilling the transmitter specification. It was configured to generate a pair of quadrature signals, where the sine wave was used for the RF signal output and both signals were used as the LO signals in the receiver. As the system clock frequency was 200 MHz, Nyquist’s theorem dictates the maximum synthesisable frequency from the IP Core was 100 MHz. However, the onboard DAC had a sampling rate of 1.6 GHz so the hardware could support

signals of up to 800 MHz. To match the DAC sampling rate of 1.6 GSa/s, eight parallel phase-offset DDS IP cores were used, each one slightly offset in phase from the previous by a fixed amount. Consequently, eight individual samples of the output signal were generated on each clock cycle. These samples were then interleaved (in time) to generate a coherent signal. This effectively raised the sampling rate by a factor of eight, to 1.6 GHz, matching that of the DAC and allowing for signals up to 800 MHz to be synthesised. This technique of generating signals with frequency above the system clock frequency is known as supersampling.

Figure 3.5 shows a Python simulation demonstrating this method for a 50 MHz output signal. Each marker type displays the samples generated from a single DDS, therefore the time between samples from the same DDS is 5 ns, matching the clock frequency of 200 MHz. Between consecutive samples from a given DDS, there are seven other interleaved samples generated from the remaining DDS cores, raising the effective sampling rate.

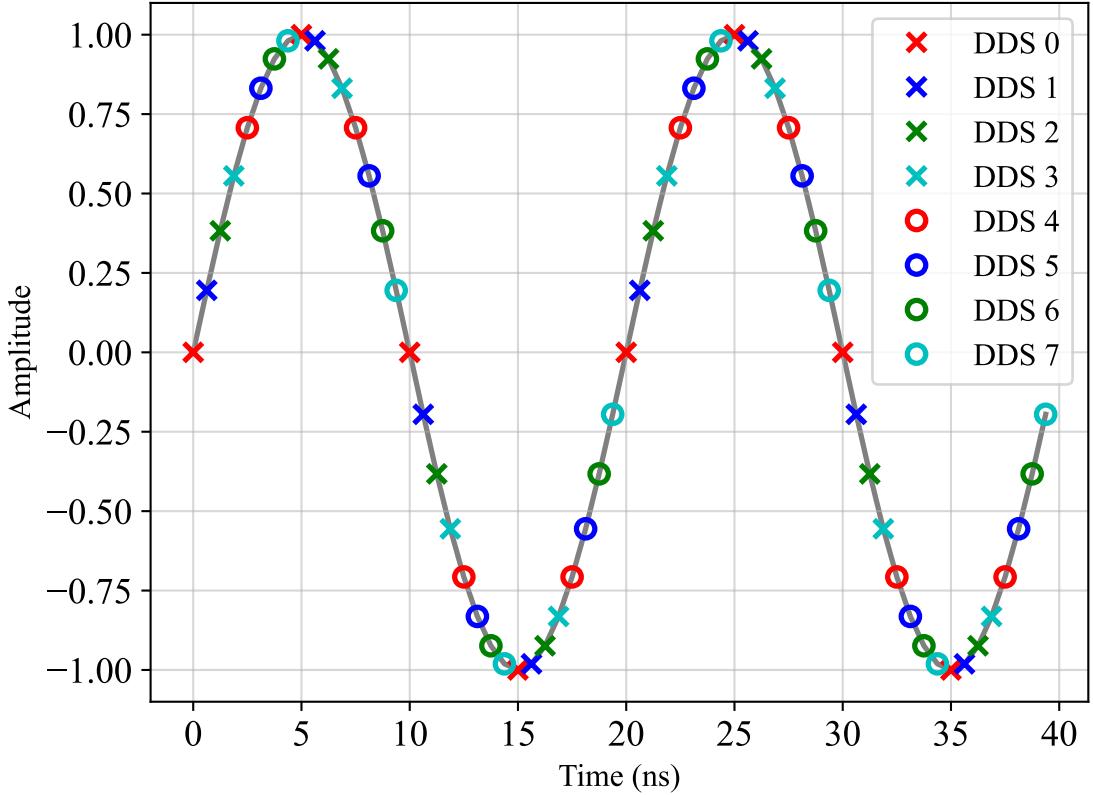


Figure 3.5: A 50 MHz sine wave generated using an 8-core DDS Python simulation. The grey line has to be superimposed in order to guide the eye.

When the synthesised frequency was smaller than the clock frequency, the phase offset, p_i , for the i^{th} DDS core is related to the phase increment by

$$p_i = \frac{i}{8} \Delta\theta(f_{\text{syn}}) \quad (3.1)$$

where $i = 0, 1, \dots, 7$, f_{syn} is the frequency of the synthesised signal and $\Delta\theta$ is the phase increment defined in equation (2.2). This equation states that each DDS core is $1/8$ of a cycle ahead of the previous.

Synthesizing super-sampled frequencies introduced modifications to the calculation of both the phase increment and phase offset. The values of both parameters were dependent

on what Nyquist zone the synthesised frequency was in. The Nyquist zone referred to the zones bounded by consecutive multiples of the system clock frequency, for example, the first zone ($K=1$) was 0 MHz - 200 MHz, the second zone ($K=2$) was 200 MHz - 400 MHz. This was defined as a numerical factor, K , as

$$K = \text{ceil}\left(\frac{f_{syn}}{f_{clk}}\right). \quad (3.2)$$

The scale factor was required because each DDS was driven by the 200 MHz system clock. If $f_{syn} > 200$ MHz, the quantity in equation (2.2) becomes greater than one. In this case $\Delta\theta > 2^N$ where N is the binary width of $\Delta\theta$. Since $\Delta\theta$ corresponds to a memory address within the LUT, binary overflow occurs. The addition of the scaling factor, K , prevents the overflow. The amount the synthesised frequency ‘overflows’ the clock frequency was given by

$$f_{over} = f_{syn} - (K - 1)f_{clk}. \quad (3.3)$$

The rescaled phase increment was then calculated by making the following transformation

$$\Delta\theta(f_{syn}) \rightarrow \Delta\theta(f_{over}) \quad (3.4)$$

where the functional form of $\Delta\theta$ remained the same as defined in equation (2.2). Binary overflow automatically occurred in the FPGA, so no extra computation was required to take place in order to perform this rescaling.

The rescaled phase offset of the i^{th} DDS core was found to take the form

$$p_i(f_{over}) = \frac{i}{8}K\Delta\theta\left(\frac{f_{over}}{K}\right) \quad (3.5)$$

where $i = 0, \dots, 7$. This expression takes into account the overflowed frequency ensuring that each DDS is $1/8^{\text{th}}$ of a cycle ahead of the previous. In the case of synthesised frequencies within the first Nyquist zone, $f_{syn} < f_{clk}$ yields $K = 1$. Equation (3.3) gives $f_{over} = f_{syn}$ and therefore the rescaled phase increment and phase offset become equal to the usual case in equations (2.2) and (3.1). The phase increment and phase offset for each DDS IP Core was

calculated within the DDS module using the frequency stored within the FPGA registers via these equations.

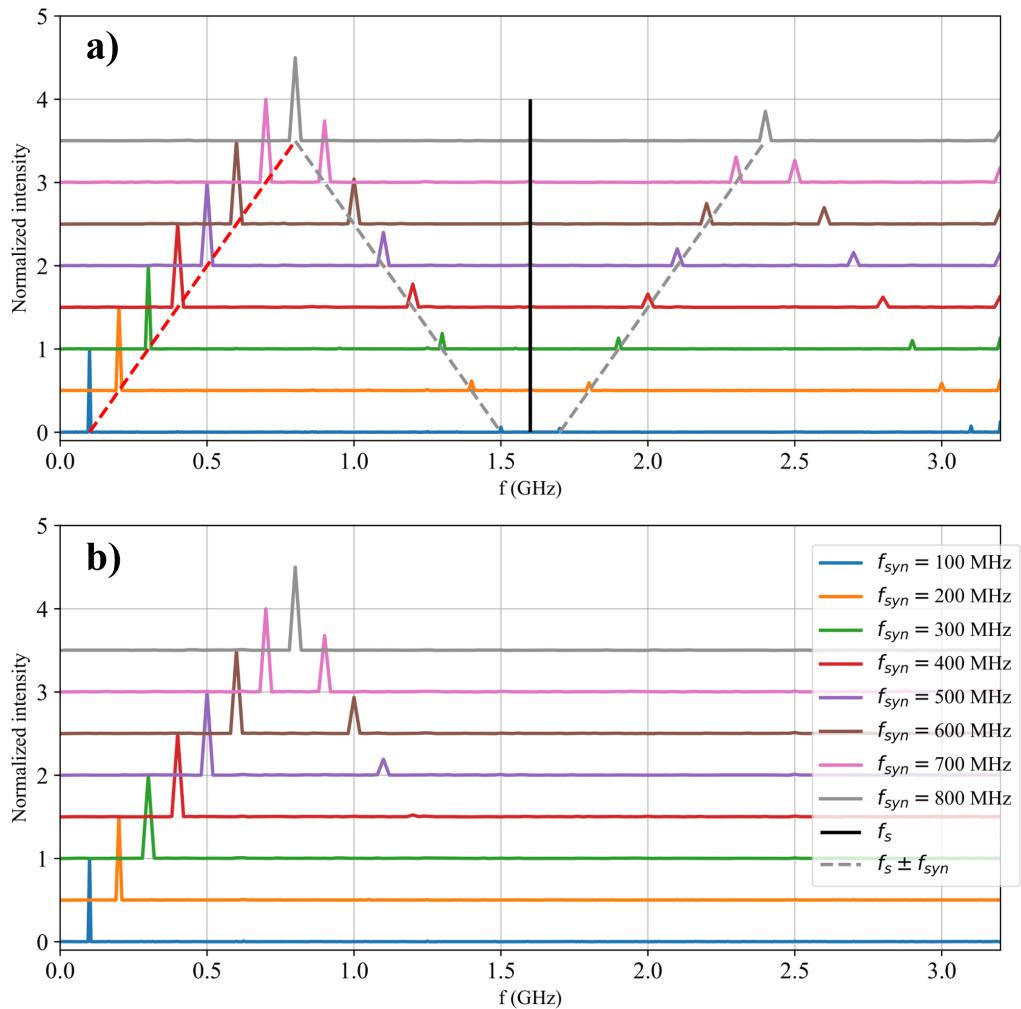


Figure 3.6: a) The spectrum of the RF signal generated by the SDR14 at varying frequencies from 100 MHz - 800 MHz. The red and grey dotted lines represent the dominant and secondary frequency components expected from the ADC response. The solid black line indicates the sampling frequency. b) The spectrum of the RF signal generated by the SDR14 at varying frequencies from 100 MHz - 800 MHz, with a reconstruction LPF used to filter the high frequency components, as described in the text.

Figure 3.6 shows a set of measurements characterizing the RF output of the FPGA. The measurements were obtained by connecting the DAC output directly to an oscilloscope, while the frequency was swept from 0 MHz - 800 MHz. These measurements were performed by bypassing the QPSK module, in order to solely evaluate the performance of the DDS module. The spectrum of the signals was obtained by a fast Fourier transform (FFT) and the intensity was normalized to the largest peak. The red dotted line indicates the frequency of the dominant peak at the specified frequency indicating the DDS module was functioning correctly.

Every spectrum in figure 3.6a contained higher frequency components which depended linearly on the synthesised frequency indicated by the grey dotted lines. These components originated from the use of a DAC to generate the analogue signal. The output of the DAC in the time domain was a series of voltage pulses with varying amplitude [61]. In the frequency domain this yields a sinc envelope resulting in a series of secondary image peaks in addition to the dominant frequency at f_{syn} . These secondary peaks appear at $f_s \pm f_{syn}, 2f_s \pm f_{syn}, \dots$ which is consistent with figure 3.6a.

The secondary peaks were filtered out using a reconstruction filter as shown in figure 3.3. A Minicircuits VLF800+ 800 MHz LPF was used for this purpose. Figure 3.6b shows the same measurements as the previous figure with the reconstruction filter. For synthesised signals with frequency less than 400 MHz, the image peaks were removed completely. For higher frequencies, the attenuation was consistent with the roll-off from the filter response.

Finally, the peak-to-peak voltage of the synthesised signal after filtering was measured as a function of frequency and shown in figure 3.7. It was calculated from the integral of the power spectral density across the whole frequency range. The voltage was roughly independent of frequency over the range 10 MHz to 800 MHz and agreed well with the data sheet value of 0.4 V.

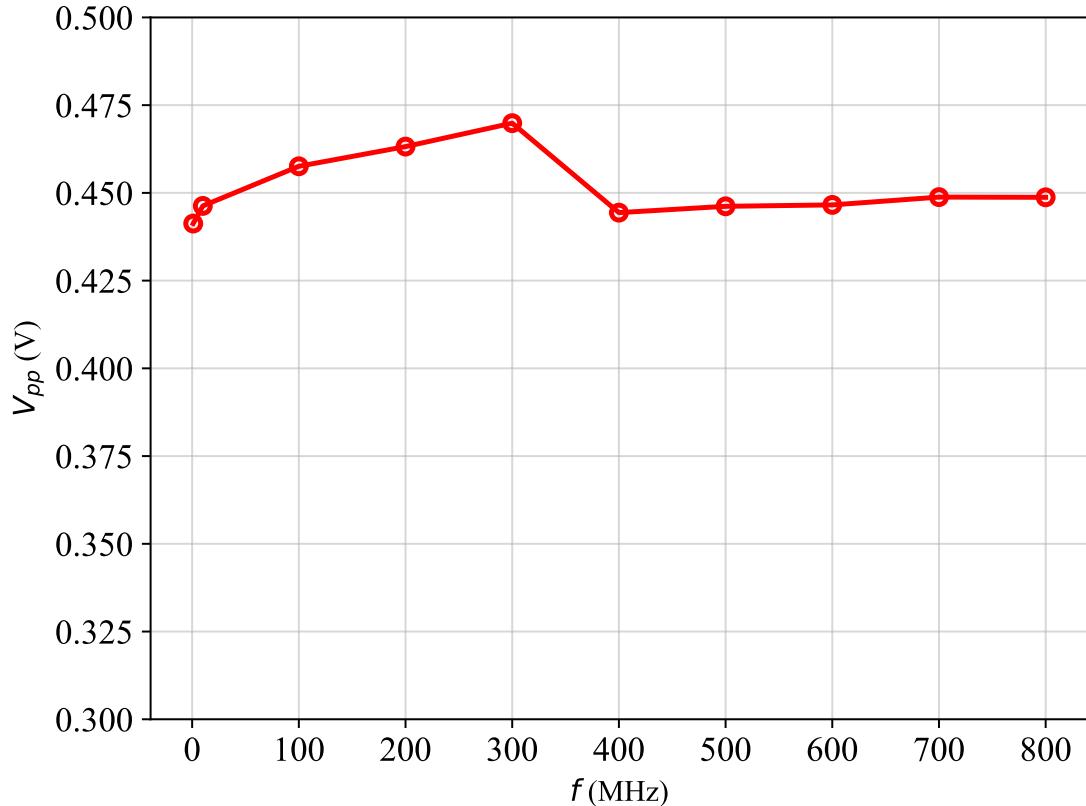


Figure 3.7: The output voltage of the SDR14 as function of frequency

Finally, the peak-to-peak voltage of the synthesised signal after filtering was measured as a function of frequency and shown in figure 3.7. It was calculated from the integral of the power spectral density across the whole frequency range. The voltage was roughly independent of frequency over the range 10 MHz to 800 MHz and agreed well with the data sheet value of 0.4V. The small deviation was expected as the comparison was made with the typical DC voltage.

QPSK

The QPSK module was responsible for applying the TX and RX phase shift to the 16-bit samples generated by the DDS module. The implementation was essentially that of a QPSK modulator described in section 2.1.2. The module took a pair of sine and cosine reference signals, generated by the DDS, and a 2-bit phase encoding for both the TX shift and RX shift as inputs. Each 2-bit phase encoding was derived from a 5-bit value stored in the device registers, the conversion between values is shown in table 3.2. The QPSK module outputted the RF signal with the TX phase shift applied and a pair of LO signals both shifted by the RX phase. The RF output was also reduced to 14-bits wide by dropping the final two least significant bits in order to match the 14-bit DAC resolution.

Table 3.2: Conversion table between the 5-bit TX phase and RX phase stored in the user registers and their 2-bit encoding forms.

5-bit encoding	2-bit encoding	Phase difference (radians)
0	11	$\pi/4$
9	10	$3\pi/4$
18	00	$5\pi/4$
27	01	$7\pi/4$

The output signal of the QPSK module can be written as

$$y_{out} = \sum_{m=0}^1 (-1)^{x[m]+1} S_m \quad (3.6)$$

where S_m is either the sine or cosine reference signal and $x[m]$ is the value of the m^{th} bit of the 2-bit phase encoding. The full behaviour implemented by 3.6 displayed in table 2.1. The signal samples generated from the DDS used a two's complement representation. Therefore, a negation of the sine and cosine samples were performed by inverting all of the bits and adding one. This was done efficiently using a logical operator as an XOR (exclusive OR)

between the sample and a 16-bit string of ones followed by a NOT operation. With these operations, equation (3.6) becomes

$$y_{out}[n] = \sum_{m=0}^1 (\sim x[m]) \oplus S_m[n] \quad (3.7)$$

where \sim and \oplus represent NOT and XOR respectively and $g[n]$ indicates the n^{th} bit of a signal g . This implementation was chosen as the bit-wise operations were performed using logic gates which reduced the number of logical components required. This was necessary as the QPSK module had to operate on 24 parallel samples every clock cycle, 8 samples from the RF output and 16 samples from the pair of LO signals.

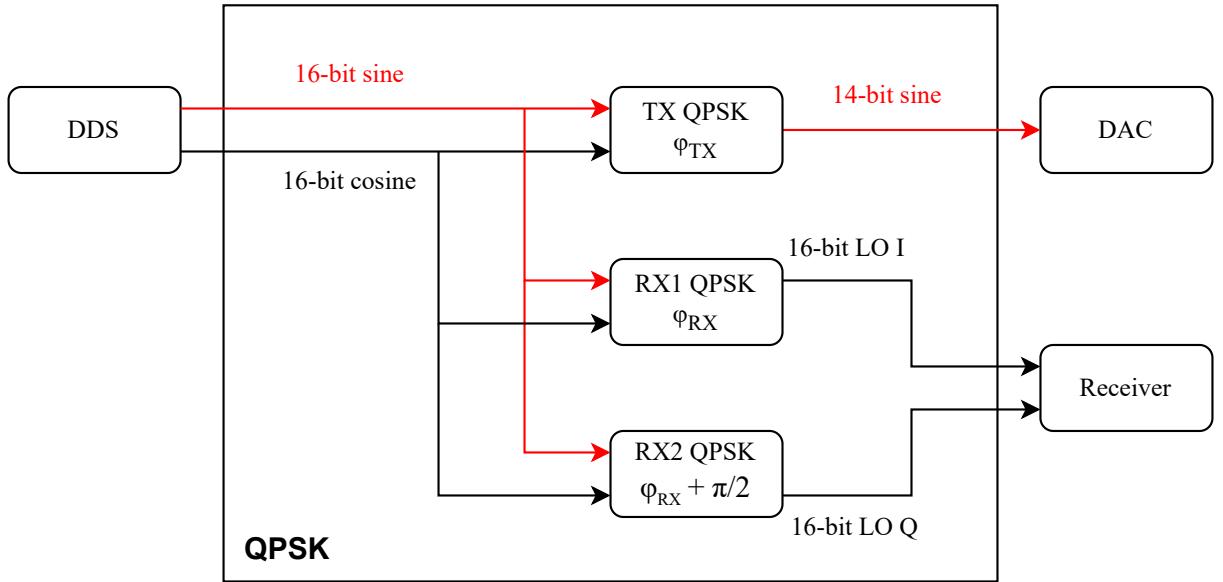


Figure 3.8: An overview of the internal structure of the QPSK module. Each QPSK instance is labelled with its corresponding phase shift.

The full internal architecture of the QPSK module is shown in figure 3.8. Due to the generation of three outputs, the QPSK module acted as wrapper module for three separate QPSK instances. Each instance performed the QPSK modulation described above. The first operation, TX QPSK, applied the TX phase shift to the RF output. The remaining operations, RX1 QPSK and RX2 QPSK, applied the RX phase shift to each LO signal

separately. As the LO signals were required to remain in quadrature after the phase shift, the phase-encoding of the RX2 QPSK operation was fixed to be always $\pi/2$ ahead of RX1 QPSK operation. Finally, the wrapper module outputted the three signals to the rest of the logic.

Figure 3.9 shows 10 MHz RF signal generated by the TX QPSK instance. It shows the output for all values of the binary phase encoding: 00, 01, 10, 11. The solid lines are sine waves plotted with the labelled phase shift. Each signal is out of phase by 90° with the neighbouring signal indicating the TX QPSK instance is functioning correctly.

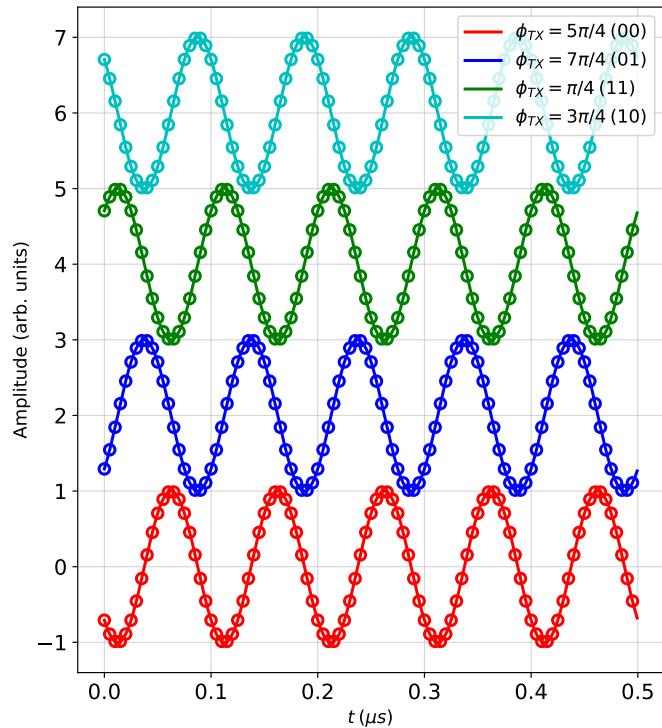


Figure 3.9: A 10 MHz RF signal output from the TX QPSK instance as a function of time. Each line is labelled with the phase shift and the corresponding binary encoding stored in the FPGA user registers.

Figure 3.10a and figure 3.10b shows the 10 MHz LO signals generated by the RX1 and RX2 QPSK instances. The signals demonstrate the correct $\pi/2$ phase-shift between each coloured signal in each subfigure. Furthermore, the LO relationship between the RX signals has been preserved after the phase shift was applied by the RX1 and RX2 QPSK instances. This is demonstrated by the identically coloured signals in the left and right panels of figure 3.10 having a $\pi/2$ shift.

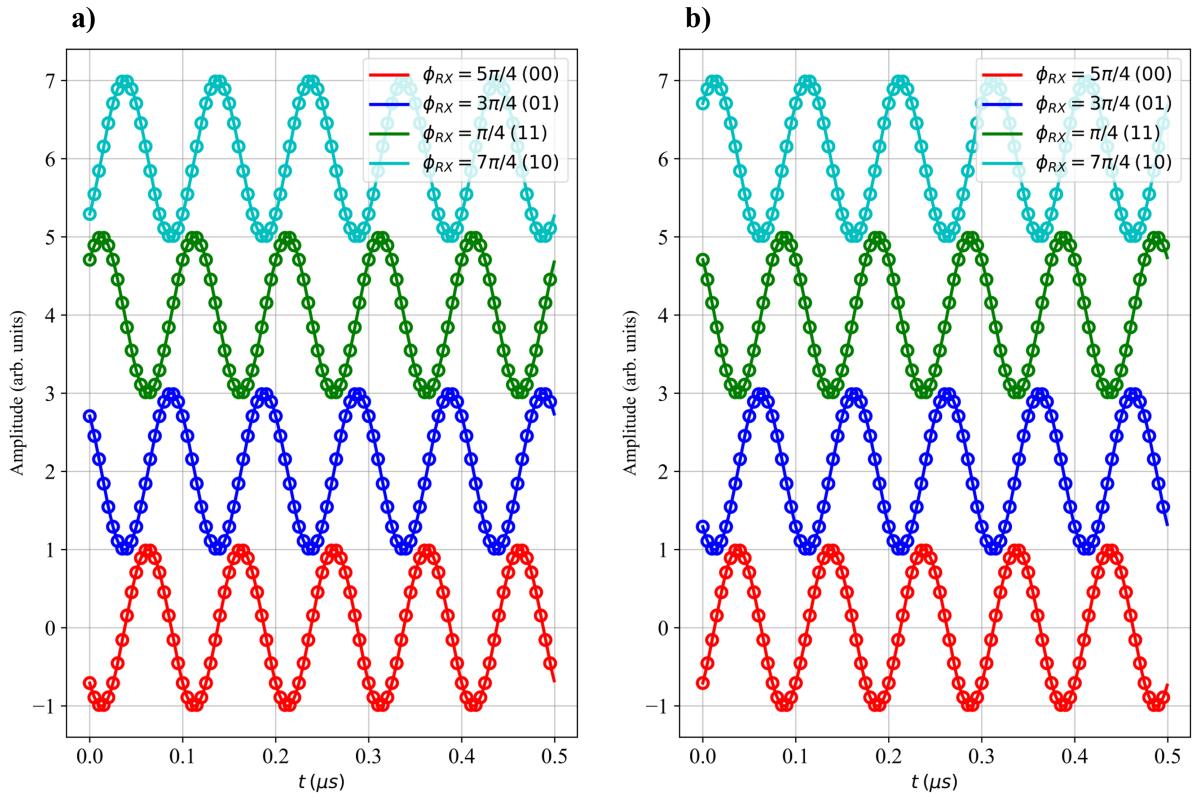


Figure 3.10: A pair of 10 MHz LO signal outputs from the RX1 and RX2 QPSK instances labelled with their corresponding RX phase. a) The in-phase (I) signal. b) The quadrature (Q) signal.

3.2.2 Pulse programmer

The pulse programmer module was responsible for the timing of each NMR experiment. The timer was reset at the start of each experiment and was incremented by 5 ns on every positive edge of the system clock. The resolution of the timer was defined by the inverse of the 200 MHz system clock. The timer value was stored in a 64-bit wide register, which allowed for a maximum experiment duration of over 500 years, clearly sufficient for this application.

TTL Pulses

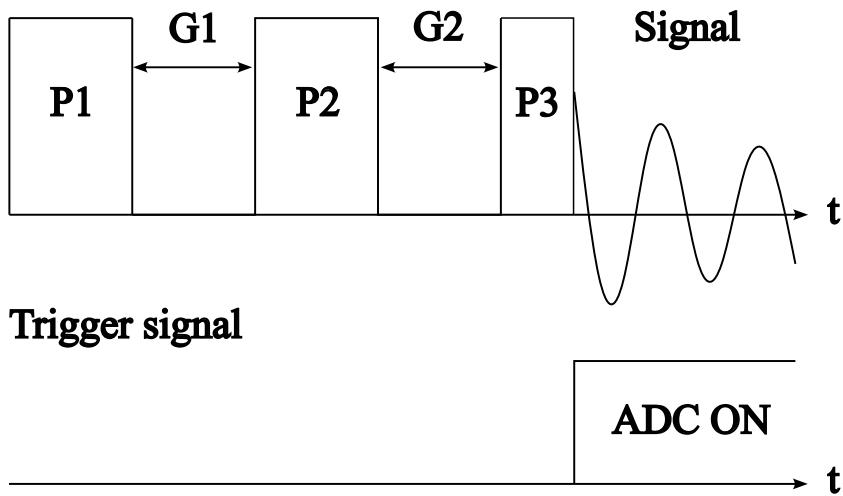


Figure 3.11: A diagram showing the output of the pulse programmer module. P1, P2 and P3 are the TTL pulse lengths. G1 and G2 are gap lengths between pulses. The trigger signal is used to begin data acquisition at the start of the NMR signal.

Figure 3.11 shows an example diagram of the pulse programmer outputs during an NMR experiment. P1, P2, P3 are TTL pulses to enable the amplifier and TX gate, whilst G1 and G2 are the gap durations between pulses. These values are stored in the user registers and shown in 3.4. The trigger signal was used to enable the RX gate and is shown on the bottom axis. This signal activated the acquisition logic and was used to protect the sensitive ADC electronics from the high-power RF pulses. The trigger signal was always located after

the final TTL pulse, independent of the total number of pulses. Additionally, the trigger could be delayed in order to capture events a long time after the final pulse.

The pulse programmer was used to generate the TTL pulses for controlling the external amplifier providing the RF pulse to the NMR circuitry. The module supported the generation of up to three TTL pulses with customisable duration. The gap duration between pulses was also variable. The values of the TTL pulse and gap duration were taken from the FPGA user registers at the start of each experiment. The TTL pulses were represented as a square wave, where a value of ‘1’ represented times where the amplifier was enabled. Programmatically, they were generated using an IF statement, comparing the value of the timer to the values from the registers. This led to the FPGA tools synthesizing this using a single multiplexer component. Figure 3.12 shows a selection of example TTL pulses sequences generated over the timescales in the specification.

The amplifier enable signal also functioned as the TX gate. The RF output signal was multiplied by the amplifier enable signal which ‘gated’ the RF output signal, and meant the RF signal was fixed to ‘0’ at times outside the TTL pulses allowing for best performance from the amplifier. Figure 3.13 shows a 213 MHz RF output signal gated by the TTL pulse signal. Figure 3.13c shows the spectrum of the gated section of the signal, showing a single 213 MHz component.

The TTL pulses also controlled the TX and RX phase shifts applied to the transmitter signals, which allowed for a different phase shift for each pulse. This was done by storing the TX and RX shifts in an array. A pointer was initialised to the first entry in the arrays and was used to designate the active TX/RX phase for a given pulse. The pointer value was incremented after each TTL pulse, stepping through the phases. The active TX/RX phase was sent to the QPSK module where it was used as an input for the TX, RX1 and RX2 QPSK instances (ϕ_{TX} and ϕ_{RX} in figure 3.8).

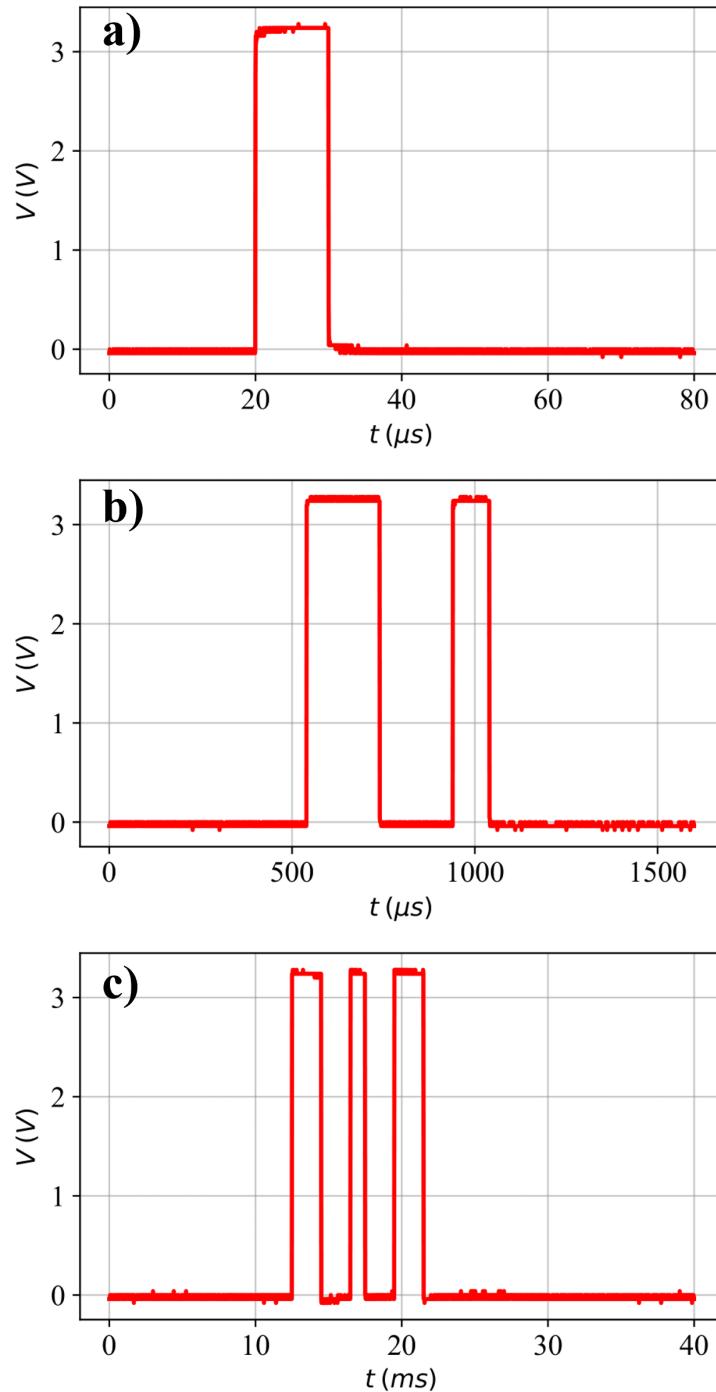


Figure 3.12: A set of TTL pulse sequences generated by the pulse programmer module with the following parameters a) $P1=20 \mu s$. b) $P1=200 \mu s$, $G1=200 \mu s$, $P2=100 \mu s$. c) $P1=2 \text{ ms}$, $G1=2 \text{ ms}$, $P2=1 \text{ ms}$, $G2=2 \text{ ms}$, $P3=2 \text{ ms}$.

