

Concepts in low-cost and flow NMR

Zur Erlangung des akademischen Grades eines
DOKTORS DER INGENIEURWISSENSCHAFTEN (Dr.-Ing.)

von der KIT-Fakultät für Maschinenbau des
Karlsruher Instituts für Technologie (KIT)
angenommene

DISSERTATION

von

MSc. Pedro Freire Silva

Tag der mündlichen Prüfung: 03.02.2022
Hauptreferent: Prof. Dr. Jan G. Korvink
Korreferent: Prof. Dr. Arno Kentgens



Concepts in low-cost and flow NMR

Ph.D Thesis of

Pedro Freire Silva

at the Department of Mechanical Engineering
Institute Microstructure Technology (IMT)

Reviewer: Prof. Dr. Jan G. Korvink

Second reviewer: Prof. Dr. Arno Kentgens

6. June 2016 – 23. September 2019

Karlsruher Institut für Technologie
Institut für Mikrostrukturtechnik
H.-v.-Helmholtz Platz 1
76344 Eggenstein-Leopoldshafen

I declare that I have developed and written the enclosed thesis completely by myself, and have not used sources or means without declaration in the text.

Karlsruhe, 1/1/2020

.....

(Pedro Freire Silva)

Abstract

Magnetic resonance (MR) has long been a staple in the fields of chemistry and medical science, prized for its ability to non-invasively probe a wide array of characteristics. Despite its successes, the technique remains a high-end solution in most industries, with other analytical chemistry approaches and imaging techniques being preferred in everyday, low-end applications. This stems from the MR need for high-precision manufacturing techniques, MR-compatible components, and complex operation and post-processing procedures. These further limit the use of MR, namely that of low-sensitivity nuclear magnetic resonance (NMR), as a transception method in other metrology applications despite its desirable and unique advantages. This body of work showcases new techniques to drastically lower the price of NMR, in generalised transception applications, by targeting the high-integration of multiple mass-manufacturable subsystems.

Starting from a target application, the measurement of liquid flow through NMR, a top-down approach was followed to best adjust to the system-like behaviour of NMR hardware, which shows a sharp price/performance dependence on how the various subsystems are integrated. A theoretical approach was initially applied to the field of flow NMR, discovering a new operational method that introduces a correlation of the NMR signal to fluid flow while maintaining the statistical significance of the measurement. This method set top-level constraints leading to the conceptualisation of a complete system that can be fully mass-manufacturable, as well as offset the usual complexity of radio-frequency control/transception to commercial Software-Defined Radio (SDR) boards.

Having chosen the measurement procedure, the most directly correlated subsystems were the RF transception and shimming coils. These were thus developed as printed circuits boards (PCBs) with all the necessary electrical components built in, achieving high spin-sensitivity and shimming power efficiency, respectively. The new design of the transception coil, beyond a previously unattainable geometrical decoupling, was able to generate the first geometrically-differential behaviour between two NMR samples, which drastically eases the performance requirements of the follow-up electronics, namely SDR ADCs. A second concept in inductive readout was investigated; the use of adaptive, modular and highly-arrayed coils on flexible PCBs. These showed the first static coil design that can be electrically-modified to work optimally at a wide range of frequencies, by adjusting the windings' connections and thus the coil's self-resonant frequency to closer to the working frequency.

Finally, several magnet geometries were evaluated and simulated, with the optimal solution for the system being a C-shaped gap magnet. The first principles methodology used to evaluate these was repurposed from magnetostatic reciprocity theory into the field of NMR, and allowed for several new insights into magnet development. These, along with a literature research into magnetic material modelling, showed how field inhomogeneity can be expected from fabrication variations and non-linear effects of the magnets used. Through this methodology, three new concepts were developed: a self-assembly pathway into the development of Halbach micromagnets, a profilometry magnet with an adjustable penetration depth and an optimal magnetisation distribution that allows for ultra-high fields in small-scale NMR.

Zusammenfassung

Die Magnetresonanz (MR) ist seit langem ein fester Bestandteil der Chemie und der Medizin und wird wegen ihrer Fähigkeit geschätzt, ein breites Spektrum von Merkmalen nicht-invasiv zu untersuchen. Trotz ihrer Erfolge bleibt die Technik in den meisten Industriezweigen eine High-End-Lösung, da andere Ansätze der analytischen Chemie und bildgebende Verfahren in alltäglichen, einfachen Anwendungen bevorzugt werden. Dies ergibt sich aus dem Bedarf an hochpräzisen Fertigungstechniken, MR-kompatiblen Komponenten und komplexen Anwendungs- und Nachbearbeitungsverfahren im Fall von MR. Diese schränken die Verwendung der MR, genauer die der Kernspinresonanz (NMR) mit geringer Empfindlichkeit, als Transzeptionsmethode in anderen metrologischen Anwendungen trotz ihrer wünschenswerten und einzigartigen Eigenschaften weiter ein. In dieser Arbeit werden neue Techniken vorgestellt, mit denen der Preis der NMR in allgemeinen Transzeptionslösungen drastisch gesenkt werden kann, indem man auf hohe Integration mehrerer in Massenfertigung herstellbarer Subsysteme abzielt.

Ausgehend von einer Zielanwendung, der Messung des Flüssigkeitsdurchflusses mithilfe der NMR, wurde ein Top-Down-Ansatz verfolgt, um die Betriebsstrategie am besten an das Systemverhalten der NMR-Hardware anzupassen. Dieses Systemverhalten zeigt eine starke Preis-/Leistungsabhängigkeit davon, wie die verschiedenen Subsysteme integriert werden. Zunächst wurde ein theoretischer Ansatz auf dem Gebiet der Durchfluss-NMR angewandt, wobei eine neue optimale Betriebsmethode entdeckt wurde, die eine Korrelation des NMR-Signals mit dem Flüssigkeitsstrom herstellt, wobei die statistische Signifikanz der Messung erhalten bleibt. Diese Methode legt Bedingungen auf höchster Ebene fest, die zur Konzeptualisierung eines vollständigen Systems führten, das sowohl vollständig in Massenproduktion hergestellt werden kann als auch die übliche Komplexität der Hochfrequenzsteuerung/-übertragung auf das Niveau kommerzieller Software-Defined Radio (SDR)-Boards senkt.

Nach der Wahl des Messverfahrens waren die am direktesten korrelierten Subsysteme die HF-Transzeptions- und Shimming-Spulen. Diese wurden daher als gedruckte Schaltungen (PCBs) entwickelt, in die alle notwendigen elektrischen Komponenten eingebaut waren, wodurch eine hohe Spinempfindlichkeit bzw. ein hoher Wirkungsgrad im Shimming erreicht wurde. Das neue Design der Transzeptionsspule - mit einer bisher unerreichten geometrischen Entkopplung - ermöglichte zum ersten Mal differentielles Auslesen von zwei NMR-Proben, wodurch die Leistungsanforderungen an die Folgeelektronik, nämlich SDR-ADCs, drastisch gesenkt wurde. Ein zweites Konzept des induktiven Auslesens wurde untersucht, nämlich die Verwendung von adaptiven, modularen Multi-Array-Spulen auf flexiblen Leiterplatten. Diese zeigten das erste statische Spulendesign, das elektrisch so modifiziert werden kann, dass es in einem breiten Frequenzbereich optimal funktioniert, indem die Anschlüsse der Wicklung und damit die Eigenresonanzfrequenz der Spule nahe an die Arbeitsfrequenz angepasst werden.

Außerdem wurden mehrere Magnetgeometrien betrachtet und simuliert, wobei die optimale Lösung für das System ein C-förmiger Spaltnagnet ist. Die erste Methodik zur Bewertung dieser Prinzipien wurde von der magnetostatischen Reziprozitätstheorie auf den Bereich der NMR übertragen, wodurch neue Erkenntnisse über die Magnetentwicklung gewonnen werden konnten. Diese, zusammen mit einer Literaturrecherche zur Modellierung magnetischer Materialien, zeigten, wie Feldinhomogenität aufgrund von Herstellungsvariationen und nichtlinearen Effekten der verwendeten Magnete geschätzt werden kann. Mit dieser Methodik wurden drei neue Konzepte entwickelt: ein Selbstmontagepfad bei der Entwicklung von Halbach-Mikromagneten, ein Profilmagnet mit einstellbarer Eindringtiefe und eine optimale Magnetisierungsverteilung, die ultrahohe Felder für Kleinskalen-NMR ermöglicht.

Nomenclature

Abbreviations

Abbreviation	Meaning
MR	Magnetic resonance
MRI	Magnetic resonance imaging
EPR	Electron paramagnetic resonance
NMR	Nuclear magnetic resonance
FID	Free induction decay
TE	Echo time
TR	Repetition time
SE	Spin echo
SNR	Signal-to-noise ratio
DOF	Degree of freedom
FT	Fourier transform
FWHM	Full width half maximum
ppm/ppb	Parts per million/billion
FS	Full scale
RF	Radio frequency
AC/DC	Alternating/Direct current
SRF	Self-resonant frequency
EMI	Electromagnetic interference
PSD	Power spectral density
LNA	Low-noise amplifier
VGA	Variable gain amplifier
ADC	Analog-to-digital converter
DAC	Digital-to-analog converter
PCB	Printed circuit board

Symbols

Symbol	Meaning	Unit
ω_0	Larmor frequency	rad/s
B_0	Static magnetic field	T
B_1	RF magnetic field	T
M_0	Rest net magnetisation	A/m
M_z	Longitudinal magnetisation	A/m
M_{xy}	Transverse magnetisation	A/m
T_1	Spin-lattice relaxation time	s
T_2	Spin-spin relaxation time	s
T_2^*	Effective relaxation time	s
Z	Impedance	Ω
R	Resistance	Ω
L	Inductance	H
C	Capacitance	F
Q	Quality factor	-
Δf	MR receiver's bandwidth	Hz
F	Noise factor	-
NF	Noise figure	dB
G	Gain	dB

Constants

Symbol	Meaning	Value	SI base unit
\hbar	Reduced Planck constant	1.055×10^{-34}	$\text{Kg m}^2 \text{s}^{-1}$
k_B	Boltzmann constant	1.381×10^{-23}	$\text{Kg m}^2 \text{s}^{-2} \text{K}^{-1}$
N_A	Avogadro number	6.022×10^{23}	-
γ_H	^1H gyromagnetic ratio	267.5×10^6	$\text{s}^{-1} \text{T}^{-1}$
γ_D	^2H gyromagnetic ratio	41.07×10^6	$\text{s}^{-1} \text{T}^{-1}$
T	Room temperature	293.2	K
μ_0	Vacuum permeability	$4\pi \times 10^{-7}$	$\text{Kg m}^1 \text{s}^{-2} \text{A}^{-2}$
ϵ_0	Vacuum permittivity	8.854×10^{-12}	$\text{A}^2 \text{s}^4 \text{Kg}^{-1} \text{m}^{-3}$

Contents

Abstract	i
Zusammenfassung	iii
Nomenclature	v
1 Introduction	1
1.1 Motivation for low-cost and flow NMR	1
1.2 Manuscript outline	7
1.3 Results and publications	8
2 Essential theory	9
2.1 Phenomenological NMR	9
2.2 Electrical instrumentation	12
2.3 Magnetostatics in NMR	15
2.4 Signal conditioning and post-processing	19
2.5 Computational methods	21
3 Flow NMR methodology	25
3.1 Net-phase flow encoding	25
3.1.1 Introduction	25
3.1.2 Front-end design and constraints	27
3.1.3 Flow encoding and measurement SNR	33
3.1.4 Conclusion	37
4 Inductive signal transception	39
4.1 Geometrically-differential stripline	39
4.1.1 Introduction	39
4.1.2 NMR detector	41
4.1.3 Electronics and operation	45

Contents

4.1.4	Differential NMR	50
4.1.5	Conclusion	52
4.2	Configurable, roll-up double-helix	53
4.2.1	Introduction	53
4.2.2	Theory	54
4.2.3	Resonator design	58
4.2.4	Adjustable performance	62
4.2.5	Conclusion	66
5	Permanent magnet design	69
5.1	Magnetostatic reciprocity	69
5.1.1	Introduction	69
5.1.2	Theory and corollaries	71
5.1.3	Modelling and result standardisation	75
5.1.4	Self-aligned magnets	77
5.1.5	NMR profilometry	80
5.1.6	High-field designs	82
5.1.7	Conclusion	84
6	Conclusion	87
	Bibliography	91
	Acknowledgments	105
	List of figures	106
	List of tables	108

1 | Introduction

1.1 Motivation for low-cost and flow NMR

This manuscript documents the various scientific advances made to better measure liquid flow by means of NMR transception, and the low-cost hardware that this task relies on. For an understanding of why one would want to pursue such an endeavour, this section offers a broad, non-technical explanation of the problem, highlighting the large gap between what can currently be obtained from portable NMR flowmeters and from research-level, high-end NMR systems, while comparing these to the abilities of flowmeters based on alternative techniques. An initial motivation is easily conveyed through the incredible insights into liquids and their multidimensional flow characteristics shown in Figure 1.1.

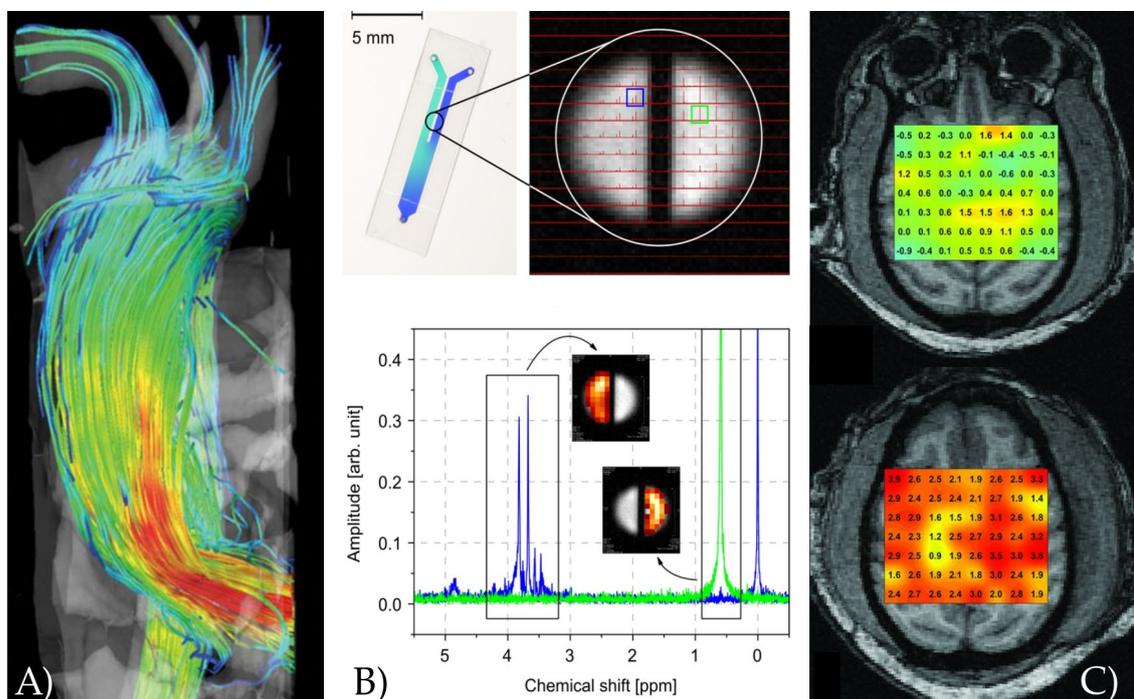


Figure 1.1: A) shows the velocity of blood flow in the aorta, obtained from [1] with permission. In B), chemical shift imaging results are overlaid on a normal MRI image of two NMR-identifiable fluids, obtained from [2] under a CC-4.0 license. An overlay of a brain temperature profile on an MRI image is depicted in C) before (top) and after a stroke, as reproduced from [3] through due use.

1 Introduction

Despite the extraordinary insights they offer, NMR flowmetry and multidimensional queries of a fluid require complex and expensive hardware, as well as weakly demanding applications, with a long lasting behaviour that can be repeatably averaged and/or sequentially probed. This has kept this wealth of information constrained to research fields, where it offers a viable experimental option to computational fluid dynamics. Applications that require high throughput, cheap/compact hardware and/or demanding flow regimes, exemplified in Figure 1.2, can only obtain modest information from the flow, which is occasionally sufficient for the challenge at hand.

These results could be obtained with portable, compact and low-cost hardware based on commercial solutions adapted from standard telecommunications and power-electronics hardware. Such approaches have seen significant success with the emergence of low-cost prototyping, proving to be a fitting solution for applications in quality control [7], non-invasive structural analysis [8] and even medical imaging applications in developing countries [9]. Solutions are mostly based on permanent magnets, coil inductors and a spectrometer board connected to a computer, all similar to the ones depicted in Figure 1.3.