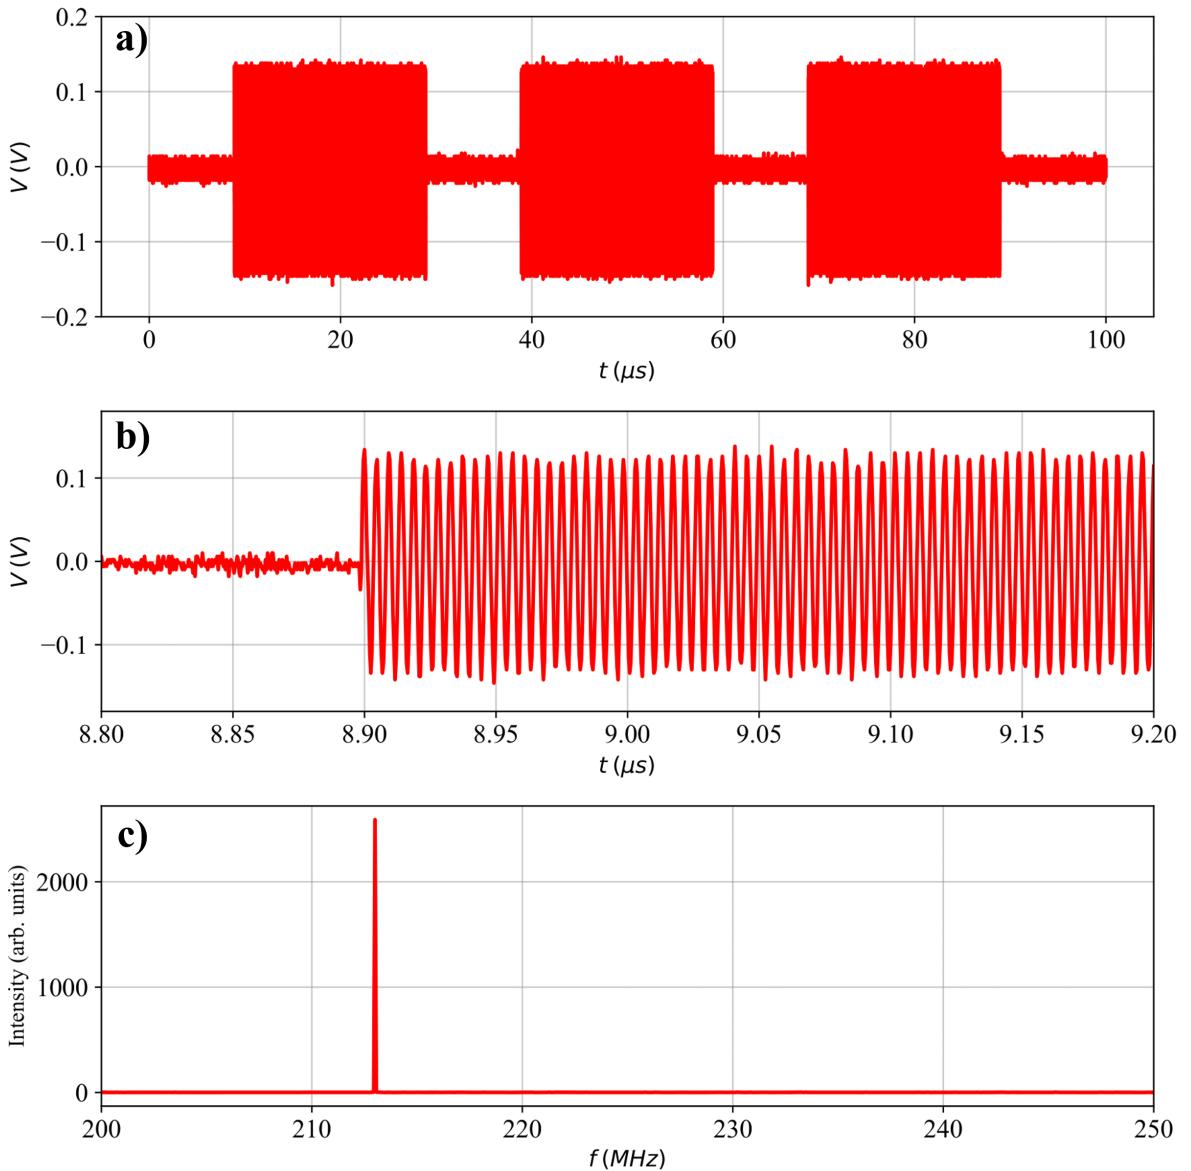


Figure 3.13: a) A 213 MHz RF output after being gated by the TTL pulse sequence. b) A zoom-in of the start of P1. c) The spectrum of the signal during P1. There is a single component at 213 MHz as expected.

Finally, the pulse programmer was used to trigger the data acquisition logic. The FPGA was configured to level trigger on an external signal provided through the TRIG IN port. This trigger signal was generated after the last TTL pulse by the pulse programmer

module and connected through the internal logic to the external trigger input. This meant the trigger signal was internal to the device. The trigger signal was also used as the RX gate which enabled the DAC. The DAC was only enabled within the trigger pulse, which ensured that the DAC was protected from the application of the high-powered RF signal to the NMR set-up. The start time and duration of the trigger pulse were customisable.

3.2.3 Receiver

The receiver section consisted of 3 modules: a quadrature mixer, and a decimation module formed of FIR LPF and downampler sub-modules (not shown in figure 3.3). It was responsible for taking the input signal from the ADC, mixing and downsampling the signal before depositing the resulting signals into the FPGA DRAM. The NMR spectrometer receiver chain was discussed in further detail in section 2.2.3. The data within the DRAM was then extracted to the host PC using the C++ API provided by SP Devices [140]. The NMR signal was digitized using the front-panel DAC with a sampling rate of 800 MHz and precision of 14 bits, generating four parallel 14-bit input samples every clock cycle. Therefore, the mixer and decimation operations had to use super-sampling techniques in order to process the parallel inputs correctly.

The quadrature mixer was implemented using a set of parallel multiplier IP cores [141]. This module had to perform eight individual 14-bit multiplication operations, with four operations coming from multiplying the NMR signal with each LO signal (I and Q) individually. A set of eight Multiplier IP cores were used to implement this process. The IP cores were used to ensure an efficient architecture for the multiplication operation and to ensure the synthesis tools made use of the onboard DSP48 modules. The IP cores generated a pipelined architecture for each multiplication operation, optimizing the timing of the module to ensure timing closure was met.

Figure 3.14 shows a schematic diagram of the output of the mixer module. A dummy input signal was generated using a function generator and the mixer outputs were sent to the FPGA DRAM, bypassing the FIR filter module. The figure shows a mixed signal with a pair of frequency components located at the sum and difference frequencies of the input and LO signals as expected. However, there is also the presence of an additional frequency component matching the frequency of the LO signals, f_{syn} . This phenomenon is known as LO leakage. For this application, the presence of the leakage signal is not significant. This is because the rest of the receiver chain only requires the difference frequency component of the mixer output. For a typical condensed matter NMR experiment, the leakage signal will be of the order 100 MHz however the difference frequency will be of the order 100 Hz. Therefore, for all operational frequencies the subsequent FIR LPF stage will always filter the leakage signal.

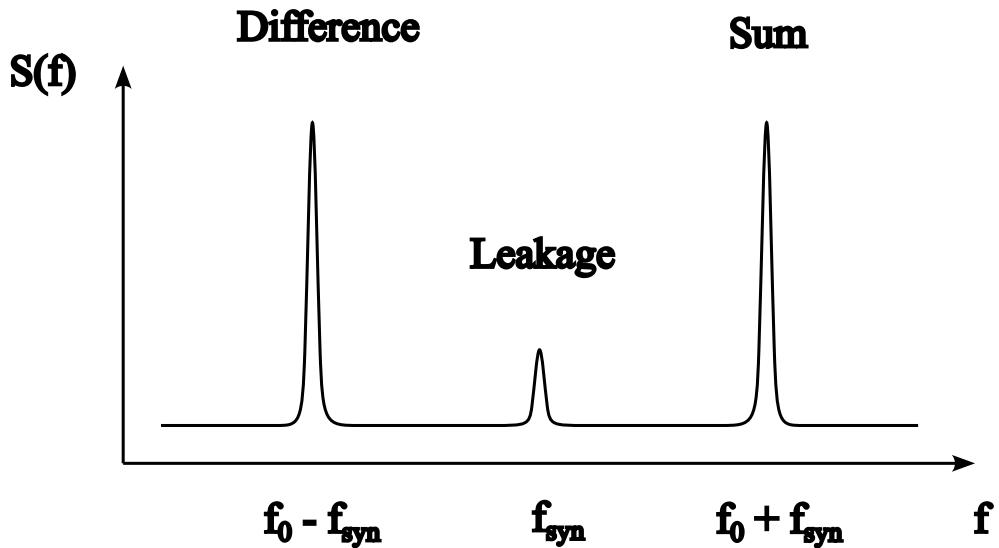


Figure 3.14: A schematic diagram of the output spectrum from the mixer module. The signal contained components at the sum and difference frequency as expected with an additional leakage frequency component at the LO frequency.

The FIR filter and downsampling modules were combined to implement a decimation operation describing in section 2.1.3. The decimation chain was required to reduce the

sampling rate of the NMR signal from 400 MHz to 12.5 MHz which was suitable for the host PC. Five separate decimation stages were implemented, each stage implemented a decimation of factor two. This allowed for halfband LPF FIR filters to be used which increased the computational efficiency of the decimation modules.

The first two decimation stages required supersampling techniques as the sample rate of the signal was larger than the system clock frequency. Consequently, the first two stages required parallel filtering of the input signal to process the set of input samples on each clock cycle. For example, in the first filter stage there were four output samples to match the four input samples on each clock cycle. A filter stage that takes in n new inputs and needs to produce n parallel outputs on each clock cycle will be referred to as n -parallel. Therefore, the first and second stages of decimation used a 4-parallel and 2-parallel FIR LPF respectively. Between each filter module, there was a downampler which removed every other sample from the filter output. For 1-parallel (conventional) filters, the FIR Compiler IP core was used for the implementation of the filter A custom approach was required for the parallel filters as the IP core didn't support parallel filtering.

The custom parallelised FIR filter design used a delay line to store a bank of input samples. The length of the delay line was required to be long enough so there were enough input samples for each parallel filter to calculate its output sample. The number of inputs required for the filter to compute its output is equal to the number of taps of the filter. A schematic diagram of a 5 tap 4-parallel filter with a 8 samples delay line is shown in figure 3.15. A set of input samples, $X[n]$ are stored within the input delay line. They are used as the input to a set of n filters producing the output samples $Y[n]$. Here $X[4n]$ indicates the first sample (in time) to be shifted into the delay line. The coloured arrows indicate the inputs used by each filter

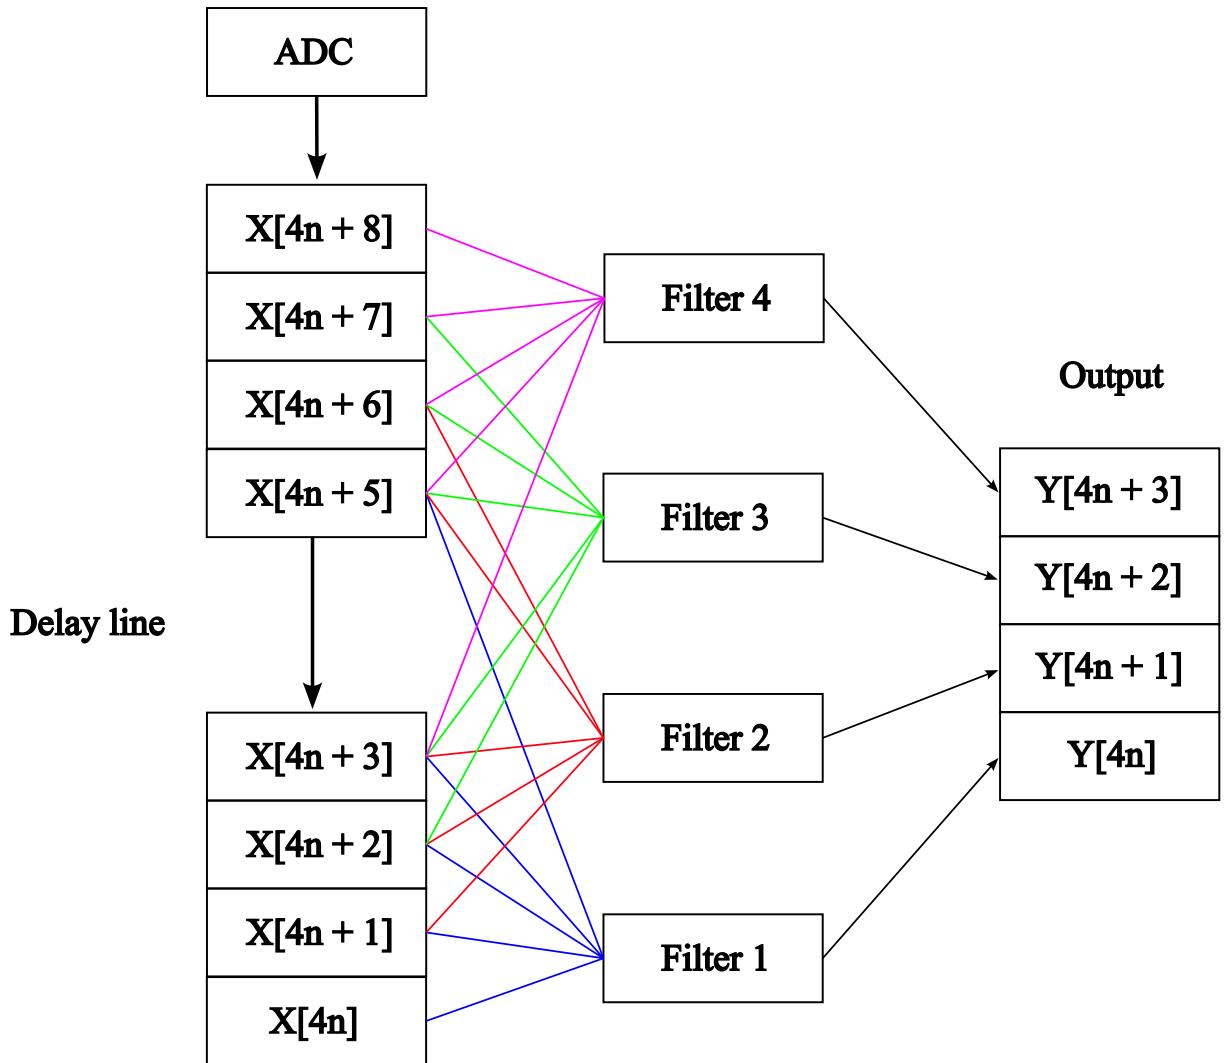


Figure 3.15: The initial state of a 5-tap 4-parallel FIR filter. The filters act on a set of input samples, $X[n]$, and produce output samples $Y[n]$. The coloured lines denote the samples being used for each separate filter operation.

Figure 3.16 indicates the same filter logic after one clock cycle has elapsed. The input samples are shifted in blocks of four during each clock cycle because it is a 4-parallel filter. The coloured arrows are fixed to specific positions in the delay line therefore the filter inputs are updated as the inputs are shifted. This produces the four new output samples for this clock cycle.

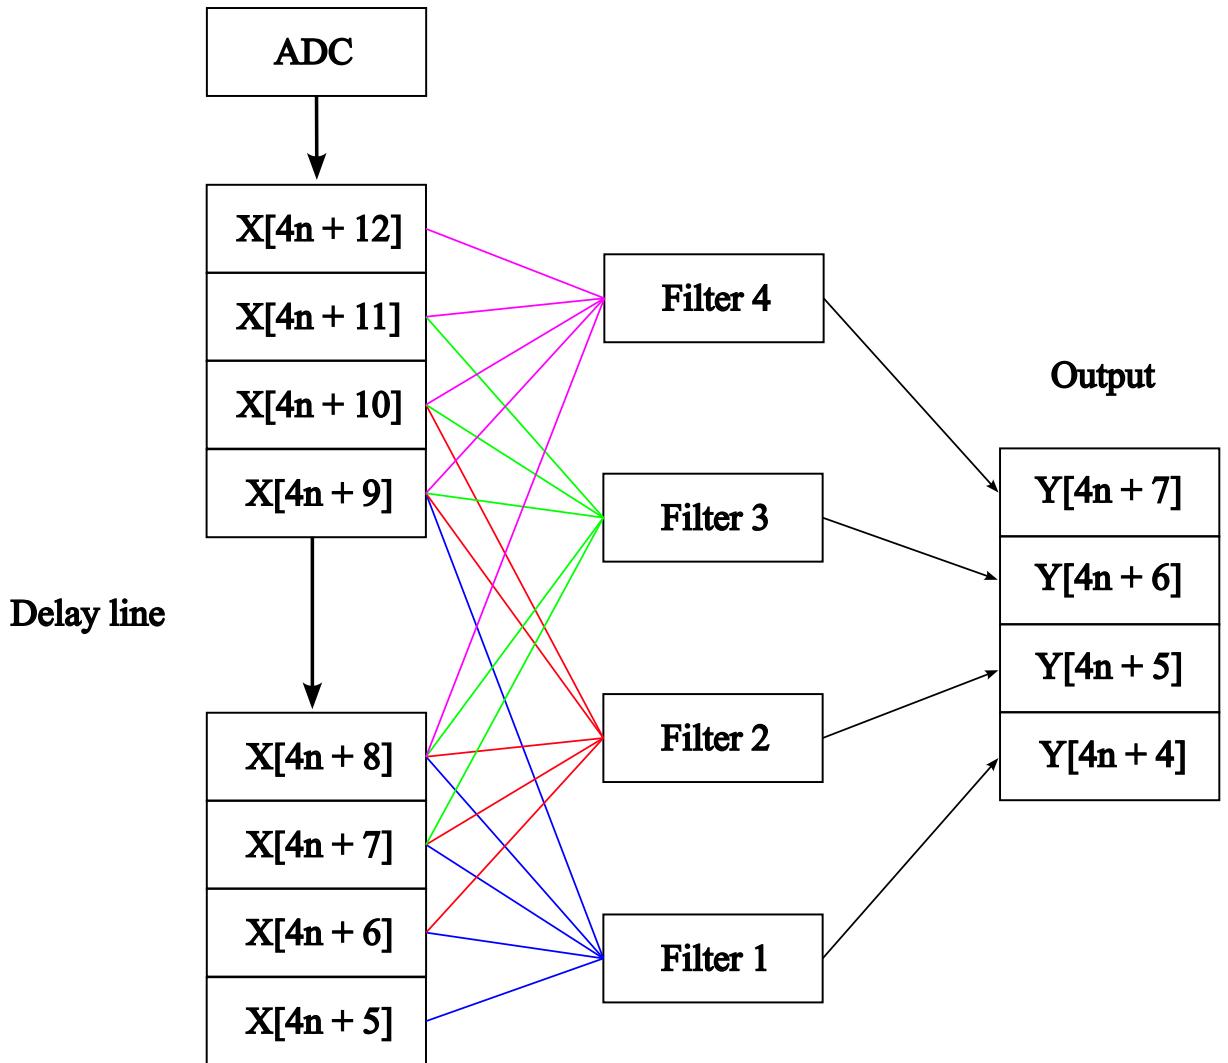


Figure 3.16: The next state of the 4-parallel FIR filter shown in figure 3.15. The inputs have been shifted by 4 samples.

As shown in the figures, the delay line must be long enough to populate all of the filter inputs at once. The set of inputs of each filter are offset by one sample to ensure all samples in the delay line are used. The length of the delay line was calculated using the number of taps and the number of parallel filters. The FPGA implementation of an FIR filter can be realised as a set of multiplications and additions. This can be demonstrated by considering equation 2.13 which states the filter output is given by a convolution operation. A convolution is represented graphically by the overlap in area when one function is translated

over the other. In this case, time passing causes new samples to replace the current samples in the delay line. This has the same effect as translating the input signal over the filter response (formed from the set of filter coefficients). Therefore the required mathematical operations are the multiplications between the input samples and filter coefficients and then a summation to generate a the filter output sample for each parallel filter (equation 2.15).

The 4-parallel FIR filter architecture was simulated in Python before implementing on the FPGA in order to verify its functionality. The simulated 4-parallel filter was a 29-tap 200MHz LPF where the taps were generated using the windowing method outlined in section 2.1.4. The delay line was simulated using a queue data structure and the FPGA clock cycles were implemented using a for loop. On each iteration of the loop, four new inputs were added to the start delay line and four samples were removed from the end. The output samples were generated only by using multiplication and addition operations to best simulate the FPGA logic. Figure 3.17 shows the result of filtering an input signal with frequency components of 150 MHz and 250 MHz in order to demonstrate a clear scenario where one component should be filtered. The 250 MHz component has been filtered to an amplitude consistent with the filter response of the 29 tap filter demonstrating the desired behaviour (more taps would fully attenuate the peak). Furthermore, the filtered signal was compared to the in-built convolution function from the SciPy library [142], which whilst not shown, produced an identical output.

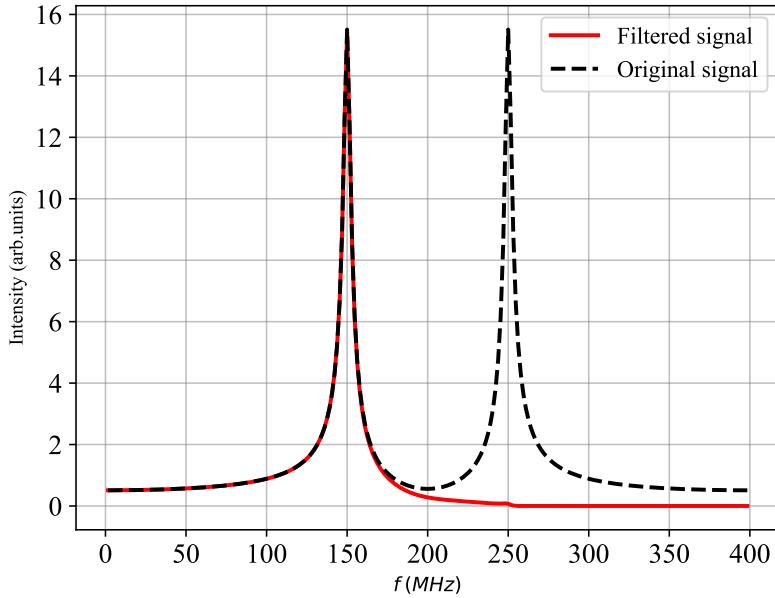


Figure 3.17: A python simulation of a 29-tap, 4-parallel 200 MHz FIR LPF applied to an input signal with frequency components of 150 MHz and 250 MHz.

The Verilog implementation of the FIR module emulated that of the Python simulation. The multiplication and addition operations were performed using IP cores to provide an efficient implementation. The module took in the input signal, coefficient values and the level of parallelisation as inputs. This meant that it was simple to reuse the modules across the design by simply changing the coefficient values on each instance of the module. The coefficient values were calculated using the windowing method explained in section 2.1.4 in a standalone Python script. In Verilog, the coefficient values were stored as 18-bit signed numbers, with 17 fractional bits to provide adequate precision.

The first filter in the chain was a 19-tap 200 MHz LPF designed to remove aliasing during the first decimation stage. This was done to decimate the sample rate from 800 MHz - 400MHz. This was followed by a downsampler that removed a pair of samples from the FIR filter output, so the next filter in the chain was a 2-parallel filter.

Figure 3.18 shows the frequency response of this filter. It was measured by sweeping the LO signal frequency from 0 MHz to 400 MHz and sending the quadrature mixer leakage signal through the filter. The filter output bypassed the downsampling module and was sent to the DRAM to be retrieved by the host PC. The response of the filter was as expected with a cutoff at 200 MHz.

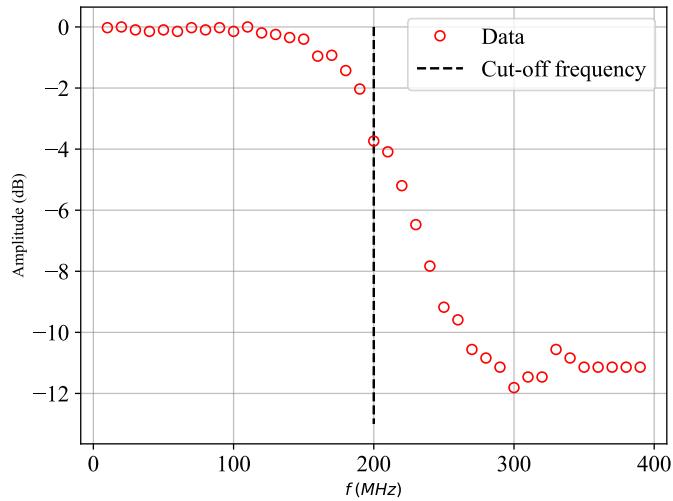


Figure 3.18: The frequency response of the 200 MHz FIR LPF implemented on the SDR14.

The filter response would be sharper by adding more taps at the cost of more computation. This was unnecessary in this case. As there were a chain of FIR filters, one for each decimation stage, any aliasing introduced from frequencies above the cutoff frequency of the earlier filters is removed by filters further along the chain. Therefore, early filters in the chain used fewer taps than the later filters, saving resources compared to using an identical number of taps for all filters.

3.3 Python control software

A Python-based graphical user interface (GUI) application was developed in order to provide the user with an interface to control the SDR14. The application was developed in line with the following aims

- Provide a simple, easy-to-use interface to allow users with little prior experience to run NMR experiments.
- Allow automation of NMR experiments so long experiments could be performed with minimal user input over the experiment duration.
- Generate a record of performed NMR experiments containing the NMR parameters and sample properties to facilitate simple repeatability.
- Provide a live plot of the acquired data throughout the experiment, allowing the user to view the status of the data in real-time.
- Integrate temperature and magnetic field control of a Quantum Design PPMS within the NMR experiments.
- Return the apparatus to a safe state when errors occur, to minimize damage to components in the system.

This application was built using the PyQt5 [143] framework providing access to a set of familiar widgets (buttons, input fields, tables) used across many applications. Furthermore this framework is compatible with all major desktop platforms allowing it to be ran on all PCs. The host PC interacted with the SDR14 via a C++ API provided by SP Devices [140]. The API was accessed using the built-in ctypes [144] package which provided access to a range of C compatible data types and calling functions from a .dll file. NumPy [145]

and PyQtGraph [146] were used for manipulation of the acquired data from the SDR14 and visualising the data respectively. In this instance PyQtGraph was used over the more popular plotting package Matplotlib [147] due to Matplotlib causing lag when plotting live data. PyQtGraph plotting was efficient enough to live-plot the incoming data, with minimal lag when plotting up to 10,000 data points per second. The MultiPyVu package [148] was used to integrate control of the PPMS into the program. This package was developed by Quantum Design and was recommended over using standard GPIB commands [149].

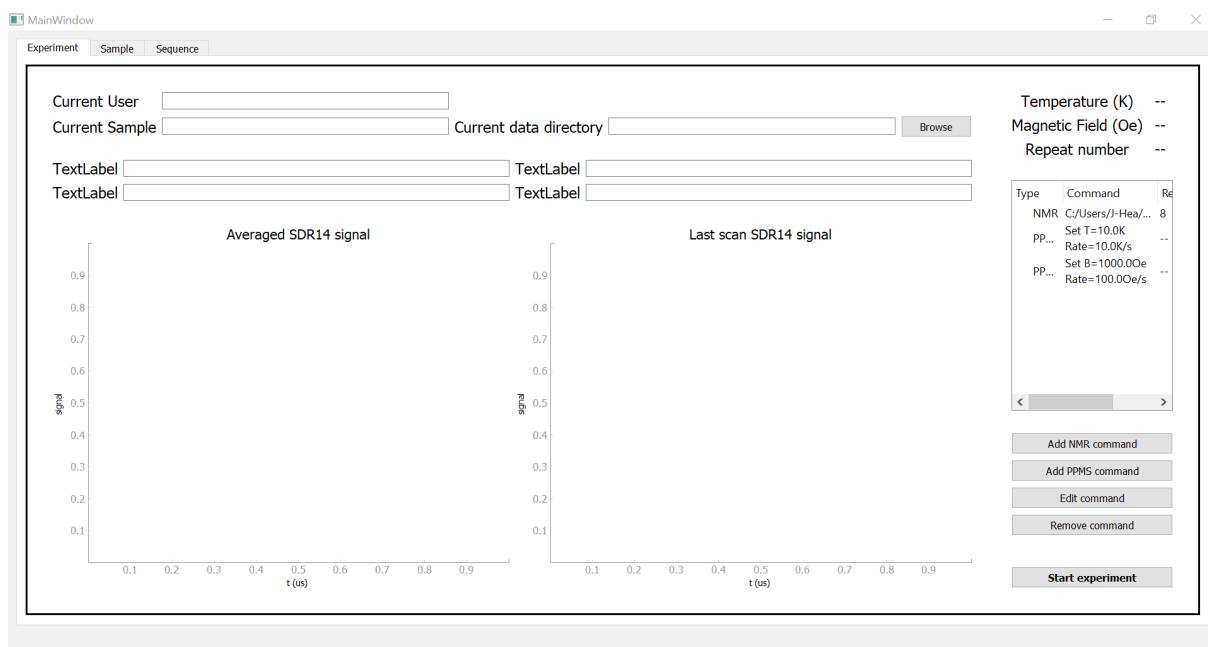


Figure 3.19: The main panel of the Python control GUI. The axes display a live-plot of acquired data during an experiment. The experimental command list is the box in the right panel. The environmental conditions are shown in the top right box, updating once a second. The buttons on the lower right side are used to modify the experiment list and start the experiment.

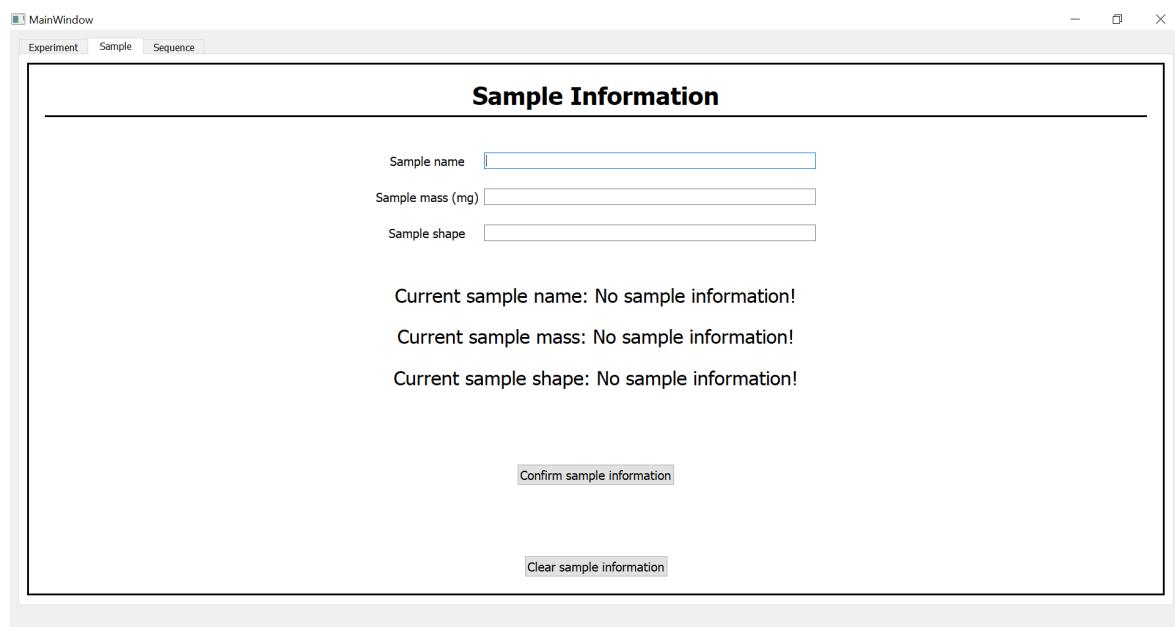


Figure 3.20: The sample panel of the Python control GUI. This form allows the user to record a sample name, mass and shape.

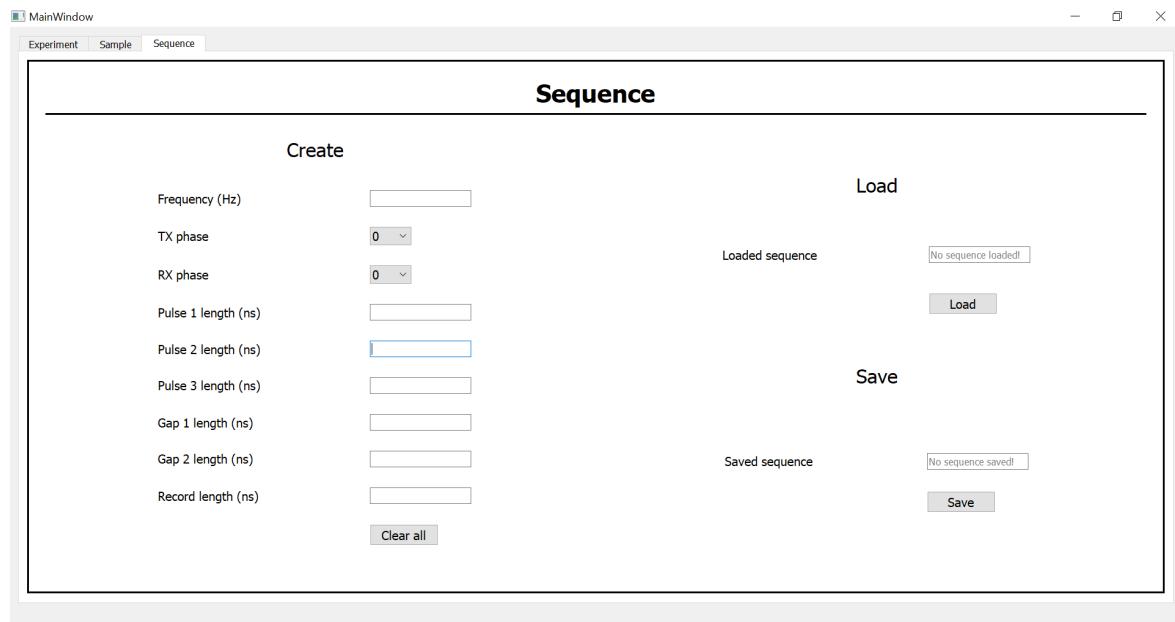


Figure 3.21: The sequence tab of the Python control GUI. This form allows for entry of parameters for an NMR command. These correspond to the parameters in table 3.1.

Figures 3.19, 3.20 and 3.21 show screenshots of the three main tabs of the GUI. The main tab displayed the current user, sample and data directory where the NMR data will be saved. The sample and data directory must be filled out before the experiment is started. The upper-right box displayed the current state of the experiment as it was running. It showed the temperature and applied magnetic field from the PPMS and the current repeat of the NMR sequence. The centre-right panel shows the command list which defined the experiment. The user adds NMR or PPMS commands to the list using the set of buttons in the bottom right corner. The NMR commands required specification of an existing sequence file generated from the sequence tab of the GUI. The PPMS commands requested the set point and change rate of either the temperature and magnetic field. When the experiment is started, each command runs from top to bottom, with the current command highlighted in green to inform the user of the experiment progress. The set of axes are used to live plot the acquired data from the SDR14 during the experiment. The right panel plots the data acquired from the latest repeat of the NMR sequence and the right panel displays the averaged signal.

The sample information tab shown in figure 3.20 was the place the user entered sample information including the sample name, mass and shape. These values were stored in the output data files so the user could easily recall the samples used during an experiment. The values entered here auto filled the sample textbox on the main display described in the previous paragraph.

Figure 3.21 shows the form used to define the NMR parameters for a NMR sequence and modifying existing sequence files. These values correspond to the variables stored in the user registers of the FPGA discussed in section 3.1. These values were validity checked in order to prevent damage to the SDR14 and keep the instrument in a predictable state. Finally, the sequence could be saved for ease of use when performing future experiments.

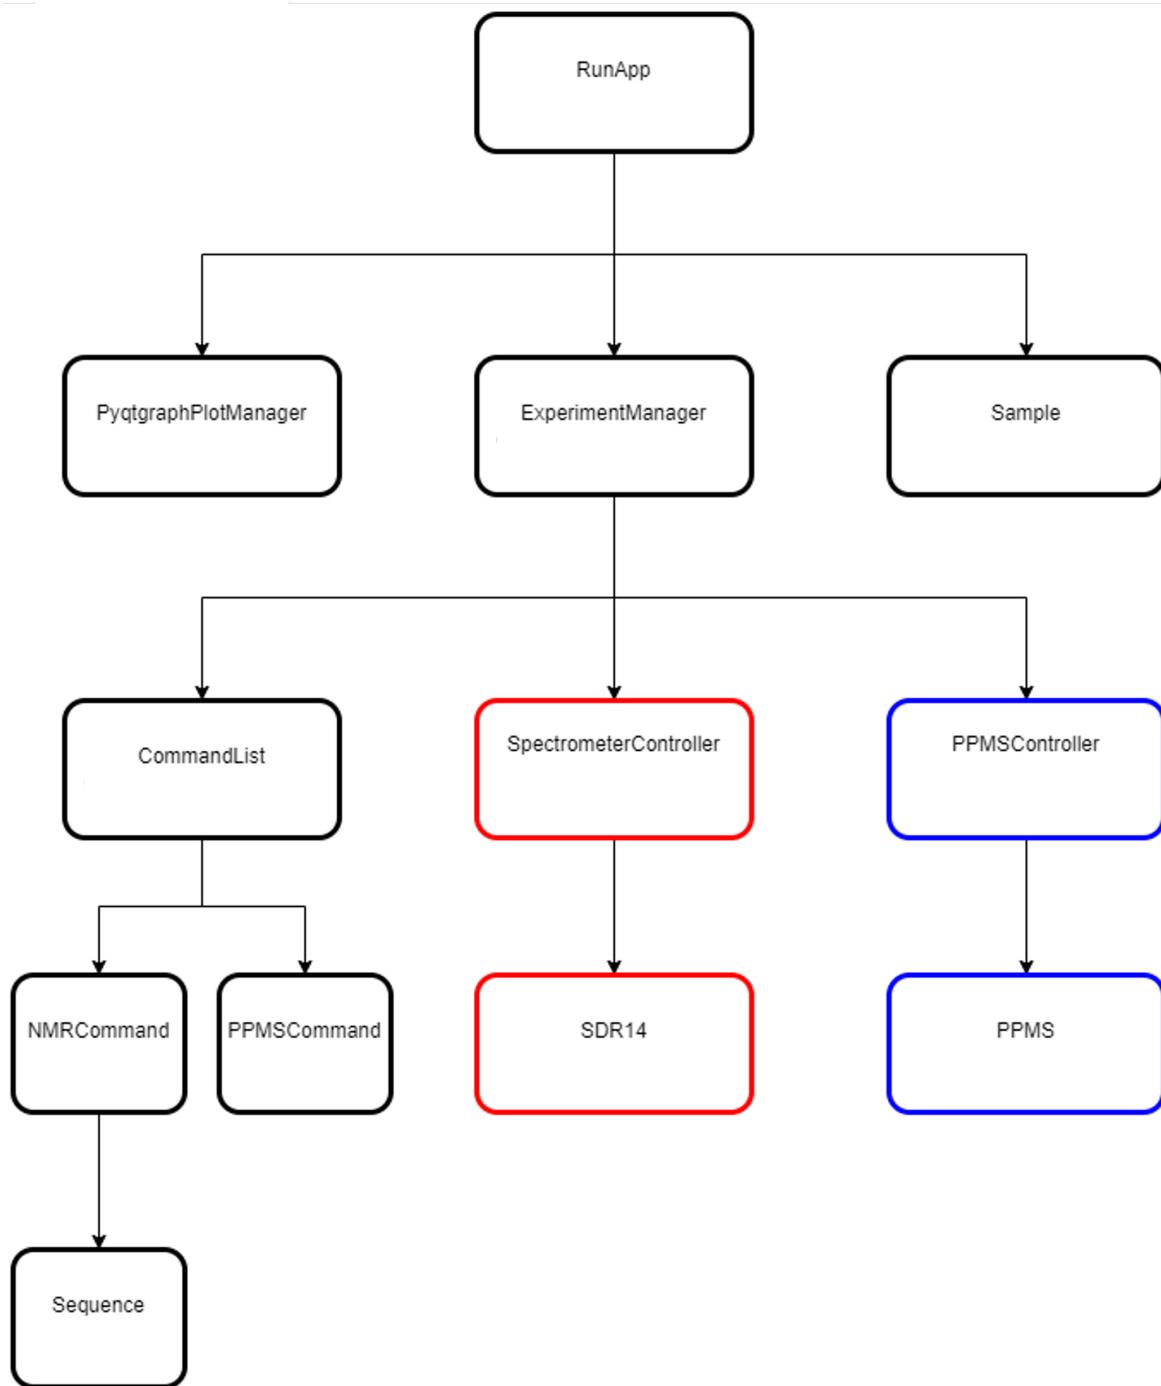


Figure 3.22: An overview of the code structure used for the Python GUI. The red and blue entities are executed on the NMR and PPMS worker threads respectively. Black entities are executed on the main thread.

A diagram showing the code structure of the application is shown in figure 3.22.

The code was developed using object-oriented programming (OOP) principles to facilitate simpler development and testing. Each box in the diagram represents a class and the arrows represent the class hierarchy of the application.

The application used multithreading to concurrently control the GUI, SDR14 and PPMS. Multithreading was required as each entity would lock its thread when operational. If a single thread were used this would lead to responsiveness issues, for example the GUI would be unresponsive during an NMR measurement as the thread would be processing the SDR14 commands. Multithreading was implemented using PyQt5's inbuilt thread support. Commands and data were sent between the threads using custom event triggers and handlers. This method ensured minimal collision between threads when cross-thread processes were run. For this application, the main thread controlled the GUI and the PPMS and SDR14 threads were designated worker threads. The main thread broadcasted events to execute commands on the worker threads. The SDR14 and PPMS worker threads were configured to only respond to events pertaining to their device. Upon finishing their processing, the worker threads would send data back to the main thread in order to be saved and plotted. The worker threads then entered a 'listener' state, waiting for the next event to arrive.

The Sample and Sequence classes were used to purely store data regarding the NMR sequence parameters and sample properties. These were developed in order to easily keep track of sample and sequence parameters throughout execution of the application. The classes contained validity checks to ensure valid values for any user input data, for example the NMR frequency must be 800 MHz at most, in order to keep the equipment in a predictable state.

The SDR14 and PPMS classes were written as wrappers around the corresponding device API. In the case of the SDR14, this was done to provide a Python interface around the C++ API. This class provided a set of complex SDR14 functions by combining a set

of lower-level functions directly from the API. For example, all FPGA registers could be reset in a single Python function call, instead of having to iterate through all registers using the lower-level functionality. The functionality in these classes was accessed using a parent SpectrometerController or PPMSController class. Essentially, these parent classes were used to combine the required properties for the multithreading architecture with the functionality within the SDR14 and PPMS child classes.

The CommandList class was responsible for maintaining the schedule during an experiment. This data structure contained a list of all commands and their order in which they would be executed during an experiment. It contained two types of command: NMRCCommand and PPMSCommand. These classes stored information relating to commands performed by either the SDR14 or the PPMS. An NMR command stored the NMR sequence parameters and number of repeats. A PPMS command stored the condition to be changed (temperature or magnetic field) as well as the set value and rate of change.

The ExperimentManager class was responsible for running the entire experiment. Every experiment was defined using a CommandList entity containing a set of NMRCCommands and PPMSCommands. This class kept track of the current command during an experiment and controlled the scheduling of commands. This involved sending the current task to its respective worker thread and listening for data from the worker threads in order to determine when the next command should be executed. Finally, it was responsible for sending data to the GUI in order to be displayed to the user. It repeated these processes until the CommandList was traversed. This class handled any mathematical processing of the data and saving.

The PyQtGraphPlotManager was used to implement all operations related to plotting the data. This included generating the plot axes and labels, setting plot colours and style in addition to plotting newly acquired data during the experiment.

3.4 NMR probe

To perform NMR experiments within the PPMS, an NMR probe that integrated with the PPMS hardware was constructed. Quantum Design provided a number of blank probes that were suitable for custom measurements [150]. In this case, they were not used as they didn't have enough space for the addition of the NMR resonance circuit and tunable capacitors, so a bespoke probe was designed and built.

Figure 3.23 shows a photograph of the entire NMR probe. The probe was designed using Autodesk Inventor 2019 and Onshape, an online CAD application. The probe consisted of three sections: the probe head, the probe body and the sample chamber.

The probe was designed in accordance with the following set of specifications

- All materials that were located within the PPMS chamber must be non-magnetic.
- The thermal conductivity of the materials must be low enough to reduce heat transfer from the probe head to the sample chamber.
- The probe head must be vacuum sealed in order to use the PPMS temperature control during an NMR measurement.
- The probe must fit within the 26 mm diameter of the PPMS magnet bore [151].
- The capacitors must be able to be adjusted whilst the PPMS chamber remained under vacuum.

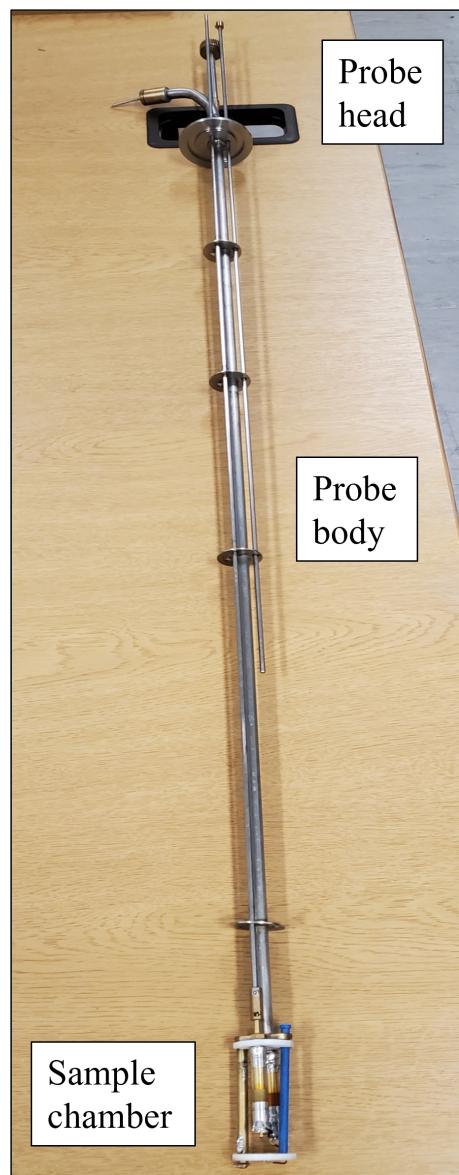


Figure 3.23: An image of the NMR probe designed for use in the QD PPMS. The labels indicate the three main sections of the probe.

The probe head used a modified KF40 flange fitting, which connected to the sample chamber access port on the top of the PPMS. The probe head was made from stainless steel was chosen due to it's thermal conductivity of 50.2 W/mK, which was lower than alternatives such as brass and aluminium which had 205.0 W/mK and 109.0 W/mK respectively [152]. As the probe head remained outside the PPMS at all times, it was at room temperature so a small thermal conductivity was a necessity. A stainless steel tube was welded to the flange, which provided structural support between the probe head and sample chamber. Additionally, the tube contained the coaxial cable that transmitted the NMR RF pulse and detected NMR signal to/from the sample chamber. A COAX CO. LTD SC-219-50/PBC-PBC [153] cable was used as it was non-magnetic and suitable for use at cryogenic temperatures. A BNC connector was soldered onto the exposed end of the coaxial cable (not shown in figure 3.23). Three stainless steel rods were inserted through the probe head to facilitate capacitor tuning for the tuning and matching capacitors and sample rotation. O-rings were inserted on the feedthrough between each rod and the KF-40 flange. A stainless-steel plate was screwed down onto the flange, compressing the set of O-rings providing a vacuum sealed interface whilst still allowing for the rods to be rotated. A picture of the probe head from two different angles is shown in figure 3.24a and figure 3.24b, with the key components labelled.

The probe body was defined as the space between the sample chamber and probe head. It contained a set of stainless-steel baffles welded to the central stainless steel tube in order to centralise the probe location with the PPMS bore and ensure alignment of the adjustment rods. Holes were drilled in the baffles and adjustment rods so liquid helium could freely circulate within the chamber.

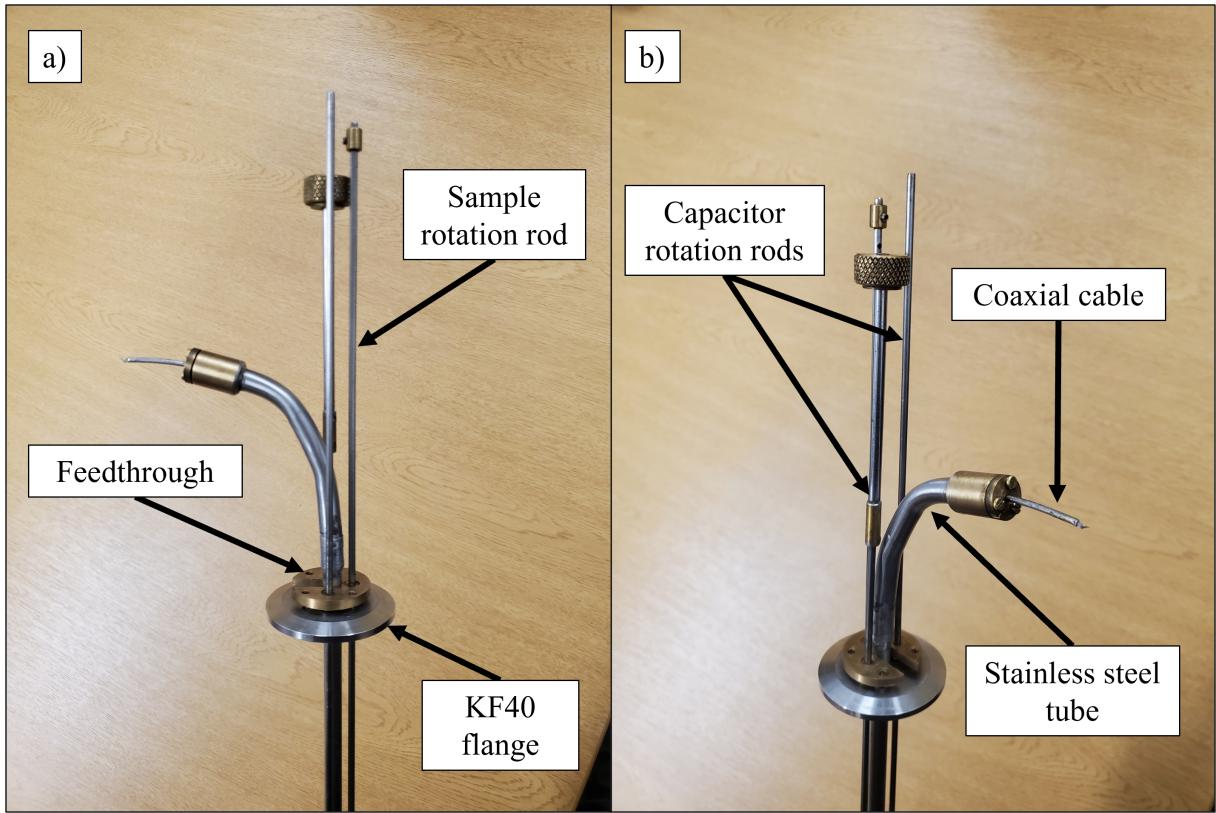


Figure 3.24: An image of the probe head from a) the left side and b) the right side. The labelled components are discussed in the text.

The sample chamber housed the NMR resonance circuitry and the sample and was constructed from components made from PLA and brass. A picture of the sample chamber is shown in figure 3.25 A brass connector was soldered to the bottom of the central stainless-steel tube. A 3D printed PLA structure was connected to the brass using a pressed fit brass rod and PLA rod. A screw and PLA cap were used to fix the brass and PLA rod in place respectively. The capacitors were suspended from the top of the PLA structure using a panel mount screw. The capacitor screw adjustment terminals were aligned with the bottom of the capacitor rotation rods and were connected using a brass and PLA interface. The brass interface provided a ground point for the resonance circuit mounted on the probe. This allowed for capacitor adjustment from the exposed rods on the probe head. The sample rotation rod was terminated just above the PLA structure since this configuration did not

support sample rotation. The sample was suspended in the NMR coil (not shown) in the area highlighted in grey in figure 3.25. One side of the coil was soldered to the capacitors and the other was soldered to the bottom of the brass support rod.

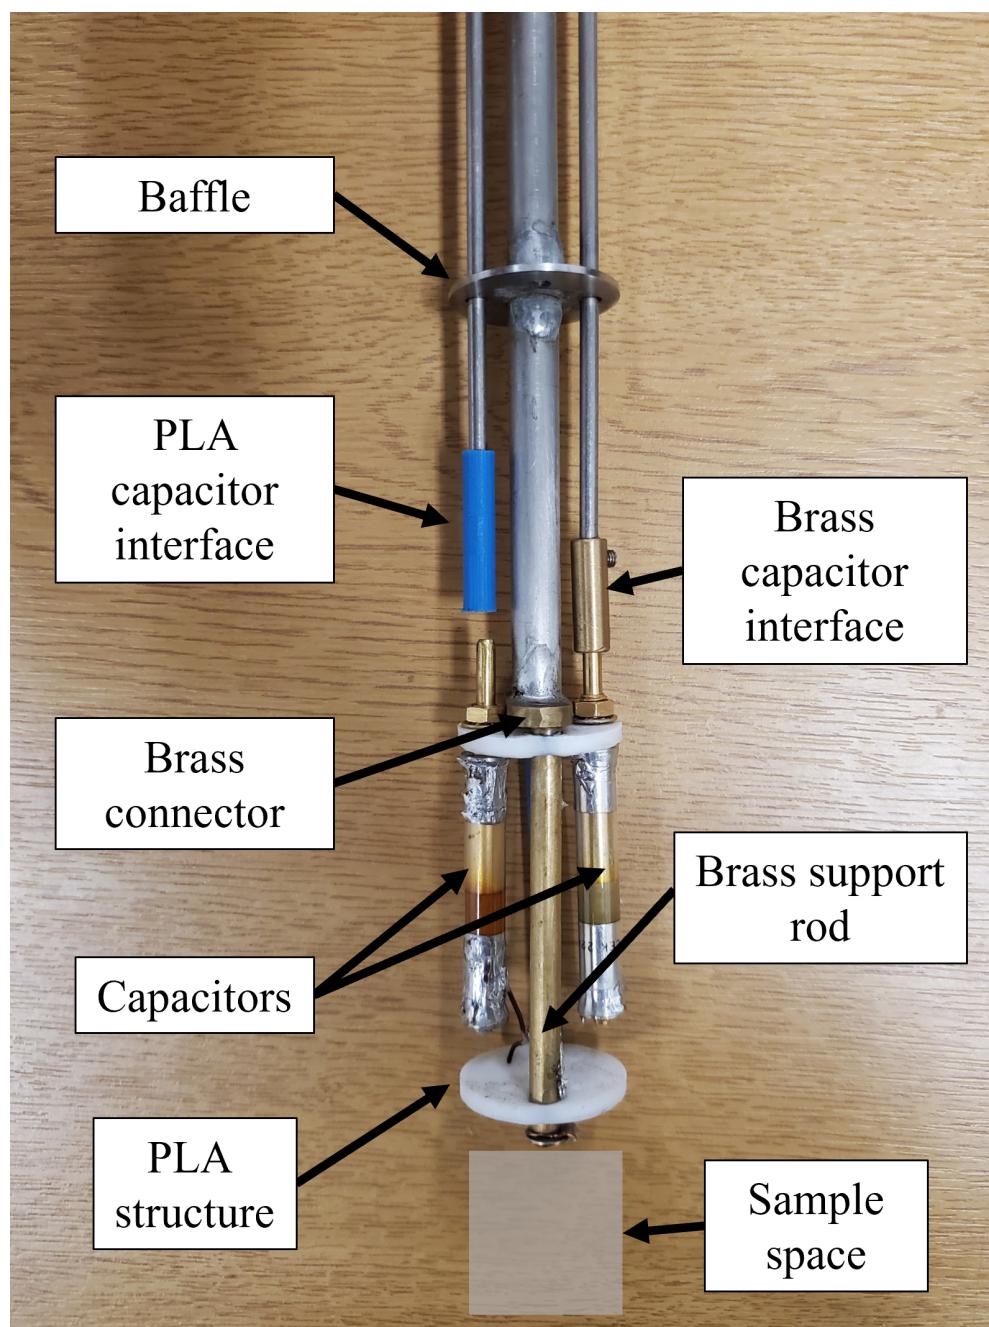


Figure 3.25: An image of the sample chamber. The region where the sample and coil are mounted is highlighted in grey.

3.5 ^{59}Co NMR

^{59}Co powder is a ferromagnetic compound with a strong internal magnetic field of approximately 22 T [154] caused by dipolar interaction between the nuclear spins and electron spins. This is due to the coexistence of face centred cubic (fcc) and hexagonal close packed (hcp) structure displayed by ^{59}Co [155]. Due to this, zero-field NMR is possible as the internal field acts as the static field B_0 . Zero-field NMR signals from ^{59}Co powder were first observed by Gossard and Portis displaying a large resonance peak at 213 MHz [156]. This peak is attributed to presence of magnetic domain walls in the fcc components of the powder sample. Additional measurements of the ^{59}Co powder spectrum have found multiple resonance peaks across the range 210 MHz - 223 MHz [155]. Each peak can be attributed to the existence of either fcc or hcp magnetic domains and domain walls [155]. Furthermore, ^{59}Co powder displays an enhanced NMR signal, allowing it to be measured at room temperature [155].

This section will present ^{59}Co NMR measurements used to verify the performance of the FPGA-based NMR spectrometer. It will include the measurements used to characterise the response of the quarter-wave transformer and NMR tank circuit. The full setup used for these measurements was equivalent to that in figure 2.24 in section 2.2.4. The measurements were performed at room temperature at zero field with the sample mounted in the NMR probe. The NMR frequency was chosen to be 213 MHz to match the peak in the spectrum and thus providing the largest signal. The test sample had a mass of 20 mg and was contained in a 5 mm gelatin capsule and sealed with teflon tape.

3.5.1 Quarter wave transformer

Figure 3.26a shows the quarter wave transformer used in the cobalt powder NMR setup. It was tuned for a frequency of 210 MHz, therefore supported frequencies in the range 190

MHz - 230 MHz and was suitable for these measurements. It was constructed using four 1N4148TR high power switching diodes [157] for the pair of crossed diodes and a short piece of RG-58 coaxial cable. The cable length was determined by connecting a longer piece of cable to a tee and the S_{11} phase shift was measured using a vector network analyser (VNA). The cable was shortened until the quarter-wave frequency reached the desired value.

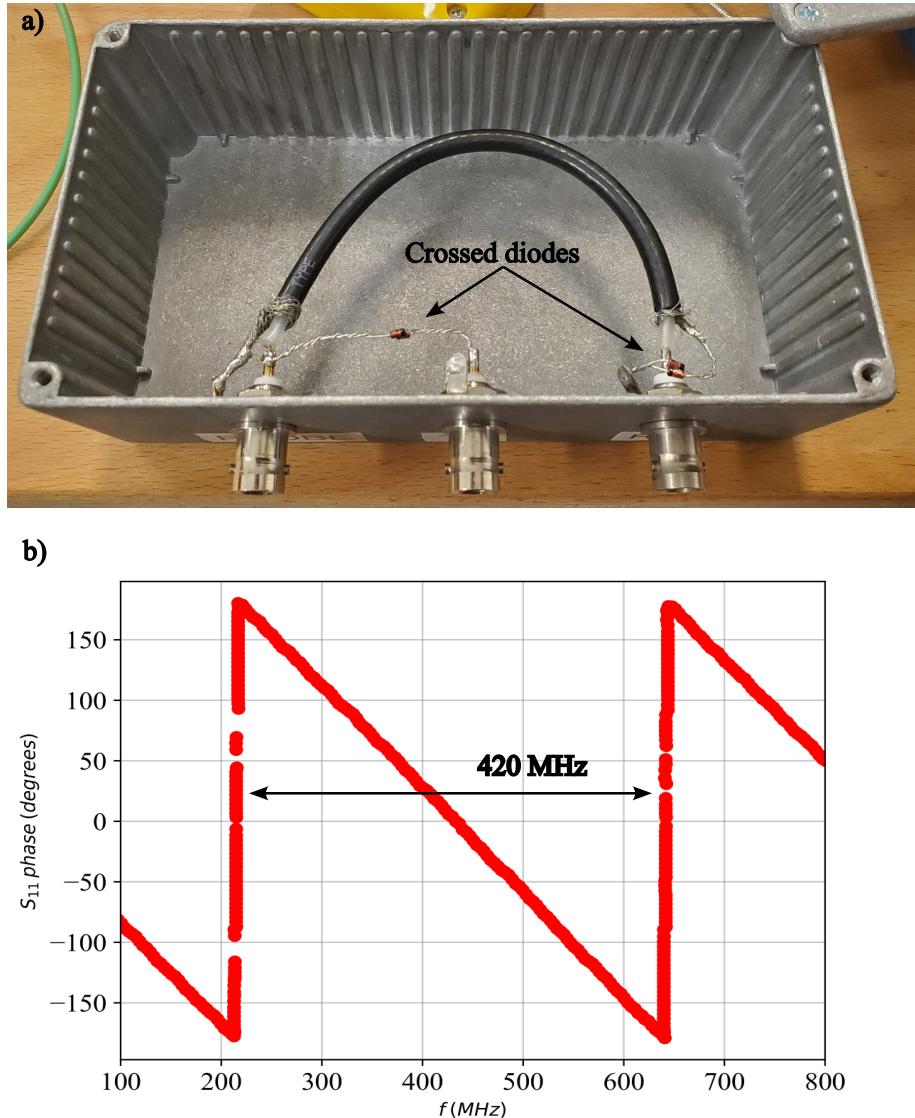


Figure 3.26: a) The 210 MHz quarter wave transformer used to obtain the ^{59}Co powder NMR signals. b) The VNA measurement of S_{11} phase shift of the quarter wave transformer shown in a).

Figure 3.26b shows the VNA S_{11} phase shift measurement for the final length of cable. The quarter wave frequency of the coaxial cable was obtained by calculating the frequency at which a 180° phase shift between the outgoing and reflected RF signal occurred. This was done by dividing the 360° phase shift of 420 MHz by 2 to obtain the desired frequency of 210 MHz.

3.5.2 Resonance circuit

The resonance circuit used on the NMR probe is shown in figure 3.27.

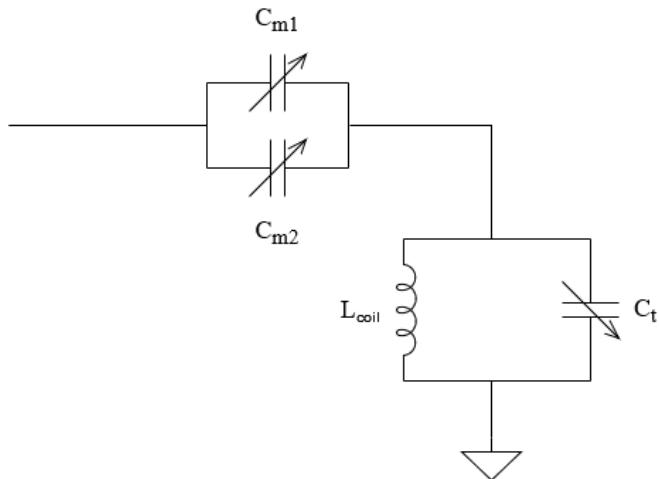


Figure 3.27: The circuit diagram for the resonance circuit used for the ^{59}Co powder NMR measurements.

All of the capacitors were Voltronics NMTM120CEK non-magnetic, glass variable trimmer capacitors with a tuneable range of 1 pF - 120 pF. It was slightly adapted from the circuit described in section 2.2.4 with the modification of a pair of matching capacitors. As the matching capacitors were placed in parallel, the total capacitance was the sum of each capacitor. A second matching capacitor was added as a single matching capacitor was found to be insufficient for impedance matching for frequencies above 200 MHz, when the circuit was mounted on the probe. The resonant circuit was easily matched using a single matching

capacitor when tested on a bench-top test setup which indicated some additional sources of parasitic capacitance from the probe setup. The coil had diameter 5 mm with 3 turns and was made by winding 22 awg single core copper wire around the straw containing the sample capsule.

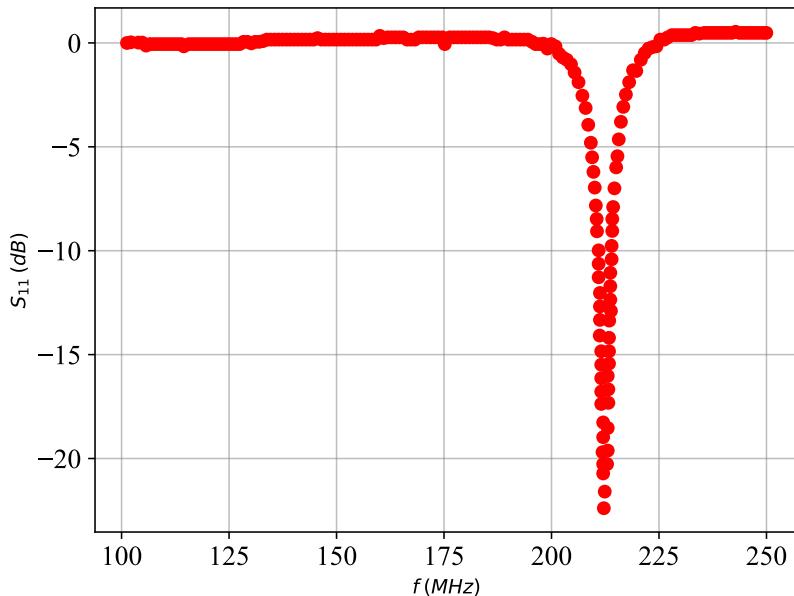


Figure 3.28: The S_{11} measurement of the tuned resonance circuit for the ^{59}Co powder NMR measurements.

Figure 3.28 shows the S_{11} transmission parameter of the probe mounted resonance circuit over the frequency range 100 MHz - 250 MHz measured using a VNA. The circuit was tuned and matched with the coil wrapped around the sample. The trough at 213 MHz indicated a well-tuned and matched circuit. At 213 MHz, $S_{11} = -23\text{dB}$ which corresponded to a 99.5% power transmission into the resonance circuit from the VNA signal generator.

The Q-factor of the tank circuit was calculated from the data shown in figure 3.28. It was calculated using the ratio

$$Q = \frac{\Delta f}{f_0} \quad (3.8)$$

where Δf is the full width at half maximum of the resonance peak and f_0 is the resonance frequency. The Q-factor of the circuit was calculated to be $Q = 84 \pm 1$. A typical value of Q for an NMR set-up is $Q > 50$ [158]. Hence, the circuit met this condition and was suitable for NMR measurements.

3.5.3 ^{59}Co spin echo

The ^{59}Co spin echo signal was acquired using a standard spin echo sequence described in section 2.2.2. The pulse sequence had parameters $f_{syn} = 213$ MHz, $\text{P1} = 1 \mu\text{s}$, $\text{G1} = 5 \mu\text{s}$ and $\text{P2} = 2 \mu\text{s}$. The signal was averaged over 128 repeats with a recycle delay of approximately 1 s between measurements. The sample had a mass of 20 mg and was contained within a gelatine capsule. Measurements were performed with the sample mounted on the custom NMR probe (outside of the PPMS) using the electronics characterised in the previous two sections.