Analysing the theory, it becomes clear that any NMR spectrometer can be used to measure flow and thus advances in both fields, compact NMR and flow NMR, move hand-in-hand. With the ultimate goal of advancing the field of compact NMR flowmetry, an initial assessment of NMR's abilities against standard flowmetry techniques is necessary. This comparison allows one to gauge whether its strengths as a scalar transception method, its non-invasive nature and its sensitivity to most liquid samples can be sufficiently leveraged to overcome its disadvantage of a low overall sensitivity, which might entail either a slow measurement, low-accuracy results or an exceedingly expensive experimental setup.

Commercial flowmeter applications [14] employ an extraordinary breadth of techniques to achieve a correlation to fluid flow, among which one can encounter thermal, vibrational, electromagnetic, optical and mechanical interactions at various scales. Due to NMR's voluminous and expensive implementation at large channel diameters, because of a necessary electro/permanent magnet, it becomes clear that NMR can best be leveraged at small to moderate scales (i.e. ID < 10 cm), despite existing solutions for large oil pipelines [15]. Simultaneously, large scale implementations are ideally suited for mechanical sensors, as they best leverage their benefits while minimising the relative effect of physical interaction with the flow.

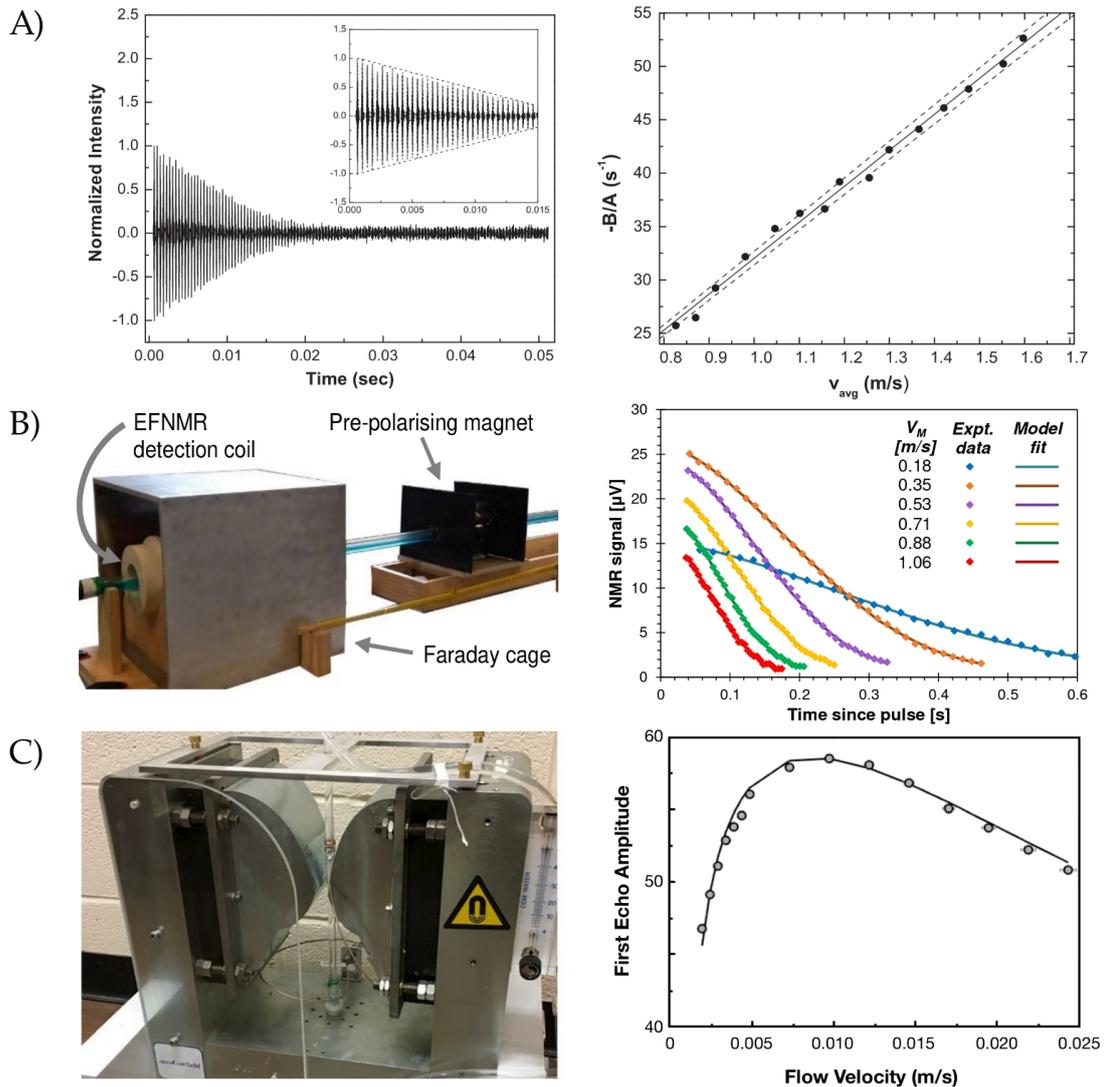


Figure 1.2: A) shows a standard NMR measurement, a CPMG sequence, and the evolution of its initial slope with increasing flow, as the measured sample leaves the measurement volume faster. Results obtained with permission from [4]. B) depicts a similar dependence, in a more simplified setup using the earth's magnetic field, obtained from [5] with permission. In C), a magnetic field gradient is applied on the flow and the amplitude of the echo of the response is correlated to flow, seen to be maximal for the time-optimal gradient moment. Results obtained from [6] with permission.

Conversely, small-scale and non-invasive applications must be handled with external transception methods, among which those based on the Coriolis effect, optical encoding, and thermal readouts have been highly successful [14]. To illustrate some ways these methods can be correlated to flow, three applications are shown in Figure 1.4 and an overview of top-performing non-invasive flowmeters is shown in Table 1.1. As seen, existing solutions offer a turndown ratio from around 100 (i.e the ratio between their highest and lowest measurable flows) up to 1000 in special cases, and can measure flow down to an accuracy of $\approx 0.5\%$. Whereas one would expect the flow velocity to be a key measurement parameter, which it remains in mass-flow measurements (e.g. Coriolis), due to an ability to scale the cross-section of a channel and thus its velocity profile, one can easily adjust it to the optimal range of the transduction technique.

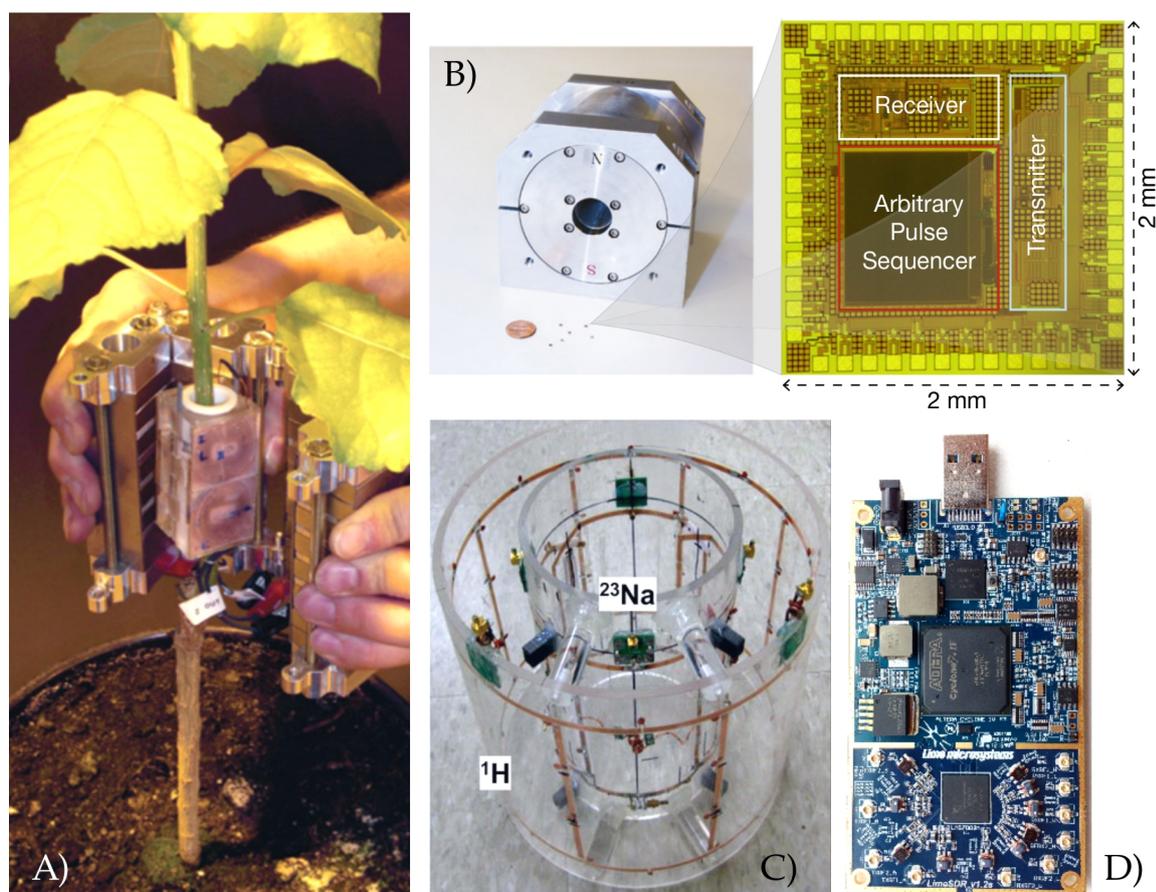


Figure 1.3: A) shows a coil and magnet combination for portable NMR on plants, as obtained from [10] with permission. B) shows how the hardware front-end can be miniaturised to be inserted in a low-cost lab magnet, as per [11] under a CCBY-4.0 license. In C), a standard birdcage coil for measuring hydrogen and sodium isotopes is shown with permission from [12]. A recent low-cost solution to acquire and control the signals is depicted in D), an FPGA-based solution, with a chip front-end, in a commercial SDR board [13], shown through due use.

Having set the two most common performance indicators in flowmetry, readout accuracy and turn-down ratio, one can now gauge from first principles whether NMR might successfully out-compete other techniques. Milli-scale NMR routinely achieves a statistical significance below 0.1% (i.e. SNR > 1000) when measuring high concentration species (e.g. the water in an aqueous solution) at normal flow repetition rates (i.e. ≈ 100 ms), and thus the flow-NMR measurement's accuracy upper bound lies beyond that of any other existing technique. However, this correlation also sets a bound on the possible NMR flowmetry applications, in which the detection volume must be large enough to be measured with sufficient statistical significance, in real-time.

Table 1.1: Comparative analysis of different flowmetry encoding and transduction methods and their respective turndown ratio and accuracy.

Configuration (transduction)	Implementation	Year [Ref.]	Fluid	Channel size (μm)	Flow range (min^{-1})	Accuracy (% FS)
Hot wire/film (resistive)	Polyamide catheter in vivo	2015 [16]	Glucose water	-	<200 mL	2.5%
Calorimetric (luminescence)	Printing and PDMS on glass	2016 [17]	Ethanol Water	135x27	<0.1 μL 100-500 μL	10% 6.6%
Calorimetric (resistive)	SU-8 on PMMA	2010 [18]	Water	2000x20	0.04-25 μL (CV)	-
Calorimetric (electric)	Post-CMOS (suspended)	2016 [19]	N_2	500x500	<100 mL	0.1%
Thermal wave (resistive)	SiRN on Si (suspended)	2014 [20]	Water	$\varnothing 20$	<500 nL	<2%
Magnetic gradient (NMR)	PVC tubing	2018 [6]	Water	$\varnothing 800$	28-10000 μL	$\approx 0.28\%$
Coriolis oscillation (capacitive)	Steel tubing	2013 [21]	Water, Isopropanol	$\varnothing 1590$	83-17000 nL	<0.83%
Optical evanescence (optical)	Plastic tubing	2017 [22]	Water	$\varnothing 150$	1.8-1000 nL	0.18%

Starting from this initial NMR accuracy, the act of acquiring and encoding the measurement must not introduce systematic errors above 0.1% of the encoded parameter. Fortunately, current low-cost electronics far outperform this limit, enabling amplitude accuracy below 0.01% (e.g. 16-bit DACs/ADCs), phase control below the millidegree range, and frequency readouts with a precision below 1 ppm (e.g. OXCO crystal).

A significant advantage of an NMR flowmeter stems from its frequency-based behaviour. Whereas the turndown ratio is often ultimately bound by a ratio of lengths (e.g. the length of two measurement sensors a certain distance apart) or times, NMR has the ability to encode flow as a modulation of its flow-independent precession. This allows for significant benefits, similar to those seen when applying such techniques in the field of telecommunications, among which a drastic increase in the upper limit of the turndown ratio. Whereas, for example, a time-of-flight technique has a measurement time strictly inversely proportional to the velocity, a phase injection into the NMR signal can be applied at a varying, optimal rate. From the sub-millidegree accuracy mentioned above and NMR's ability to encode up to thousands of degrees, one finds the upper bound of the turndown ratio to lie at six to seven orders of magnitude, far above any other technique.

1 Introduction

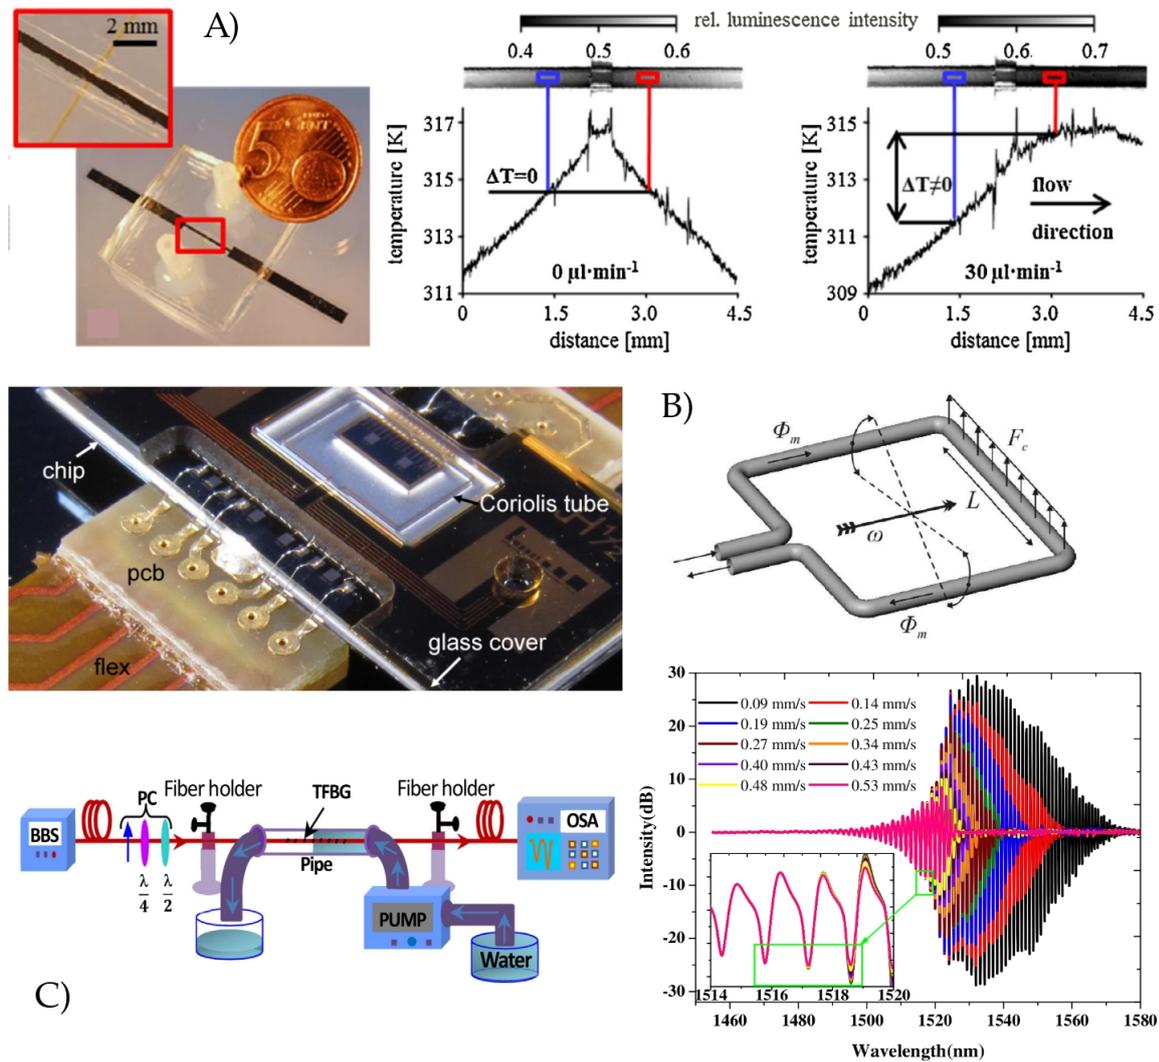


Figure 1.4: A) shows the variation in thermal profile with flow, as measured through thermoluminescence. Temperature difference is usually acquired with Pt-based meanders to measure the local temperature. Results obtained from [23] with permission. B) shows an image and operating scheme of a small-scale Coriolis flowmeter, as obtained from [21] under a CCBY-4.0 license, highlighting the Coriolis deflection appearing when exciting with force F a channel with a flowing fluid. C) depicts the change in the transfer function of a tilted Bragg grating sustaining flow-dependent cladding modes, as obtained from [22] with permission.