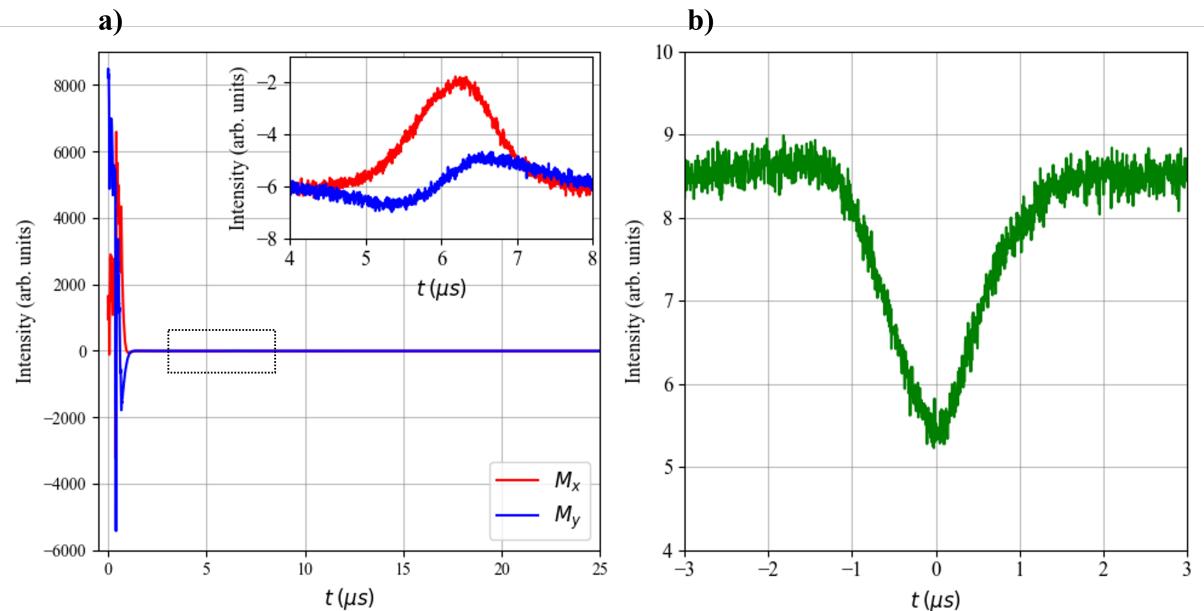


Figure 3.29: a) The raw NMR signal obtained for the ^{59}Co powder sample. The red and blue lines indicate the M_x and M_y component respectively. Inset) A zoom-in of the dotted region. b) The magnitude of the spin echo signal obtained from the raw data.

Figure 3.29a shows the raw data for the M_x and M_y components of the measured signal. This shows the full record of acquired data after averaging. This figure demonstrates the presence of the ‘leak-through’ from the RF excitation pulse on the results. The presence

of this signal then limits the minimum viable value for $G1$, as this high-power signal would fully obscure the echo signal. The inset of figure 3.29a shows a zoomed in region around $t = 6 \mu\text{s}$. The lineshapes of M_x and M_y show the characteristic absorption and dispersive shapes indicating the presence of the spin echo. Furthermore, this signal appears at the correct temporal position for a spin echo. Assuming the magnetisation begins to recover exactly in the centre of the second pulse, the expected echo position is given by $G1 + P2/2 = 6 \mu\text{s}$. Figure 3.29b shows the magnitude of the complex FID signal calculated using equation (2.38) after processing the raw data. No additional phase-shifting of the signal was performed in this step. The echo signal was shifted in time in order for $t = 0$ to be defined as the start of the echo signal. Finally, the FID was Fourier transformed to produce the spectrum shown in figure 3.30. This measured spectrum agreed well with the literature measurements shown in [156], demonstrating the performance of the FPGA-based spectrometer.

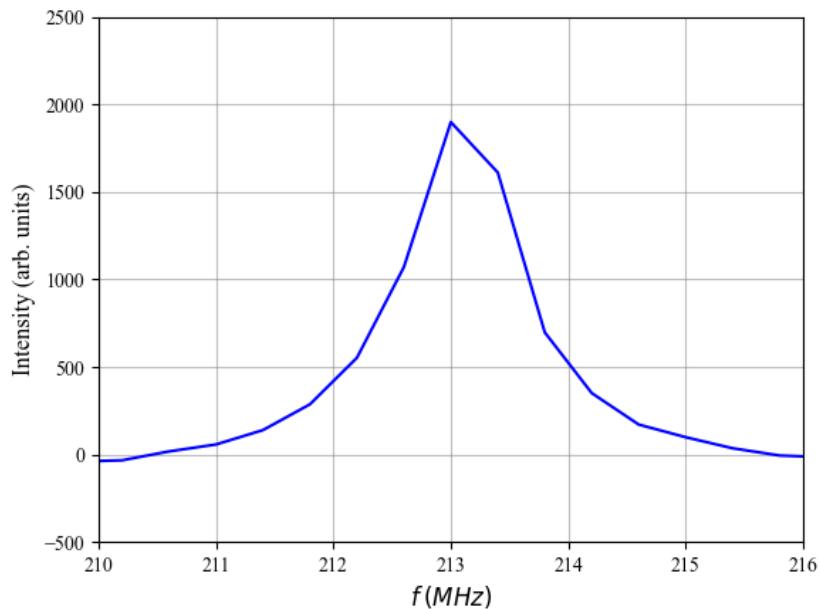


Figure 3.30: The ^{59}Co spectrum at measured using the FPGA-based NMR spectrometer. The spectrum was calculated from the raw data shown in figure 3.29 by performing the processing steps described in the main text.

3.6 Future work and conclusion

This section will describe a number of improvements that could be implemented to the FPGA-based NMR spectrometer setup. They will be divided into improvements to the spectrometer and Verilog source code, the Python GUI and the NMR probe.

3.6.1 Verilog

Firstly, the number of repeats for each NMR experiment should be handled onboard the SDR14 instead of by the host PC. This could be done by storing the number of repeats within the FPGA registers. The pulse programmer would then contain a counter that is incremented after every repeat until the desired amount were performed. In this case, downsampling and live streaming the data to the host PC would be required in order to prevent the DRAM from filling up during an experiment. If the DRAM fills up, new data would start erasing old records. This could also be implemented by pausing the experiment every set number of repeats and emptying the contents of the DRAM to the host PC. Then the data could be batch processed by the PC and plotted to the GUI and saved. This improvement would fully decouple the functionality of the SDR14 from the Python GUI increasing the portability of the device allowing it to be controlled by any custom software.

Secondly, the pulse programmer could generate a user defined number of TTL pulses. Currently, the restriction of up to three TTL pulses prevents the user from performing more advanced NMR sequences such as the CPMG sequence for measuring T_2 . This could be implemented by storing a number of pulses in the user registers. This would require modification to the triggering logic to ensure the data acquisition is performed at the correct times.

3.6.2 Python control software

The GUI could also support a number of signal processing techniques on the raw NMR data. Currently, the GUI only shows the raw NMR data and the running averaged data. This functionality could be extended to allow the user to perform a number of common NMR data processing tasks with the data. These include phase correction, Fourier Transformation, echo location and automatic peak finding. This would allow the user to quickly process the data during the experiment reducing the number of separate scripts in order to process the data.

3.6.3 NMR Probe

The main improvement to the NMR probe would be to implement sample rotation whilst the probe remained inside the PPMS allowing for the angular dependence of the NMR signals to be measured. A design for single-axis rotation of the sample was investigated during this project but was never tested inside the PPMS. This design used a pair of gears to rotate a horizontal rod within the sample, to which the sample was secured, to achieve single-axis rotation. However, there this introduced various issues such as gear backlash and limiting the space within the probe for the coil. A possible design for a single-axis probe could replicate the rotational probe provided by Quantum Design with the PPMS however this will require modifications to support the matching and tuning capacitors. The literature contains various designs for a probe that supports two-axis rotation using a set of worm gears [159], [160]. These probes are typically used in magnets with a bore radius of up to twice that of the PPMS. The small bore radius of the PPMS severely limits the amount of space available for the sample and gears. Furthermore, manufacturing the small gears required that won't wear out is a nontrivial task however may be possible using specialist machines or 3D printing.

Secondly, the probe does not currently contain a method to independently measure or control the sample temperature. The probe was designed to be used with the QD PPMS and use the PPMS systems for control of the environment. Therefore, the temperature and magnetic field values measured by the PPMS were used on the Python control GUI. The PPMS measures the temperature by taking the weighted average of three thermometers located in different locations in the sample chamber [151]. To provide a more accurate temperature measurement, a thermometer must be placed as close as possible to the sample. The probe would have to be slightly modified through the addition of a small platform to suspend the sample. This platform could contain an embedded Cernox thermometer (or similar) to provide a more accurate temperature reading at the sample position. The thermometer wiring could then be routed through the central stainless-steel tube alongside the coaxial cable, by using small teflon baffles inside the tube to ensure the two wires remained separated for the length of the tube. The probe head would be modified to provide a port for the thermometer wiring and the signal would be digitised and transmitted to the host PC.

3.6.4 NMR measurements

The spectrometer could be further tested by measuring the full ^{59}Co powder spectrum over the frequency range 210 MHz - 220 MHz. Additional measurements of the relaxation time constants, T_1 and T_2 could be performed in order to test the functionality of different pulse sequence. Temperature dependence of these measurements could be measured using the PPMS environment. After fully characterising the ^{59}Co powder, a new sample such as $\text{NiI}_2(3,5\text{-lut})_4$ could be measured in order to verify the experimental set-up can detect small NMR signals. Furthermore, measurements on a different compound would require the use of the PPMS sample environment which would demonstrate the functionality of both the probe and Python control GUI PPMS integration.

3.6.5 Conclusion

In conclusion, this chapter detailed the design, implementation and verification of the FPGA-based NMR spectrometer, which met the required specifications. An explanation of the logic and test outputs of each Verilog module was presented. The custom-built Python control software was introduced and explained. The design of a bespoke NMR probe, compatible with the Quantum Design PPMS was presented. Finally, the NMR electronics were characterised and a ^{59}Co powder spin echo signal was measured successfully using the set-up.

Chapter Four

Spin Chain

This chapter will provide a description of the spin-1 chain sample $\text{NiI}_2(3,5\text{-lut})_4$. Additionally, experimental and numerical methods will be introduced. Magnetisation, heat capacity and NMR results will be presented and discussed under varying temperature and magnetic field conditions to probe the sample throughout its phase diagram. Finally, future measurements will be proposed in order to further characterise the low temperature TLL state.

4.1 Sample information and methods

This section will introduce the spin-1 chain compound $\text{NiI}_2(3,5\text{-lut})_4$. It will provide an overview of its crystal structure and physical parameters. Next, the experimental methods used to measure the magnetisation, heat capacity and NMR results in section will be discussed. Finally, a brief description of the quantum Monte-Carlo (QMC) method for numerical simulations will be provided.

4.1.1 $\text{NiI}_2(3,5\text{-lut})_4$ sample information

$\text{NiI}_2(3,5\text{-lut})_4$ is a molecule-based AFM $S=1$ chain that belongs to the $P4/nnc$ crystallographic space group [48]. The structure of $\text{NiI}_2(3,5\text{-lut})_4$ is shown in figure 4.1a and 4.1b. It consists of Ni-I chains running parallel with the c axis where each Ni ion is separated by $c=9.9783\text{\AA}$. Neighbouring chains are offset by $c/2$. Each Ni ion is surrounded by 28 C sites split across four ligands each containing seven sites. As the sample is three-dimensional, there are two potential exchange pathways: intra-chain and inter-chain denoted by J and J_{\perp} respectively. The directions of these exchange pathways are shown in figure 4.1b. The ligands separate the chains and minimise the inter-chain exchange interaction, J_{\perp} [48]. The intra-chain exchange interaction, J , is dominant as it is facilitated by the I ions between the Ni ions. Therefore, the sample has a high degree of one-dimensionality and the Ni ions form an $S = 1$ chain.

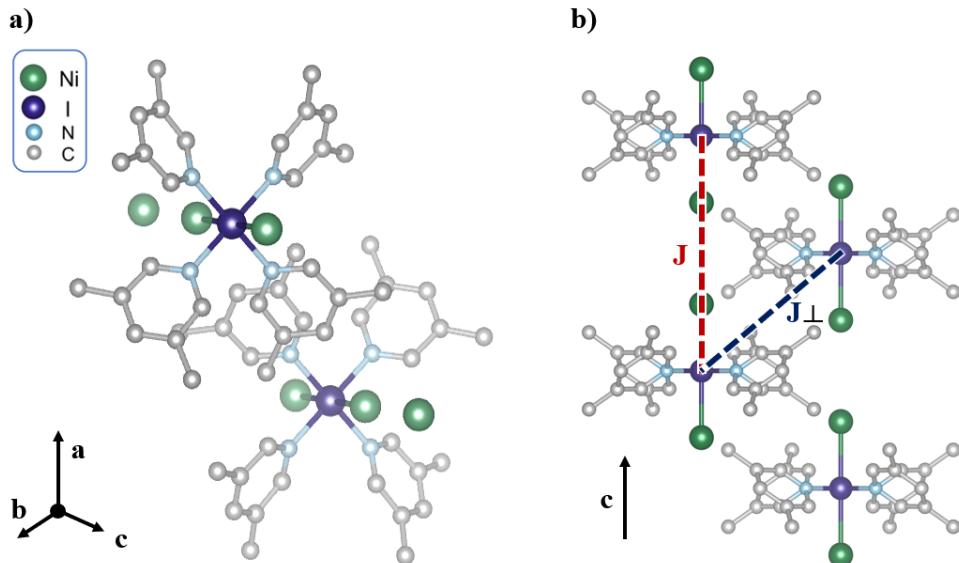


Figure 4.1: a) A diagram showing the structure of $\text{NiI}_2(3,5\text{-lut})_4$. The chains run along the c direction. b) The chain viewed along the c direction. Pairs of lutidine ligands overlap when viewed from this angle. The dashed lines indicate the direction of the J and J_{\perp} intra/inter chain interactions.

The compound has been shown to be described by the following AFM $S = 1$ chain Hamiltonian

$$H = D \sum_i (\mathbf{S}_i^z)^2 + J \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{S}_j + J_{\perp} \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{S}_{j\perp} + \mu_B \mu_0 \sum_i \mathbf{H} \cdot \mathbf{g} \cdot \mathbf{S}_i \quad (4.1)$$

Table 4.1 shows a list of the key parameters of $\text{NiI}_2(3,5\text{-lut})_4$ presented in [48]. This sample was chosen for the following measurements as it has an accessible H_c for typical condensed matter laboratory equipment. Furthermore, it demonstrates good agreement with theory with the ratio $\frac{\Delta_0}{J} \approx 0.41$ and has a small ratio $\frac{D}{J}$ removing any anisotropic effects [48].

Table 4.1: Parameter values for $\text{NiI}_2(3,5\text{-lut})_4$. All values taken from [48].

Property	Value
a (Å)	12.0019
b (Å)	12.0019
c (Å)	9.9783
J (K)	17.5
D (K)	-1.2
$\frac{D}{J}$	-0.07
$\frac{\Delta_0}{J}$	0.40
\bar{H}_c (T)	4.4
H_s^{\parallel} (T)	46.0
H_s^{\perp} (T)	50.7

A powder sample of $\text{NiI}_2(3,5\text{-lut})_4$ was used to collect the measurements in this section. The sample was synthesised for us in collaboration with the Superconductivity and Magnetism group from the University of Warwick [161] and was used for all of the experiments performed in this thesis. Details of the synthesis process and sample characterisation can be found in the following [48], [162], [163].

4.1.2 Experimental

This section introduces the experimental hardware and techniques used to measure the magnetisation, heat capacity and NMR results presented in section 4.2.

Magnetisation

Magnetisation measurements were measured using a Quantum Design Magnetic Property Measurement System 3 (MPMS3) [164]. This machine utilised a sensitive Superconducting Quantum Interference Device (SQUID) to perform high precision magnetisation measurements over a wide range of temperature and magnetic fields.

The base system provides a controllable temperature range of 1.8 K - 400 K. Employing the oven or iQuantum Helium-3 insert options extends the available temperature range to 0.5 K-1000 K. The magnetic field is generated using a superconducting magnet and ranges from 0 T - 7 T. The magnet is located within a shielded mu-metal dewar and cooled using liquid nitrogen and helium. A superconducting switch allows for rapid charging and discharging of the magnet at maximum rate of 700 Oe/s. These conditions allow for the magnetisation and the magnetic susceptibility, $\chi = M/H$, to be precisely measured as a function of temperature and applied magnetic field.

The bore of the superconducting magnet contains a second order gradiometer formed from pairs of clockwise and anticlockwise superconducting loops. The gradiometer measures the second derivative of the magnetic field. This screens the SQUID measurement from noise caused by uniform magnetic fields and linear magnetic field gradients from external sources. The sample is fixed to a sample rod and suspended within the gradiometer using a linear motor. The motor is used to move the sample along the length of the gradiometer. A current is generated within the gradiometer coils as a result of the change in flux due to

the magnetic moment of the sample as it is swept through the coils. The gradiometer is inductively coupled to the SQUID which provides a sensitive current to voltage conversion, producing a spatially varying voltage signal.

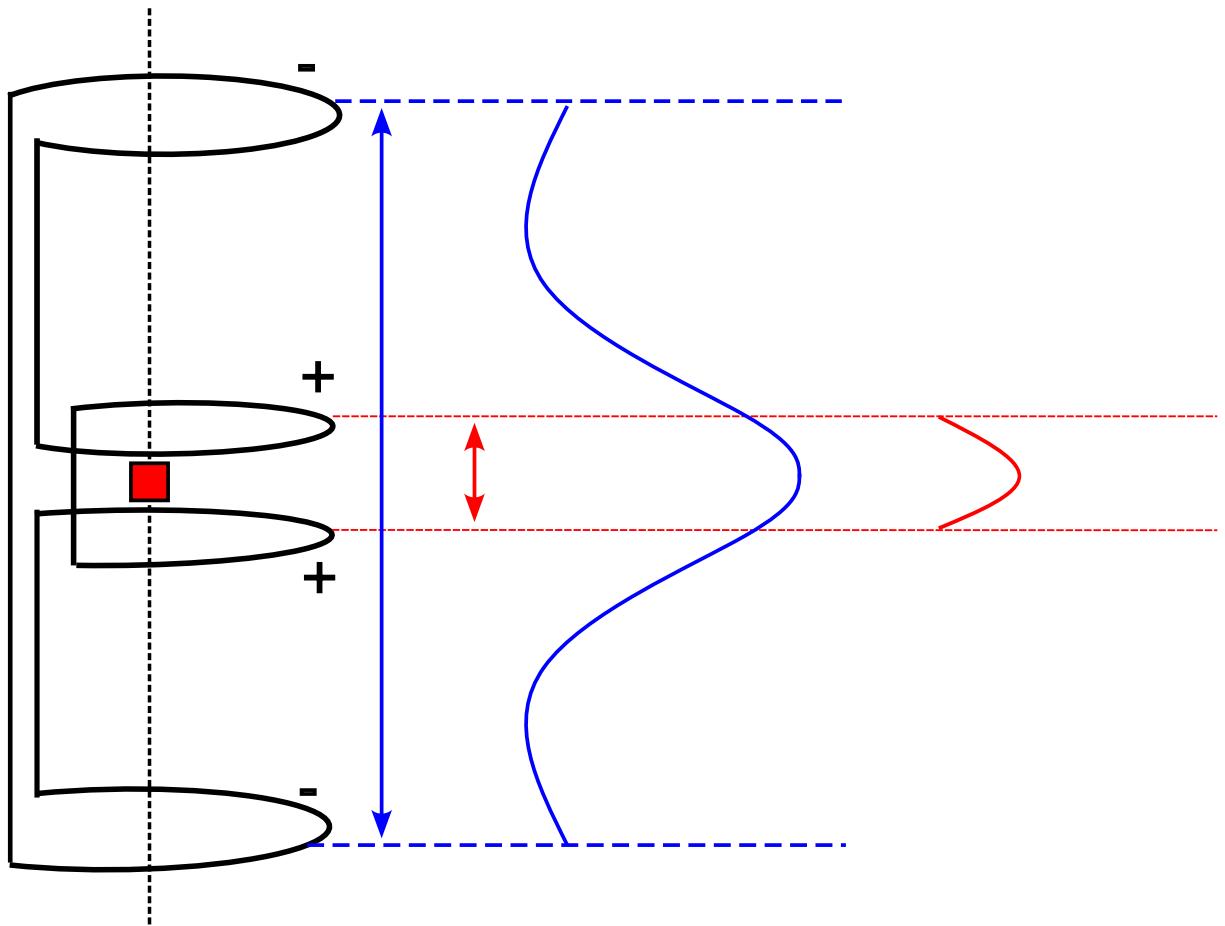


Figure 4.2: A diagram to show the SQUID measurement setup. The sample is represented by the red square and is moved along the dotted line by the linear motor. The coil winding direction is denoted by the + and - symbols. The blue curve shows the SQUID response generated using the DC scan mode. The red curve shows the response from the SQUID-VSM operating mode. The arrows denote the range of movement during each mode of operation.

The MPMS3 supports two methods of measuring the magnetisation: DC scan and SQUID-VSM (vibrating sample magnetometry). The DC scan mode sweeps the sample along the full length of the gradiometer. This generates the dipole signal shown in figure 4.2. The magnetisation is extracted by fitting the voltage curve to the dipole equation. The software generates two fits: the free centre fit and fixed centre fit. The free centre fit includes the centre position of the sample as a fitting parameter and provides better fits if the sample position drifts compared to the obtained value during the initial calibration.

The SQUID-VSM mode operates like a traditional VSM with the addition of a SQUID to the set-up. The sample is vibrated at the centre of the gradiometer at a fixed frequency, ω_0 . This produces an approximately parabolic voltage signal shown in figure 4.2. In a traditional VSM set-up this voltage would directly correspond to the induced voltage in the pickup coils due to the moving magnetic moments within the sample. The use of the SQUID as a current to voltage converter provides the improved sensitivity for the SQUID-VSM over traditional VSM. The SQUID voltage is then amplified and digitized before being presented to the user.

Heat capacity

The heat capacity is a quantity that relates an input of energy, ΔQ to the temperature change of the material ΔT

$$\Delta Q = mC\Delta T \quad (4.2)$$

where C is the heat capacity at constant pressure and m is the mass of the sample. Heat capacity measurements were performed using a Quantum Design Physical Property Measurement System (PPMS) [151], [165]. These measurements can be performed within a temperature range of 2 K - 300 K and magnetic field range of 0 T - 9 T.

To perform the measurements, the sample is mounted within a puck that fits inside

the bore of the PPMS magnet. A diagram of the puck is shown in figure 4.3. The puck contains a platform with an embedded thermometer and heater to measure the temperature change of the sample. The platform is suspended from the puck frame by a set of conductive wires. The wires provide electrical connections to the components on the platform and a thermal link from the platform to the puck frame, which acts as a thermal bath. The thermal conductance of the wire is such that the platform can reach a thermal equilibrium with the frame within the duration of the measurement. The sample is fixed to the platform using a small amount of Apiezon N grease to ensure sufficient thermal contact between platform and heater.

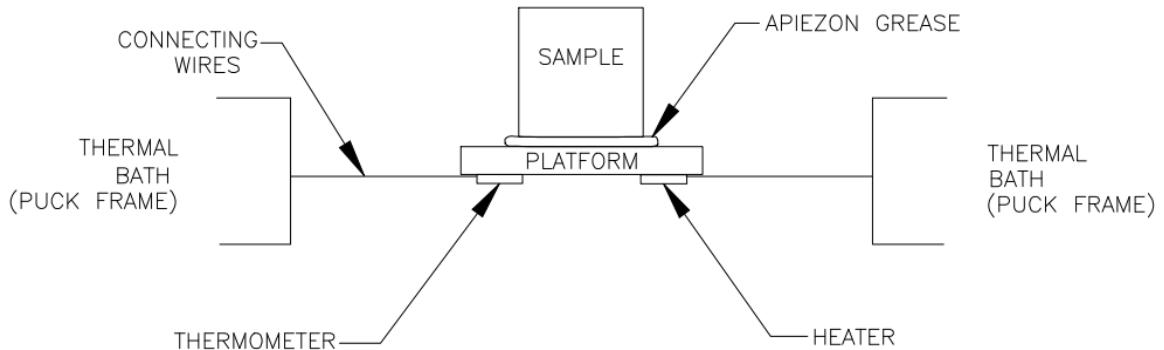


Figure 4.3: A schematic diagram of the heat capacity puck provided by QD. Reproduced from [165].

The measurements are performed by applying a pulse of heat at fixed power to the sample and letting the sample cool. The pulse causes the sample temperature to rise from a known initial temperature. When the pulse ends, the sample temperature relaxes back to the temperature of the puck frame. The temperature is recorded throughout the heating and cooling phases. This method of heat capacity measurement is known as a relaxation method [165], [166].

The change in temperature of the sample and platform can be related to the heat capacity of the sample by treating the sample, platform and puck frame (thermal bath) as

separate entities. This model is used to account for non-ideal thermal coupling between sample and platform. Figure 4.4 shows a diagram of the entities within the system. K_s and K_p are the sample-platform and platform-thermal bath thermal couplings. $P(t)$ is the pulse of heat applied to the sample. The heat flow between these entities is described by the following system of differential equations

$$C_p \frac{dT_p}{dt} = P(t) - K_p(T_p(t) - T_b) + K_s(T_s(t) - T_p(t)) \quad (4.3)$$

$$C_s \frac{dT_s}{dt} = -K_s(T_s(t) - T_p(t)) \quad (4.4)$$

where C_p , T_p are the heat capacity and temperature of the platform, C_s , T_s are the heat capacity of the sample and T_b is the temperature of the thermal bath. Equation (4.3) and (4.4) describe the temperature change of the platform and sample separately. The sign of each term represents if energy is being transferred to or leaked from the entity. Solving the differential equations for T_p (the temperature measured by the platform thermometer) yields

$$T_p(t) = -(Ae^{-t/\tau_1} + Be^{-t/\tau_2}) \quad (4.5)$$

where A and B are constants, τ_1 is the relaxation time constant between the platform and puck frame and τ_2 is the relaxation time constant between the sample and platform. The sample heat capacity has the form

$$C_s = K_p \tau_1 \left(1 - \frac{K_p \tau_2}{C_p}\right) + K_p \tau_2 - C_p \quad (4.6)$$

These equations are known as the ‘Two tau model’ [165].

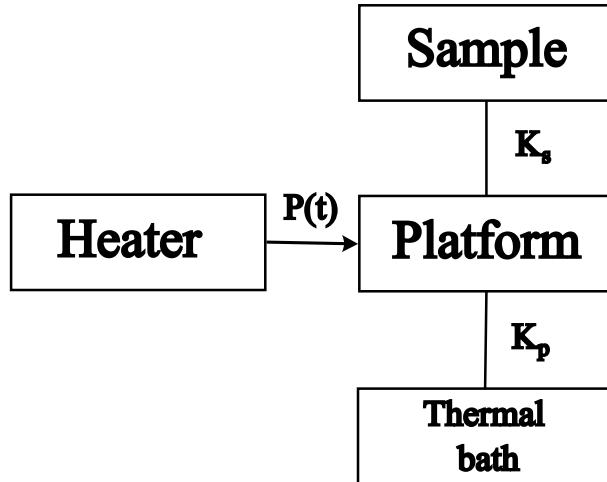


Figure 4.4: A schematic diagram of the heater, sample, platform and thermal bath entities in the two-tau model. K_s and K_p denote the sample-platform and platform-thermal bath respectively. $P(t)$ represents the pulse of heat applied to the platform by the heater.

In practice, to measure the sample heat capacity with the PPMS an additional set of ‘addenda’ measurements must be performed. The addenda measurements are used to measure the heat capacity of the empty platform and uses a simpler model to calculate C_p . The addenda measurements are be performed under the same temperature and magnetic field conditions as the sample measurements. Once both the sample and addenda measurements have been performed the software calculates C_s by fitting to the two-tau model.

NMR

The room temperature NMR measurements were measured using an NMR technique known as magic angle spinning (MAS) NMR. NMR spectra on solid-state powder samples generally display broad lines within their spectra originating from the rotational distribution of molecules within the sample [167]. The spectra obtained from these samples are known as ‘powder patterns’. The linewidth provides information about the local electronic environment of the nuclear spins such as the dipole fields described in section 2.2.5. If a spectrum

contains multiple closely spaced broad NMR lines, there is a loss in resolution due to overlap from all lines. Therefore, individual lines cannot be resolved and information pertaining to individual nuclear sites cannot be extracted.

MAS NMR is a technique used to remove the orientation dependent interactions within the NMR interaction Hamiltonian. The main feature of the technique is rotating the sample at a frequency in the kHz range and tilting the sample at an angle of 54.74° (the ‘magic angle’) from the direction of the static field B_0 [167]. The sample rotation is achieved by inserting the sample into a rotor and the system is driven by the injection of air into the sample chamber to spin the rotor [168]. A schematic diagram of a MAS setup is shown in figure 4.5.

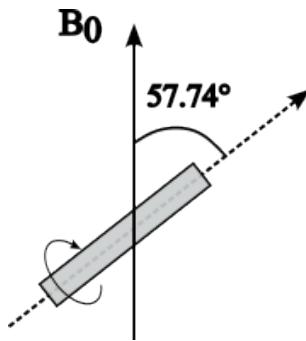


Figure 4.5: A diagram showing a MAS NMR setup. The sample is contained within the sample holder (grey square) and is suspended at the ‘magic angle’ from the static field.

The magic angle is calculated by considering the dipole interaction Hamiltonian in equation 2.45. If the interaction is small compared to the Zeeman interaction, the secular approximation can be made, dropping higher order terms from the interaction [169]. Applying the secular approximation and writing equation 2.45 in polar coordinates yields

$$H_{dip} = \frac{-\mu_0 \gamma_I \gamma_S \hbar^2}{16\pi^3 r^3} (3 \cos^2(\theta) - 1) I_z S_z \quad (4.7)$$

where γ_I , γ_S are the gyromagnetic ratio of the nuclear spin and electron spin respectively. r and θ are the distance and angle between the spins. The magic angle can be found by

considering $3\cos^2\theta - 1 = 0$, where the dipole interaction vanishes. Therefore, MAS has the effect of averaging the rotationally dependent terms of the interaction Hamiltonian to zero reducing or fully removing the line-broadening from the NMR lineshape [167].

Figure 4.6a) and 4.6b) show a schematic diagram of the powder pattern and MAS NMR spectrum respectively. Rotating the sample introduces ‘spinning sidebands’ into the spectrum at multiples of the spinning frequency f_{rot} . If a sample has many NMR peaks, this results in a complex MAS NMR spectrum as each peak will generate a set of spinning sidebands. The spinning sidebands are identified by examining the frequency difference between pairs of peaks.

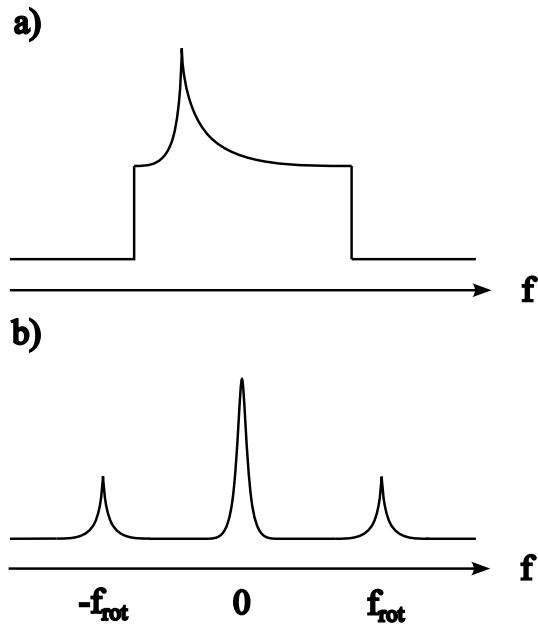


Figure 4.6: a) An example powder pattern spectrum caused by line broadening. b) The same spectrum measured using MAS NMR. Spinning sidebands are introduced at multiples of the rotational frequency f_{rot} .

4.1.3 Simulation

This section will present a basic introduction to Quantum Monte-Carlo (QMC) simulations, a common numerical method used to simulate many body systems in condensed matter physics. It will conclude with an overview of the Algorithms and libraries for physics simulations (ALPS) software package, which was used for the QMC calculations performed in section 4.2.

Quantum Monte-Carlo

Quantum Monte Carlo (QMC) simulations encompass a large range of computational approaches to study quantum many-body systems [170]. They have been used across various fields of physics which contain the many-body problem including condensed matter physics and astrophysics [171]. Within condensed matter, they have been used to study phase transitions [172], lattice defects [173] and Bose-Einstein Condensates [174]. There are many flavours of QMC algorithms such as variational Monte Carlo [175] and Path Integral Monte Carlo [176] due to some algorithms favouring specific systems. At the simplest level, all QMC algorithms share the application of classical Monte Carlo approaches to a quantum problem [177]. As a result, the general idea is to randomly sample a subset of states within the full Hilbert space of the system. Provided these states are sampled carefully, reasonable estimates of observables such as the magnetisation and heat capacity can be computed [170]. Consequently, QMC methods can be used to simulate large systems of the order of 1000 particles [170], [177].

Monte Carlo methods refer to methods that make heavy use of random numbers to estimate statistical averages [178]. This technique can be used to numerically solve multi-dimensional integrals. The integral, I , of a d -dimensional integral of a function $f(\mathbf{x})$ over a

volume V can be written as

$$I = \int_V d^d \mathbf{x} f(\mathbf{x}) = \frac{\int_V d^d \mathbf{x} f(\mathbf{x})}{V} V = \bar{f}_V V \quad (4.8)$$

where \bar{f}_V is the average of $f(\mathbf{x})$ over the volume. Therefore, a possible way to calculate \bar{f}_V is to randomly draw samples uniformly distributed within V and calculate the average in the usual way.

$$I \approx \frac{\sum_{i=1}^N f(\mathbf{x}_i)}{N} V \quad (4.9)$$

where N is the number of random samples and \mathbf{x}_i is the i th random sample. The error in I is proportional to $N^{-\frac{1}{2}}$ (the standard error on the mean), so provided N is large enough, a reasonable numerical value for I can be computed [178].

Monte Carlo integration is particularly relevant in condensed matter physics when considering the average of an observable in a quantum system. The partition function of a system with N particles is

$$Z_N = \sum_{states} e^{-E_{state}/k_B T} \quad (4.10)$$

where E is the energy of the state and T is the temperature [170]. The Monte Carlo approach to calculate observables in the system is to randomly generate a set of states from within the ensemble of states of the system. The average value of an observable, O , is then given by

$$\langle O \rangle = \frac{\sum_{i=1}^N O_i e^{-E_i/k_B T}}{\sum_{i=1}^N e^{-E_i/k_B T}} \quad (4.11)$$

where O_i is the value of the observable in state i .

The states that are randomly sampled during this process must be distributed in the same way as the physical distribution of states in the system. For a system at equilibrium at a finite temperature, the probability of observing a state i , p_i , is given by the Boltzmann distribution:

$$p_i = e^{-E_i/k_B T} \quad (4.12)$$

Therefore, if a state is selected from an underlying uniform distribution (i.e. each state has an equal chance of being chosen), the system will not be modelled correctly. A simple example can be shown by considering an $S = \frac{1}{2}$ ferromagnetic Ising model system. The ground state below the critical temperature consists of all spins aligned in the $+z$ or $-z$ direction so these states have the dominant Boltzmann weights. However, the chance of sampling one of the ground states becomes vanishingly small if all states have an equal probability to be sampled since the state space grows exponentially. To resolve this issue, a bias has to be added to the sampling process in order to sample states from a Boltzmann distribution, this is known as ‘importance sampling’.

This is commonly performed by using a Markov Chain. This is a process that defines a set of probabilistic rules that transform an input state to an output state. The most famous algorithm to map a state sampled from a uniform distribution to a Boltzmann distribution is the Metropolis algorithm [170]. For the Ising model it is as follows:

1. Sample a state from a uniform distribution and calculate the energy, E .
2. Flip a random spin and calculate the energy of the new state, E_{new} .
3. If E_{new} is lower than E , accept the new state.
4. If E_{new} is greater than E , accept the new state with a probability $p = e^{\frac{E_{new}-E}{k_B T}}$.

It can be shown that repeated applications of the Metropolis algorithm will lead to a Boltzmann distribution and at this point the simulation is considered valid [170]. Additionally, once a valid state has been obtained, further applications of the algorithm will always produce another a valid state [170]. Therefore, the value of an observable can now be computed for each valid step of simulation. The final consideration is the error on a computed observable is the standard deviation of the mean. This requires that each measurement is statistically uncorrelated, which is clearly not the case when using this algorithm. A simple

way to uphold this requirement is to only compute the observable every P steps, where P is large enough that the autocorrelations in the Markov chain have decayed [170].

The main limitation of QMC methods is known as the ‘sign problem’. This is caused by the existence of negative or more generally, complex off-diagonal elements in the Hamiltonian of quantum systems [179]. These elements can result in the generation of negative probability distributions, which are unphysical. The problem isn’t present in classical Monte Carlo as the Hamiltonian will always be diagonal. The existence of the sign problem is a result of the choice of basis for a given system. Generally this problem can be mitigated at the cost of an exponential increase in runtime or can be removed entirely by performing a basis transformation on the Hamiltonian [179].

Algorithms and Libraries for Physics Simulations

The Algorithms and Libraries for Physics Simulations (ALPS) project [180] is an open-source software project used to simulate quantum systems. It was designed to provide an accessible interface to a collection of popular numerical algorithms without having to implement them from scratch. ALPS includes a range of QMC algorithms as well as other simulation methods such as exact diagonalisation (ED), density matrix renormalisation group (DMRG) and dynamical mean field theory (DMFT). A python package, `pyalps`, is included with the installation to allow for auto-generation of simulation parameters, performing batch simulations and loading and plotting the outputs using simple Python scripts.

An ALPS simulation requires a parameter file which is a text file containing details about the lattice and Hamiltonian to be simulated, as well any environmental conditions such as temperature and magnetic field. An excerpt of a Hamiltonian file is shown in figure 4.7a. The ‘Sitebases’ tag is used to specify local quantum numbers and operators such as S^+ , S^- along with how the operators modify the quantum numbers. The ‘Parameter’ tag specifies

any additional parameters of the model such as magnetic field or exchange interaction. The ‘Hamiltonian’ tag is then used to combine the sitebases and parameters to construct the model. The lattice is also specified using an XML format and this contains information about the geometry of the system. An excerpt of a lattice file is shown in figure 4.7b. The ‘Vertex’ tag is used to specify the position of sites within the system and specify the type of particle residing on that site. The ‘Edge’ tag is used to specify an edge connecting two vertices, allowing for interaction pathways between sites. Combining these files allows for the simulation of a range of custom quantum models on custom lattice geometries.

a)

```

<MODELS>
<SITEBASIS name="spin">
  <PARAMETER name="local_spin" default="local_s"/>
  <PARAMETER name="local_S" default="1/2"/>
  <QUANTUMNUMBER name="S" min="local_spin" max="local_spin"/>
  <QUANTUMNUMBER name="Sz" min="-S" max="S"/>
  <OPERATOR name="Splus" matrixelement="sqrt(S*(S+1)-Sz*(Sz+1))">
    <CHANGE quantumnumber="Sz" change="1"/>
  </OPERATOR>
  <OPERATOR name="Sminus" matrixelement="sqrt(S*(S+1)-Sz*(Sz-1))">
    <CHANGE quantumnumber="Sz" change="-1"/>
  </OPERATOR>
  <OPERATOR name="Sz" matrixelement="Sz"/>
</SITEBASIS>

```

b)

```

<GRAPH name="4-site mixed" vertices="4">
  <VERTEX id="1" type="0"/>
  <VERTEX id="2" type="1"/>
  <VERTEX id="3" type="0"/>
  <VERTEX id="4" type="1"/>
  <EDGE type="0" source="1" target="2"/>
  <EDGE type="0" source="2" target="3"/>
  <EDGE type="0" source="3" target="4"/>
  <EDGE type="0" source="4" target="1"/>
  <EDGE type="1" source="1" target="3"/>
  <EDGE type="1" source="2" target="4"/>
</GRAPH>

```

Figure 4.7: a) The Hamiltonian file specifying the ‘spin’ model in ALPS. b) The lattice file specifying a square lattice with interactions along the edges and main diagonals in ALPS.

By default, ALPS includes a number of commonly used lattices and models. All QMC simulations performed in this thesis used the ‘spin’ model which specified a Heisenberg model with a magnetic field and anisotropy factors with the ‘chain lattice’ geometry which specified a length N 1D chain. QMC simulations were performed with the ‘directed loop’ QMC algorithm which was recommended for spin models in a magnetic field.

4.2 Results and Discussion

This section will present the magnetisation, heat capacity and NMR results on a powder sample of $\text{NiI}_2(3,5\text{-lut})_4$.

4.2.1 Magnetisation

Low-field magnetization

Low-field magnetization measurements of a $\text{NiI}_2(3,5\text{-lut})_4$ were performed using a QD MPMS3 SQUID-VSM over a temperature range 2 K - 300 K with $\mu_0 H = 0.1$ T. These measurements were performed using a 16 mg powder sample. The powder was tightly packed into the standard MPMS3 plastic pill sample capsule sealed with epoxy resin in order to prevent sample motion during the measurement. The capsule was inserted into a 5 mm diameter drinking straw and was connected to the sample rod. The drinking straw was chosen to minimise any background signals from the sample holder. The low-field measurements were used to characterise additional contributions from the sample such as the background contribution. This was done using two approaches: a Curie-Weiss fit and a Pade Approximation.

The Curie-Weiss method involved fitting the susceptibility data to the following Curie-Weiss law [181]

$$\chi(T) = \frac{C}{(T - \theta_{CW})} + \chi_0 \quad (4.13)$$

where C is the Curie constant, θ_{CW} is the Curie-Weiss temperature and χ_0 is a temperature independent contribution to the susceptibility. Furthermore, the Curie constant is related to the g -factor of the compound by [181]

$$C = \frac{N_A \mu_0 \mu_B^2 g^2 S(S+1)}{3k_B} \quad (4.14)$$

where S is the spin.

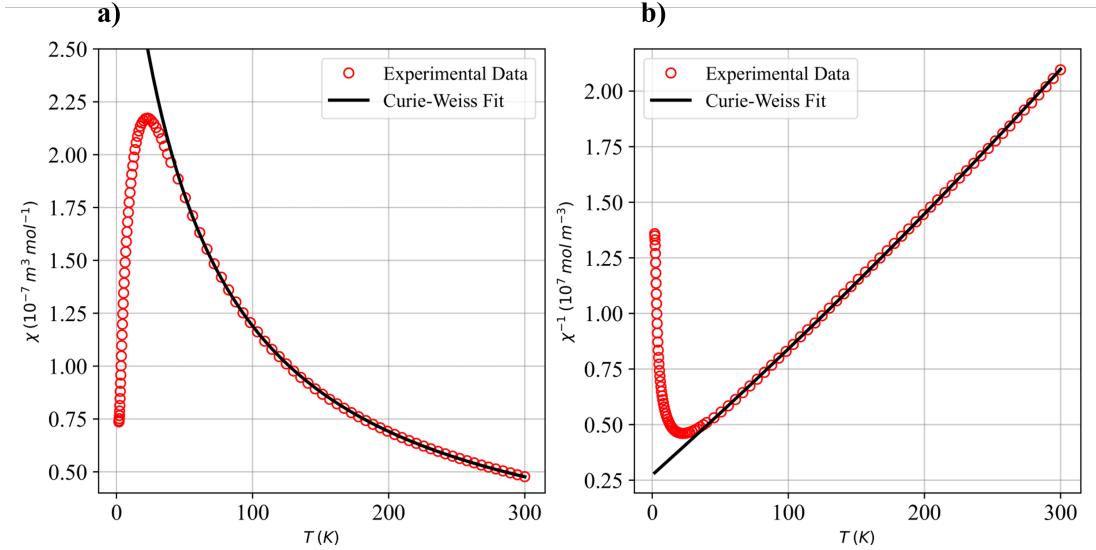


Figure 4.8: a) Susceptibility of the $\text{NiI}_2(3,5\text{-lut})_4$ powder sample at $\mu_0 H = 0.1$ T as a function of temperature. The solid line denotes a Curie-Weiss fit. b) Inverse susceptibility of the $\text{NiI}_2(3,5\text{-lut})_4$ powder sample at $\mu_0 H = 0.1$ T as a function of temperature.

The Curie-Weiss law was used to characterise the high-temperature regime of the low-field magnetisation data. Figure 4.8 shows the low-field molar susceptibility measurements of $\text{NiI}_2(3,5\text{-lut})_4$. The molar susceptibility was calculated using the sample mass, $m=16$ mg and the molar mass of $\text{NiI}_2(3,5\text{-lut})_4$, $M_m = 741.11098 \text{ g/mol}$. The data show a typical response for a AFM chain displaying a peak at the Neel temperature, T_N , around 20 K. The high temperature ($T > T_N$) behaviour shows a paramagnetic response and follow the Curie-Weiss law in equation (4.13). The low temperature ($T < T_N$) behaviour shows a sharp downturn in susceptibility caused by the build-up of antiferromagnetic interactions.

The high temperature data (from 50 K - 300 K) were fit to equation (4.13) yielding the fitting parameters $\theta_{CW} = -50.7$ K, $C = 1.9 \times 10^{-5} \text{ Km}^3 \text{ mol}^{-1}$ and $\chi_0 = -6.1 \times 10^{-9} \text{ m}^3 \text{ mol}^{-1}$. An estimate of the g -factor was calculated using equation (4.14) as $g_{powder} = 2.5$ which was larger than the literature value of $\text{NiI}_2(3,5\text{-lut})_4$ is $g_{powder} = 2.197$ [48]. Finally,

$\theta_{CW} \approx 2T_N$ indicated a small degree of frustration within the compound [181].

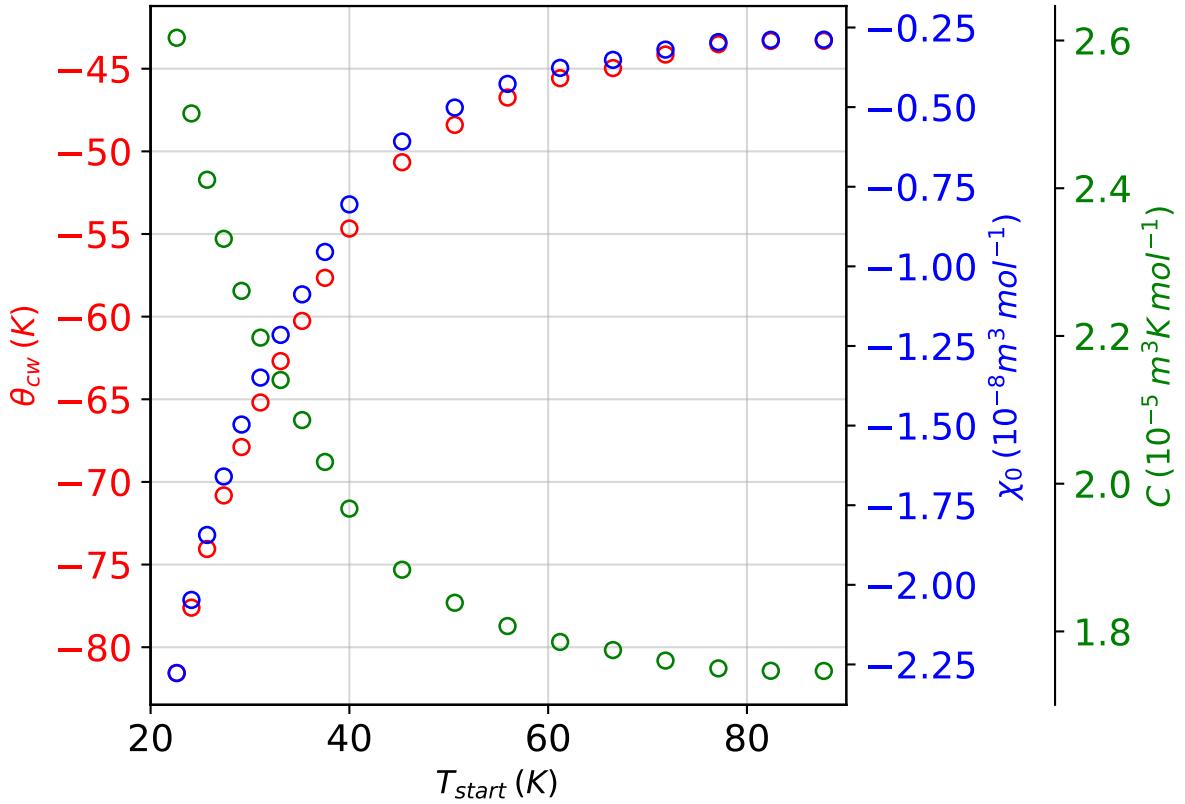


Figure 4.9: The fitting parameters obtained by fitting the low-field susceptibility data to a Curie-Weiss law. The x -axis represents the start temperature of the fitting region. The end of the region was always 300 K.

Figure 4.9 shows the variation of the Curie-Weiss fitting parameters as a function of temperature. This demonstrates the crossover between the AFM behaviour and paramagnetic behaviour region clearly. Above 50 K, the fitting parameters begin to plateau indicating small variations in the fits within this region. Below 50 K, there is a large variation of the parameters as a function of temperature as the AFM correlations begin to dominate the paramagnetic behaviour, causing the data to deviate from the Curie-Weiss law.

The second approach was to fit the data over the whole range by making use of a Padé Approximation which was used to characterise the impurity contribution to the susceptibility. The expected temperature dependency of a spin chain was given by [182]

$$\chi(T) = (1 - x)\chi_{chain} + 2x\chi_{imp} + \chi_0 \quad (4.15)$$

where the first term is the susceptibility from the chain, the second term accounts for the existence of impurities within the sample and the final term is the same temperature independant term found in equation (4.13). The variable x is the fraction of impurity fraction. The impurity term can be written as [182]

$$\chi_{imp}(T) = \frac{N_A \mu_B^2 g_{imp}^2 S(S+1)}{3k_B(T - \theta_{imp})}. \quad (4.16)$$

The Padé Approximation was used to obtain the following form for χ_{chain} [183]

$$\chi_{chain} = \frac{N_A \mu_B^2 g_{imp}^2 S(S+1)}{3k_B T} \exp\left(\frac{-\Delta J}{k_B T}\right) \frac{1 + \sum_{i=1}^m A_i \left(\frac{J}{k_B T}\right)^i}{1 + \sum_{j=1}^n B_j \left(\frac{J}{k_B T}\right)^j} \quad (4.17)$$

where A_i and B_j are the set of dimensionless coefficients found in table 4.2

Table 4.2: A table showing the values of the coefficients A_i and B_i calculated using a Padé Approximation. Values taken from [183].

i	A_i	B_i
1	0.67855	1.6
2	1.2698	2.6533
3	0.65478	2.5159
4	0.14123	1.6783
5	0.087773	0.41951
6	-9.1750×10^{-5}	0.041205

Figure 4.10 shows the low temperature regime of the low-field susceptibility data. The data showed a small upturn at the lowest temperatures that was attributed to the

presence of $S = 1/2$ end chain moments (impurities) within the sample. These are present due to lattice defects within the sample breaking the chains of $S = 1$ moments. When a chain is broken, $S = 1/2$ moments are introduced at each end. The data was fitted to equation (4.15) in order to characterise the contribution of the end chain moments to the susceptibility. x , θ_{imp} and χ_0 were treated as fitting parameters and the values $x = 0.0041$, $\theta_{imp} = 0.9$ K and $\chi_0 = 0.222 \text{ } m^3\text{mol}^{-1}$ were obtained. The fit agreed well with the data except at the lowest temperatures. This indicated that measurements at lower temperatures may be required in order to correctly characterise the impurity contribution. This is because the temperature dependence of the $S = 1/2$ moments follows a $1/T$ dependence[109] so it will rapidly increase as the temperature is lowered, becoming the dominant contribution to the measured susceptibility.

Quantum critical scaling

As outlined in section 2.3, the quantum critical region around H_c at the transition between the gapped Haldane phase and TLL phase is described by a $z = 2$ dynamical exponent. The magnetisation in this regime follows the scaling relaxation [110], [126]

$$M(T, H) = N_A g \mu_B \sqrt{\frac{2k_B T \Delta}{c^2}} \mathcal{M} \left(\frac{g \mu_0 \mu_B (H - H_c)}{k_B T} \right) \quad (4.18)$$

where \mathcal{M} is a scaling function defined as [126]

$$\mathcal{M}(x) = \frac{1}{\pi} \int_0^\infty \frac{1}{e^{y^2-x} + 1} dy. \quad (4.19)$$

This equation implies magnetization curves measured at different magnetic fields will collapse on to the same curve when the quantity $M(T, H)/(N_A g \mu_B \sqrt{2k_B T}/2c)$ is plotted against $g \mu_0 \mu_B (H - H_c)/k_B T$. Figure 4.11 shows a demonstration of quantum critical scaling around H_c for $\text{NiI}_2(3,5\text{-lut})_4$ using the parameters found in table 4.1 and the relation $c = 2JS$.

In-field magnetisation measurements were performed around the critical field. H_c from

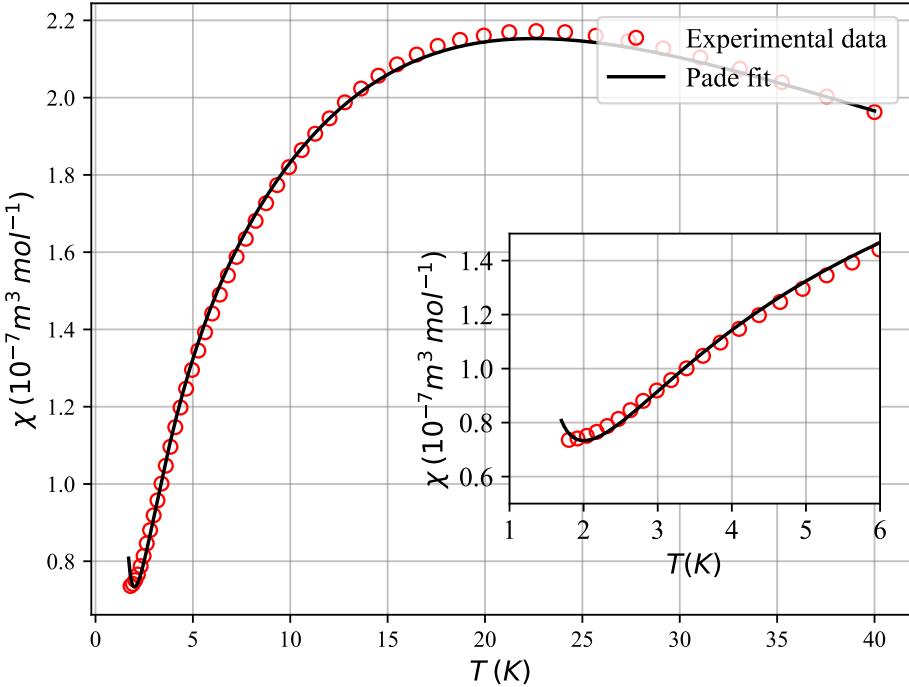


Figure 4.10: The low temperature region of the low field susceptibility data showing the fit to equation (4.15). Inset) A zoom in of the low temperature region, a small upturn in the data can be seen at the lowest temperatures.

0 K to 40 K. Figure 4.12a) shows the raw magnetisation measurements of $\text{NiI}_2(3,5\text{-lut})_4$ from 3.8 T to 5.0 T measured using SQUID-VSM. The data was processed by subtracting the impurity and background magnetisation present in equation (4.15) to obtain the inherent magnetisation of the chain. The impurity magnetisation contribution was modelled using the $S = 1/2$ Brillouin function [109]

$$M_{imp}(T, H) = g\mu_B N_A S x \tanh\left(\frac{Sg\mu_B H}{k_B T}\right) \quad (4.20)$$

where x is the impurity concentration found from the Padé approximation fitting in the previous section. The background magnetisation, M_0 , originating from χ_0 was temperature independent and calculated and subtracted using $M_0 = \chi_0 H$ at each field.

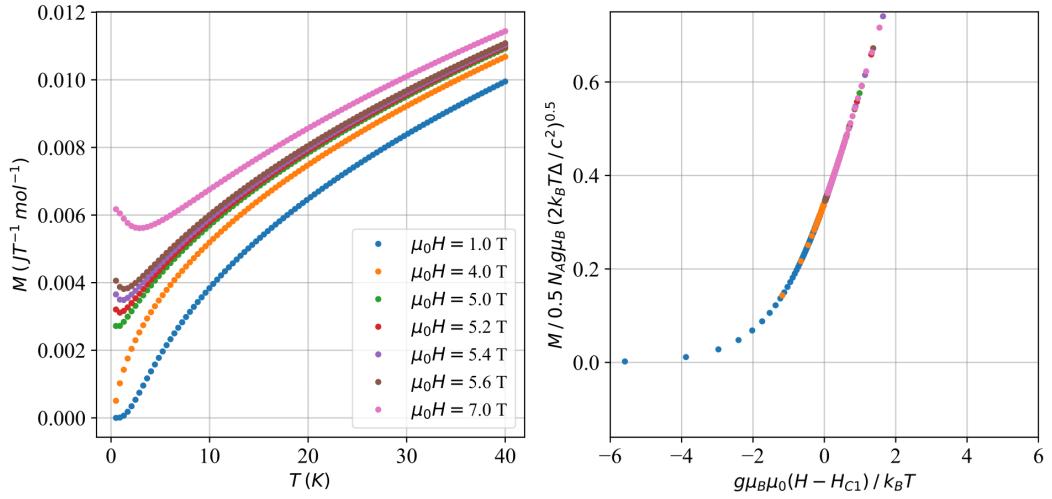


Figure 4.11: a) The theoretical magnetisation calculated using equation (4.18) for $NiI_2(3,5-lut)_4$. b) Demonstration of magnetisation scaling by plotting the rescaled variables described in the text.

Figure 4.12b shows the experimental magnetisation data with the impurity and background contribution subtracted, plotted in terms of the scaling variables. The results do not agree with the theoretical scaling behaviour. The fanning-out of the scaled data represents an over subtraction of the impurity and background contribution of the magnetisation. The main reason for this is that the scaled variables are sensitive to the impurity and background subtraction of the magnetisation. As outlined previously, the Padé fits used to obtain x and χ_0 did not fit the data well at low temperature which means the impurity contribution was not characterised well. Additionally, the upper temperature bound on quantum critical region at H_c has not previously been established for $NiI_2(3,5-lut)_4$. This means that the magnetisation measurements are not probing the quantum critical region and therefore will not demonstrate scaling. In this case, magnetisation measurements down to the mK regime would be required in order to demonstrate quantum critical scaling in $NiI_2(3,5-lut)_4$.

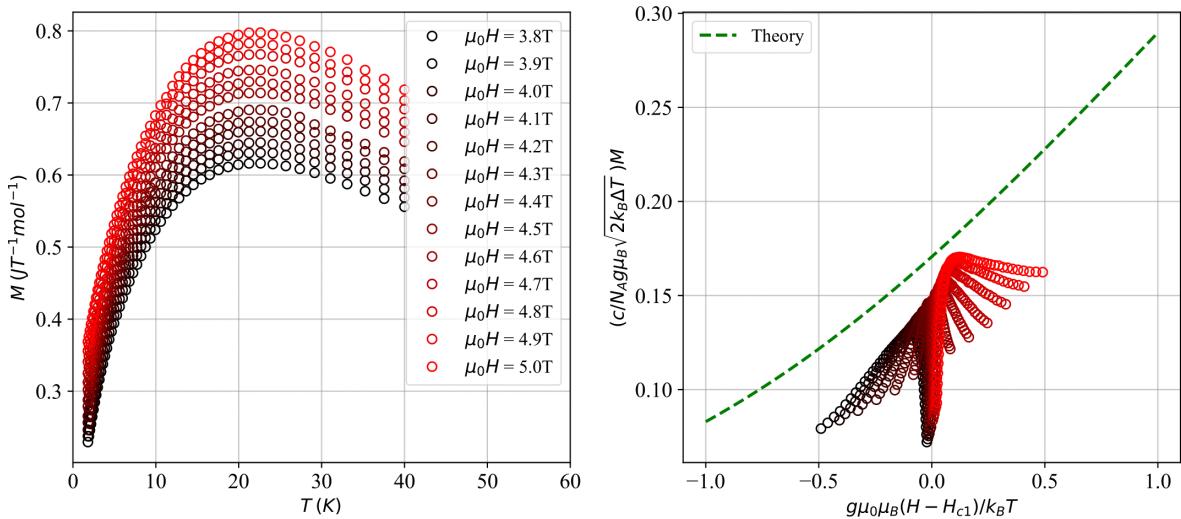


Figure 4.12: a) In-field magnetisation measurements of NiI₂(3,5-lut)₄ from 3.8 T to 5.0 T. b) The data from a) plotted in terms of the scaling variables discussed in the text. The green dashed line is theoretical curve from equation (4.18).

4.2.2 Heat capacity

Calculation of the low temperature specific heat capacity

When analysing low temperature specific heat capacity of a sample, an additional contribution from the nuclear energy levels within the system will be present within the measured data. This section provides the calculations that examined the effect of the nuclear heat capacity over the temperature range of the heat capacity measurements found in the following section.

For a compound, the total contribution to the heat capacity is calculated by summing the contribution from each unique nuclear species within the compound. The contribution

from a nuclear species i is given by [184]

$$C_{nuc}^i(T, H) = \frac{N_i(\gamma_i \mu_0 H \hbar)^2}{k_B T^2} \left(\frac{1}{4} \sinh^{-2} \left[\frac{x}{2} \right] - \frac{(2I_i + 1)^2}{4} \sinh^{-2} \left[(2I_i + 1) \frac{x}{2} \right] \right) \quad (4.21)$$

where $x = g_N \mu_N \mu_0 H / k_B T$, N_i is the number of nuclei of species i in 1 mole of $\text{NiI}_2(3,5\text{-lut})_4$, γ_i is the gyromagnetic ratio, I_i is the nuclear spin number, g_N is the nuclear g-factor and μ_N is the nuclear magneton. The natural abundance of the nuclear species was included when calculating C_{nuc}^i for isotopes that have a non-zero spin quantum number. The total nuclear contribution to the heat capacity was then calculated by summing over all elements, i

$$C_{nuc} = \sum_i C_{nuc}^i \quad (4.22)$$

Table 4.3: Parameters used to calculate the nuclear contribution to the heat capacity for $\text{NiI}_2(3,5\text{-lut})_4$ where NA is the natural abundance.

Nuclei	I	$\gamma_N (\text{rad s}^{-1} \text{T}^{-1})$	g_N	Moles in 1 mole of $\text{NiI}_2(3,5\text{-lut})_4$	NA
1H	$\frac{1}{2}$	2.68×10^8	5.586	36	1
13C	$\frac{1}{2}$	6.732×10^7	1.405	28	0.01109
14N	1	1.93×10^7	0.403	4	0.99636
61Ni	$\frac{3}{2}$	-2.39×10^7	-0.50	1	0.0114
127I	$\frac{5}{2}$	5.379×10^7	1.125	2	1

Table 4.3 shows the parameters used to calculate the nuclear contribution to the specific heat for $\text{NiI}_2(3,5\text{-lut})_4$. Figure 4.13 demonstrates a calculation of the nuclear contribution to the specific heat capacity at 14 T of $\text{NiI}_2(3,5\text{-lut})_4$, showing the contribution from each species of nuclei. For all fields, the total nuclear contribution is dominated by the hydrogen component due to its large gyromagnetic ratio and 1 mole of $\text{NiI}_2(3,5\text{-lut})_4$ containing a large amount of hydrogen atoms.

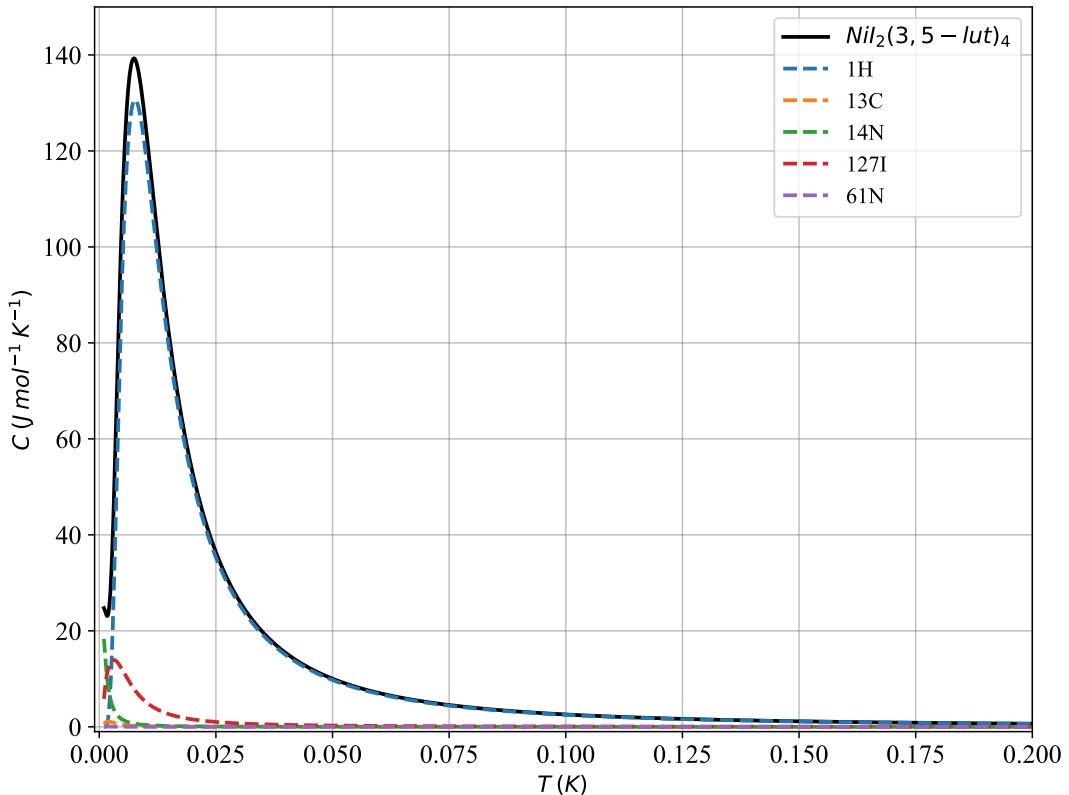


Figure 4.13: The nuclear contribution to the specific heat for $\text{NiI}_2(3,5-\text{lut})_4$ at 14 T. The dashed lines indicate the contribution from each of the constituent nuclei and the black line shows the total contribution using equation (4.22).

The total low temperature heat capacity was then calculated by combining the nuclear contribution from equation (4.22) with the TLL heat capacity in equation 2.67. Figure 4.14 shows an example calculation of the TLL heat capacity at fields 6 T, 9 T and 14 T. These values were chosen to probe the situations where the compound was near the TLL phase boundary (6 T), far from the TLL phase boundary (14 T) and an intermediate field (9 T). At 14 T, the compound should be far into the TLL state and any effects from the crossover regime from gapped to TLL should not be present. As expected, the gradient decreases with increasing field due to field dependence of the Luttinger parameter u which increases with

applied field.