While these first principles considerations imply a promising upper limit for the performance of NMR flowmetry, actual implementations will introduce significant non-idealities which may severely reduce its performance, leaving a research direction rife with scientific potential and able to reach far beyond the current state of the art results shown in Figure 1.2 or even those from the competing techniques in Table 1.1.

1.2 Manuscript outline

The structure of this manuscript, "*Concepts in low-cost and flow NMR*", aims to convey the results achieved in different areas of low-cost and flow NMR by improving on low-performing, expensive or overly complex implementations of each module of an NMR spectrometer/flowmeter. For an initial overview of the techniques and applications in the state of the art, a brief non-technical introduction to current techniques was made in Chapter 1. A comprehensive overview of the sub-field and thorough comparison to the literature are then made in the introduction of each of the technical chapters: 3, 4 and 5. These introduce a clear subdivision of the work presented in the manuscript, allowing each of the advances to be independently gauged and repurposed into any other field within NMR.

A theoretical background is introduced in Chapter 2, having been carefully developed with the explicit goal of being as short as possible, while allowing for the complete understanding of the presented results/techniques. It should enable the reader to tackle this multidisciplinary manuscript without a cumbersome in-depth dive into the various subjects. The text assumes an undergraduate-level knowledge in the fields of circuit theory, electromagnetism and signal processing. Chapter 2 introduces five thematic sections presented from a phenomenological point-of-view; giving emphasis to the final theoretical results, rather than the understanding of the ultimate fundamentals of each field. To allow for a generalisation of the theory shown, formal statements are left in a simplified but abstract form, allowing for the application of the differential equations and integrals to the appropriate conditions. These results are then made application-dependent in the following technical chapters. Similarly, to allow for an easier read of this manuscript, natural constants and coefficients stated in the nomenclature section are directly used without further explanation in the text, but application-specific variables are made clear in the appropriate location.

Chapters 3, 4 and 5 contain the newly developed concepts in the NMR subfields of flow encoding, inductive readout and NMR permanent magnets. These are given in a self-enclosed manner, with the appropriate literature comparison, application-specific theory, experimental results and references in a single chapter. This structure mimics that of the publications each of these chapters is based on, as cited at the beginning of each. When possible, experimental results are given in their reduced form, either normalised to a scaling law, as an efficiency coefficient or as a general figure of merit. This approach allows for an easier repurposing of the techniques shown to other setups, as well as a statement of each concept that is independent of the technology used.

1 Introduction

The manuscript concludes with Chapter 6, offering an overview of the results obtained, as well as a roadmap for their integration in a complete, self-sufficient prototype. Given the necessary development of the each module through a top-down approach, as NMR behaviour is particularly interlinked through thermal non-equilibria or EMI, for example, several design concerns will be rehashed so as to be avoided in continuing work. The chapter finalises with an evaluation of the results obtained, under the light of the proposed goal of advancing compact and flow NMR, so as to narrow the gap to high-end NMR solutions and eventually outperform other flowmetry techniques in select applications.

1.3 Results and publications

The results presented in detail in this manuscript are part of an overarching project including the following publications, either already published or in preparation/submission. Only the author's contribution is detailed here and one must thus consult the following documents for the remaining information.

- **P. Silva**, M. Jouda and J. Korvink, “Geometrically-differential NMR in a stripline front-end” published in *Journal of Magnetic Resonance* [24]
- **P. Silva**, S. Torres, M. Jouda, D. Mager and J. Korvink, “A multi-purpose, rolled-up, double-helix resonator” published in *Journal of Magnetic Resonance* [25]
- **P. Silva**, M. Jouda and J. Korvink, “Magnetostatic reciprocity for MR magnet design” published in *Magnetic Resonance Ampere* [26]
- **P. Silva**, M. Jouda, M. Alinaghian Jouzdani, M. Condesso and J. Korvink, “Net-phase flow NMR for compact applications” *in preparation*
- J. Franco, **P. Silva**, M. Jouda and J. Korvink, “A C++ API for the development of SDR-based NMR” *in preparation*
- Y. Liu, **P. Silva**, M. Jouda, N. MacKinnon and J. Korvink, “A differential sensor for NMR spectroscopy” *in preparation*
- S. Nimbalkar, E. Fuhrer, **P. Silva**, T. Nguyen, M. Sereno, S. Kassegne and J. Korvink “Glassy Carbon Microelectrodes Minimize Induced Voltages, Mechanical Vibrations and Artifacts in MRI” published in *Nature Microsystems and Nanoengineering* [27]
- M. Jouda, E. Fuhrer, **P. Silva**, J. Korvink, and N. MacKinnon “Automatic adaptive gain for magnetic resonance sensitivity enhancement” published in *Analytical Chemistry* [28]

2 | Essential theory

2.1 Phenomenological NMR

Magnetic resonance is an intrinsically quantum phenomenon, arising as a consequence of nucleon spin-states with differing energies, while in the presence of a magnetic field [29, 30, 31]. However, given its low coupling to other interactions, due to the vastly different energy quanta involved, it is possible to state it in a semi-classical formulation. This allows one to fully characterise the nuclear magnetic behaviour through two simple equations defining the net magnetisation of matter and its vectorial response to an any external magnetic interaction. The use of a complete quantum formulation of a spin system is nonetheless possible and well-understood [30], and even required for some applications.

As mentioned, the appearance of a magnetic field, B , on a nucleon with magnetic moment μ , creates an energy splitting of $E_m = -\mu \cdot B$, commonly referred to as the Zeeman splitting. This splitting naturally favours lower energy states and thus magnetic moments that are co-linear with the applied magnetic field. These are defined by $\mu = \gamma I$, which creates a discrete projection of the spin angular momentum I along B , taking the values $-\hbar m, \hbar(-m + 1), \dots, \hbar(m - 1), \hbar m$. Both the gyromagnetic ratio γ , and the magnetic quantum number I (a discrete state of the vector quantity with the same name, I , the spin angular momentum), are determined by the filling of atomic orbitals, and thus constants that depend on the atomic number of the isotope measured. These simplifications can be put together for a scalar form of the Zeeman energy, $E_m = -m\gamma\hbar B$, for which only unitary changes of the spin number m are statistically probable. This means that the energy of the photon necessary for a spin transition is given by $\gamma\hbar B$, corresponding to a wave angular frequency of $\omega_{\text{NMR}} = \gamma B$.

From statistical mechanics [32], it is possible to derive the population statistics of the decoupled multi-state system, by considering the high-temperature approximation of the Fermi-Dirac distribution for nucleons, which are fermions, the well-known Boltzmann distribution.

2 Essential theory

One can then compute the partition function, the sum of probability of all available states, and thus compute the net magnetisation of an ensemble:

$$\overline{M} = N\gamma\hbar \frac{\sum_{m=-I}^I m e^{-E_m/k_B T}}{\sum_{m=-I}^I e^{-E_m/k_B T}} \stackrel{E_m \ll k_B T}{\approx} \frac{N\gamma^2 \hbar^2 I(I+1)}{3k_B T} B \quad (2.1)$$

The multiplication by N carries the approximation that this is done for N independent or weakly coupled spins, an approximation that holds in room-temperature liquid NMR. Equation 2.1 thus sets that as long as there is a magnetic field applied on an ensemble, there will be a net magnetisation to be measured.

Any magnetic moment or its resulting net effect (e.g. the magnetisation over a certain volume) observes a torque given by $\tau = \mu \times B$ while under a magnetic field, defining its time dependence dL/dt . This torque gives rise to the well-known Larmor precession around a magnetic field at a frequency of γB . At the same time that the spins precess, they are exchanging energy with other nucleons and with the lattice. Interaction with the lattice will naturally revert the spin states back to a thermal equilibrium through a relaxation frequency T_1^{-1} , as defined by the appropriately named spin-lattice relaxation time. The interaction with other local magnetic dipoles will introduce a dephasing effect which is observed through a relaxation frequency of the signal of T_2^{-1} , as defined by the spin-spin relaxation time. When a strong DC magnetic field is applied in the z direction, these considerations can be put together into the differential equation defining the evolution of a spin in a vectorial model, the Bloch equation:

$$\frac{d}{dt} \begin{pmatrix} M_x \\ M_y \\ M_z \end{pmatrix} = \gamma \begin{pmatrix} -1/\gamma T_2 & B_z & -B_y \\ -B_z & -1/\gamma T_2 & B_x \\ B_y & -B_x & -1/\gamma T_1 \end{pmatrix} \begin{pmatrix} M_x \\ M_y \\ M_z \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ M_0/T_1 \end{pmatrix} \quad (2.2)$$

Due to its complexity, Equation 2.2 is best solved analytically in a referential rotating at the NMR angular frequency, a change equivalent to removing the precessing behaviour (i.e $B_z = 0$), and by defining $\mathbf{M}_{xy} = M_x + iM_y$ and $\mathbf{M}_z = M_z$. Beyond this simplification, one can split its solution for the periods of spin-excitation ($\Delta t \ll T_1, T_2$) and signal acquisition ($\mathbf{B}_{xy} = 0$):

$$\frac{d}{dt} \begin{pmatrix} \mathbf{M}_{xy} \\ \mathbf{M}_z \end{pmatrix} = \frac{i\gamma}{2} \begin{pmatrix} \mathbf{B}_{xy} \mathbf{M}_z \\ \mathbf{B}_{xy}^* \mathbf{M}_{xy} - \mathbf{B}_{xy} \mathbf{M}_{xy}^* \end{pmatrix} \quad (2.3)$$

$$\frac{d}{dt} \begin{pmatrix} \mathbf{M}_{xy} \\ \mathbf{M}_z \end{pmatrix} = \begin{pmatrix} -T_2^{-1} \mathbf{M}_{xy} \\ -T_1^{-1} (\mathbf{M}_z - \mathbf{M}_{z0}) \end{pmatrix} \quad (2.4)$$

Equation 2.3 can now be easily interpreted as a standard rotation if one considers an on-resonance magnetic field, B_1^+ , along the x direction and oscillating at the frequency of the rotating referential (i.e. $B_1^+ = B_{xy} = 1$). This response is what enables a rotation of the net magnetisation away from equilibrium and thus starts-off the NMR precession and measurement process, usually with a 90° pulse/rotation away from the initial/rest state of $M_z = M_{z0}$, $M_{xy} = 0$.

Simultaneously, Equation 2.4, starting from an ideally excited state $M_z = 0$, $M_{xy} = M_{z0}$, immediately awards the time evolution that is the measurement baseline of any NMR/MRI experiment now back in the laboratory frame of reference:

$$M_{xy} = s_{\text{NMR}}(t) = M_{z0} e^{i\omega_{\text{NMR}}t} e^{-t/T_2} \quad (2.5)$$

$$M_z = M_{z0}(1 - e^{-t/T_1}) \quad (2.6)$$

The long-lasting, periodic evolution of the NMR signal in Equation 2.5 implies a well resolved frequency response directly proportional to the value of γB_0 . This propriety of the NMR response is of paramount importance as it opens the door to the fields of NMR spectroscopy, investigating small changes in the value of γ , and MR imaging, by introducing a spatial variation of B_0 . The frequency response is analytically given through the Fourier transform of a right-decaying exponential as:

$$S_{\text{NMR}}(\omega) = \mathfrak{F}[s_{\text{NMR}}(t)] = \int_0^\infty s_{\text{NMR}}(t) e^{-i\omega t} = \frac{T_2^{-1} - i(\omega - \omega_{\text{NMR}})}{T_2^{-2} + (\omega - \omega_{\text{NMR}})^2} \quad (2.7)$$

Equation 2.7 is what gives rise to the narrowband ($T_2 \gg \omega_{\text{NMR}}^{-1}$) peaks seen in NMR spectra. Its complex nature can be split into real and imaginary parts, respectively the absorption and dispersion responses, or into magnitude and phase. The latter split is also what gives rise to the two information channels employed pixel-wise in MRI methods, but which are outside of the scope of this manuscript.

Overall, one can now understand the fundamental reason as to why the nuclear magnetic response exists and, from Equation 2.1, what its amplitude is. From Equation 2.2, one can deduce how to manipulate the underlying NMR statistics, represented by a net magnetisation, to obtain a readout signal before a return to thermal equilibrium. This very same measured signal, acquired with the methods to be introduced in Section 2.2, can then be interpreted as a time-evolving or frequency-based response through Equations 2.6 and 2.7.

2.2 Electrical instrumentation

Given the ability to manipulate the net magnetisation with an AC magnetic field at the Larmor frequency, as seen in Equation 2.3, a common detection method is to use the very same coil to readout the voltage induced on it by the precessing magnetisation. In an alternating manner, this coil is used for spin manipulation and detection, as explored experimentally later, but developed with the main goal of an optimised signal reception.

When implementing an optimised receiver, one aims to optimise the voltage induced by the spins on some coil. However, it is also possible to state its signal as a function of the field per unit current, B_i , generated by the coil on a volume containing the sample [33, 34], which is a much simpler and thus preferable approach. Simultaneously, the induced signal voltage ξ , will carry voltage noise n_V , and together make up the SNR of the inductive readout of a sample of volume V_S , through:

$$\text{SNR}_{\text{NMR}} = \xi/n_V = n_V^{-1} \int (\partial/\partial t)(-B_i \cdot M_{xy}) dV_S \quad (2.8)$$

Having defined the measurement's quality through Equation 2.8, a natural follow-up is to state B_i and n_V as a function of sensor-dependent quantities, given that the observable magnetisation M_{xy} and its state are already defined by Equations 2.1 and 2.2, respectively.

In an inductor with an inductance L , it is possible to equate the energy stored inside the magnetic field created by a current I on a medium of permeability μ , and the electrical work done to generate it [34] through $L I^2 = \int (H \cdot B) dV_{\mathbb{R}^3}$. The later can be simplified using the ratio of the magnetic energy placed inside V_S to the total energy, the filling factor η , into $\mu_0^{-1} \eta V_S B^2$. This relation, along with the thermal white noise in a dissipative conductive medium at temperature T , allows for the statement of the two sensor-dependent variables to be optimised for:

$$B_i \approx \sqrt{\mu \eta L / V_S} \quad (2.9)$$

$$n_V = \sqrt{4k_B T R} \quad (2.10)$$

The previous results introduce an explicit optimisation functional to be used in the development of sensors for inductive NMR transception, relying solely on mechano-electrical metrics: an inductance, resistance and the efficient sample filling of the sensor.

Despite a clear formulation for the SNR of the NMR measurement, Equation 2.8 does not account for the need to acquire the information in the transceiver. The digitisation of such a signal, due to its small amplitude, requires several steps to guarantee that the SNR is degraded as little as possible, either through the effect of coupling of external noise sources or from the noisy behaviour of the various electronic components.

As an initial pre-amplification step, given the narrowband nature of the NMR response, a simple, passive solution is to use a resonant amplification system [35]. This can be achieved using a resonant cavity, in air or in a low loss-tangent dielectric, or, most commonly, through an LC circuit resonant tuned to the NMR frequency. This system is characterised by its quality factor, usually dominated by the quality factor of the inductor, of $Q = \omega L/R$. It introduces an $\sqrt{(Q^2 + 1)}$ amplification of the voltage and the noise, thereby conserving SNR, along with an impedance change to $R_L(Q^2 + 1) + 0j$, at resonance.

This system, which shows an increasingly pronounced wave behaviour as transmission lengths reach its signal's wavelength, should be impedance-matched with its load through the condition $Z_{LC} = Z_{LNA}^*$, so as to be able to transfer power to next stage in the reception chain, usually a 50Ω cable connected to an RF low-noise amplifier (LNA). To avoid an SNR-detrimental introduction of more components, the impedance change during amplification can be turned into a matching network, utilising the already implemented capacitor off-resonance and a series capacitor/inductor to compensate the marginal reactance into 50Ω. Beyond this initial amplification, multiple other variable gain amplifiers (VGA) can be used to further amplify the signal as desired. The introduction of several amplification stages then drastically reduces the effect of posterior noise or interference effects. This is analytically shown in Friis' [35] formula for the noise factor in terms of signal (V), noise (N) and voltage gain (G) at the coil, the LNA and in the rest of the chain:

$$\text{SNR}_{\text{analog}} = \sqrt{\frac{V_{\text{NMR}}^2}{N_{\text{NMR}}^2 + N_{\text{LC}}^2 + \left(\frac{N_{\text{LNA}}}{G_Q}\right)^2 + \left(\frac{N_{\text{rest}}}{G_Q G_{\text{LNA}}}\right)^2}} \quad (2.11)$$

Through the mentioned techniques, one is able to amplify the signal beyond its small amplitude and up to any desired value, at a commercially-attained noise factor slightly above that of an RF LNA, ≈ 1.06 . At this point, by matching the dynamic range of the signal and the ADC at play, one can successfully digitise the signal and post-process the data.