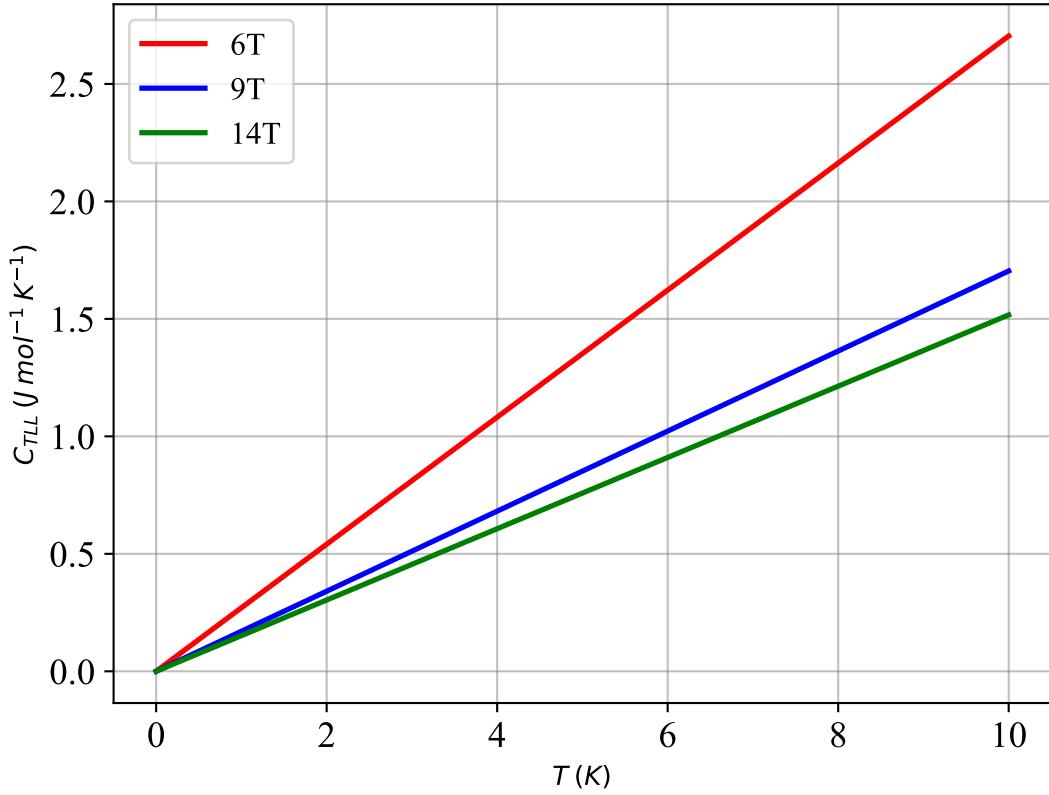


Figure 4.14: An example calculation of the TLL heat capacity of $\text{NiI}_2(3,5\text{-lut})_4$ at 6 T, 9 T and 14 T using equation (2.67).

Figure 4.15a shows each component of the total low temperature heat capacity of $\text{NiI}_2(3,5\text{-lut})_4$ at 9 T, where the solid black line indicates the total heat capacity. Figure 4.15b shows the total heat capacity calculated at 6 T, 9 T and 14 T of $\text{NiI}_2(3,5\text{-lut})_4$. Therefore, for temperatures less than 1 K, the dominant component is the nuclear component which would have to be subtracted to correctly obtain the TLL specific heat. However, for measurements above 1 K, the effect of the nuclear heat capacity is negligible compared to the linear TLL heat capacity. It should be noted that this analysis only holds when the compound is within the TLL phase which also has a temperature ceiling of approximately

$0.05J = 0.875$ K [185]. In conclusion, over the temperature range 2 K - 10 K where the heat capacity measurements in next section were performed, the effect of the nuclear heat capacity was minimal.

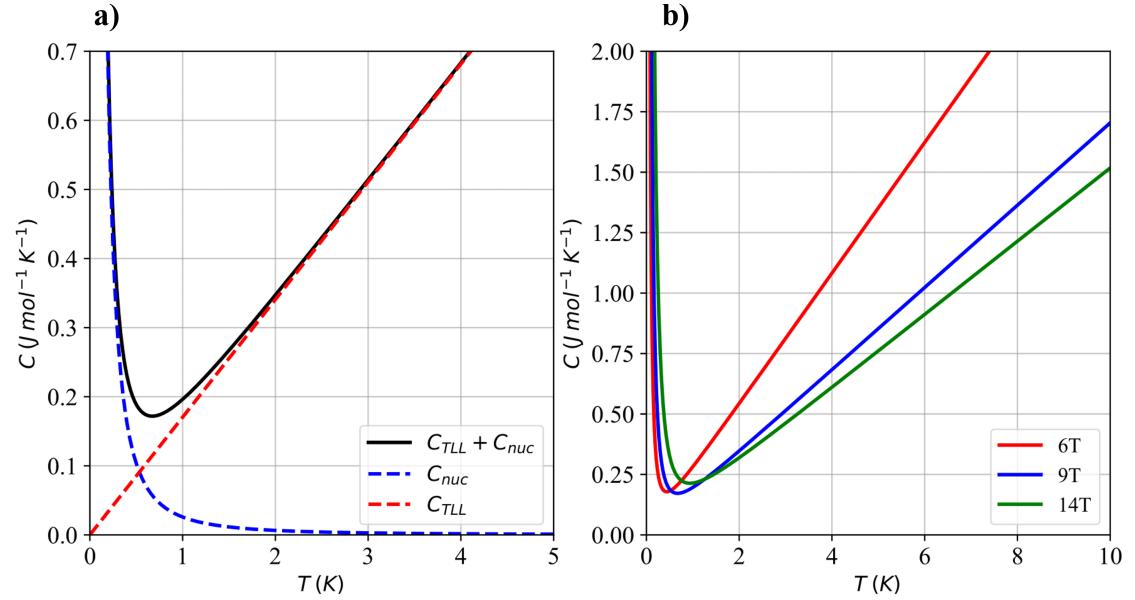


Figure 4.15: a) The total low temperature heat capacity for $\text{NiI}_2(3,5\text{-lut})_4$ at 9 T. The black line indicates the total heat capacity and the red and blue dashed lines indicate the TLL and nuclear component respectively. b) The total calculated low temperature heat capacity for $\text{NiI}_2(3,5\text{-lut})_4$ at 6 T, 9 T and 14 T.

Zero-field measurements

The measured heat capacity contained contributions from both the lattice and magnetic moments in the compound. The field-independent lattice heat capacity was characterised using zero-field heat capacity measurements taken over the temperature range 2 K - 300 K. The magnetic heat capacity was then obtained by subtracting the lattice heat capacity. The characterization of the lattice heat capacity was performed using two methods: Debye-Einstein fitting and comparison to QMC simulations. This section will present the procedures

performed on the zero-field heat capacity measurements in order to obtain the lattice heat capacity.

The Debye-Einstein method consisted of fitting the measured heat capacity data to a Debye-Einstein model of the form

$$C_{lat}(T) = \frac{3A_D}{x_D^3} \int_0^{x_D} \frac{x^4 e^x}{(e^x - 1)^2} dx + \sum_{i=1}^n A_{E,i} \frac{\theta_{E,i}^2}{T^2} \frac{e^{(\theta_{E,i}/T)}}{[e^{(\theta_{E,i}/T)} - 1]^2} \quad (4.23)$$

where $x = \hbar\omega/k_B T$, $x_D = T/\theta_D$, A_D and $A_{E,i}$ are the Debye and Einstein constants and θ_D , $\theta_{E,i}$ are the Debye and Einstein temperatures respectively [186]. Additionally, n is the number of Einstein terms. For $T \ll T_D$ The first term is the Debye term which encodes the characteristic T^3 low temperature behaviour into the model. The remaining n Einstein terms are responsible for the high temperature behaviour.

Table 4.4: Fitting parameters used to characterise the lattice contribution to the the heat capacity using the Debye-Einstein model.

Parameter	Set 1	Set 2
A_D (K)	151	217
θ_D (K)	277	124
$A_{E,1}$ (K)	260	220
$\theta_{E,1}$ (K)	138	-692
$A_{E,2}$ (K)	276	223
$\theta_{E,2}$ (K)	826	-1458
$A_{E,3}$ (K)	4991	170
$\theta_{E,2}$ (K)	3195	-226

Table 4.4 shows the fitting parameters obtained by fitting the zero-field heat capacity data to equation (4.23). Two sets of parameters were generated from the fits using slightly different approaches. Set 1 included n , the number of Einstein terms, as a fitting parameter which yielded a value of $n = 3$ in addition to the rest of the parameters. Set 2 didn't use

n as a fitting parameter, instead a fixed value of $n = 3$ was used due to it being the fitted value from set 1. This was done in order to reduce the number of free parameters within the model in order to improve the fit.

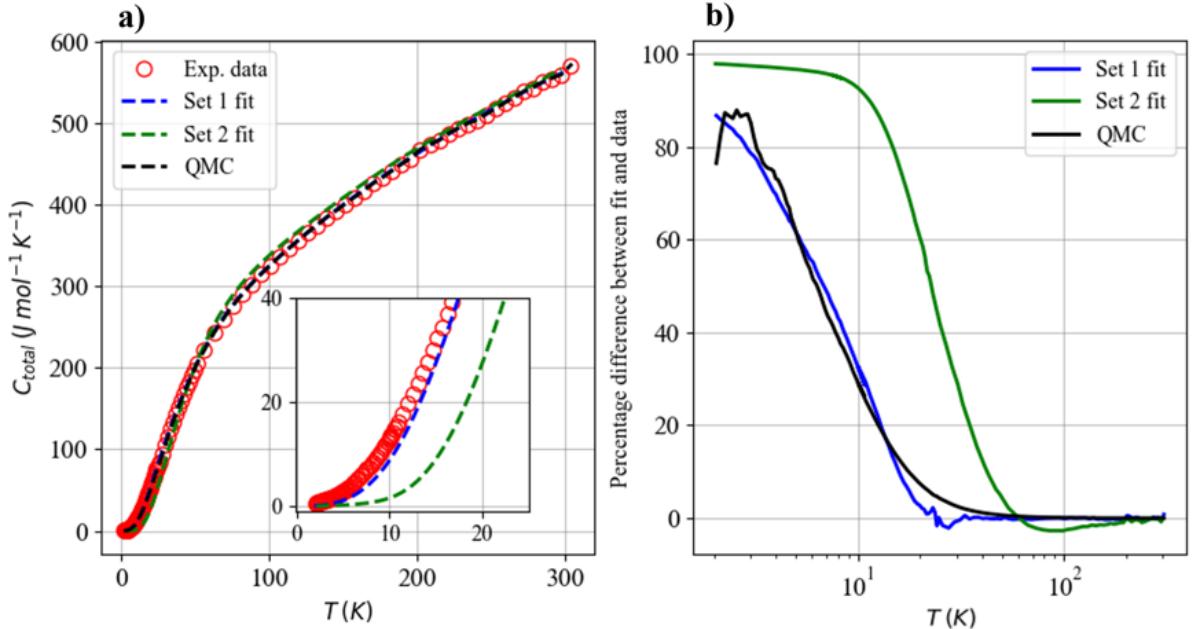


Figure 4.16: a) The zero-field heat capacity of $\text{NiI}_2(3,5\text{-lut})_4$ as a function of temperature. The circles denote the measured data, the blue and green dotted lines indicate the Debye-Einstein fits using parameter set 1 and set 2 respectively. inset) Zoomed into section highlighting the low temperature zero-field heat capacity. b) Percentage difference between the Debye-Einstein fits, QMC method and the experimental data.

Figure 4.16a shows the zero-field heat capacity data with the Debye-Einstein fits using parameter set 1 and set 2 overlaid. In the high temperature regime, both fits match the data well. However, the magnetic specific heat contribution is negligible in the high temperature regime so fit quality in the low temperature regime was examined. The inset of figure 4.16a shows a clear deviation of the set 1 parameter fit from the low temperature data relative to the set 2 parameter fit. Figure 4.16b shows the percentage difference between the fits and the experimental data as a function of temperature. For temperatures less than 10 K, the

fit performed using parameter set 2 provided a closer fit to the data. Consequently, for the following calculation of the magnetic specific heat, the lattice specific heat contribution was modelled using parameter set 2.

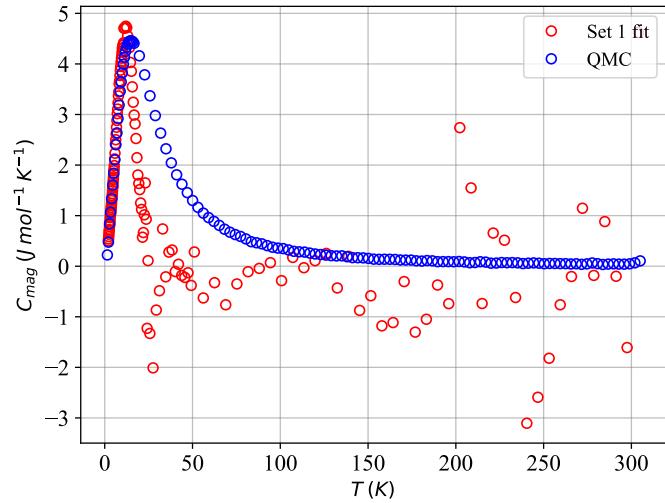


Figure 4.17: The zero-field magnetic heat capacity of $\text{NiI}_2(3,5\text{-lut})_4$ as a function of temperature. The red circles represent the magnetic heat capacity calculated by subtracting the lattice heat capacity modelled using parameter set 2 and the blue circles indicate the result of a QMC simulation.

Figure 4.17 shows the magnetic heat capacity obtained using the Debye-Einstein fit with parameter set 2 and the QMC method. In the Debye-Einstein case this was calculated by subtracting the fit discussed in the previous section from the experimental data. However, this resulted in many negative-valued magnetic heat capacity data points with a large amount of scatter around zero. A negative heat capacity was unphysical and a smooth decrease to zero was expected for this quantity as temperature was increased. This indicated that the Debye-Einstein fitting was unsuitable for this analysis leading to the adoption of a QMC-based method to characterise the lattice contribution of the specific heat.

The QMC approach involved performing a zero-field QMC simulation in order to

obtain the magnetic heat capacity curve shown in figure 4.17. The simulated curve was interpolated using cubic spline interpolation to produce a continuous function. Finally, the interpolated QMC magnetic heat capacity was subtracted from the experimental data to characterise the lattice heat capacity. The performance of this method compared to the Debye-Einstein fits is shown in figure 4.16b. It generated a lattice contribution of a similar profile as the parameter set 2 fit at low temperatures. However, as it was generated via a simulation this ensured the magnetic heat capacity remained non-negative across the entire temperature range. Therefore, the lattice heat capacity generated using the QMC method was used to calculate the in-field magnetic specific heat discussed in the following section.

In-field measurements

This section will present the in-field heat capacity measurements measured from 0 T - 9 T over the temperature range 2 K - 10 K.

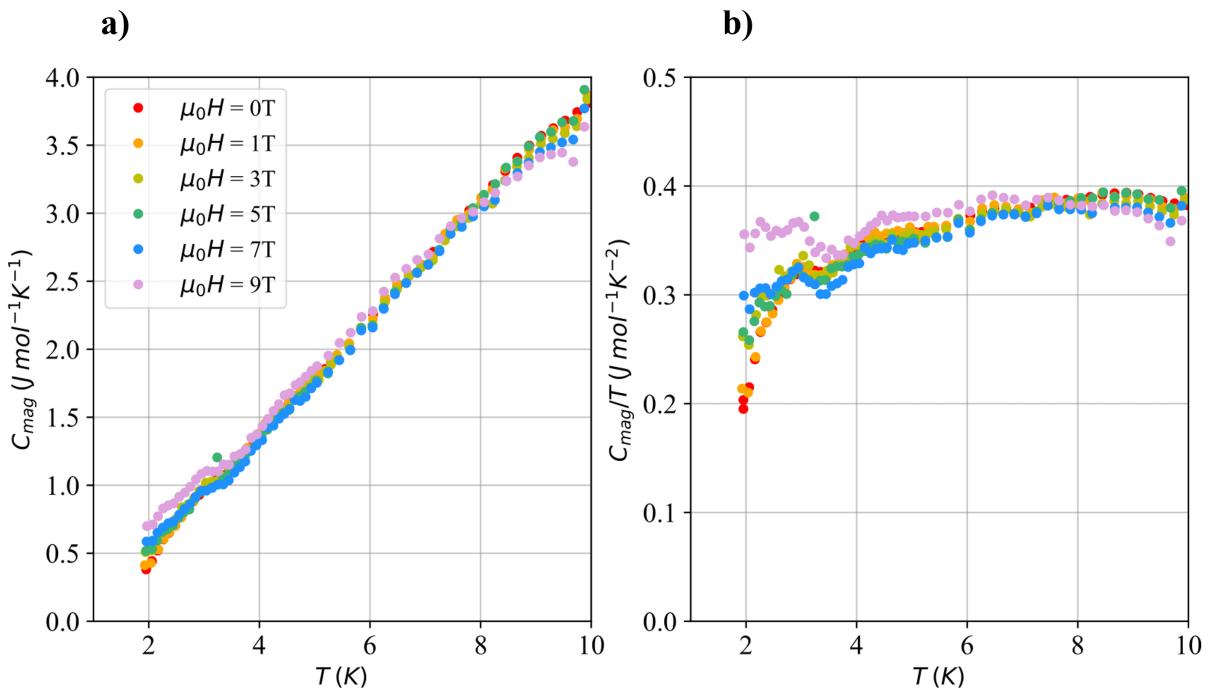


Figure 4.18: a) The in-field magnetic heat capacity of $\text{NiI}_2(3,5\text{-lut})_4$ from 0 T - 9 T. b) The in-field magnetic heat capacity divided by the temperature, to accentuate the low temperature behaviour

Figure (4.18) shows the infiel magnetic heat capacity measurements over the range 1.8 K - 10 K under applied fields 0 T - 9 T. The magnetic heat capacity was calculated by subtracting the lattice heat capacity obtained in the previous section from the measured data. Above 4 K, the data is independent of applied field and all follow a similar curve. Below 4 K, there is a systematic upturn as a function of field for the 5 T, 7 T and 9 T data. This data was obtained in conditions above the estimated value of $\mu_0 H_c \approx 4.4$ T and within a temperature range of the order 1 K which would place the sample in the vicinity of the TLL phase. Therefore, this upturn could signal the transition from the gapped phase to the TLL phase.

4.2.3 ^{13}C NMR

Dipole field calculations

The dipolar field contribution to the Knight shift was calculated by considering the effect of the ^{28}Ni electron spins on each of the 28 ^{13}C sites. By expanding the dot products in equation 2.45, the dipole Hamiltonian can be written in the following form

$$\mathcal{H}_{dip} = -\frac{\mu_0}{4\pi} \gamma_e \gamma_n \hbar^2 \sum_j \mathbf{I} \cdot A_{j,dip} \cdot \mathbf{S}_j^T \quad (4.24)$$

where \mathbf{S}^T denotes the transpose of \mathbf{S} and $A_{j,dip}$ is the dipole interaction tensor for the interaction between ^{13}C nuclear spin, \mathbf{I} , and the j^{th} ^{28}Ni electron spin. The tensor $A_{j,dip}$ encodes the directional dependence of the dipole interaction and has the explicit form

$$A_{j,dip} = \frac{1}{r_j^5} \begin{bmatrix} r_j^2 - 3x_j^2 & -3x_j y_j & 3x_j z_j \\ -3y_j x_j & r_j^2 - 3y_j^2 & 3y_j z_j \\ 3z_j x_j & -3y_j z_j & r_j^2 - 3z_j^2 \end{bmatrix} \quad (4.25)$$

where x_j , y_j and z_j are the distance between the ^{13}C nuclear spin and the ^{28}Ni electron spin at site j . The dipole tensor was then found for a given nuclear spin summing over all electron spins j

$$A_{dip} = \sum_j A_{j,dip} \quad (4.26)$$

The dipole interactions were simulated by generating all ^{28}Ni sites within a sphere of radius r_{max} using the crystallographic vectors. Then the quantity $A_{j,dip}$ was calculated for a given ^{13}C site by iterating through all ^{28}Ni sites within the sphere. Equation (4.25) has a r_j^{-3} dependence so ^{28}Ni sites located far from a ^{13}C site would generate a smaller contribution to the dipole field at the ^{13}C site. Therefore, provided r_{max} was large enough, the value of $A_{j,dip}$ would converge. A value of $r_{max} = 150$ Å was chosen for all calculations. This value was obtained by calculating A_{dip} using a range of r_{max} values. Figure 4.19 shows the convergence

of the diagonal components of A_{dip} for $r_{max} \approx 150 \text{ \AA}$. The total dipole tensor for the ^{13}C site was then calculated by equation (4.26). This process was repeated for each ^{13}C site to yield 28 values for the dipole tensor.

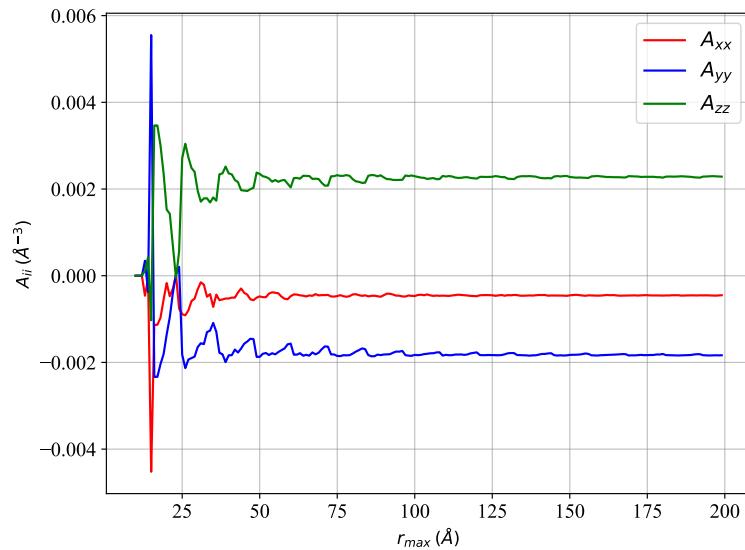


Figure 4.19: The diagonal components of the dipole tensor, A_{dip} , plotted as a function of r_{max} . The components converge for $r_{max} \approx 150 \text{ \AA}$.

The output of the dipole interaction simulation is shown in figure 4.20b. This shows the diagonal components of A_{dip} for all 28 ^{13}C sites, labelled in figure 4.20a. The calculation yielded seven unique dipole tensors demonstrating the symmetry of the ligand structures of $\text{NiI}_2(3,5\text{-lut})_4$. The colours of the ^{13}C sites in figure 4.20a correspond to equivalent sites in each individual ligand. For example, site 0, 13, 14, and 27 all had an identical A_{dip} . Furthermore, within any given lutidine ligand, there were four unique average values of the diagonal components of A_{dip} . For example, within ring 1, sites 0 and 6, 1 and 5, 2 and 4, and site 3 result in four different average values. This indicates that the powder spectrum of $\text{NiI}_2(3,5\text{-lut})_4$ should contain four NMR lines with each line corresponding to a set of sites that produce the same average value of A_{dip} . An example labelling of the equivalent ^{13}C sites on a ligand is shown in inset of figure 4.22c.

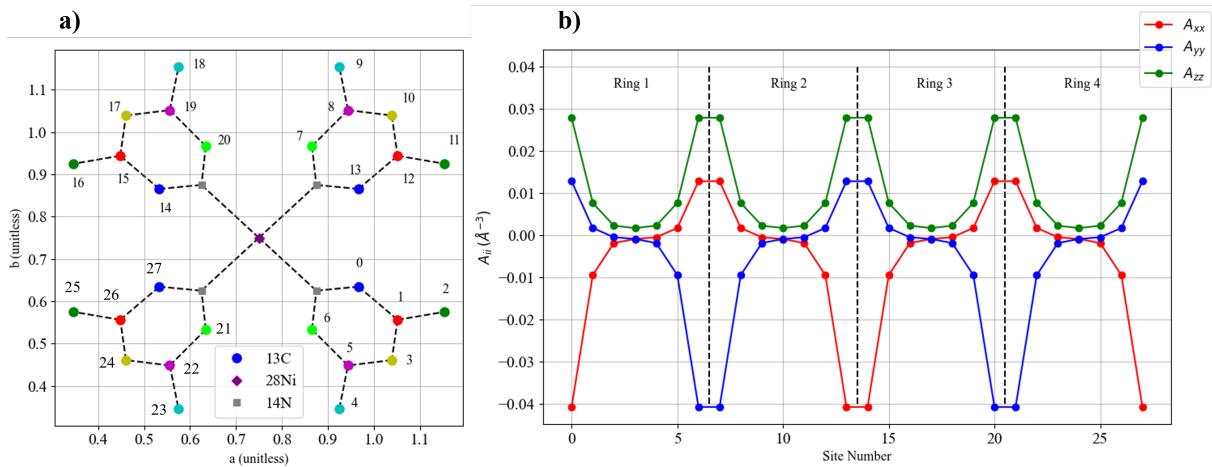


Figure 4.20: a) A diagram of a ^{28}Ni site in $\text{NiI}_2(3,5\text{-lut})_4$ viewed in the a - b plane. ^{13}C sites are numbered and their colour denotes equivalent sites in each ligand. b) The diagonal components of A_{dip} for each ^{13}C site. The site number corresponds to the numbering in a). The lines guide the eye.

^{13}C NMR spectra

The ^{13}C NMR spectra of $\text{NiI}_2(3,5\text{-lut})_4$ was measured experimentally using two techniques: MAS NMR and a conventional spin-echo sequence. The MAS NMR was performed at a fixed temperature of 300 K in a static magnetic field of 11.4 T and a spinning frequency of $f_{rot} = 20$ KHz. The spectrum was averaged over 204,800 repeats with a recycle delay of 750 ms, in order for the signal to fully recover after each acquisition. Figure 4.21 shows the MAS ^{13}C NMR spectrum as a function of NMR shift. Four peaks originating from the sample were identified at 400 ppm, 70 ppm, 0 ppm and -120 ppm respectively, in agreement with the simulations performed in the previous section. The identification of four peaks in the spectrum was expected from the symmetry of the 3,5-lutidine molecule. The peaks labelled with a \star represent spinning sidebands introduced from the 20 kHz spinning frequency. The ‘true’ peaks were identified by considering each individual peak and matching up where the spinning sidebands were located for that peak. The spinning sideband frequencies for the

labelled set of four ‘true’ peaks matched the set of peaks present in spectrum most accurately.

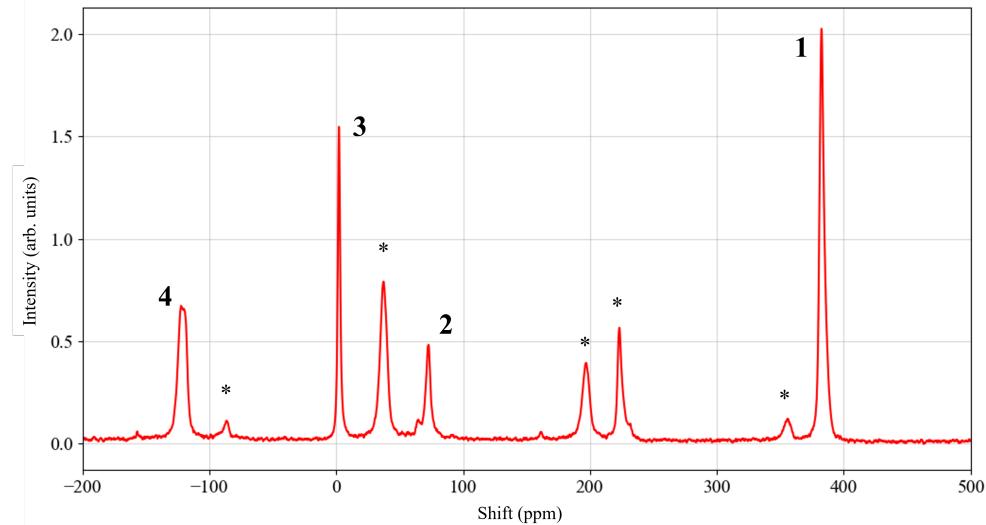


Figure 4.21: The MAS ^{13}C NMR spectra of $\text{NiI}_2(3,5\text{-lut})_4$ at 300 K. The peaks labelled with a number indicate ‘true’ peaks. These labels are consistent with the labels in figure 4.22. The peaks labelled with a \star indicate spinning sidebands.

The conventional NMR results were performed over a temperature range of 2K - 220 K in a static magnetic field of 7.075T, using a spin-echo sequence averaged over 2048 scans with a 150ms recycle delay. Figures 4.22a and 4.22b show the ^{13}C NMR spectra as a function of temperature. There are two distinct peaks present at all temperatures. The central, non-shifting, peak at 0 ppm remains at a constant shift as a function of temperature indicating that it is decoupled from the magnetic interactions with the sample. The shifting peak position increases from 1000 ppm at 300 K, to a maximum shift of 1500 ppm at approximately 15 K before rapidly decreasing to 600 ppm at the lowest temperatures. Figure 4.22d plots the peak position of the shifting and non-shifting peak against temperature. The magnetic susceptibility measured at 7 T is overlaid. The shifting peak position matches the general shape of the susceptibility over the full temperature region indicating that the shifting peak is strongly coupled to the internal magnetic interactions.

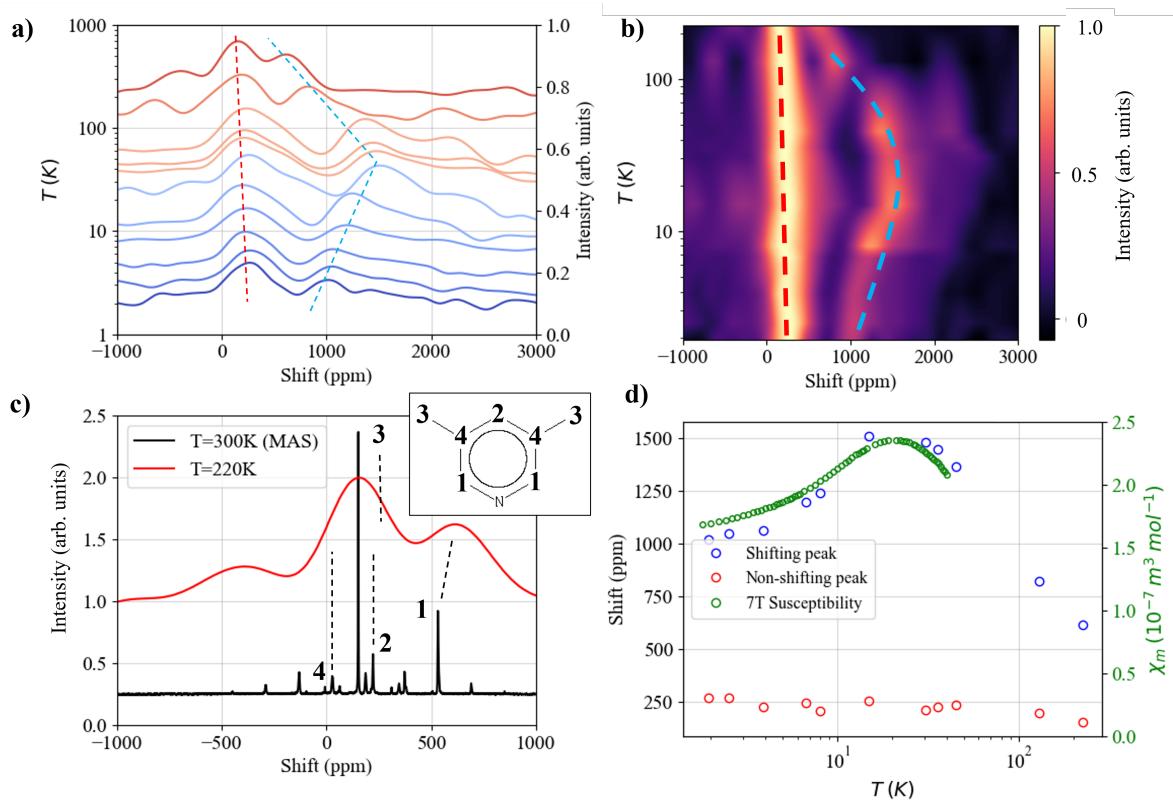


Figure 4.22: a) ^{13}C NMR spectra as a function of temperature measured using the spin echo sequence measured in a magnetic field of 7.075 T. Dotted lines are to guide the eye. b) Colour map generated from the data in a) highlighting the trajectory of each peak as a function of temperature. The dotted lines are to guide the eye to show the movement of the peak as a function of temperature. c) A comparison of the 300 K MAS NMR spectrum with the 220 K spin-echo signal. The peaks are numbered according to figure 4.21. The dotted lines indicated equivalent peaks between each spectrum. inset) The molecular structure of 3-5 lutidine, labelled with the ^{13}C sites found in the spectra. d) The NMR shift of the shifting and non-shifting peaks as a function of temperature overlaid with magnetic susceptibility measured at 7 T. Error bars have been omitted as they are too small to represent.

Figure 4.22c shows the comparison between the 300 K MAS ^{13}C NMR spectrum and the 220 K spin echo spectrum. The MAS spectrum peak numbering is equivalent with figure 4.21. The MAS spectrum was shifted from that in figure 4.21 due to a different choice in

reference frequencies. After the shift was applied, both spectra show good agreement. The dotted lines indicate matching peaks between the spectra. Peak 1 was attributed to the shifting peak whereas peaks 2, 3 and 4 were assigned to the non-shifting peak. This allowed the peaks to be matched with the ^{13}C sites on the 3,5 lutidine ligand shown on the inset of figure 4.22c. This was done by considering the effect of the dipole interaction on each of the ^{13}C sites. The ^{13}C sites in position 1 are closest to the ^{28}Ni site, hence experience a larger dipole interaction. Therefore, this causes strong coupling at these sites leading to the shifting behaviour shown in the temperature dependent spectra. Sites 2,3 and 4 are all located much farther from the ^{28}Ni site so they are less coupled to the dipole interactions and therefore form the non-shifting peak. Furthermore, sites 2,3 and 4 are all positioned at a similar distance from the ^{28}Ni site, which results in their NMR lines being closely spaced away from the line generated from site 1.

T_1 relaxation time

^{13}C T_1 measurements were performed in order to further probe the ^{28}C sites using both MAS NMR and conventional NMR. Both measurements were obtained using the saturation recovery sequence described in section 2.2.2.