Due to the limited discretisation step of the ADC, quantisation noise is introduced in the final analog step. This approximately white noise can be partially reduced by ADC processing gain, through which an oversampling ratio ($\text{OSR} = f_{\text{ADC}}/2f_{\text{NMR}}$) of 4 introduces another equivalent bit of precision, or 6 dB in SNR [36]. Assuming a signal much larger than the least significant bit (LSB) in an N-bit ADC and uncorrelated with f_{ADC} , a valid assumption as the spectrometer and the atomic/NMR clocks are independent, one can define the maximum achievable SNR, the signal to quantisation noise ratio, SQNR, as:

$$\text{SQNR}_{\text{ADC}} = [6.02N + 1.76 + 10\log_{10}(\text{OSR})] \text{ dB} \quad (2.12)$$

2 Essential theory

Despite SNR being defined in terms of voltages in NMR, as in Equations 2.8 and 2.11, it is stated in terms of power in Equations 2.12 and 2.14 to resemble standard RF formulations.

When the previously obtained SQNR is below the measurement's SNR, as expected from Equation 2.8, it degrades the overall result and thus heterodyne detection must be employed. This widespread detection scheme for high-frequency signals relies on the mixing of the two signals, the NMR response and a local oscillator (LO), and the low-pass filtering (LPF) of their response:

$$\sin(\omega_{NMR}t) e^{-it\omega_{LO}} \stackrel{LPF}{\approx} \frac{-i}{2} e^{it(\omega_{LO}-\omega_{NMR})} + \frac{\sin[(\omega_{LO}+\omega_{NMR})t] + i \sin[(\omega_{LO}-\omega_{NMR})t]}{2} \quad (2.13)$$

The resulting signal maintains spectral coherence around the centre frequency $\omega_{LO} - \omega_{NMR}$, usually zero, but in a shifted frequency band, while filtering out all spectral components outside of the bandwidth set by the LPF. While this approach removes the dominant effect of the quantisation noise (through a larger OSR) and removes out-of-band noise, which could be aliased over the signal during digitisation, it relies on a well resolved LO frequency.

The LO reference frequency of the system originates from an in-built clock source, usually a crystal or an oscillator, which carries a certain error. The most common performance metric for oscillators is the maximum drift with temperature, critical in telecommunications, but mostly negligible in NMR due to the far higher drift of the magnetisation intensity of any permanent magnet and an usual possibility to post-process the deviation away. This is routinely done by locking into an NMR-clock, either a secondary dedicated NMR system, with the same field drift, or the drifting behaviour of a dominant NMR species in the sample, such as a solvent.

However, an oftentimes neglected metric that directly correlates to the performance of the measurement is the instantaneous clock jitter. Usually defined by its Allan variance over a period of time, Υ , jitter introduces a voltage error corresponding to a measurement before or after the expected sampling moment thus introducing an equivalent voltage noise and degrading SNR, beyond its direct impact on spectral resolution. For a signal at the Nyquist frequency, the upper bound of the SNR is lowest and can be stated as [37]:

$$\text{SNR}_{\text{jitter}} = \left[10 \log_{10} \left(\frac{1}{\omega^2 \Upsilon^2} \right) \right] \text{ dB} \quad (2.14)$$

With an overview of the process, one can see that having defined the initial and maximal SNR of the measurement with Equation 2.8, one must simply design a sensor, amplification and acquisition RF-pathway that does not introduce any of the SNR-limiting effects of Equations 2.11, 2.12 and 2.14.

2.3 Magnetostatics in NMR

Posing the hardest engineering challenge in MR, the implementation of a strong but low-cost magnetic field at a sufficiently-high homogeneity is of paramount importance, requiring the implementation of either an electromagnet, a superconducting design or a permanent magnet setup. Given the scope of this manuscript, portable and low-cost NMR hardware, only solutions based on permanent magnets are considered here.

So far, Section 2.1 has defined the MR response under the assumption that a static and homogeneous magnetic field is applied over all the measured spins, which then give out an ideal response measured through the methods in Section 2.2. However, in real implementations, a non-homogeneous magnetic field means that each spin will evolve under different experimental conditions and thus have a differing precession frequency. After all the spins are coherently excited, they continually acquire a constant marginal phase according to their deviation from the mode intensity of the magnetic field, $\gamma\Delta B_0 t$, eventually becoming fully dephased. As their contribution to the overall signal sums destructively, the net MR response decays faster, resulting in a larger bandwidth in the frequency domain. Assuming the inhomogeneity widens the FWHM of the frequency peak while maintaining a Lorentzian-like shape, one can state their effect as an additional decay rate/frequency, $(\gamma\Delta B_0)^{-1}$, beyond the expected $(T_2)^{-1}$, for a total decay rate defined as $(T_2^*)^{-1}$ [31].

From this new exponential decay rate, the dependence of the magnetisation and the induced voltages on the magnetic field intensity [33] and the optimised timing of an averaged MR experiment based on the Ernst angle [38], one knows all the necessary values to judge a magnetic field implementation. Relying on the the Lorentzian deviation of the magnetic field, ΔB_0 , its intensity, B_0 , the T_1 of the desired sample and the repetition rate of the experiment T_R , one can state a measurement's SNR per unit square root of time depending on whether it is single-frequency/high-inhomogeneity (sSNR_t, $T_R \approx T_2^*$, $(\gamma\Delta B_0)^{-1} \gg T_2^{-1}$) or frequency-resolved (fSNR_t, $T_R \approx T_2$) measurement as:

$$\text{sSNR}_t \propto B_0^{7/4} e^{-(\gamma\Delta B_0 T_1)^{-1}} \quad (2.15)$$

$$\text{fSNR}_t \propto B_0^{7/4} (1 + T_2 \gamma \Delta B_0)^{-0.5} \quad (2.16)$$

Any magnetic field solution in MR will thus have the goal of maximising the optimisation functional defined in Equations 2.15/2.16 for its desired read-out methodology and sample of interest.

2 Essential theory

Knowing the goal of an MR magnet, the first step after choosing a prospective magnet geometry is to analyse it through either the sum of the contributions of each part of the structure or through the solution of Gauss' equation, $\nabla \cdot (\mu_0 H + B_r) = 0$, depending on whether several small magnets or a continuous structure are employed.

Whereas Gauss' equation must be numerically solved with the methods discussed in Section 2.5, a unitary magnetic dipole μ aligned with the z -axis can have its axially-symmetric magnetic field analytically defined [39] in polar coordinates $\{\rho, \theta\}$ as:

$$\vec{B}_d = \frac{\mu_0}{4\pi} \frac{|\vec{m}_z|}{(\rho^2 + z^2)^{3/2}} [1.5\sin(2\theta)\hat{e}_\rho + (3\cos^2(\theta) - 1)\hat{e}_z] \quad (2.17)$$

At this point, knowing how one can obtain the magnetic field of a magnetisation (M), magnetic moment ($\mu = MV$) or remanent field ($B_R = \mu_0 M$) distribution in space, as well as the worth of such a field for an MR experiment, one must now go over how these field sources can be generated from a given material.

Magnetic materials relevant for MR applications can generally be divided into soft and hard magnetic materials, depending on whether they show a hysteresis effect after being magnetised [40]. Most materials show a non-hysteretic behaviour, being either paramagnetic, diamagnetic or, if they show a pronounced magnetisation response, soft-ferromagnetic. However, given that these materials do not retain a magnetic alignment after the removal of the magnetic field, H , one must resort to the use of hard-ferromagnetic materials to maintain a long-lasting magnetic field without continuous power dissipation. These materials, after a strong electromagnetic pulse is applied and removed, retain ordered magnetic dipoles and thus a remanent magnetisation, B_R , as shown in Figure 2.1.

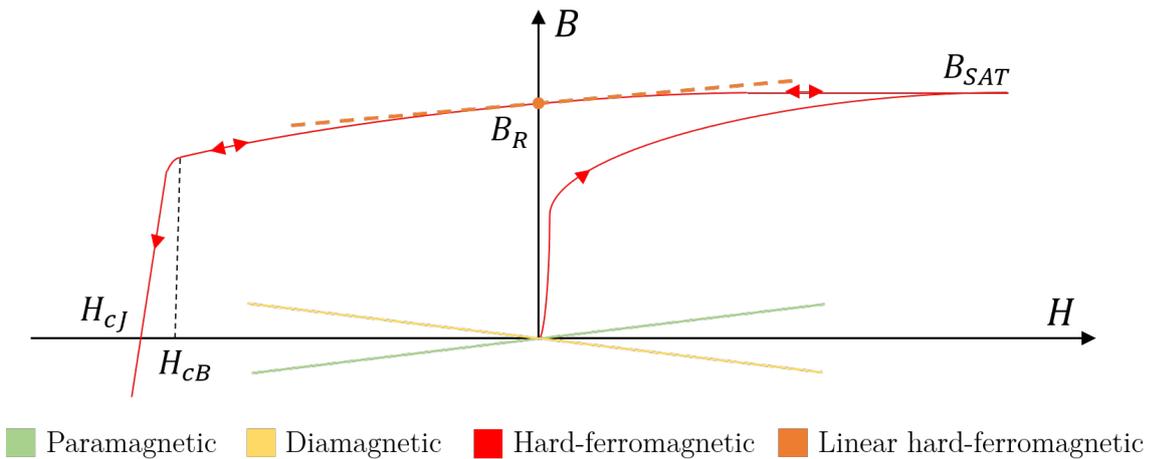


Figure 2.1: Schematic H-B curves of the collinear response of standard magnetic materials when a magnetic field H is applied and removed. Values for the saturation field, B_{Sat} , remanent field, B_R , demagnetisation coercivity, H_{cB} , and intrinsic coercivity, H_{cJ} are shown schematically. The derivative of the curves depicts the permeability along the magnetisation axis, μ_{\parallel} .

Due to the increasing availability of high H_{cB} materials, their robustness to the assembly process and ease of modelling, most designs now use solely NdFeB alloys, despite the potential of a higher B_{SAT} in soft-ferromagnetic materials.

NdFeB, and similar materials, can be characterised through a constant remanent field, a linear response along their direction of magnetisation, μ_{\parallel} , and another in all perpendicular directions, μ_{\perp} , thus mimicking a linear hard-ferromagnetic material [41]. Their response can be modelled according to the following constitutive relation:

$$\begin{pmatrix} B_x \\ B_y \\ B_z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ B_R \end{pmatrix} + \begin{pmatrix} \mu_{\perp} & 0 & 0 \\ 0 & \mu_{\perp} & 0 \\ 0 & 0 & \mu_{\parallel} \end{pmatrix} \cdot \begin{pmatrix} H_x \\ H_y \\ H_z \end{pmatrix} \quad (2.18)$$

Despite the clear preference in literature for simplified models, which naturally allow for all the modelling benefits of the principle of superposition, the use of non-linear responses might hold a largely undiscovered potential. Similar simplifications were preferred in the early years of electrical circuits and up until the introduction of the transistor, which showed a controllable saturation behaviour on its I-V curve. Similarly, several magnetic materials show a saturation behaviour of their H-B curve, but are disregarded in MR as it becomes complicated to model their response as either a constant or linear perturbation. To showcase their potential, a new concept was tested; that of magnetic parts with transistor-like properties shown here to homogenise a magnetic field.

A promising magnetic alloy is $\text{Fe}_{0.5}\text{Ni}_{0.5}$, which presents a high initial permeability, the highest B_{SAT} among nickel-iron alloys, and a sharp saturation of its H-B profile [42]. Magnetically, this translates to a small 'resistance' to the passage of magnetic flux at low flux intensities and a large resistance past a certain value. Such a response mimics that of an electrical limiter, which shows a low resistance below a certain threshold after which current is shunted away from the path of interest. This effect can be ideally repurposed into magnetic field design, by introducing a physical pathway for the passage of flux, up to a certain value, at which the flux will prefer other pathways. If this transition can be made sudden enough, it can be used to homogenise the magnetic field as seen in the comparison shown in Figure 2.2.

The comb structure is able to 'limit' the magnetic field intensity, setting it at the value of the saturation magnetic field of the material multiplied by the packing factor of the combs, and equally 'distribute' it along the gap down to the curvature effect of propagation in air. For visualisation, the comb's pitch was made large, with a reduction leading to a larger height at a ppm homogeneity. This result contrasts with the standard response of a bulk soft-ferromagnet, which would introduce a curvature of the field in itself, further emphasising the inhomogeneous tendency of end-effects.

2 Essential theory

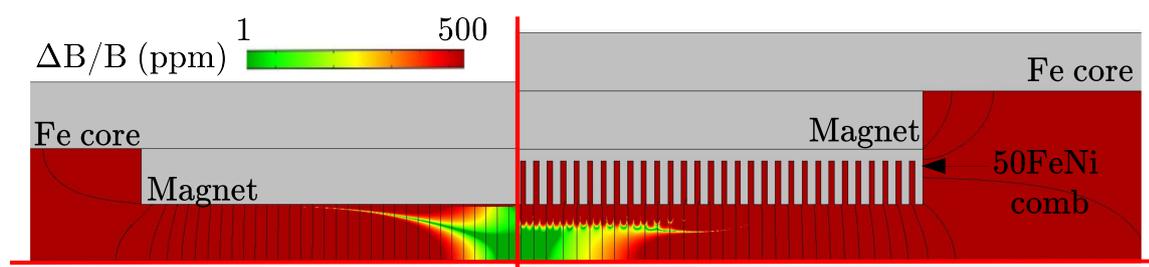


Figure 2.2: Comparison of the magnetic field homogeneity in two designs using the same magnet (NdFeB) and core (Fe) material but with the option of a field-limiting magnetic part (right). Comb structure is made from 50FeNi and was simulated with a large pitch for visual purposes. Red lines denote symmetries used when simulating.

These results segue well into to the last effect in magnetic materials that is routinely seen in MR experiments, a thermal drift [40]. A core tenet of any MR application is the correlation between a frequency/magnetic field and a chemical interaction (spectroscopy) or location in space (imaging) and therefore any drifting of this correspondence will cause a loss of information. This becomes particularly severe when considering that useful spectroscopic resolution often goes below the ppm range, which corresponds to a thermal drift below 10 mK for most materials and only 1 mK for NdFeB. Any thermal drift above these limits during an acquisition will effectively display a convolution of the sharp peak with the drift profile, blurring any chemical resolution. Fortunately, along with tracking and predictive methods, the use of field 'limiting' techniques, similar to the ones shown in Figure 2.2, means that one can use composite structures with a net-zero thermal drift [43], opening up a pathway towards solving this problem.

Unlike the difficult design, characterisation and fabrication procedures of MR magnets, the *macroscopic* theory concerning these can be easily summed up as some distribution of material reacting according to Equation 2.18, that generates a magnetic field according to Maxwell's equation or Equation 2.17, with the end goal of optimising the quality of the measurement, as defined per Equations 2.15/2.16.

2.4 Signal conditioning and post-processing

Having established how to manipulate, acquire and digitise the NMR signal, one must now set a concrete figure of merit derived from the raw data, which can then be experimentally employed. This final procedure allows the non-specialist to use NMR as a tool, observing the resulting chemical spectrum, MRI image or scalar output.

An initial step in post-processing most NMR data is the zero-filling of the receive buffer, which increases the smoothness of the frequency spectrum, and a multiplication by one or more of the several well-known windowing functions (f), an apodisation procedure [29]. As the digital data received is the result of hardware downconversion, explained in Section 2.2, two I/Q or Real/Imaginary buffers should be used:

$$s(t) = \prod f_i(t) (Re_{ADC+0}(\theta, t) + iIm_{ADC+0}(\theta, t)) \quad (2.19)$$

In NMR spectroscopy, the ultimate goal is the statistically significant identification of all relevant frequency peaks. When it comes to apodisation, this presents a trade-off between methods favouring a better quantified peak through a better SNR, and methods awarding a narrow bandwidth in the frequency domain, so as to not hide neighbouring peaks. However, when using NMR as a scalar transception channel, frequency-resolution is no longer a concern, and one should therefore apply an exponential filter. This windowing function optimally improves SNR by giving a lower weighing factor to the decayed/low-SNR part of the signal, at the cost of frequency resolution [44].

Computing the amplitude of the NMR response, s_A , one of the core scalar results, becomes as trivial as seen in Equation 2.20. The amplitude of the signal, now DC after hardware downconversion, can then be obtained by the Fourier transform at zero frequency of the time-domain signal, from its start moment, t_{dead} , up to the end of the measurement, T_{meas} :

$$s(\omega)_A = \left| \int_{t_{dead}}^{T_{meas}} e^{-i\omega t} [Re_{ADC}(\theta, t) + iIm_{ADC}(\theta, t)] e^{-t/T_2^*} dt \right| \quad (2.20)$$

This SNR increase, awarded by the exponential windowing function, directly improves the statistical significance of experimental correlations to the amplitude of the NMR signal, as well as to its phase, a central result in phase-encoding methods [45]:

$$\sigma_\phi = \text{SNR}_{|NMR|}^{-1} \quad (2.21)$$

Knowing the standard deviation (i.e. error) of the phase, through Equation 2.21, one must then define what is considered to be the scalar phase signal, s_ϕ . This is commonly set as a series differential measurement between the natural response of the system, which

2 Essential theory

includes the phase relative to the spectrometer, and the response of the system with an additional encoding sequence. This differential phase can be averaged to a single scalar value carrying still all the benefits of apodisation:

$$s(\omega)_\phi = \arg \left(\int_{t_{\text{dead}}}^{T_{\text{meas}}} e^{-i\omega t} [Re_{\text{ADC}}(\theta, t) + iIm_{\text{ADC}}(\theta, t)] e^{-t/T_2^*} dt \right) - \arg \left(\int_{t_{\text{dead}}}^{T_{\text{meas}}} e^{-i\omega t} [Re_{\text{ADC}}(\theta, t) + iIm_{\text{ADC}}(\theta, t)] e^{-t/T_2^*} dt \right)_{\text{ref}} \quad (2.22)$$

Having established clear scalar metrics in Equations 2.20 and 2.22, one can now pose multiple research questions. Can sample evaporation cause a decrease in s_A ? Can fluid flow be correlated to s_ϕ ? The answer to these is a clear yes, with more complicated tests requiring simply the appropriate redesign of the baseline encoding, read-out or post-processing steps shown in the sections above.

A necessary final step, is to define how the repetition time of an experiment affects the spins' conditioning, as it is often necessary to run the same experiment multiple times, be it to evaluate the time evolution of some response, the behaviour of various samples or to simply average results. Given the non-negligible T_1 duration that the system requires to return to maximum signal availability, it becomes clear that a trade-off between running multiple experiments faster and waiting for an optimal relaxed state of the system must be reached. For this reason, the results in the previous equations were stated as a function of their flip angle θ , which must be optimised.

This was researched by Ernst [38], which analytically set the ideal excitation angle, θ_E , for a desired repetition time of an experiment, T_R , and the sample's spin-lattice relaxation time, T_1 , when wishing to correlate the results of sequential experiments in the same sample:

$$\theta = \theta_E = \arccos(e^{-T_R/T_1}) \quad (2.23)$$

Whereas the result in Equation 2.23 serves as a good baseline, and even the optimal angle for SNR-limited applications in stationary samples, use cases requiring T_1/T_2 imaging contrasts or fast flowing samples require further optimisation [46].

2.5 Computational methods

Four theoretical sections in, it has become clear how to acquire an MR response; what one is actually measuring, how one sets up the magnetic environment for an experiment, how the underlying quantity is transduced into an electrical signal and how this signal is post-processed into the final desired information. The remaining question is how one can physically implement these steps into a working prototype.

In low-cost implementations, the electronics have become trivial with the commercialisation of out-of-the-box SDR solutions, which can simply be programmed for low-field NMR operation. However, the development of suitable RF front-end designs, magnetic pre-polarisers/precession magnets, sample channels and thermal outlets remains an open problem with multiple possible solutions. These can be experimentally built and tested, but this effort would be inordinate when compared to the currently favoured approach of simulation and optimisation *in-silico*.

The workhorse of current simulation approaches is the use of Finite Element Modelling techniques. The method, when applied to physical problems, starts with the choice of the differential equation(s) ruling the phenomenon. For solving, this equation is restated in a variational form using a support basis function, which takes into account the discretisation of the nD space used. This discretisation, called a mesh, turns an analysis of continuous space to a sparse algebraic operation, with an error which can be estimated. Nowadays, this multistep process can be achieved with out-of-the-box commercial software such as COMSOL, Ansys or Abaqus, open-source toolboxes such as Moose or CalculiX and even mathematical computing frameworks such as Matlab and Mathematica.

Regardless of the tool being used, the initial step in simulating a scenario is to choose the differential equation ruling the physical evolution of the system. In this manuscript, Gauss' law was used for magnetostatic problems (Equation 2.24), Ampère's/Poisson's law for low-frequency/DC EM fields (Equation 2.25), the single-frequency wave equation for high-frequency EM fields (Equation 2.26) and the Navier-Stokes/continuity equations for fluid flow (Equation 2.27). These were later solved for the underlying field of interest shown below in bold, using the appropriate boundary conditions at the outer surface (zero value of \mathbf{E} , \mathbf{H} , \mathbf{V} and \mathbf{u}).

$$\nabla \cdot (\boldsymbol{\mu} \cdot \mathbf{H} + B_r) = 0 \quad (2.24)$$

$$\nabla \times \mathbf{H} = -\sigma \nabla \mathbf{V} \quad \wedge \quad \nabla^2 \mathbf{V} = 0 \quad (2.25)$$

$$\nabla \times \boldsymbol{\mu}_r^{-1} (\nabla \times \mathbf{E}) - k_0^2 (\epsilon_r - \frac{j\sigma}{\omega \epsilon_0}) \cdot \mathbf{E} = 0 \quad (2.26)$$

$$\rho(\mathbf{u} \cdot \nabla) \mathbf{u} = \nabla \cdot [-\mathbf{p} + \boldsymbol{\mu} (\nabla \mathbf{u} + (\nabla \mathbf{u})^T)] \quad \wedge \quad \nabla \cdot \mathbf{u} = 0 \quad (2.27)$$

2 Essential theory

A published example of the above approach concerns the phenomena involved in MRI in conducting brain electrodes [27]. In the case below, a research question was posed: should glassy-carbon microelectrodes replace standard platinum ones and what are the effects at play?

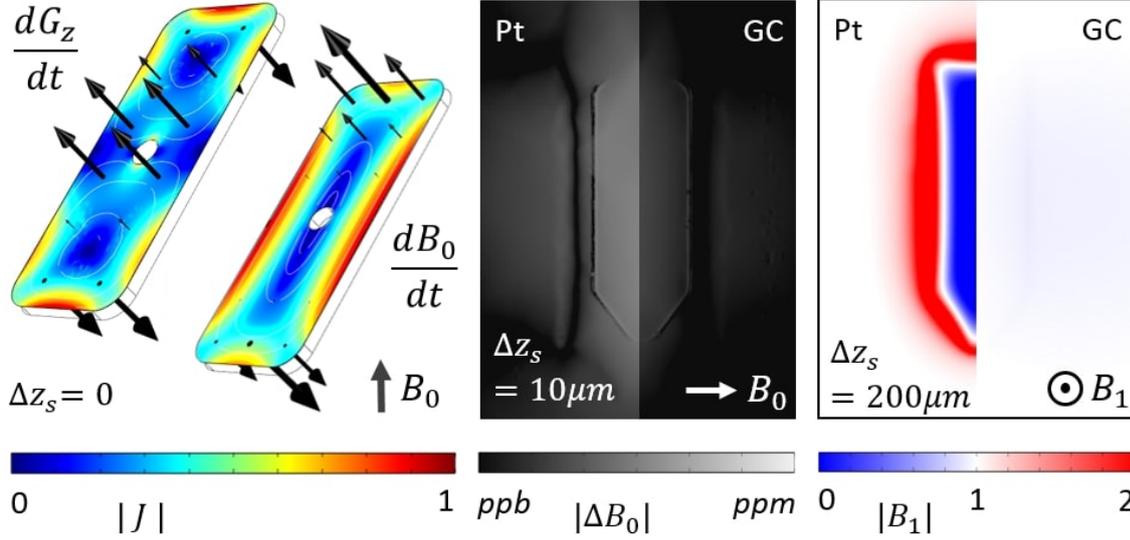


Figure 2.3: Comparison of the behaviour of two GC and Pt 20x10 mm electrodes. On the left, the induced currents and forces generated by each effect of a unitary gradient's slew rate are shown, in the centre, the B_0 relative-inhomogeneity map is depicted and, on the right, the normalised B_1 inhomogeneity map. Results are for a water phantom in a 1.5T field, shown at the relevant distances to the electrode's surface, Δz . Image repurposed from [27] with permission.

From the raw simulation results, Figure 2.3 could be generated, showcasing an the exact comparison that leads to a clear answer to the question posed. It can immediately be seen that a reduced susceptibility mismatch ($\Delta\sigma_{H_2O-GC} \approx 0.04\Delta\sigma_{H_2O-Pt}$) and drastically reduced area conductivity ($\sigma_{t_{GC}} = 0.005 \sigma_{t_{Pt}}$) lead to a significant reduction in artefacts arising from B_1 and B_0 inhomogeneity.

Simultaneously, vibrational effects could be studied. When applying a current/force-generating B_0 gradient as $B_G(t) = t\vec{G} \cdot \vec{\Delta r}$ on the electrodes, the mechanical response of the microelectrode can be split into two separate contributions by considering that $\vec{\Delta r} = \vec{\Delta}_{CG-CE} + \vec{\Delta}(r)_{E-CE}$. The first concerns the response to a spatially varying field around the electrode's centre (CE), $\vec{\Delta}(r)_{E-CE}$, and the second represents the uniform response due to the fixed displacement from the centre of the gradient field (CG), $\vec{\Delta}_{CG-CE}$, which imposes an equivalent effect of dB_0/dt . Both these effects could be easily observed in simulation through the torque-generating currents in the electrodes, leading to a more optimized design of future electrode shapes. Similarly, a decreased conductivity could be seen to lead to a reduction in the mechanical response coming from gradient switching, further emphasising the benefits of carbon electrodes over platinum.

The methods presented in this section show how one can avoid a large manufacturing and testing effort by running experiments in a computer, evaluating structures until one obtains the desired behaviour. Having this ability as a baseline, and assuming a well-defined optimisation functional can be defined (e.g. those in Equations 2.8 or 2.15), a number of methods to optimise the structure become available. These can be broadly divided into shape and topology optimisation, depending on whether one explicitly parameterises the structure being researched or not [47]. More specifically, whereas shape optimisation methods take structure-defining coefficients, such as a length and width of a box or the coefficients defining a curve, and optimise for their value, topology optimisation works on a completely unconstrained region in space, discretised with a geometrical mesh, which may or may not contain the end structure.

While shape optimisation is directly accessible to numerical optimisation, such as gradient-based or genetic methods, topology optimisation requires a framework connecting the mesh to a structure, the most common of which are the density, topological derivative and level-set methods. Given the breadth of existing numerical method solvers and their mesh, geometry and physics dependent optimality, these are considered to be outside of the scope of this manuscript.

3 | Flow NMR methodology

Not five years after the discovery of NMR itself, in 1946, did it become evident that the measurement experienced a very pronounced motional sensitivity and could thus be used to measure flow velocity [48]. Nevertheless, and despite interim advances in state of the art of magnetic resonance and flow visualisation, these techniques are not in widespread use in industry [49], where they could be used to better control and characterise inline flow processes. The work in this section was developed targeting this very same opportunity. Motivated by the high upper bound for the performance of an NMR flowmeter, seen in Section 1, the knowledge that an NMR flowmeter can be repurposed from any low-cost spectrometer, and with all the necessary theory from Section 2, one can now develop an NMR flowmetry *method* targeting low-cost applications/hardware.

3.1 Net-phase flow encoding

The contents of this section are based on the author's contribution to the material in "Net-phase flow NMR for compact applications", in preparation.

3.1.1 Introduction

NMR, being a nuclear interaction, can be used to non-invasively study a variety of liquid compounds, constituting one of the few available techniques when wishing to measure biological, corrosive or high-temperature fluids. When considering the wealth of behaviours that fluid flow can display, the task becomes more challenging and competing solutions based on transception techniques without spatial discretisation/encoding become less desirable. Fortunately, NMR has already proven to be able to measure laminar, turbulent or even multiphase behaviour in liquids, aerosols, slurries or emulsion colloids, all in a non-invasive manner [50]. The approaches used to achieve this are generally divided into imaging techniques, for unknown or complex scenarios, or one-dimensional responses, for known behaviours (e.g. total flow in a range known to be in a laminar regime).

3 Flow NMR methodology

Imaging methods are able to generate a wealth of 4D information on the flow [51], which can further be made species-selective or even temperature and pressure dependent. However, these results are complex to achieve and require either high-performance instrumentation or simplistic flow conditions, so as to offset the SNR loss that comes with spatial division of the flow field into voxels, a major limitation in NMR. This loss is intrinsically linked to a reduced statistical certainty in measuring what is effectively a subset of all the sources of information. Furthermore, imaging encoding and readout often require moderate to long acquisition times [52], depending on the desired spatial resolution, which prove to be limiting for fast flow rates or when high probing rates are necessary.

Another approach to flow analysis, preferable in low-cost/low-sensitivity setups, is based on the reduction of obtained information as, in constrained, repeatable and calibrated flow conditions, the use of imaging techniques awards little extra information. From first principles, it is self-evident that laminar flow is preferable in an engineering measurement (e.g. total flow), as it maintains flow behaviour (e.g. dual phase separation) and introduces a lower pressure drop. Similarly, if the flow channel is pre-engineered and well-characterised, it allows for a repeatable laminar flow profile and a behaviour consistent with previous calibrations. Given these presuppositions, one can thus correlate the flow measurement to a single one-dimensional response to flow. This approach has been used to encode the NMR signal's magnitude with the flow information, with or without time resolution.

Single-dimension measurements are embodied through either local or non-local setups/techniques. Local techniques [53, 54] use a single coil where either the sensitive volume is saturated and the measured signal originates from the fresh spins entering the volume, or the sensitive volume is excited and the signal decay is exacerbated as excited spins flow away. Non-local methods are based in different readouts of a sample excited with one coil and detected with another one downstream, in a time-of-flight manner [55]. It becomes clear that the aforementioned techniques intrinsically suffer from two large limitations: their dimension-linked response and a sub-optimal acquisition. The former arises from a general need to have a significant magnitude change within the remanence time in the coil, for example, which ultimately limits the measurable range of a setup. Simultaneously, the acquisition procedure requires either a significant wait time (e.g. the time of flight), or the acquisition of a signal going from a high to a low amplitude, which degrades the overall SNR of the flow measurement.

In an effort to merge the benefits of local, *flow-adjustable* phase-encoding procedures used for flow-imaging and the low-information/statistically-significant *one dimensional response* of portable systems, a simplified system is proposed in Section 3.1.2. Its technical implementation mimics a lower dimensional flow and thus enables the use of the new flow-encoding method introduced in Section 3.1.3. Despite its control requirements, this approach has a promising upper bound ending with a measurement precision matching that of the NMR measurement itself. Section 3.1.4 then dissects the strengths of the technique and discusses its implementation and possible improvements.

3.1.2 Front-end design and constraints

A critical initial step in developing any operational method is considering the hardware it will run on, as well as any necessary constraints to its experimental environment. For this reason, and with the goal of achieving low-cost NMR-based flowmetry, an intrinsically SNR-limited solution, a top-down approach to the design of the system was taken, aiming to average as many flowing spins together as possible, so as to maximise flow/spin sensitivity.

An assembly was thus developed to have a synergistic behaviour between its flow channel, which will generate the flow profile, the sensitivity profile of the detector and the volume-selective shim-set. Furthermore, as these will need to be placed inside of a strong but cheaply-attained magnetic field, another spatial correlation will be further introduced. This section will focus on the development of sub-components found to be constraining at the system level, leaving the development of standard electronics (e.g. current sources or digital electronics), magnet design details (e.g. materials, safety or fabrication) and thermal management undiscussed. The end assembly, shown in Figures 3.1 and 3.6, could be manufactured and assembled using only commercial parts and fabrication techniques and a standard electronics/mechanical workshop.

Any systems development procedure must rely on reducing the search-space of design possibilities through an evaluation of their relative impact on the overall objective. Given the goal of a low-cost/complexity implementation and the large effort associated with Halbach magnets, as discussed in Section 5, the initial choice was that of using a C-shaped gap magnet, whose easy fabrication and underlying planar integration are ideally compatible with mass-manufacturable, PCB-based solutions. This immediately constrained the search-space, ruling out techniques used with cylindrical bores, a decision shown in Sections 4.1 and 5 to not be overly performance-limiting.

3 Flow NMR methodology



Figure 3.1: Photograph of the assembly of the complete prototype fabricated. Picture includes the front-end for NMR signal transception, a LimeSDR spectrometer with a USB connection, a bi-planar set of shims/gradients, a board containing controllable current sources and a C-shaped gap magnet with a thermal management system. A disposable PMMA planar insert, containing the flowing sample, is shown entering the user-accessible front of the setup, on the right.