Figure 4.23a shows the obtained spectra used to calculate T_1 . Each line represents a different value of τ , the pulse separation. The recovery behaviour is clearly demonstrated in this figure as the amplitude of the 0 ppm peak is still increasing as τ increases which indicates a slow relaxation time. Conversely, the peak at -120 ppm reached a constant amplitude with small values of τ which indicated a fast relaxation time. Figure 4.23b shows the recovery curves used to obtain T_1 for each peak. The intensity of a given peak was calculated by integrating the area under the peak. T_1 was then obtained through fitting to equation (2.35). The values of T_1 for each peak are shown in table 4.5

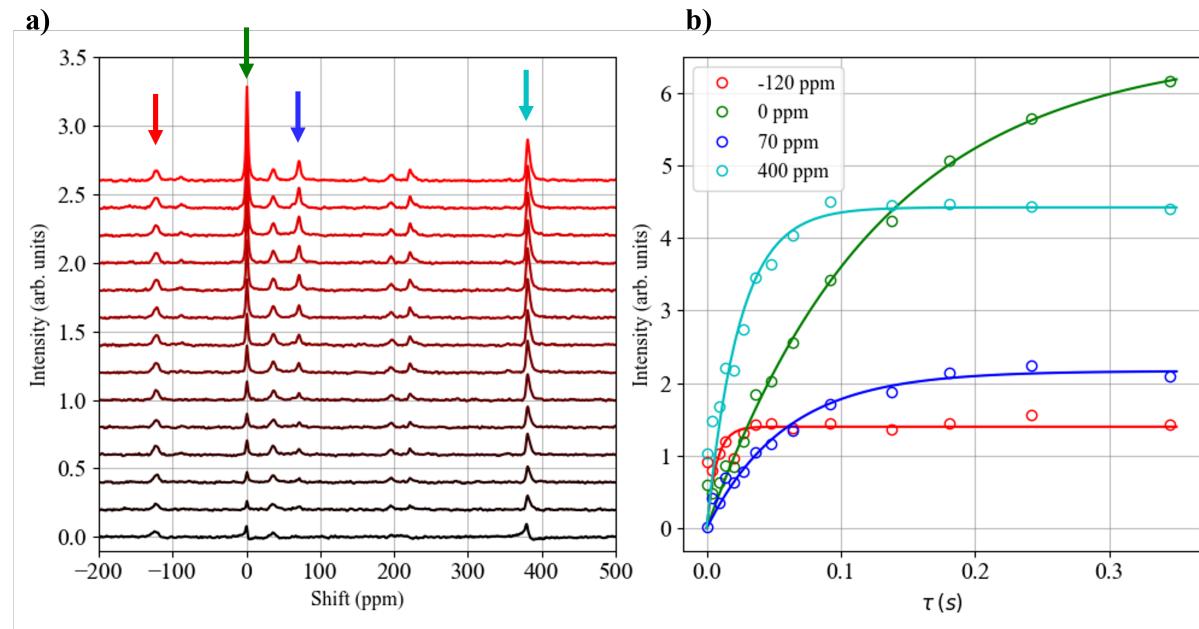


Figure 4.23: a) ^{13}C T_1 spectra obtained from MAS NMR at 300 K using the saturation recovery sequence, the arrows indicate the peaks used to measure T_1 . b) Magnetisation recovery curves calculated for a peak located at the indicated NMR shift. Error bars have been omitted as they are too small to represent.

The 0 ppm peak shows the largest value of T_1 . This agrees well with the ^{13}C site labelling from the previous section. This peak was attributed to the farthest ^{13}C site from the ^{28}Ni spin due to its weak coupling to the internal magnetism. As the site is only weakly coupled, it experiences minimal fluctuations from the dipolar fields and therefore the relaxation time is long. Furthermore, the site 1 peak at 400 ppm showed a relatively small T_1 value indicating strong coupling to fluctuations from the internal fields agreeing well with the discussion in the previous section. The shortest T_1 was obtained for the -120 ppm peak corresponding to site 4. This doesn't agree with the expected relationship with the dipolar field coupling. This difference was attributed to the existence of strong longitudinal fluctuations of the internal field which cannot be characterised using this method.

Table 4.5: ^{13}C T_1 values for the set of highlighted peaks in figure 4.23.

Peak shift (ppm)	T_1 (ms)
-120	7 ± 1
0	127 ± 1
70	58 ± 1
400	24 ± 1

The temperature dependence of the ^{13}C T_1^{-1} inverse time constant was measured using conventional spin echo NMR and the results are shown in figure 4.24. The low temperature data was expected to follow a power law behaviour of $T_1^{-1} \propto T^\alpha$ as outlined in section 2.3. The green region indicates the expected power law behaviour for an attractive TLL state bounded by the lines $\alpha = -0.5$ and $\alpha = 0.75$ corresponding to $K = 1$ and $K = 2$ respectively. The measured data does not lie within the expected region with an exponent of approximately $\alpha=1.25$ resulting in $K = -2$ using equation (2.69). This indicates that upper temperature boundary of the TLL has not been crossed by these measurements, consistent with the heat capacity measurements in the previous section, and $\text{NiL}_2(3,5\text{-lut})_4$ displays some intermediate crossover state down to roughly 1 K.

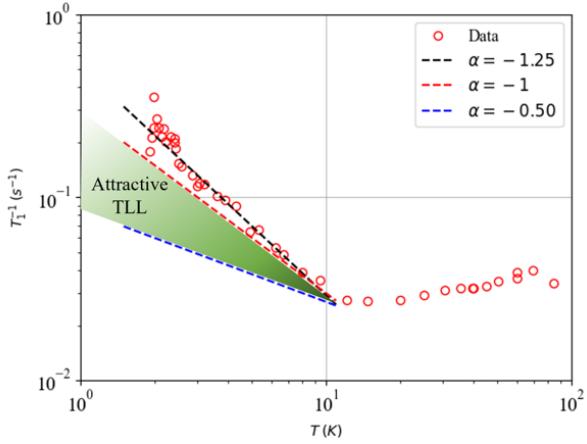


Figure 4.24: The temperature dependence of ^{13}C T_1^{-1} . The dotted lines indicate power laws with varying exponents. The green region indicates the expected power law for an attractive TLL.

4.3 Conclusion and Future Work

In conclusion, this chapter presented the magnetisation, heat capacity and ^{13}C NMR measurements performed on $\text{NiI}_2(3,5-\text{lut})_4$. The magnetisation results were characterised using a Curie-Weiss and a Padé Approximation fit. Quantum critical scaling of the magnetisation around $H_c = 4.4$ T was not successful due to the poor characterisation of the low temperature $S = \frac{1}{2}$ impurity contributions. Additionally, heat capacity measurements were presented in order to calculate the TLL parameters however no signatures of the TLL state were found over the measured temperature regime. One further extension to these measurements would be to obtain a theoretical value of u and K through heat capacity measurements in the mK regime and comparing them to their expected values from DMRG calculations [187]. Then ^{13}C NMR results were presented demonstrating the effectiveness of MAS NMR techniques for use in condensed matter physics. The NMR lines in the spectra were attributed to specific ^{13}C sites within $\text{NiI}_2(3,5-\text{lut})_4$ by using a combination of temperature dependent NMR spectra and T_1 measurements. Finally, the temperature dependence of the ^{13}C T_1 time constant was used to verify $\text{NiI}_2(3,5-\text{lut})_4$ does not demonstrate an attractive TLL state

down to 1 K and instead hosts an unexpected crossover state between the high temperature state and expected TLL state.

This investigation could be extended by performing the magnetisation, heat capacity and NMR measurements down to the mK regime. This would allow for characterisation of the impurity contribution of the magnetisation due to it becoming the dominating component at low temperatures. Furthermore, it would generate a larger range for the expected T -linear behaviour in the heat capacity to be observed. However, the nuclear spin contribution would have to be carefully calculated and subtracted due to its growing contribution at lower temperatures. Finally, the ^{13}C NMR measurements would indicate a crossover to the attractive TLL regime as expected from the phase diagram by showing the expected power law behaviour. Once the temperature boundary for the TLL state was fully established, the TLL can be fully characterised through determination of the parameters K and u by following the approach of [110], [187].

Chapter Five

Conclusion

In conclusion, Chapter 3 introduced the design and implementation of the FPGA-based NMR spectrometer. The Verilog code structure was introduced and the DSP methods used within each module were described. The functionality of the spectrometer was verified through the testing of each Verilog module. Python control software to control the spectrometer was developed. The design and construction of a bespoke NMR probe compatible with the Quantum Design PPMS was discussed. The spectrometer was tested by successfully measuring the spin echo signal from a test sample of ^{59}Co powder using an NMR frequency of 213 MHz. Characterisation of the electronics used to perform the experiment were included. This yielded an NMR spectrum consistent with spectrum from the literature.

Chapter 4 detailed the magnetisation, heat capacity and ^{13}C NMR measurements on the isotropic Haldane chain sample $\text{NiI}_2(3,5\text{-lut})_4$. The zero-field molar susceptibility data was fit to a Curie-Weiss law obtaining a g -factor of $g = 2.5 \pm 0.1$, which deviated from the reported value of $g = 2.197$. A Pade approximation-based fit was used to characterise the impurity contribution to the magnetisation. This was partially successful, providing a good fit in the high temperature regime however poorly characterised the impurity-dominated low temperature data. Quantum critical scaling of the in-field magnetisation around $H_c = 4.4$

T was performed using magnetisation data obtained at 3.8 T - 5.0 T which didn't follow the theoretical behaviour expected from a $z = 2$ quantum critical point, in part due to the poor Pade approximation fit.

The heat capacity behaviour at low temperatures was modelled by considering both the nuclear contribution and TLL contribution. Zero-field heat capacity measurements were performed from 2 K - 300 K in order to characterise the lattice contribution to the specific heat. This was done through two approaches: a Debye-Einstein fit and QMC approach. The in-field magnetic heat capacity was calculated by subtracting the lattice contribution for in-field heat capacity data over the field range 1 T - 9 T. The data with $H > H_c$ showed the appearance of distinct low temperature shoulder occurring at around 3 K. However, there were no signatures of the expected TLL state found within $\text{NiI}_2(3,5\text{-lut})_4$ in these measurements.

The temperature dependence of the ^{13}C NMR spectrum and longitudinal relaxation time constant T_1 were measured using MAS NMR at 300K and conventional spin echo experiments from 220 K - 2 K. Four ^{13}C NMR lines were resolved using the MAS NMR data at shifts of -120 ppm, 0 ppm, 70 ppm and 400 ppm, demonstrating the usefulness of MAS NMR in condensed matter experiments. MAS NMR ^{13}C T_1 measurements yielded values of (7 ± 1) ms, (127 ± 1) ms, (58 ± 1) ms, (24 ± 1) ms respectively. These lines were then matched to the line-broadened spin echo measurements by considering the crystallographic site positions within the lutidine ligand and the T_1 values. Two broad peaks were measured in the spin-echo NMR spectra and the NMR shift of the peak as a function of temperature was measured. One peak was found to have close to zero shift, indicating small coupling to the internal magnetic environment, whereas the other peak showed an NMR shift that mirrored the qualitative shape of the susceptibility measurements, indicated a strong coupling to the internal magnetic interactions. The ^{13}C T_1 measurements yielded a power law dependence on temperature with an approximate exponent of $\alpha = -1.25$. This value is outside the range

for the attractive TLL state expected for an isotropic Haldane chain, indicating the chain was not in a TLL state.

References

- [1] F. Bloch, “Nuclear induction,” *Phys. Rev.*, vol. 70, no. 7-8, pp. 460–474, 1946. DOI: [10.1103/PhysRev.70.460](https://doi.org/10.1103/PhysRev.70.460).
- [2] F. Bloch, W.E.Hansen, and M.Packard, “The Nuclear Induction Experiment,” *Phys. Rev.*, vol. 70, no. 5, pp. 474–485, 1946.
- [3] E. M. Purcell, H. C. Torrey, and R. V. Pound, “Resonance Absorption by a Nuclear Magnetic Moments in a Solid,” *Phys. Rev.*, vol. 69, no. 1-2, pp. 37–38, 1947. DOI: [10.1103/PhysRev.69.37](https://doi.org/10.1103/PhysRev.69.37).
- [4] C. Berthier, M. Horvatić, M. H. Julien, H. Mayaffre, and S. Krämer, “Résonance magnétique nucléaire en champs magnétiques intenses: application à la physique de la matière condensée,” *Comptes Rendus Physique*, vol. 18, no. 5-6, pp. 331–348, 2017. DOI: [10.1016/j.crhy.2017.09.009](https://doi.org/10.1016/j.crhy.2017.09.009).
- [5] R. Tycko, “NMR at low and ultralow temperatures,” *Acc. Chem. Res.*, vol. 46, no. 9, pp. 1923–1932, 2013. DOI: [10.1021/ar300358z](https://doi.org/10.1021/ar300358z).
- [6] R. Řezníček, V. Chlan, and J. Haase, “Nmr and ab initio study of gallium metal under pressure,” *Phys. Rev. B*, vol. 99, no. 12, p. 125121, 2019. DOI: [10.1103/PhysRevB.99.125121](https://doi.org/10.1103/PhysRevB.99.125121).

[7] H. Kühne, H. H. Klauss, S. Grossjohann, *et al.*, “Quantum critical dynamics of an S=1/2 antiferromagnetic Heisenberg chain studied by C 13 NMR spectroscopy,” *Phys. Rev. B*, vol. 80, no. 4, pp. 1–5, 2009. DOI: [10.1103/PhysRevB.80.045110](https://doi.org/10.1103/PhysRevB.80.045110).

[8] W. G. Clark, “Pulsed nuclear resonance apparatus,” *Rev. Sci. Instrum.*, vol. 35, no. 3, pp. 316–333, 1964. DOI: [10.1063/1.1718815](https://doi.org/10.1063/1.1718815).

[9] H. Y. Chen, Y. Kim, P. Nath, and C. Hilty, “An ultra-low cost NMR device with arbitrary pulse programming,” *J. Magn. Reson.*, vol. 255, pp. 100–105, 2015. DOI: [10.1016/j.jmr.2015.02.011](https://doi.org/10.1016/j.jmr.2015.02.011).

[10] K. Takeda, “OPENCORE NMR: Open-source core modules for implementing an integrated FPGA-based NMR spectrometer,” *J. Magn. Reson.*, vol. 192, no. 2, pp. 218–229, 2008. DOI: [10.1016/j.jmr.2008.02.019](https://doi.org/10.1016/j.jmr.2008.02.019).

[11] K. Takeda, “A highly integrated FPGA-based nuclear magnetic resonance spectrometer,” *Rev. Sci. Instrum.*, vol. 78, no. 3, pp. 0–6, 2007. DOI: [10.1063/1.2712940](https://doi.org/10.1063/1.2712940).

[12] P. Hemnani, A. K. Rajarajan, G. Joshi, and S. V. G. Ravindranath, “FPGA Based RF Pulse Generator for NQR/NMR Spectrometer,” *Procedia Computer Science*, vol. 93, pp. 161–168, 2016. DOI: [10.1016/j.procs.2016.07.196](https://doi.org/10.1016/j.procs.2016.07.196).

[13] M. Pikulski, T. Shiroka, H. R. Ott, and J. Mesot, “A firmware-defined digital direct-sampling NMR spectrometer for condensed matter physics,” *Rev. Sci. Instrum.*, vol. 85, no. 9, 2014. DOI: [10.1063/1.4896351](https://doi.org/10.1063/1.4896351).

[14] J. Keeler, “Understanding NMR spectroscopy,” in Second Edition, NY, USA: John Wiley & Sons, 2011, pp. 3–6, 3–8, 4–5, 4–6, 5–11, 5–12, ISBN: 8126545887.

[15] E. Ising, “Beitrag zur theorie des ferromagnetismus,” *Zeitschrift für Physik*, vol. 31, 1 1925. DOI: <https://doi.org/10.1007/BF02980577>.

[16] L. Onsager, “Crystal statistics. i. a two-dimensional model with an order-disorder transition,” *Phys. Rev.*, vol. 65, no. 3-4, pp. 117–149, 1944. DOI: [10.1103/PhysRev.65.117](https://doi.org/10.1103/PhysRev.65.117).

[17] A. S. T. Pires, “The Heisenberg model,” in *Theoretical Tools for Spin Models in Magnetic Systems*, Bristol, UK: IOP Publishing, 2021, pp. 1–16. DOI: [10.1088/978-0-7503-3879-0ch1](https://doi.org/10.1088/978-0-7503-3879-0ch1).

[18] J. P. Renard, L. P. Regnault, and M. Verdaguer, “Haldane quantum spin chains,” in *Magnetism: Molecules to Materials I*, New Jersey, USA: John Wiley & Sons, Ltd, 2001, pp. 49–93. DOI: <https://doi.org/10.1002/3527600841.ch2>.

[19] C. Rovira, “Molecular spin ladders,” *Chem. Eur. J.*, vol. 6, no. 10, pp. 1723–1729, 2000. DOI: [https://doi.org/10.1002/\(SICI\)1521-3765\(20000515\)6:10<1723::AID-CHEM1723>3.0.CO;2-Z](https://doi.org/10.1002/(SICI)1521-3765(20000515)6:10<1723::AID-CHEM1723>3.0.CO;2-Z).

[20] T. Barnes, E. Dagotto, J. Riera, and E. S. Swanson, “Excitation spectrum of heisenberg spin ladders,” *Phys. Rev. B*, vol. 47, no. 6, pp. 3196–3203, 1993. DOI: [10.1103/PhysRevB.47.3196](https://doi.org/10.1103/PhysRevB.47.3196).

[21] Q. Li, H. Li, J. Zhao, H.-G. Luo, and Z. Y. Xie, “Magnetization of the spin- $\frac{1}{2}$ heisenberg antiferromagnet on the triangular lattice,” *Phys. Rev. B*, vol. 105, no. 18, p. 184418, 2022. DOI: [10.1103/PhysRevB.105.184418](https://doi.org/10.1103/PhysRevB.105.184418).

[22] S. R. White and A. L. Chernyshev, “Neél order in square and triangular lattice heisenberg models,” *Phys. Rev. Lett.*, vol. 99, no. 12, p. 127004, 2007. DOI: [10.1103/PhysRevLett.99.127004](https://doi.org/10.1103/PhysRevLett.99.127004).

[23] L. Tao, “A continuous family of fully frustrated heisenberg models on the kagome lattice,” vol. 133, no. 4, p. 47001, 2021. DOI: [10.1209/0295-5075/133/47001](https://doi.org/10.1209/0295-5075/133/47001).

[24] B. Thielemann, C. Rüegg, H. M. Rønnow, *et al.*, “Direct observation of magnon fractionalization in the quantum spin ladder,” *Phys. Rev. Lett.*, vol. 102, no. 10, p. 107204, 2009. DOI: [10.1103/PhysRevLett.102.107204](https://doi.org/10.1103/PhysRevLett.102.107204).

[25] D. Schmidiger, P. Bouillot, T. Guidi, *et al.*, “Spectrum of a magnetized strong-leg quantum spin ladder,” *Phys. Rev. Lett.*, vol. 111, no. 10, p. 107202, 2013. DOI: [10.1103/PhysRevLett.111.107202](https://doi.org/10.1103/PhysRevLett.111.107202).

[26] K. Tetsuya, H. Katsuhiro, and I. Katsunori, “Magnetism of the ni₂₊ kagome lattice in rb₂ni₃s₄,” *J. Magn. Magn. Mater.*, vol. 177-181, pp. 591–592, 1998. DOI: [https://doi.org/10.1016/S0304-8853\(97\)00515-5](https://doi.org/10.1016/S0304-8853(97)00515-5).

[27] N. Oba, C. Michioka, M. Kato, K. Yoshimura, and K. Mibu, “Physical properties of the novel jarosite-type compound nafes₃(seo₄)₂(oh)₆,” *J. Phys. Chem. Sol.*, vol. 66, no. 8, pp. 1438–1441, 2005. DOI: <https://doi.org/10.1016/j.jpcs.2005.05.027>.

[28] E. Nocerino, C. Witteveen, S. Kobayashi, *et al.*, “Nuclear and magnetic spin structure of the antiferromagnetic triangular lattice compound lictre2 investigated by usr, neutron and x-ray diffraction,” *Sci. Rep.*, vol. 12, no. 1, pp. 1438–1441, 2022. DOI: <https://doi.org/10.1038/s41598-022-25921-9>.

[29] T. M. Brill and J. P. Boucher, “Excitations in haldane spin chains: Electron spin resonance in nеп,” *Phys. Scr.*, vol. 1994, no. T55, p. 156, 1994. DOI: [10.1088/0031-8949/1994/T55/027](https://doi.org/10.1088/0031-8949/1994/T55/027).

[30] F. D. M. Haldane, “Nonlinear field theory of large-spin heisenberg antiferromagnets: Semiclassically quantized solitons of the one-dimensional easy-axis néel state,” *Phys. Rev. Lett.*, vol. 50, no. 15, pp. 1153–1156, 1983. DOI: [10.1103/PhysRevLett.50.1153](https://doi.org/10.1103/PhysRevLett.50.1153).

[31] O. V. Maximov, S. V. Streltsov, and A. N. Vasiliev, “Long range ordered, dimerized, large-d and haldane phases in spin 1 chain compounds,” *Crit. Rev. Solid State Mater. Sci.*, vol. 46, no. 4, pp. 371–383, 2021. DOI: [10.1080/10408436.2020.1852911](https://doi.org/10.1080/10408436.2020.1852911).

[32] P. B. Johnson, J. A. Rayne, and S. A. Friedberg, “Magnetic properties of csnicl3 and rbnicl3,” *J. Appl. Phys*, vol. 50, no. B3, pp. 1853–1855, 2008. DOI: [10.1063/1.327143](https://doi.org/10.1063/1.327143).

[33] I. A. Zaliznyak, S.-H. Lee, and S. V. Petrov, “Continuum in the spin-excitation spectrum of a haldane chain observed by neutron scattering in csnicl3,” *Phys. Rev. Lett.*, vol. 87, no. 1, p. 017202, 2001. DOI: [10.1103/PhysRevLett.87.017202](https://doi.org/10.1103/PhysRevLett.87.017202).

[34] A. N. Vasiliev, O. L. Ignatchik, M. Isobe, and Y. Ueda, “Long range néel order in the quasi-one-dimensional vanadium-based $s=1$ pyroxenes $(\text{lina})\text{v}(\text{si},\text{ge})_2\text{o}_6$,” *Phys. Rev. B*, vol. 70, no. 13, p. 132415, DOI: [10.1103/PhysRevB.70.132415](https://doi.org/10.1103/PhysRevB.70.132415).

[35] Y. Narumi, R. Sato, K. Kindo, and M. Hagiwara, “Magnetic property of an $s=1$ antiferromagnetic dimer compound,” *J. Magn. Magn. Mater.*, vol. 177-181, pp. 685–686, 1998. DOI: [https://doi.org/10.1016/S0304-8853\(97\)00549-0](https://doi.org/10.1016/S0304-8853(97)00549-0).

[36] Y. Narumi, K. Kindo, M. Hagiwara, *et al.*, “High-field magnetization of $s=1$ antiferromagnetic bond-alternating chain compounds,” *Phys. Rev. B*, vol. 69, no. 17, p. 174405, 2004. DOI: [10.1103/PhysRevB.69.174405](https://doi.org/10.1103/PhysRevB.69.174405).

[37] N. Tsyrulin, C. D. Batista, V. S. Zapf, *et al.*, “Neutron study of the magnetism in $\text{ni}\text{cl}_{24}\text{sc}(\text{nh}_2)_2$,” *J. Phys. Condens. Matter*, vol. 25, no. 21, p. 216008, 2013. DOI: [10.1088/0953-8984/25/21/216008](https://doi.org/10.1088/0953-8984/25/21/216008).

[38] T. Takeuchi, H. Hori, T. Yosida, *et al.*, “Magnetization process of haldane materials tmn_{11} and nina_2 ,” *J. Phys. Soc. Jpn*, vol. 61, no. 9, pp. 3262–3266, 1992. DOI: [10.1143/JPSJ.61.3262](https://doi.org/10.1143/JPSJ.61.3262).

[39] T. Goto, T. Ishikawa, Y. Shimaoka, and Y. Fujii, “Quantum spin dynamics studied by the nuclear magnetic relaxation of protons in the haldane-gap system $(\text{ch}_3)_4\text{nni}(\text{no}_2)_3$,” *Phys. Rev. B*, vol. 73, no. 21, p. 214406, 2006. DOI: [10.1103/PhysRevB.73.214406](https://doi.org/10.1103/PhysRevB.73.214406).

[40] S. Tomonaga, “Remarks on bloch’s method of sound waves applied to many-fermion problems,” *Prog. Theor. Exp. Phys.*, vol. 5, no. 4, pp. 544–569, 1950. DOI: [10.1143/ptp/5.4.544](https://doi.org/10.1143/ptp/5.4.544).

[41] J. M. Luttinger, “An exactly soluble model of a many-fermion system,” *J. Math. Phys.*, vol. 4, no. 9, pp. 1154–1162, 1963. DOI: [10.1063/1.1704046](https://doi.org/10.1063/1.1704046).

[42] R. Egger and A. O. Gogolin, “Effective low-energy theory for correlated carbon nanotubes,” *Phys. Rev. Lett.*, vol. 79, no. 25, pp. 5082–5085, 1997. DOI: [10.1103/PhysRevLett.79.5082](https://doi.org/10.1103/PhysRevLett.79.5082).

[43] H. Ishii, H. Kataura, H. Shiozawa, *et al.*, “Direct observation of tomonaga–luttinger-liquid state in carbon nanotubes at low temperatures,” *Nature*, vol. 426, no. 6966, pp. 540–544, 2003. DOI: [10.1038/nature02074](https://doi.org/10.1038/nature02074).

[44] K.-I. Imura, K.-V. Pham, P. Lederer, and F. Piéchon, “Conductance of one-dimensional quantum wires,” *Phys. Rev. B*, vol. 66, no. 3, p. 035313, 2002. DOI: [10.1103/PhysRevB.66.035313](https://doi.org/10.1103/PhysRevB.66.035313).

[45] C.-Y. Hou, A. Rahmani, A. E. Feiguin, and C. Chamon, “Junctions of multiple quantum wires with different luttinger parameters,” *Phys. Rev. B*, vol. 86, no. 7, p. 075451, 2012. DOI: [10.1103/PhysRevB.86.075451](https://doi.org/10.1103/PhysRevB.86.075451).

[46] X. G. Wen, “Chiral luttinger liquid and the edge excitations in the fractional quantum hall states,” *Phys. Rev. B*, vol. 41, no. 18, pp. 12 838–12 844, 1990. DOI: [10.1103/PhysRevB.41.12838](https://doi.org/10.1103/PhysRevB.41.12838).

[47] Y. Wang, V. Ponomarenko, Z. Wan, *et al.*, “Transport in helical luttinger liquids in the fractional quantum hall regime,” *Nature Communications*, vol. 12, no. 1, p. 5312, 2021. DOI: [10.1038/s41467-021-25631-2](https://doi.org/10.1038/s41467-021-25631-2).

[48] R. C. Williams, W. J. A. Blackmore, S. P. M. Curley, *et al.*, “Near-ideal molecule-based haldane spin chain,” *Phys. Rev. Res.*, vol. 2, no. 1, p. 013082, 2020. DOI: [10.1103/PhysRevResearch.2.013082](https://doi.org/10.1103/PhysRevResearch.2.013082).

[49] F. Nebeker, *Fifty Years of Signal Processing*. New Jersey, USA: IEEE Signal Processing Society, 1998, pp. 1–3. [Online]. Available: <https://signalprocessingsociety.org/uploads/history/history.pdf>.

[50] E. Davies, “Chapter 1 - vision, the challenge,” in *Computer Vision: Algorithms and Applications*, Fifth Edition, Berlin, Germany: Springer Science and Business Media, 2018, pp. 1–15.

[51] J. Meng, J. Zhang, and H. Zhao, “Overview of the speech recognition technology,” *2012 Fourth International Conference on Computational and Information Sciences*, pp. 199–202, 2012. DOI: [10.1109/ICCIS.2012.202](https://doi.org/10.1109/ICCIS.2012.202).

[52] L. E. Frenzel, “Chapter 5 - electronic circuits: Digital: Practically everything is digital these days,” in *Electronics Explained: Fundamentals for Engineers, Technicians, and Makers*, Second Edition, Oxford, UK: Newnes, 2018, pp. 134–135. DOI: [10.1016/B978-0-12-811641-8.00005-9](https://doi.org/10.1016/B978-0-12-811641-8.00005-9).

[53] A. Bensky, “Chapter 5 - signal generation and transmitters,” in *Short-range Wireless Communication*, Third Edition, Amsterdam, The Netherlands: Elsevier, 2019, pp. 139–141. DOI: [10.1016/C2017-0-02356-X](https://doi.org/10.1016/C2017-0-02356-X).

[54] P. R. Dale, S. W. Fardo, V. Chandra, and B. W. Fardo, “7.3.2 crystal-controlled clocks,” in *Electronic Digital System Fundamentals*, Second Edition, River Publishers, 2023, pp. 182–183, ISBN: 978-87-7022-739-1. [Online]. Available: <https://app.knovel.com/mlink/khtml/id:kt013E1E63/electronic-digital-system/crystal-controlled-clocks>.

[55] T. J. Roush, “Chapter 5 - signal sampling and distortion,” in *Wireless Receiver Architectures and Design*, First Edition, Amsterdam, The Netherlands: Elsevier, 2014, pp. 263–303. DOI: [10.1016/b978-0-12-378640-1.00005-6](https://doi.org/10.1016/b978-0-12-378640-1.00005-6).

[56] G. Brookshear and D. Brylow, “Computer science: An overview,” in Thirteenth edition, London, UK: Pearson Education, 2019, pp. 70–75, ISBN: 9781292263441.

[57] Xilinx Inc., *DDS Compiler*, 2021. [Online]. Available: https://www.xilinx.com/products/intellectual-property/dds_compiler.html#overview.

[58] G. Ali, “Chapter 7 - passband digital transmission,” in *Introduction to Digital Communications*, Amsterdam, The Netherlands: Elsevier, 2016, pp. 340–342. DOI: <https://doi.org/10.1016/C2012-0-06171-6>.

[59] A. Balarabe, R. Ahmad, Z. L. Hassan, and M. Aliyu, “A Comprehensive Analysis of M-ary PSK and M-ary QAM Schemes,” *Asian Journal of Research in Computer Science*, vol. 7, no. 3, pp. 63–71, 2021. DOI: [10.9734/ajrcos/2021/v7i330183](https://doi.org/10.9734/ajrcos/2021/v7i330183).

[60] M. Steer, “2.13: Phase Shift Keying Modulation,” in *Microwave and RF Design I - Radio Systems*, CA, USA: LibreTexts Engineering, 2023. [Online]. Available: [https://eng.libretexts.org/Bookshelves/Electrical_Engineering/Electronics/Microwave_and_RF_Design_I_-_Radio_Systems_\(Steer\)](https://eng.libretexts.org/Bookshelves/Electrical_Engineering/Electronics/Microwave_and_RF_Design_I_-_Radio_Systems_(Steer)).

[61] S. Smith, *The Scientist’s and Engineer’s Guide to Digital Signal Processing*, First Edition. CA, USA: California Technical Publishing, 1997, pp. 40, 41, 46, 47, 96–98, 177–178, 320. [Online]. Available: <http://www.dspguide.com/>.

[62] *Mathematical basics of band-limited sampling and aliasing*, Accessed: 2023-10-20, 2005. [Online]. Available: <https://pdfserv.maximintegrated.com/en/an/TUT3628.pdf>.

[63] A. Oppenheim, R. Schafer, and J. Buck, “Sampling of Continuous-Time Signals,” in *Discrete Time Signal Processing*, Second Edition, NJ, USA: Prentice-Hall Inc., 1998, p. 149, ISBN: 0137549202.

[64] M. Parker, *Digital Signal Processing 101*, Second Edition. Oxford, UK: Newnes, 2017, pp. 65–66, 75–82. DOI: [10.1016/b978-0-12-811453-7.00008-1](https://doi.org/10.1016/b978-0-12-811453-7.00008-1).

[65] V. Schmid and M. B. Anke, “Chapter 3 - subband coding and wavelet transform,” in *Pattern Recognition and Signal Analysis in Medical Imaging*, Second Edition, Amsterdam, The Netherlands: Elsevier, 2014, pp. 71–111. DOI: [10.1016/C2012-0-00347-X](https://doi.org/10.1016/C2012-0-00347-X).

[66] R. Oshana, “Chapter 7 - overview of dsp algorithms,” in *Overview of DSP Algorithms*, First Edition, Oxford, UK: Newnes, 2012, pp. 113–127. DOI: [10.1016/B978-0-12-386535-9.00007-X](https://doi.org/10.1016/B978-0-12-386535-9.00007-X).

[67] W. van Drongelen, “Chapter 18 - filters: Digital filters,” in *Signal Processing for Neuroscientists*, Second Edition, Amsterdam, The Netherlands: Elsevier, 2018, pp. 345–354. DOI: [10.1016/B978-0-12-810482-8.00018-7](https://doi.org/10.1016/B978-0-12-810482-8.00018-7).

[68] R. Mortimer, “Functional series and integral transforms,” in *Mathematics for Physics Chemistry*, Fourth Edition, Amsterdam, The Netherlands: Elsevier, 2013, pp. 136–138, ISBN: 0124158099.

[69] I. A. Sulaiman, H. M. Hassan, M. Danish, M. Singh, P. K. Singh, and M. Rajoriya, “Design, comparison and analysis of low pass fir filter using window techniques method,” *Mater. Today*, vol. 49, pp. 3117–3121, 2022. DOI: [10.1016/j.matpr.2020.10.952](https://doi.org/10.1016/j.matpr.2020.10.952).

[70] R. M. Lewitt, “Multidimensional digital image representations using generalized kaiser-bessel window functions,” *J. Opt. Soc. Am. A*, vol. 7, no. 10, 1990. DOI: [10.1364/josaa.7.001834](https://doi.org/10.1364/josaa.7.001834).

[71] D. Tufts, D. Rorabacher, and W. Mosier, “Designing simple, effective digital filters,” *IEEE Transaction on Audio and Electroacoustics*, vol. 18, no. 2, pp. 142–158, 1970. DOI: [10.1109/TAU.1970.1162094](https://doi.org/10.1109/TAU.1970.1162094).