The initial step in dimensioning the magnet was to take into consideration the PCB stack shown in Figure 3.6, for shimming and RF transception, which required a gap of 15 mm to allow for a low-cost implementation in 1-2 mm thick PCBs and two ≈ 4 mm thick sample holders. From this top-level constraint, a simulation in a commercial FEM solver (COMSOL Multiphysics, COMSOL AB) was used to set acceptable dimensions for the rest of the prototype by optimising for a trade-off between commercial magnet availability, field intensity and field homogeneity. The end result, shown in Figure 3.2, was dominated by the chosen N48 NdFeB disk magnets (60 mm OD, 10 mm height, SuperMagnetMan, USA), reaching a field strength of 0.70 T (compared to 0.73 T in simulation) and simulated inhomogeneity on the VOI (set as $3 \times 10 \times 10$ mm) of ≈ 1400 ppm.

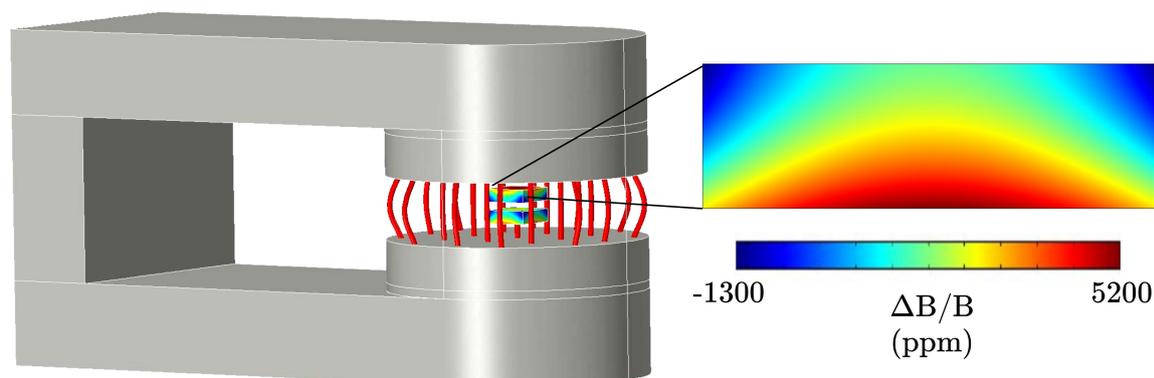


Figure 3.2: CAD model of the magnetically-active components in the NMR spectrometer. The two VOIs have their inhomogeneity profiles highlighted and the magnetic field lines in the gap are shown in red. Materials used are Stainless Steel 405 and N48 NdFeB with a remanent field of 1.38 T.

Having established the magnetic environment to be overlaid onto the flow, one must now develop the flow channel shaping its profile. If one wishes to encode a relationship between the point-wise flow properties (i.e. velocity) and the NMR response, any variation of the velocity profile will translate to summing/averaging unequal magnitudes, phases or frequencies, which will inevitably lead to a loss in signal and thus SNR. To avoid this, one attempted to approximate a 2D flow profile which could be stated as $v(z)$. In a laminar regime, this behaviour would mimic the evolution of a 2D lamina with the same flow-encoded response moving along the detector. To perfectly achieve this, one would need to make the width and length of the channel infinitely large, but approximate solutions were found to be feasible.

To test the validity of the 2D flow approximation, the expected behaviour was studied by solving the Navier-Stokes equation for incompressible flow and a no-slip boundary condition at the walls. The velocity profile on the VOI, shown along with the entire flow channel in Figure 3.3, was calculated to have an average relative velocity standard deviation of 9.4% across an ideally constant-velocity 10x10 mm planar lamina. While the deviation found could be reduced by further widening the channel, the marginal effect was found to be small and not worth further complicating the insertion of a wider system into the gap.

Following standard planar fabrication techniques, the flow channel in Figure 3.3 was fabricated using three PMMA pieces laser-cut and glued together with UV-curable glue (Dymax 1187-M, Dymax Europe GmbH) and found able to withstand the maximum pressure used, 8 bar. The total length of the flow measurement chamber should consist of the length of the VOI, 10 mm, and twice the entrance length necessary for the flow to develop. The width should be increased as much as possible, within the constraints, so as to best approximate a 2D flow profile along the width of the VOI, 10 mm. The flow chamber was thus set at a size of 27x19x3 mm, making use of all the available space within the measurement coil, shown in Figure 3.6.

This channel could be interfaced with a flow system based on a pressure control mechanism (OB1 MkIII+, Elveflow). With two ports, 0-200 mBar and 0-8 Bar, a pressure stability of 0.006% FS, which is well below the precision of common flow measurements, one could consider any achieved flow to remain steady-state for the duration of the measurement. Given the variable hydraulic resistance of the setup, due to flexible tubing, a high-precision Coriolis flowmeter (mini CORI-FLOW M13, Bronkhorst) can be further used to benchmark the flow results over a water flow of 1-2000 g/h, at a precision of 0.2%.

3 Flow NMR methodology

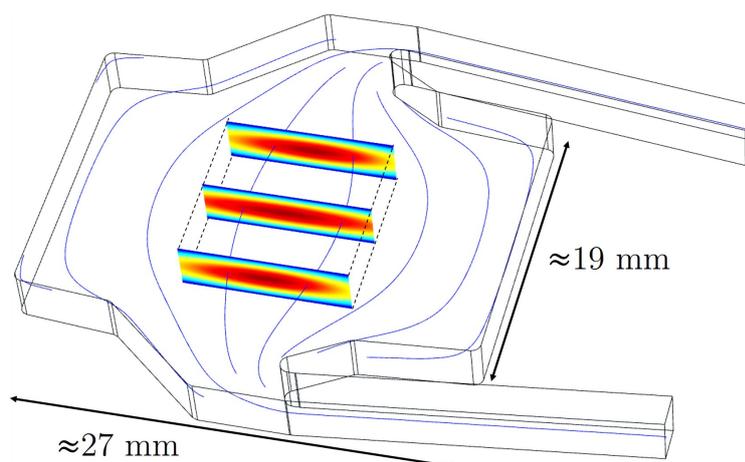


Figure 3.3: Schematic of the flow channel and the streamlines of the fluid flow. In colour, the normalised profile of the velocity along the channel is shown on three cross-sections of the VOI.

Having generated a volume where the velocity approximately matches the desired profile, one must now *select* this volume as the source of signal. This was possible through a specific design of the detection coil. The use of a stripline geometry allowed for a large volume enclosed by the measurement coil while at the same time selecting the desired VOI as a source of signal, unlike what is possible in common resonator geometries. The system and the electronics interfacing the low-cost spectrometer to be used, a LimeSDR, are further detailed in Section 4.1. The narrow, most sensitive section of the stripline was numerically optimised for an optimal total response coming from the VOI, resulting in a 7 mm width and 10 mm length, and is seen to have a higher sensitivity to more central spins, which further emphasises the volume that best represents the desired 2D flow, as seen in Figure 3.3. Excitation on the entire flow volume can be done with a decoupled transmission coil with a flat, solenoid-like design.

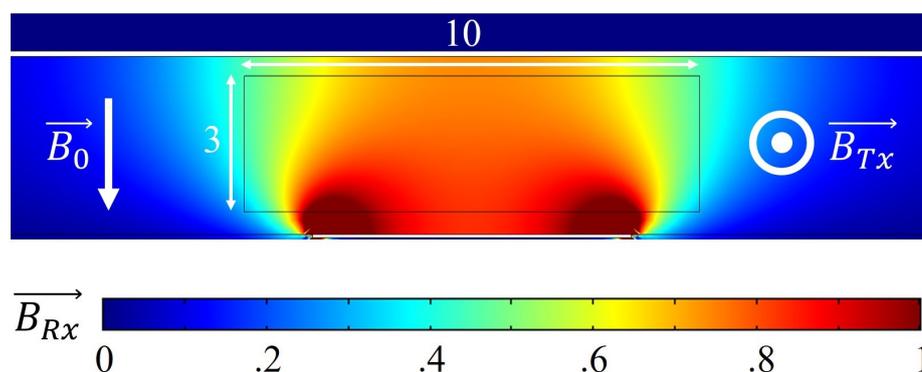


Figure 3.4: Normalised sensitivity profile of the 7 mm wide stripline optimised for detection on the VOI. Directions of the main precession field, B_0 , and the RF excitation field, B_{Tx} , are shown for clarity. The stripline's detection field is in-plane, and perpendicular to both B_{Tx} and B_0 .

The whole measurement setup was designed to enable flow-encoding sequences in a permanent magnet setup, and thus a final component was necessary to achieve this; a set of shimming/gradient coils. The design of the shim system obtained was intrinsically linked to the choice of a gap magnet and the placement of the VOI at the centre of the gap. This not only means a local behaviour of the field (i.e. VOI length of 10 mm over a 60 mm gap) but also a point of rotational symmetry, which implies a local extremum of the field intensity. These two features imply only small high-order spatial components of the field are to be expected and thus a second-order shim set will suffice.

Standard active shimming techniques employ a set of coils capable of producing an orthogonal basis of the B_z component of the magnetic field, assuming a local approximation of $\Delta|B| \approx \Delta B_z$, most commonly through a set of solid spherical harmonics [56]. This set of coils is obtained using the target field method (TFM) [57] which numerically solves the inverse problem associated with the Biot-Savart law, finding the current distribution that best generates the desired field distribution on the VOI.

The continuous current distribution awarded by the method can then be discretised into PCB-manufacturable copper tracks by taking the level set of the stream functions and connecting them together in series [58]. The implementation in a bi-planar coil geometry is based on the work developed in [59] and was used to design seven shimming/gradient coils, whose simulated performance is quantified in Table 3.1.

Table 3.1: Simulated performance of the seven spherical harmonic (SH) shims as given by their position (z), the standard deviation of the generated field on the VOI relative to the SH (E), the power dissipated per T/m or T/m² (P_d), the total resistance (R), the inductance (L) [60] and the inductive rise time (t_{LR}). Results obtained using the code from Miguel Condesso [61].

SH	z (mm)	E (ppm)	P_d (mW)	R (Ω)	L (μ H)	t_{LR} (μ s)
y	7.42	881	9.72	4.1	12	2.9
x	7.09	936	8.60	4.0	11	2.8
z	6.87	658	4.59	3.1	9.4	3.0
yz	6.56	617	18.6	1.4	2.6	0.5
xz	6.34	659	15.8	1.4	2.6	0.5
xy	6.04	756	396	2.4	3.9	0.6
$2z^2 - r^2$	5.82	327	24.9	3.2	6.3	0.5

3 Flow NMR methodology

To experimentally confirm the performance of the shims, the flow channel shown was filled with water and placed inside a commercial NMR system (ICON, Bruker BioSpin) with its shim system disabled. Using a much larger gap magnet further reduced high order inhomogeneity components coming from the underlying B_0 , allowing one to observe the purity of the shims. Using the shim system, it was possible to drastically improve the unshimmed decay time ($\text{FWHM}_{no\ shim} \approx 1/T_2^* = 990\text{Hz} = (1\text{ ms})^{-1}$) to a far longer decay ($\text{FWHM}_{shim} = 28\text{ Hz} = (36\text{ ms})^{-1}$). While this adjustment was manually achieved and could be further perfected, the obtained results, through showing a reduction in inhomogeneity of two orders of magnitude, imply a shim purity below the 1% range, down to the simulated $\approx 0.1\%$.

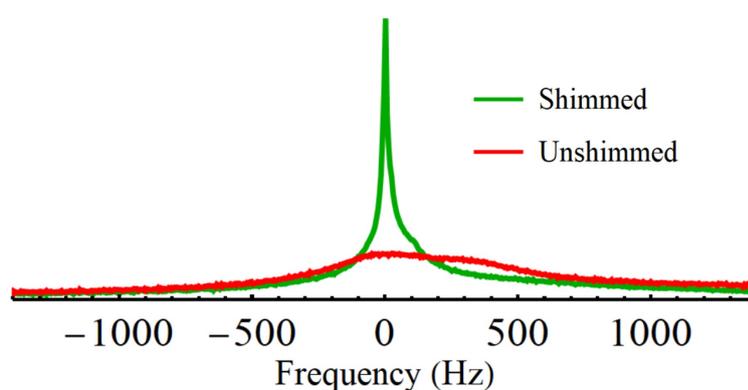


Figure 3.5: Magnitude of the Fourier transform of two equal and sequential experiments with and without the shims enabled.

The shims required for the linewidth correction in Figure 3.5 had a total power dissipation of 6 mW when correcting for a linewidth of 990 Hz. When scaling this value up to the expected inhomogeneity of the gap magnet used in the prototype, one obtains a power dissipation at the 250 mW order of magnitude. This heating power corresponds to an adiabatic heating of the water in the flow channel of $\approx 8\text{ mK}/100\text{ ms}$, a negligible value even in a worst case scenario (i.e. slow flow rates and poor heat management).

Each of the aforementioned parts, devised to be assembled together, were further tested in the ICON imaging system, with a FLASH imaging sequence being used to measure the selection effect of the three components developed. At short TE values (1.6 ms), a mostly proton-density contrast allows one to see how the major signal contribution is derived from proximity to the stripline, with its outline and the flow channel being visible in Figure 3.6. As the spins dephase (TE=20 ms), only the spins on the VOI remain coherent, due to selective behaviour of the shim system, showing a larger signal contribution from the spins at the centre of the VOI. It then becomes clear that the combined weighing of the stripline's sensitivity profile and selective shims will help further emphasise the effect of the wide flow channel, approaching the desired 2D flow.

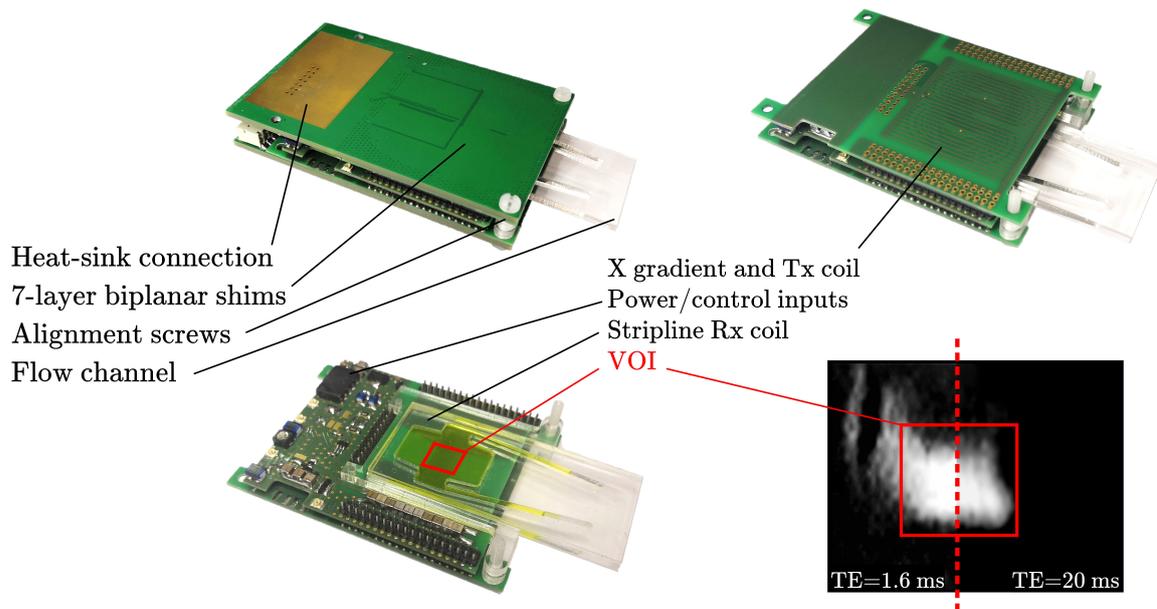


Figure 3.6: Sequential disassembly steps of the measurement setup used for the experiment, with the final photo showing the VOI imaged highlighted in red. MRI images show the same area, mirrored, in two sequential FLASH experiments with different TE.

3.1.3 Flow encoding and measurement SNR

Having proven that a near-2D behaviour can be locally obtained, one must now decide on an encoding procedure that can optimally leverage this ability in any equivalent, volume-selective device. As a prospective method, phase encoding was found to be a promising technique to measure flow, as it introduces information in an often disregarded dimension of 2D low-cost NMR, its phase. Unlike magnitude-encoding methods which, by definition, require a change in magnitude away from the maximum [4, 5, 6], degrading overall SNR, phase can be registered at the maximum signal intensity. A phase encoding approach was thus used to encode flow, as an $\mathbb{R} \rightarrow \mathbb{R}$ function, to best achieve results in SNR-limited applications.

The underlying logic in achieving correlation between phase and velocity is based on generating a gradient of some response along the flow direction. In two steps, some effect is applied to the spins and afterwards removed, for a total null effect. However, if the spins flow along the gradient, whatever was applied will differ from whatever is removed, thus leaving a remaining effect. This is analytically described in Equation 3.1 by considering the incremented phase that is encoded into a spin ensemble after a symmetric bipolar gradient G_v is applied, for a total time T_{enc} , on a fluid element moving at velocity $v_0 + a_0 t$.

3 Flow NMR methodology

$$\begin{aligned}
\phi_{T_{enc}} &= \int_0^{T_{enc}} \gamma(B_0 + G(t) \cdot r(t)) dt \\
&= \int_0^{T_{enc}} \gamma G(t) \cdot (r_0 + v_0 t + 0.5 a_0 t^2 + \dots) dt + \phi_{B_0} \\
&\approx \int_0^{T_{enc}/2} \gamma G_0 \cdot (r_0 + v_0 t + 0.5 a_0 t^2) dt + \int_{T_{enc}/2}^{T_{enc}} \gamma(-G_0) \cdot (r_0 + v_0 t + 0.5 a_0 t^2) dt + \phi_{B_0} \\
&\approx \gamma G_0 T_{enc}^2 (2v_0 + a_0 T_{enc}) + \phi_{B_0} = k_{enc} \bar{v} + \phi_{B_0}
\end{aligned} \tag{3.1}$$

The simplest, modular encoding that achieves a phase injection with spin refocusing is a gradient spin echo sequence, with small modifications to account for effects in portable systems, shown in Figure 3.7. It should employ a correctly engineered 180° pulse, as pulses in inhomogeneous systems have a pronounced shape/power-dependent tilt-angle distribution across the sample, due to the excitation-linewidth and peak-linewidth being in the same order of magnitude. The use of a double unipolar gradient, allowed by the inversion pulse, enables the use of single-quadrant current sources, as well as the rephasing of the signal into an echo, partially avoiding the decay stemming from the inhomogeneity of a portable system. It should be further noted that for long/narrowband pulses, Equation 3.1 must receive a linear correction term into $T_{enc}^2 + T_{enc} \Delta T_{G1-G2}$, instead of T_{enc}^2 , to account for the flow between gradient lobes.

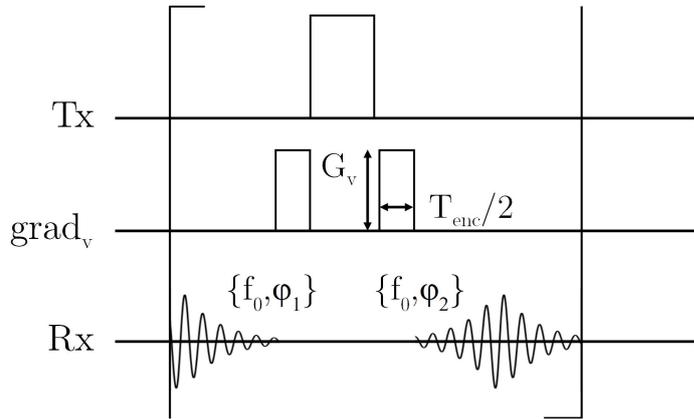


Figure 3.7: Pulse sequence block can be applied during any readout to encode flow, establishing $k_{enc} = \gamma G_v T_{enc}^2 / 4$ to obtain the phase difference $\phi = \phi_2 - \phi_1$. Tx line shows the 180 degree inversion pulse applied.

The acquisition of any two signals within the same measurement procedure allows for a robustness to variable flow conditions, magnetic field drifts and altering matching conditions, as two sequential phase profiles can be measured in two conditions as similar as possible. At the same time, its magnitude-conservative and frequency-independent response, allow for its integration into any other spectroscopic method, enabling enhanced flow characterisations.

The encoding method shown above, while well-resolved for each spin, will showcase an altered response when measured without imaging techniques (i.e. spatial discretisation) in a real flow profile. Ideally, the response of plug flow would match directly a phase-shifted response, without any convolution with some velocity profile, but generating plug flow is not achievable in generalised flow setups. Most standard applications employ a cylindrical channel, which performs sub-optimally due to the sharp velocity/phase distribution ($v = 2\bar{v}(1 - r^2/r_0^2)$) introduced by the no-slip conditions at the channel's boundary. Attempting to reduce this limitation, one could remove wall effects by measuring only the centre of the channel, to the detriment of SNR, or by extending the channel width sufficiently beyond the measurement region, thereby achieving an approximately 2-dimensional flow profile ($v = 1.5\bar{v}(1 - z^2/z_0^2)$). The net-phase of the signal consisting of the sum of all spins, each encoded with a profile/velocity-dependent phase $k_{enc}v$, is shown in Figure 3.8.

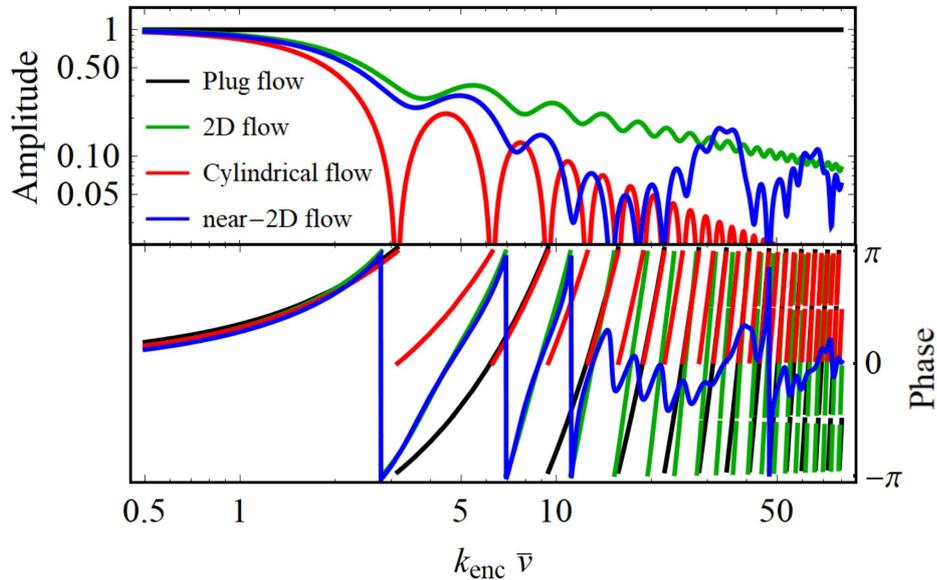


Figure 3.8: Magnitude and phase response of the NMR signal from a phase-encoded fluid with average velocity \bar{v} . Results shown for four flow conditions: $v = \bar{v}$, $v = 1.5\bar{v}(1 - z^2/z_0^2)$, $v = 2\bar{v}(1 - r^2/r_0^2)$ and the flow profile inside the VOI in Figure 3.3.

3 Flow NMR methodology

After establishing the ability to measure a fluid across a large encoding range ($k_{enc}\bar{v}$), the natural follow-up is to define the optimisation functional for the desired measurement. Wanting to have a high precision measurement of flow, Q , the natural functional will be the signal to noise ratio of the measurement, defined as Q/σ_Q . In any well-defined channel with a cross-section A , one can trivially state $Q/\sigma_Q = \bar{v}/\sigma_{\bar{v}}$, along with $\phi = \Phi(v)$ and $\bar{v} = V(v)$. Put together, it follows that $Q = A\bar{v} = AV(\Phi^{-1}(\phi))$, and therefore the SNR of the flow measurement can be stated as a function of the measurement's phase SNR, as long as Φ^{-1} is locally well defined.

Simultaneously, one can precisely state the standard deviation of a phase measurement as SNR^{-1} [45] and the standard deviation of the phase difference between an initial reference and a second measurement as $\sigma_{\phi_2-\phi_1} = \sigma_{\phi_{2-1}} = (\sigma_{\phi_2}^2 + \sigma_{\phi_1}^2)^{0.5}$. Acquiring a reference at a baseline SNR_0 and a second one at $M \times \text{SNR}_0$, due to the amplitude loss after encoding shown in Figure 3.8, one can simplify the previous dependency into $\sigma_{\phi_{2-1}} = \text{SNR}^{-1} [(M^2+1)/M^2]^{0.5}$, leading to:

$$\text{SNR}_Q = \frac{Q}{\sigma_Q} = \frac{A\bar{v}}{A\sigma_{\bar{v}}} = \frac{k_{enc}\bar{v}}{k_{enc}\frac{d\bar{v}}{d\phi_{2-1}}\sigma_{\phi_{2-1}}} = k_{enc}\bar{v} \text{SNR}_0 \sqrt{\frac{M^2}{M^2+1}} \frac{d\phi_{2-1}}{d(k_{enc}\bar{v})} \quad (3.2)$$

This formula leads to the well-known closed-form for the SNR of a flow measurement in a plug flow regime ($\phi = k_{enc}v$, $\bar{v} = v$, $M = 1$), usually approximated well within a single imaging voxel, of:

$$\text{SNR}_Q = \frac{\bar{v}}{\sigma_{\bar{v}}} = \frac{k_{enc}\bar{v}}{\sigma_{\phi_{2-1}}} = k_{enc}\bar{v} \frac{\text{SNR}_0}{\sqrt{2}} \quad (3.3)$$

From the statement in Equation 3.2 and the dependencies shown in Figure 3.8, or coming from any other flow regime, one can now state and compare the performance of net-phase NMR flowmetry measurements in varying channels:

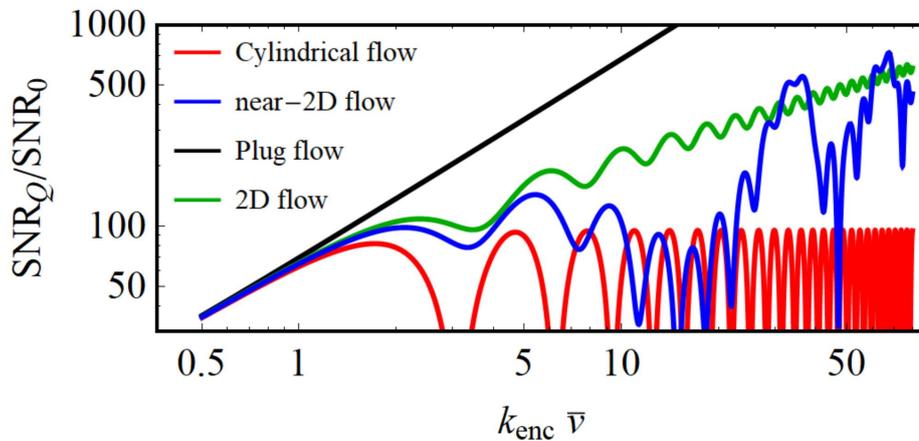


Figure 3.9: Dependence of the flow measurement's SNR on an increasing flow encoding intensity $k_{enc}\bar{v}$, for the four flow conditions introduced above.

3.1.4 Conclusion

In this section, and stemming from the simple first-principles theoretical results shown in Sections 3.1.2 and 3.1.3, three overarching conclusions could be drawn.

Initially, one could see how phase-based encoding methods achieve an SNR of the flow measurement *above that of the NMR measurement itself* and that, among them, near-2D, net-phase encoding methods can drastically outperform current magnitude-based flow encoding solutions. At the same time, on the more practical side, one could show that these can be implemented using commercial techniques, at a low cost. This was demonstrated by the design assembled, shown in Figure 3.1, which could be built for under 1000€, using a standard electronics and mechanical workshop, and requiring only two weeks of implementation effort.

Together, these advances depict how the phase-based encoding approach shown in Figure 3.8, based on engineered flow profiles, open the door to the simplification of NMR flowmeters while showing several advantages when compared to magnitude-based methods:

- *Improved SNR*; the results in Figure 3.9 introduce two new clear advantages past the state of the art. Initially, they depict a drastic SNR advantage of highly phase-wrapped encodings versus standard $-\pi$ to π approaches. Simultaneously, they show how the engineering of a near-2D flow allows for a larger SNR when compared to cylindrical flow, all while remaining fully compatible with standard NMR methods.
- *Large locally-bijective domain*; similarly to what is expected from an ideal plug flow, but very differently from what would be achievable in cylindrical flow, the 2D/near-2D flows show a continually non-zero amplitude of their response. This propriety, while trivial at first look, as one could adjust k_{enc} for a non-zero value of the response of cylindrical flow, means that one continually has a locally-bijective correspondence of phase to velocity. At high wrapping numbers, where a small change in the average velocity can mean a large behaviour change, one does not run the risk of a missed measurement, which allows for the continuous tracking of velocity/wrapping numbers, even in highly dynamic flow.
- *Velocity independence*; the statement of the theoretical framework as a function of $k_{enc}\bar{v}$ instead of simply \bar{v} allows for an encoding process that is fully velocity independent as long as k_{enc} can be adequately adjusted. Such an ability contrasts with the low-turndown ratio of magnitude-based methods, which do not have this freedom. This adjustment further allows one to leverage the increasing statistical significance of measurements at large $k_{enc}\bar{v}$ up to the upper and lower bounds of k_{enc} or up until secondary effects become limiting.

3 Flow NMR methodology

At low velocities, the inability to generate a sufficient k_{enc} , due to a limited current supply, thermal concerns or a T_{enc} rivalling the measurement time or T_2^* , will impart a small, non-statistically significant phase, and thus flow cannot be discerned. At the polar opposite, high speeds, or equivalently low-SNR cases, will disallow phase tracking, thus forcing the reduction k_{enc} to a regime in which one can bijectively track the evolution of the velocity.

Given the motivations above, several future developments can be expected in the field. Promising directions include the shaping of the top and bottom plates of the flow chamber (to better mimic a 2D-like flow profile), the use of a slice-selection step to measure only certain lamina (to obtain an SNR-preferred behaviour approaching that of 1D flow) and the exploration of ultra-high wrapping regimes (for an increased SNR of the flow measurement).

4 | Inductive signal transception

A core part of any NMR spectrometer, whether used for flowmetry or otherwise, is its transduction system, which converts the magnetic response of spins into an electrically measurable signal. Even though systems relying on optical [62], electrical [63], magneto-optical [64] and mechano-optical [65] transduction have also proven successful, the large majority still relies on the inductive readout of the AC magnetic field, using the base methods laid out in Section 2.2. The two following sections, 4.1 and 4.2, showcase two low-cost inductive systems developed from scratch to target gap and Halbach magnets, the two geometries most often used in low-cost NMR.

4.1 Geometrically-differential stripline

The contents of this section are based on the author's contribution to the material published in "Geometrically-differential NMR in a stripline front-end" [24] with the permission of Elsevier.

4.1.1 Introduction

Within inductive solutions, and as NMR progressed into low-cost applications while attempting to maintain performance levels comparable to standard imaging and spectroscopic systems, several approaches for the implementation of the front-end have been proposed. These have often featured a significant leap in at least one aspect of the technique with all other concepts held constant, thus resulting either in a low level of integration, or low versatility.

In this section, a different approach is followed, proposing a complete front-end that is conceived through a top-down approach based on the previous chapters' requirements and whose subsystems are of low manufacturing complexity, commercially available and downscalable.

4 Inductive signal transception

Currently, a first barrier to NMR simplification is the acquisition and digitisation of RF signals, for which there are no simple commercial solutions [66]. For this reason, solutions have historically been based on microprocessors [67], FPGAs [68], and more recently, commercial Software Defined Radio boards [69]. The former require significant technical knowledge in RF design and digital hardware, thereby excluding many potential non-specialist users. The approach shown here was therefore to build upon a commercially available open-source SDR board, the LimeSDR (Lime Microsystems, UK), as a back-end to be connected to the resonator described in this section in detail. The board offers two channels with transmission/reception radiofrequency control, so that a complete spectrometer solution can be realised. Details beyond its integration with the proposed front-end can be found in literature [70].

A second major bottleneck is the availability of suitably high-performance magnetic fields as required for spin polarisation of the sample and spin precession during measurement. Permanent magnet systems have evolved considerably, due to both new design approaches and magnetic materials, and have most commonly been implemented as cylindrical Halbach arrays [71] due to their high field strength, low fringe field, and acceptable homogeneity, as discussed further in Section 5. However, these arrangements have a limited uniform volume, and require a complex assembly procedure and extensive characterisation of its components, which are unsuitable for low-cost implementations. In addition, Halbach field patterns are strongly susceptible to variations in geometry and temperature. In contrast, and more suitably for miniaturised integration, gap magnets suffer less from these issues, they require only two planar magnetic layers, so that their design requires virtually no specialised approach, and their geometry features multidirectional access at an acceptable field strength compromise.

The pre-selection of these subsystems shifts the development towards finding an NMR transceiver that uses a planar construction principle and does not suffer from the proximity to conductive magnet material. In addition, it must contain suitable interface electronics to correct for the power and noise limitations introduced by low-cost spectrometers, so as not to compromise the measurement's operational versatility and signal-to-noise ratio. The remainder of Section 4.1 discusses the optimal design of the transceiving sensor, presenting signal amplification and experiment control schemes, and concludes with results obtained from a microfluidic sample holder chip that showcase a new measurement method in NMR. Sections 4.1.5 wraps up with an overview of the achieved results and some suggestions for their integration current experiments in planar, low-field NMR systems.