[72] D. J. Goodman and M. J. Carey, “Nine digital filters for decimation and interpolation,” *IEEE Transactions on acoustics, speech and signal processing*, vol. 29, no. 2, pp. 121–126, 1977. DOI: [10.1109/TASSP.1977.1162921](https://doi.org/10.1109/TASSP.1977.1162921).

[73] E. Monmasson and M. N. Cirstea, “Fpga design methodology for industrial control systems - a review,” *IEEE Transactions on Industrial Electronics*, vol. 54, no. 4, pp. 1824–1842, 2007. DOI: [10.1109/TIE.2007.898281](https://doi.org/10.1109/TIE.2007.898281).

[74] C. Vincenzo, S. Andrea, and S. Pesce, “Chapter thirteen - on-board implementation,” in *Modern Spacecraft Guidance, Navigation, and Control*, Amsterdam, The Netherlands: Elsevier, 2023, pp. 685–712. DOI: [10.1016/C2020-0-03563-2](https://doi.org/10.1016/C2020-0-03563-2).

[75] M. Swarup, S. Steve, and B. Mohamad, “Chapter 4 - signal processing hardware,” in *Implantable Biomedical Microsystems*, Amsterdam, The Netherlands: Elsevier, 2015, pp. 73–81. DOI: [10.1016/C2013-0-12618-9](https://doi.org/10.1016/C2013-0-12618-9).

[76] C. Maxfield, “Chapter 9 - HDL-Based Design Flows,” in *The Design Warrior’s Guide to FPGAs: Devices, Tools and Flows*, Oxford, UK: Newnes, 2004, pp. 153–177. DOI: <https://doi.org/10.1016/B978-075067604-5/50010-6>.

[77] V. J. Coli, “Fpga design technology,” *Microprocessors and microsystems*, vol. 17, no. 7, pp. 383–389, 1993. DOI: [https://doi.org/10.1016/0141-9331\(93\)90060-K](https://doi.org/10.1016/0141-9331(93)90060-K).

[78] IEEE Computer Society, “IEEE Standard for Verilog Hardware Description Language.,” *IEEE Std 1364-2005 (Revision of IEEE Std 1364-2001)*, pp. 1–590, 2006. DOI: [10.1109/IEEESTD.2006.99495](https://doi.org/10.1109/IEEESTD.2006.99495).

[79] IEEE Computer Society, “IEEE Standard for VHDL Language Reference Manual,” *IEEE Std 1076-2019*, pp. 1–673, 2019. DOI: [10.1109/IEEESTD.2019.8938196](https://doi.org/10.1109/IEEESTD.2019.8938196).

[80] S. Harris and D. Harris, “3 - sequential logic design,” in *Digital Design and Computer Architecture*, MA, USA: Morgan Kaufmann, 2016, pp. 108–171. DOI: [10.1016/C2013-0-14352-8](https://doi.org/10.1016/C2013-0-14352-8).

[81] D. Harris, “5 - clocking,” in *Skew-Tolerant Circuit Design*, MA, USA: Morgan Kaufmann, 2001, pp. 143–160. DOI: [10.1016/B978-1-55860-636-4.X5000-X](https://doi.org/10.1016/B978-1-55860-636-4.X5000-X).

[82] J. Hennessy and D. Patterson, “Fundamentals of quantitative design and analysis,” in *Computer Architecture: A Quantitive Approach*, Fifth Edition, MA, USA: Morgan Kaufmann, 1995, pp. 43–46, ISBN: 012383872X.

[83] O. Cadenas and G. Megson, “A clocking technique for fpga pipelined designs,” *Journal of Systems Architecture*, vol. 50, no. 11, pp. 687–696, 2004. DOI: [10.1016/j.sysarc.2004.04.001](https://doi.org/10.1016/j.sysarc.2004.04.001).

[84] M. H. Levitt, “2. Magnetism,” in *Spin Dynamics: Basics of Nuclear Magnetic Resonance*, Second Edition, NY, USA: John Wiley & Sons, 2008, pp. 23–26, ISBN: 0471489220.

[85] N. Curro, “Nuclear magnetic resonance as a probe of strongly correlated electron systems,” in *Strongly Correlated Systems: Experimental Techniques*, Berlin, Germany: Springer Berlin, Heidelberg, 2015, pp. 5–6, 14–15, 21. DOI: [10.1007/978-3-662-44133-6](https://doi.org/10.1007/978-3-662-44133-6).

[86] B. Bransden and C. J. Joachain, “14. Quantum Statistics,” in *Quantum Mechanics*, Second Edition, NJ, USA: Prentice Hall, 2000, pp. 641–657, ISBN: 0582356911.

[87] C. Bengs and M. H. Levitt, “A master equation for spin systems far from equilibrium,” *J. Magn. Reson.*, vol. 310, p. 106 645, 2020. DOI: <https://doi.org/10.1016/j.jmr.2019.106645>.

[88] P. Wittek, “4 - quantum computing,” in *Quantum Machine Learning*, Amsterdam, The Netherlands: Elsevier, 2014, pp. 41–45. DOI: [10.1016/C2013-0-19170-2](https://doi.org/10.1016/C2013-0-19170-2).

[89] F. Bloch, “Nuclear induction,” *Physica*, vol. 17, no. 3-4, pp. 272–281, 1951. DOI: [10.1016/0031-8914\(51\)90068-7](https://doi.org/10.1016/0031-8914(51)90068-7).

[90] C. P. Slichter, “Principles of magnetic resonance,” in Third Edition, Berlin, Germany: Springer Berlin, Heidelberg, 1990, pp. 36–38. DOI: [10.1007/978-3-662-09441-9](https://doi.org/10.1007/978-3-662-09441-9).

[91] G. B. Chavhan, P. S. Babyn, B. Thomas, M. M. Shroff, and E. M. Haacke, “Principles, Techniques, and Applications of T2*-based MR Imaging and Its Special Applications,” *Radiographics*, vol. 29, no. 5, pp. 1433–1449, 2009. DOI: [10.1148/rg.295095034](https://doi.org/10.1148/rg.295095034).

[92] P. Thureau, G. Mollica, F. Ziarelli, and S. Viel, “Chapter Five - Solid-State 1H NMR Studies of Homonuclear Dipolar Couplings,” in *Annual Reports on NMR Spectroscopy*, Amsterdam, The Netherlands: Elsevier, 2014, pp. 217–249. DOI: <https://doi.org/10.1016/B978-0-12-800184-4.00005-9>.

[93] E. L. Hahn, “Spin echoes,” *Phys. Rev.*, vol. 80, no. 4, pp. 580–594, 1950. DOI: [10.1103/PhysRev.80.580](https://doi.org/10.1103/PhysRev.80.580).

[94] H. Y. Carr and E. M. Purcell, “Effects of diffusion on free precession in nuclear magnetic resonance experiments,” *Phys. Rev.*, vol. 94, no. 3, pp. 630–638, 1954. DOI: [10.1103/PhysRev.94.630](https://doi.org/10.1103/PhysRev.94.630).

[95] S. Meiboom and D. Gill, “Modified Spin-Echo Method for Measuring Nuclear Relaxation Times,” *Rev. Sci. Instrum.*, vol. 29, no. 8, pp. 688–691, 2004. DOI: [10.1063/1.1716296](https://doi.org/10.1063/1.1716296).

[96] P. Morris and A. Price, “3.02 - Image Contrast and Resolution in MRI,” in *Comprehensive Biomedical Physics*, Amsterdam, The Netherlands: Elsevier, 2014, pp. 21–36. DOI: <https://doi.org/10.1016/B978-0-444-53632-7.00302-6>.

[97] E. Fukushima and S. B. W. Roeder, “Experimental pulse NMR : A nuts and bolts approach,” in MA, USA: Addison Wesley Longman Publishing Co, 1981, pp. 310, 388–392, 400–403, 408–411, ISBN: 0201104032.

[98] D. W. Sinkovits and M. S. Conradi, “Improved rejection of transmitter noise: A convenient scheme with resonant crossed diodes,” *J. Magn. Reson.*, vol. 171, no. 1, pp. 11–14, 2004. DOI: <https://doi.org/10.1016/j.jmr.2004.07.015>.

[99] C. Li and D.-P. Schreurs, “Chapter 1 - Fundamentals of microwave engineering,” in *Principles and Applications of RF/Microwave in Healthcare and Biosensing*, Amsterdam, The Netherlands: Elsevier, 2017, pp. 1–52. DOI: <https://doi.org/10.1016/B978-0-12-802903-9.00001-1>.

[100] D. J. Shanefield, “Chapter 11 - filters and resonance,” in *Industrial Electronics for Engineers, Chemists, and Technicians*, NY, USA: William Andrew Publishing, 2001, pp. 117–128. DOI: <https://doi.org/10.1016/B978-081551467-1.50012-2>.

[101] R. E. Walstedt, “2. Introduction to NMR Studies of Metals, Metallic Compounds, and Superconductors,” in *The NMR probe of High-Tc Materials*, First Edition, Berlin, Germany: Springer Berlin, Heidelberg, 2008, pp. 22–25. DOI: [10.1007/978-3-540-75565-4](https://doi.org/10.1007/978-3-540-75565-4).

[102] H. Takeda, M. Itoh, and H. Sakurai, “Local magnetic properties of the antiferromagnetic metal NaV₂O₄ with frustrated double-chain structure investigated by NMR and magnetization measurements,” *Phys. Rev. B*, vol. 86, no. 17, pp. 1–8, 2012. DOI: [10.1103/PhysRevB.86.174405](https://doi.org/10.1103/PhysRevB.86.174405).

[103] H. Takeda, H. Yasuoka, M. Yoshida, *et al.*, “V51 NMR study on the S=1/2 square lattice antiferromagnet K₂V₃O₈,” *Phys. Rev. B*, vol. 100, no. 5, pp. 1–12, 2019. DOI: <https://doi.org/10.1103/PhysRevB.100.054406>.

[104] S. Kambe, H. Sakai, Y. Tokunaga, *et al.*, “¹¹⁵In nqr study with evidence for two magnetic quantum critical points in dual ce site superconductor ce₃pt₁₁,” *Phys. Rev. B*, vol. 101, no. 8, pp. 1–5, 2020. DOI: [10.1103/PhysRevB.101.081103](https://doi.org/10.1103/PhysRevB.101.081103).

[105] J. Seliger, V. Žagar, and R. Blinc, “¹⁴N NQR study of the structural phase transitions in NH₄NO₃,” *Zeitschrift für Physik B Condensed Matter*, vol. 77, pp. 439–443, 1989. DOI: [10.1007/BF01453795](https://doi.org/10.1007/BF01453795).

[106] G. Jeschke, “5.2: Dipole-dipole interaction,” in *Electron Paramagnetic Resonance*, CA, USA: LibreTexts Chemistry, 2024. [Online]. Available: [https://chem.libretexts.org/Bookshelves/Physical_and_Theoretical_Chemistry_Textbook_Maps/Electron_Paramagnetic_Resonance_\(Jenschke\)](https://chem.libretexts.org/Bookshelves/Physical_and_Theoretical_Chemistry_Textbook_Maps/Electron_Paramagnetic_Resonance_(Jenschke)).

[107] K. Nawa, M. Takigawa, M. Yoshida, and K. Yoshimura, “Anisotropic spin fluctuations in the quasi one-dimensional frustrated magnet LiCuVO₄,” *J. Phys. Soc. Jpn*, vol. 82, no. 9, p. 094709, 2013. DOI: [10.7566/JPSJ.82.094709](https://doi.org/10.7566/JPSJ.82.094709).

[108] S. Li, S. Kawai, Y. Kobayashi, and M. Itoh, “Magnetic excitation and local magnetic susceptibility of the excitonic insulator Ta₂NiSe₅ investigated by Se ⁷⁷NMR,” *Phys. Rev. B*, vol. 97, no. 16, pp. 1–9, 2018. DOI: [10.1103/PhysRevB.97.165127](https://doi.org/10.1103/PhysRevB.97.165127).

[109] S. Blundell, *Magnetism in Condensed Matter*, First Edition. Oxford, UK: Oxford University Press, 2001, pp. 25, 85, 92, ISBN: 0198505914.

[110] D. Schmidiger, “Physics of a strong-leg quantum spin ladder,” PhD Thesis, ETH Zürich, 2014. [Online]. Available: <https://www.research-collection.ethz.ch/handle/20.500.11850/98354>.

[111] K. Binder, “Theory of first-order phase transitions,” *Rep. Prog. Phys.*, vol. 50, no. 7, p. 783, 1987. DOI: [10.1088/0034-4885/50/7/001](https://doi.org/10.1088/0034-4885/50/7/001).

[112] L. P. Kadanoff, W. Götze, D. Hamblen, *et al.*, “Static Phenomena near critical points: Theory and Experiment,” *Rev. Mod. Phys.*, vol. 39, no. 2, 1967. DOI: [10.1103/RevModPhys.39.395](https://doi.org/10.1103/RevModPhys.39.395).

[113] M. Vojta, “Quantum Phase Transitions,” *Rep. Prog. Phys.*, vol. 66, pp. 2069–2110, 2003. DOI: [10.1088/0034-4885/66/12/R01](https://doi.org/10.1088/0034-4885/66/12/R01).

[114] S. Haravifard, Z. Yamani, and B. D. Gaulin, “Chapter 2 - Quantum Phase Transitions,” in *Neutron Scattering - Magnetic and Quantum Phenomena*, Amsterdam, The Netherlands: Elsevier, 2015. DOI: <https://doi.org/10.1016/B978-0-12-802049-4.00002-6>.

[115] A. Carollo, D. Valenti, and B. Spagnolo, “Geometry of quantum phase transitions,” *Phys. Rep.*, vol. 838, pp. 1–72, 2020. DOI: <https://doi.org/10.1016/j.physrep.2019.11.002>.

[116] A. Zheludev, “Quantum critical dynamics and scaling in one-dimensional antiferromagnets,” *J. Exp. Theor. Phys.*, vol. 131, no. 1, pp. 34–45, 2020. DOI: [10.1134/S1063776120070183](https://doi.org/10.1134/S1063776120070183).

[117] M. A. Continentino, “Quantum scaling in many-body systems,” *Phys. Rep.*, vol. 239, no. 3, pp. 179–213, 1994. DOI: [10.1016/0370-1573\(94\)90112-0](https://doi.org/10.1016/0370-1573(94)90112-0).

[118] I. Bose, “Quantum magnets: a brief overview,” *arXiv:cond-mat:0107399*, 2001. [Online]. Available: <https://arxiv.org/pdf/cond-mat/0107399.pdf>.

[119] T. Hewitt, “Phase diagram of the anisotropic heisenberg spin ladder,” PhD Thesis, University of Kent, 2017. [Online]. Available: <https://kar.kent.ac.uk/68462/>.

[120] F. Mila and F. C. Zhang, “On the origin of biquadratic exchange in spin 1 chains,” *Eur. Phys. J. B*, vol. 16, pp. 7–10, 2000. DOI: [10.1007/s100510070242](https://doi.org/10.1007/s100510070242).

[121] M. Perfetti and J. Bendix, “The Multiple Faces, and Phases, of Magnetic Anisotropy,” *Inorg. Chem.*, vol. 58, no. 18, pp. 11 875–11 882, 2019. DOI: [10.1021/acs.inorgchem.9b00636](https://doi.org/10.1021/acs.inorgchem.9b00636).

[122] J. des Cloizeaux and J. J. Pearson, “Spin-wave spectrum of the antiferromagnetic linear chain,” *Phys. Rev.*, vol. 128, no. 5, pp. 2131–2135, 1962. DOI: [10.1103/PhysRev.128.2131](https://doi.org/10.1103/PhysRev.128.2131).

[123] N. D. Mermin and H. Wagner, “Absence of ferromagnetism or antiferromagnetism in one- or two-dimensional isotropic heisenberg models,” *Phys. Rev. Lett.*, vol. 17, no. 22, pp. 1133–1136, 1966. DOI: [10.1103/PhysRevLett.17.1133](https://doi.org/10.1103/PhysRevLett.17.1133).

[124] M. P. Nightingale and H. W. J. Blöte, “Gap of the linear spin-1 heisenberg antiferromagnet: A monte carlo calculation,” *Phys. Rev. B*, vol. 33, no. 1, pp. 659–661, 1986. DOI: [10.1103/PhysRevB.33.659](https://doi.org/10.1103/PhysRevB.33.659).

[125] E. S. Sørensen and I. Affleck, “Large-scale numerical evidence for bose condensation in the $s=1$ antiferromagnetic chain in a strong field,” *Phys. Rev. Lett.*, vol. 71, no. 10, pp. 1633–1636, 1993. DOI: [10.1103/PhysRevLett.71.1633](https://doi.org/10.1103/PhysRevLett.71.1633).

[126] S. Sachdev, T. Senthil, and R. Shankar, “Finite-temperature properties of quantum antiferromagnets in a uniform magnetic field in one and two dimensions,” *Phys. Rev. B*, vol. 50, no. 1, pp. 258–272, 1994. DOI: [10.1103/PhysRevB.50.258](https://doi.org/10.1103/PhysRevB.50.258).

[127] O. Golinelli, T. Jolicœur, and R. Lacaze, “Haldane gaps in a spin-1 heisenberg chain with easy-plane single-ion anisotropy,” *Phys. Rev. B*, vol. 45, no. 17, pp. 9798–9805, 1992. DOI: [10.1103/PhysRevB.45.9798](https://doi.org/10.1103/PhysRevB.45.9798).

[128] S. Hu, B. Normand, X. Wang, and L. Yu, “Accurate determination of the gaussian transition in spin-1 chains with single-ion anisotropy,” *Phys. Rev. B*, vol. 84, no. 22, p. 220402, 2011. DOI: [10.1103/PhysRevB.84.220402](https://doi.org/10.1103/PhysRevB.84.220402).

[129] T. Giamarchi, “PECULIARITIES OF $D = 1$,” in *Quantum Physics in One Dimension*, Online Edition, Oxford, UK: Oxford University Press, 2003, pp. 1–3. DOI: [10.1093/acprof:oso/9780198525004.003.0001](https://doi.org/10.1093/acprof:oso/9780198525004.003.0001).

[130] F. D. M. Haldane, “‘Luttinger liquid theory’ of one-dimensional quantum fluids. I. Properties of the Luttinger model and their extension to the general 1D interacting spinless Fermi gas,” *J. Phys. C: Solid State Phys.*, vol. 14, no. 19, p. 2585, 1981. DOI: [10.1088/0022-3719/14/19/010](https://doi.org/10.1088/0022-3719/14/19/010).

[131] M. Jeong, D. Schmidiger, H. Mayaffre, *et al.*, “Dichotomy between Attractive and Repulsive Tomonaga-Luttinger Liquids in Spin Ladders,” *Phys. Rev. Lett.*, vol. 117, no. 10, pp. 20–22, 2016. DOI: [10.1103/PhysRevLett.117.106402](https://doi.org/10.1103/PhysRevLett.117.106402).

[132] R. M. Konik and P. Fendley, “Haldane-gapped spin chains as luttinger liquids: Correlation functions at finite field,” *Phys. Rev. B*, vol. 66, no. 14, p. 144 416, 2002. DOI: [10.1103/PhysRevB.66.144416](https://doi.org/10.1103/PhysRevB.66.144416).

[133] J. S. Möller, T. Lancaster, S. J. Blundell, *et al.*, “Quantum-critical spin dynamics in a tomonaga-luttinger liquid studied with muon-spin relaxation,” *Phys. Rev. B*, vol. 95, no. 2, p. 020 402, 2017. DOI: [10.1103/PhysRevB.95.020402](https://doi.org/10.1103/PhysRevB.95.020402).

[134] *SDR14*, <https://www.spdevices.com/products/hardware/14-bit-digitizers/sdr14>, Accessed: 2020-09-20.

[135] *Virtex-6 Family Overview*, https://www.xilinx.com/support/documentation/data_sheets/ds150.pdf, Accessed: 2023-11-26, 2015.

[136] *M9005A PXIe Chassis*, <https://www.keysight.com/en/pd-2745421-pn-M9005A/pxie-chassis-5-slot-3u-with-integrated-system-module?cc=GB&lc=eng&jmpid=zzfindm9005a>, Accessed: 2023-11-26.

[137] *ISE Design Suite*, <https://www.xilinx.com/products/design-tools/ise-design-suite.html>, Accessed: 2023-11-26.

[138] *ISim User Guide UG660 (v14.1)*, 2012. [Online]. Available: https://www.xilinx.com/content/dam/xilinx/support/documents/sw_manuals/xilinx14_1/plugin_ism.pdf.

[139] *AMD CORE Generator System*, <https://www.xilinx.com/products/design-tools/coregen.html>, Accessed: 2023-11-26.

[140] Teledyne SP Devices, *Teledyne SP Devices ADQAPI Reference Guide*, 2022. [Online]. Available: https://www.spdevices.com/en-us/Products_/Documents/Software%20Landing%20Page/14-1351_ADQAPI_Reference_Guide.pdf.

[141] Xilinx Inc., *Xilinx DS255 Multiplier v11.2*, 2011. [Online]. Available: https://docs.xilinx.com/v/u/en-US/mult_gen_ds255.

[142] *SciPy*, <https://scipy.org/>, Accessed: 2023-11-26.

[143] *PyQt5 5.15.10*, <https://pypi.org/project/PyQt5/>, Accessed: 2023-11-26.

[144] *ctypes — A foreign function library for Python*, <https://docs.python.org/3/library/ctypes.html>, Accessed: 2023-11-26.

[145] *NumPy*, <https://numpy.org/>, Accessed: 2023-11-26.

[146] *PyQtGraph*, <https://www.pyqtgraph.org/>, Accessed: 2023-11-26.

[147] *Matplotlib: Visualization with Python*, <https://matplotlib.org/>, Accessed: 2023-11-26.

[148] *MultiPyVu 2.1.4*, <https://pypi.org/project/MultiPyVu/>, Accessed: 2023-11-26.

[149] Quantum Design, *MultiPyVu: Controlling QD Cryostats Using Python*, 2023. [Online]. Available: <https://www.qdusa.com/siteDocs/appNotes/1070-220.pdf>.

[150] Quantum Design, *Multi-function probes*. [Online]. Available: https://www.qdusa.com/siteDocs/productBrochures/1084-500_PPMS_MFP_Types.pdf.

[151] Quantum Design, *Physical Property Measurement System Hardware Manual*, 2008. [Online]. Available: https://web.njit.edu/~tyson/PPMS_Documents/PPMS_Manual/1070-150%20Rev.%20B5%20PQ%20%20PPMS%20Hardware.pdf.

[152] H. D. Young and R. A. Freedman, “Temperature and Heat,” in *University Physics*, 13th edition, London, UK: Pearson, 2012, p. 571, ISBN: 9780321762184.

[153] *SC-219/50-PBC-PBC*, <http://www.coax.co.jp/en/product/sc/219-50-pbc-pbc.html>, Accessed: 2023-11-26.

[154] M. Manjunatha, G. S. Reddy, K. J. Mallikarjunaiah, R. Damle, and K. P. Ramesh, “Determination of phase composition of cobalt nanoparticles using ^{59}Co internal field nuclear magnetic resonance,” *J Supercond Nov Magn*, vol. 32, pp. 1557–1947, 2019. DOI: <https://doi.org/10.1007/s10948-019-5083-7>.

[155] A. S. Andreev, O. B. Lapina, and S. V. Cherepanova, “A new insight into cobalt metal powder internal field ^{59}Co nmr spectra,” *Appl Magn Reson*, vol. 45, pp. 1009–1017, 2014. DOI: <https://doi.org/10.1007/s00723-014-0580-0>.

[156] A. C. Gossard and A. M. Portis, “Observation of nuclear resonance in a ferromagnet,” *Phys. Rev. Lett.*, vol. 3, no. 4, pp. 164–166, 1959. DOI: [10.1103/PhysRevLett.3.164](https://doi.org/10.1103/PhysRevLett.3.164).

[157] *1n4148tr*, Accessed: 2024-08-24, 2024. [Online]. Available: <https://www.mouser.co.uk/ProductDetail/onsemi-Fairchild/1N4148TR?qs=i4Fj9T/oRm%252BOzV8sfXrhvQ%3D%3D&srsltid=AfmBOoqDibS5ojxsznxBSylk7qxcLxCL5lBwjltvB66w6pLhcpXFiXsN>.

[158] G. M., A. D., H. M., S. Y., and M. S., “Analytical models of probe dynamics effects on nmr measurements,” *J. Magn. Reson.*, vol. 327, p. 106975, ISSN: 1090-7807. DOI: <https://doi.org/10.1016/j.jmr.2021.106975>.

[159] T. Cvitanić, M. Lukas, and M. S. Grbić, “Two-axis goniometer for single-crystal nuclear magnetic resonance measurements,” *Rev. Sci. Instrum*, vol. 90, no. 4, p. 043903, 2019. DOI: [10.1063/1.5075521](https://doi.org/10.1063/1.5075521).

[160] T. Shiroka, F. Casola, J. Mesot, W. Bachmann, and H.-R. Ott, “A two-axis goniometer for low-temperature nuclear magnetic resonance measurements on single crystals,” *Rev. Sci. Instrum*, vol. 83, no. 9, p. 093901, 2012. DOI: [10.1063/1.4748861](https://doi.org/10.1063/1.4748861).

[161] *Superconductivity and Magnetism*, <https://warwick.ac.uk/fac/sci/physics/research/condensedmatt/supermag/>, Accessed: 2024-08-10.

[162] W. J. A. Blackmore, “Ground states in low-dimensional quantum magnets,” PhD Thesis, University of Warwick, 2018. [Online]. Available: https://wrap.warwick.ac.uk/id/eprint/128299/1/WRAP_Theses_Blackmore_2018.pdf.

[163] W. J. A. Blackmore, S. P. M. Curley, R. C. Williams, *et al.*, “Magneto-structural correlations in ni₂₊ halide…halide ni₂₊ chains,” *Inorg. Chem.*, vol. 61, no. 1, pp. 141–153, 2022. DOI: [10.1021/acs.inorgchem.1c02483](https://doi.org/10.1021/acs.inorgchem.1c02483).

[164] Quantum Design, *Magnetic property measurement system user manual*, 2016. [Online]. Available: https://www.icmr.ucsbd.edu/sites/default/files/mrl_docs/instruments/1500-100%20Rev.%20F1%20MPMS%203%20Users%20Manual.pdf.

[165] Quantum Design, *Physical property measurement system heat capacity option user's manual*, 2010. [Online]. Available: https://web.njit.edu/~tyson/PPMS_Documents/PPMS_Manual/1085-150%20Heat%20Capacity.pdf.

[166] G. Ventura and L. Risegari, “12 - measurements of heat capacity,” in *The Art of Cryogenics*, Amsterdam, The Netherlands: Elsevier Science, 2008, pp. 267–288. DOI: <https://doi.org/10.1016/B978-008044479-6.50013-3>.

[167] T. Polenova, R. Gupta, and A. Goldbourt, “Magic angle spinning nmr spectroscopy: A versatile technique for structural and dynamic analysis of solid-phase systems,” *Anal. Chem.*, vol. 87, no. 11, pp. 5458–5469, 2015. DOI: [10.1021/ac504288u](https://doi.org/10.1021/ac504288u).

[168] *Cpmas probes*, Accessed: 2024-18-16, 2024. [Online]. Available: <https://www.bruker.com/en/products-and-solutions/mr/nmr/probes/Solids-Probes.html>.

[169] L. J. Mueller, “Tensors and rotations in NMR,” *Concepts in Magnetic Resonance Part A*, vol. 38A, no. 5, pp. 221–235, 2011. DOI: <https://doi.org/10.1002/cmr.a.20224>.

[170] A. W. Sandvik, “Computational studies of quantum spin systems,” *AIP Conference Proceedings*, vol. 1297, no. 1, pp. 135–338, 2010. DOI: [10.1063/1.3518900](https://doi.org/10.1063/1.3518900).

[171] I. Tews, “Quantum Monte Carlo methods for astrophysical applications,” *Front. Phys.*, vol. 8, p. 153, 2020. DOI: [10.3389/fphy.2020.00153](https://doi.org/10.3389/fphy.2020.00153).

[172] K. Sengupta, “Mott phases and quantum phase transitions of extended bose–hubbard models in 2+1 dimensions,” *Physica A: Statistical Mechanics and its Applications*, vol. 384, no. 1, pp. 115–121, 2007. DOI: <https://doi.org/10.1016/j.physa.2007.04.123>.

[173] W. D. Parker, J. W. Wilkins, and R. G. Hennig, “Accuracy of quantum monte carlo methods for point defects in solids,” *Phys. Status Solidi B*, vol. 248, no. 2, pp. 267–274, 2011. DOI: <https://doi.org/10.1002/pssb.201046149>.

[174] W. Tu, E. Moon, K. Lee, W. E. Pickett, and H. Lee, “Field-induced bose-einstein condensation and supersolid in the two-dimensional kondo necklace,” *Commun. Phys.*, vol. 5, no. 130, pp. 2399–3650, 2022. DOI: <https://doi.org/10.1038/s42005-022-00913-3>.

[175] C. Umrigar, K. Wilson, and J. Wilkins, “Optimized trial wave functions for quantum monte carlo calculations,” *Phys. Rev. Lett*, vol. 60, no. 17, pp. 1719–1722, 1988. DOI: [10.1103/PhysRevLett.60.1719](https://doi.org/10.1103/PhysRevLett.60.1719).

[176] R. Chen and K. E. Schmidt, “Path-integral quantum monte carlo calculations of light nuclei,” *Phys. Rev. C*, vol. 106, no. 4, p. 044327, 2022. DOI: [10.1103/PhysRevC.106.044327](https://doi.org/10.1103/PhysRevC.106.044327).

[177] J. Kolorenč and L. Mitas, “Applications of quantum monte carlo methods in condensed systems,” *Rep. Prog. Phys.*, vol. 74, no. 2, p. 026502, 2011. DOI: [10.1088/0034-4885/74/2/026502](https://doi.org/10.1088/0034-4885/74/2/026502).

[178] M. H. Kalos and P. A. Whitlock, “Chapter 7 - statistical physics,” in *Monte Carlo Methods*, Second Edition, NJ, USA: John Wiley & Sons, 2008, pp. 149–157. DOI: <https://doi.org/10.1002/9783527626212.ch7>.

[179] D. Hangleiter, I. Roth, D. Nagaj, and J. Eisert, “Easing the monte carlo sign problem,” *Sci. Adv.*, vol. 6, no. 33, eabb8341, 2020. DOI: [10.1126/sciadv.abb8341](https://doi.org/10.1126/sciadv.abb8341).

[180] F. Alet, P. Dayal, A. Grzesik, *et al.*, “The ALPS project: open source software for strongly correlated systems,” *J. Phys. Soc. Jpn.*, vol. 74, pp. 30–35, 2005. DOI: [10.1143/JPSJS.74S.30](https://doi.org/10.1143/JPSJS.74S.30).

[181] S. Mugiraneza and A. M. Hallas, “Tutorial: A beginner’s guide to interpreting magnetic susceptibility data with the curie-weiss law,” *Commun. Phys.*, vol. 5, no. 95, 2022. DOI: [10.1038/s42005-022-00853-y](https://doi.org/10.1038/s42005-022-00853-y).

[182] A. K. Bera, B. Lake, A. T. M. N. Islam, B. Klemke, E. Faulhaber, and J. M. Law, “Field-induced magnetic ordering and single-ion anisotropy in the quasi-one-dimensional haldane chain compound $\text{srni}_2\text{v}_{208}$: A single-crystal investigation,” *Phys. Rev. B*, vol. 87, no. 22, p. 224423, 2013. DOI: [10.1103/PhysRevB.87.224423](https://doi.org/10.1103/PhysRevB.87.224423).

[183] J. M. Law, H. Benner, and R. K. Kremer, “Padé approximations for the magnetic susceptibilities of heisenberg antiferromagnetic spin chains for various spin values,” *J. Phys. Condens. Matter*, vol. 25, no. 6, p. 065601, 2013. DOI: [10.1088/0953-8984/25/6/065601](https://doi.org/10.1088/0953-8984/25/6/065601).

[184] F. Pobell, “Solid Matter at Low Temperatures,” in *Matter and Methods at Low Temperatures*, Berlin, Germany: Springer Berlin Heidelberg, 2007, pp. 33–93. DOI: [10.1007/978-3-540-46360-3_3](https://doi.org/10.1007/978-3-540-46360-3_3).

[185] Y. Maeda, C. Hotta, and M. Oshikawa, “Universal temperature dependence of the magnetization of gapped spin chains,” *Phys. Rev. Lett.*, vol. 99, no. 5, p. 057205, 2007. DOI: [10.1103/PhysRevLett.99.057205](https://doi.org/10.1103/PhysRevLett.99.057205).

[186] D. M. Ogris and E. Gamsjäger, “Heat capacities and standard entropies and enthalpies of some compounds essential for steelmaking and refractory design approximated by

debye-einstein integrals,” *Calphad*, vol. 75, p. 102345, 2021. DOI: <https://doi.org/10.1016/j.calphad.2021.102345>.

[187] P. Bouillot, C. Kollath, A. M. Läuchli, *et al.*, “Statics and dynamics of weakly coupled antiferromagnetic spin- $\frac{1}{2}$ ladders in a magnetic field,” *Phys. Rev. B*, vol. 83, no. 5, p. 054407, 2011. DOI: [10.1103/PhysRevB.83.054407](https://doi.org/10.1103/PhysRevB.83.054407).

[188] M. Jeong, H. Mayaffre, C. Berthier, D. Schmidiger, A. Zheludev, and M. Horvatić, “Attractive tomonaga-luttinger liquid in a quantum spin ladder,” *Phys. Rev. Lett.*, vol. 111, no. 10, p. 106404, 2013. DOI: [10.1103/PhysRevLett.111.106404](https://doi.org/10.1103/PhysRevLett.111.106404).

[189] V. V. Kuzmin and P. Nacher, “Signal feedback applications in low-field NMR and MRI,” *J. Magn. Reson.*, vol. 310, p. 106622, 2020. DOI: [10.1016/j.jmr.2019.106622](https://doi.org/10.1016/j.jmr.2019.106622).