4.1.2 NMR detector

Starting from state of the art solutions, and so as to achieve optimal detection, a stripline detector was selected to be used as a standalone receive-only channel. Geometrically, the stripline offers a planar design that is suitable for use in narrow gap magnets and electrically it showcases a high detection sensitivity and ambient noise shielding. This reasoning derives from the results reported by van Bentum *et al.* [72] showing an increased sensitivity over that of a solenoid by a factor of $\sqrt{2}$. Technical advantages were offered by the ease of integration with PCB-like planar fabrication, a central and symmetric placement of the sensitive volume, and facile and modular fabrication/assembly.

The NMR signal amplitude is known to be directly proportional to the excitation magnetic field B_1 , conveniently defined through the principle of reciprocity [33]. For a coil with a unitary field of \mathbf{B}_i at the reference spectrometer phase, temperature T and resistance R_{ω_0} , and a sample with a precessing magnetisation \mathbf{M} , one can simplify Equation 2.8 into:

$$\text{SNR} = \frac{\int \frac{\partial}{\partial t} \mathbf{M} \cdot \mathbf{B}_i dV_{\text{sample}}}{\sqrt{4\pi k_B T \Delta f R_{\omega_0}}} \propto \frac{B_i}{\sqrt{R_{\omega_0}}} \quad (4.1)$$

Equation 4.1 served as the objective function for the geometrical optimisation of the stripline's signal from two sample volumes, one above and one below the strip, of 10x10x3 mm. Optimisation was performed for a Larmor frequency of 44.8 MHz, corresponding to a commercial 1.05 T system (ICON, Bruker Biospin), using an electromagnetic wave propagation FEM simulator (COMSOL Multiphysics, COMSOL AB). Given the skin depth at the working frequency, and the well-defined target volume, only the width of the narrow strip was left as a degree of freedom.

The implemented stripline, connected in parallel on both PCB layers, as is clear in Figure 4.1, has a narrow 7 mm wide path, for optimal sensitivity, and is 10 mm in length, widening rapidly up to a width of 34 mm. This achieved a near optimal total sensitivity while still maintaining an acceptable homogeneity of signal contribution from the target volume. Shielding/return paths were made with two additional PCBs and their Cu layers, spanning 34 mm × 43 mm.

The implemented design consisted of a 1 mm-thick PCB (Multi Circuit Boards) with 70 μm thick top and bottom Cu layers. The inter-PCB connections were made using several pins (TMM/CLT, SAMTEC) connected in parallel to lower their contact resistance.

Some pins were left unused so as to be re-purposed to conduct other signals (shims, gradients, excitation, etc.). The total stack thickness was 11.1 mm, allowing for two sample gaps of 4 mm and 3.8 mm, and could be trivially reduced to below 10 mm by using thinner dielectric layers.

4 Inductive signal transception

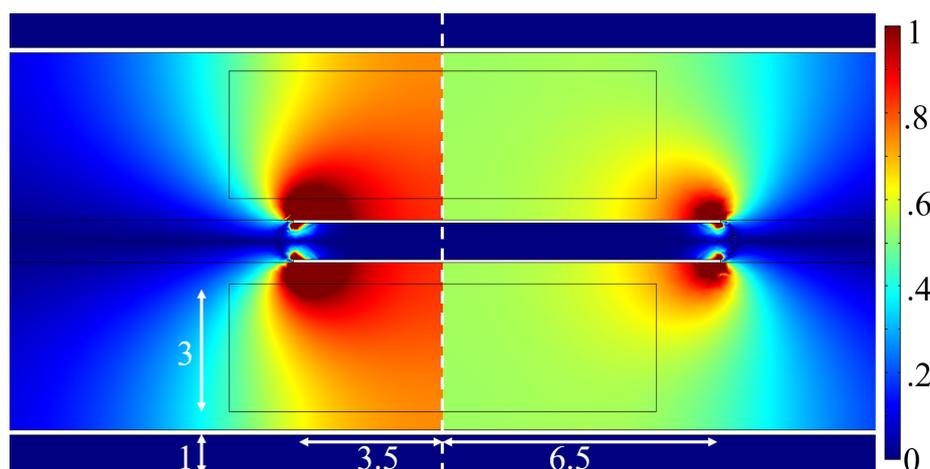


Figure 4.1: Normalised sensitivity map of the stripline over the sample volume, according to Equation 4.1. The B_1 excitation direction is parallel to the stripline's plane and perpendicular to its axis, as shown before in Figure 3.4. The cross section shown is identified in Figure 4.2 as a red dotted line. The optimised detection-only design used is shown on the left and compared to a standard transceiving coil, with a homogeneous excitation field ($\overline{\Delta|B_i|} \approx 5\%$) but lower overall sensitivity, on the right. All dimensions are in mm.

Much like for coaxial cables, striplines have little interaction with the volume outside of the shielding. This is clear from the simulation results in Figure 4.1, where the field outside is negligible (dark blue region) and thus, by reciprocity, the signal is shielded from outside interference before amplification. So as not to disturb this effect, the coil was made resonant using capacitors placed inside the shield. With the top PCB removed, thereby converting the arrangement to a microstrip configuration, one could compare the coupling of the transceiver to a sniffer coil that was symmetrically placed once below (shielded) and then above (unshielded) the centre PCB. The effective shielding was measured to be approximately 27 dB. A second critical effect that disadvantages the use of most detectors in a gap magnet is the generation of eddy currents in nearby metallic surfaces. These would in turn counter the generated B_1 field, and thus significantly reduce the detector's sensitivity. The stripline design is immune to this effect as only a negligible magnetic field amplitude is created outside of the shielding surface.

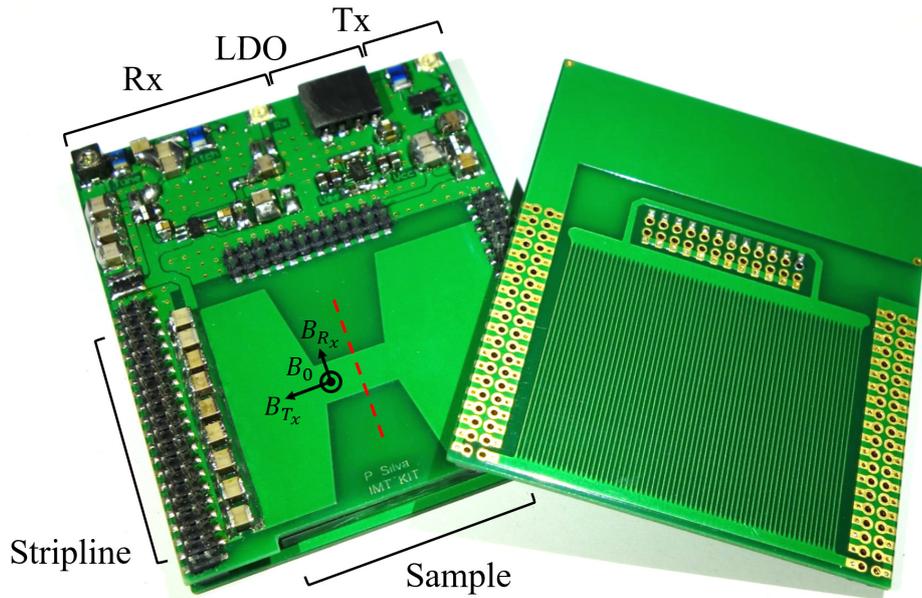


Figure 4.2: Image of the NMR detector assembly with the centre PCB, containing all of the electronics, and the bottom PCB, on the left. The top PCB is shown disassembled on the right for clarity. The resonator design is symmetric and $68 \text{ mm} \times 58 \text{ mm} \times 11.1 \text{ mm}$ in size.

Due to the very low inductance and resistance of the stripline, its lumped parameter values had to be acquired indirectly by measuring it in a resonant circuit, using an impedance analyser (E4991A, Agilent Technologies) whose results were fitted to an LC lumped-model (Advanced Design System, Keysight). Several capacitors were used in parallel (non-magnetic 1111E Series, Johanson Technology), allowing for high quality factors (Q), which, along with the varactors (BB659, Infineon), created a net capacitive system with 1.68 nF and $2.1 \text{ m}\Omega$ ($Q_C > 1000$). From this, a resonator with $L=7.7 \text{ nH}$ and $R_L=14 \text{ m}\Omega$ ($Q_L = 154$) was obtained, which would correspond to an amplified resistance, at resonance, of $R=335 \Omega$. Due to the finite resistance of the capacitor/varactor system, these values were found to be degraded to an effective resonator quality factor of $Q_{LC} = 133$ and a resistance of 290Ω .

As the end goal of an NMR detector is to achieve a high sensitivity to the enclosed precessing magnetisation, several steps were taken to insure a reliable measurement of its end figure-of-merit, so as to better optimise it. A small capillary ($ID=0.8 \text{ mm}$, $L=8 \text{ mm}$) was used as a sample to reduce any influence from radiation damping, field, and excitation inhomogeneity, and drifts on the measurement.

4 Inductive signal transception

A nutation spectrum was measured to maximise the tilt angle on the sample inside the capillary, using only the stripline as a transceiving coil and no other components (switch, varactors, etc.). This signal corresponded to that of a total ^1H magnetisation from distilled water of 3.5 mA m^{-1} , as calculated following a procedure reported by Hoult *et al.* [33], and as measured with the sample placed along the length of the stripline in a region with 0.8 relative sensitivity. The FID was recorded, and had its beginning fitted to the observed exponential decay, allowing the voltage to be correctly extrapolated to $t \rightarrow 0$, by removing the effect of the 1 ms dead-time and thus any voltage contribution from ringing. The maximum Q-amplified stripline voltage was thus $27.8 \mu\text{V}$.

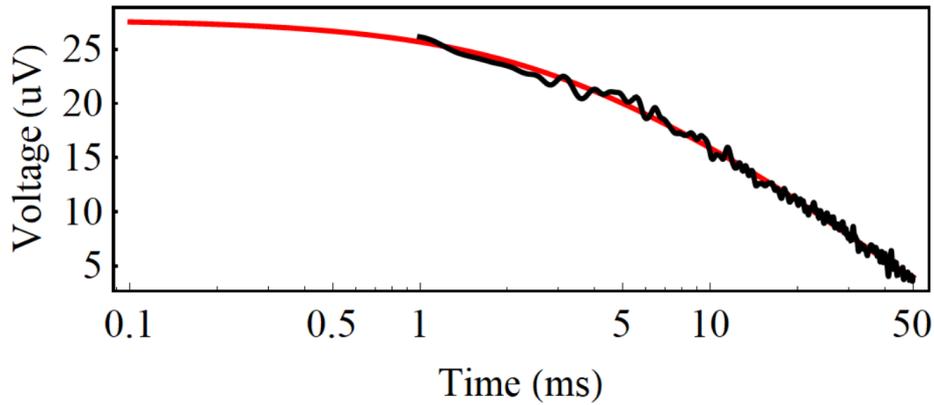


Figure 4.3: Envelope of the FID voltage recorded at the stripline, exponential least-squares fit to the points shown and subsequent extrapolation to $t \rightarrow 0$ s, in logarithmic scale.

The isolated, non-tuned LC tank was characterised ($G_Q \approx 9 \ll 133$), the external LNA's gain was removed ($G_{\text{ICON}} = 29.3 \text{ dB}$), and the voltage was input-referred to the unamplified stripline's voltage, obtaining 266 nV and consequently a maximum sensitivity of $5.7 \times 10^{-4} \text{ TA}^{-1} \Omega^{-0.5}$. This represents a universal figure of merit for the detector which can be adapted to compute other more intuitive metrics for performance, such as the limit of detection (LOD). Unfortunately, the LOD requires assumptions on the B_0 , the pulse sequence, and the measurement time used, which means its extrapolation power is fairly limited. However, adapting the reasoning and methodology in Badilita *et al.* [73], one can infer the minimum volume of water that could be measured in a single experiment, at the B_0 field used, as $968 \text{ pL Hz}^{-0.5}$.

4.1.3 Electronics and operation

To be able to transfer and later digitise the NMR signal with an analogue-to-digital converter (ADC), without sacrificing SNR, it must first be amplified to a sufficient power. The signal-reception circuit found to be the most straightforward, while achieving sufficient performance, was a high impedance acquisition of the signal across the tuned stripline. Using this technique, one does not need to match the coil's impedance, electrically tuning the resonator with a varactor instead. This is a monotonically convergent procedure, that can be automated using a software-controlled voltage, and that works for all frequencies relevant for permanent-magnet systems, as seen below in Figure 4.5.

The accessible tuning range depended on the ratio of variable to fixed capacitance, and was designed to allow for 1 MHz changes due to the interest in ^1H spectra within the B_0 range of temperature-drifting NdFeB magnets. Due to the high quality factor of the varactors, only a small degradation of the total Q was observed, as seen before. An LC notch filter was placed in the path of the biasing voltage, V_{Tune} , and the switch lines, V_{Switch} , to avoid in-band signal/noise leakages, and to filter out out-of-band noise. The filter's Q factor was chosen to match the same 1 MHz range of interest.

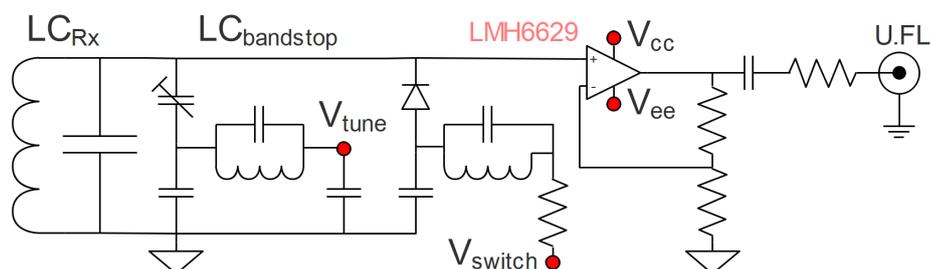


Figure 4.4: Schematic of the circuit used for tuning and switching of the detection coil, and for signal amplification.

At resonance, an inductor in parallel with a capacitor sustains an impedance transformation and achieves an amplified real impedance of $R_L(Q^2 + 1) \approx R_L Q^2$. This impedance change goes along with a Q -fold amplification of the signal, here $Q = 133$, conserving SNR. The resistance value can differ considerably depending on the size, geometry and quality of the NMR coil. An additional active amplification was possible due to the measured OPAMP gain of 24.6 dB, for a total of 68 dB. When compared to standard 50Ω matching, this coil provides a further amplification of 15 dB to 290Ω . The importance of having a large Q amplification is clear when considering the results presented in Figure 4.5, and the adaptation of Friis' formula, introduced in 2.11, and shown here manipulated into the

4 Inductive signal transception

noise factor (F) in terms of voltage density/gain:

$$F = 1 + \left(\frac{N_{\text{LNA}}}{N_{\text{NMR}} \cdot Q_{\text{LC}}} \right)^2 + \left(\frac{N_{\text{rest}}}{N_{\text{NMR}} \cdot Q_{\text{LC}} \cdot G_{\text{LNA}}} \right)^2 \quad (4.2)$$

Much as the Q amplification in the denominator reduces the impact of the LNA noise, the pre-amplification ($Q_{\text{LC}}G_{\text{LNA}}$) is able to render the SNR immune to degradation from further signal processing steps, thereby making the NMR acquisition a straightforward process, with complete acquisition and digitisation possible using an SDR board.

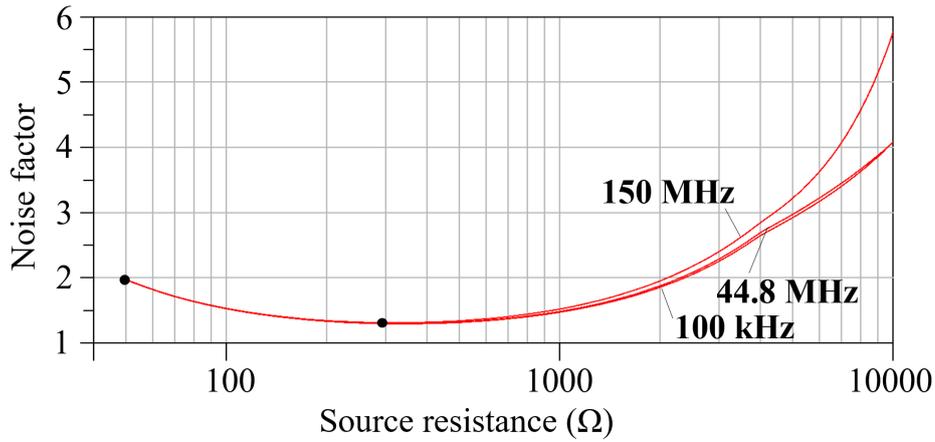


Figure 4.5: Simulated noise factor F as a function of source resistance across an LMH6629 OPAMP, in a non-inverting topology, with an 18Ω resistance to ground. At low R_{source} , F is dominated by the voltage noise of the OPAMP, and at high R_{source} by its input impedance and current noise [74].

From the simulated results for the noise factor of an NMR source with an R_{Source} resistance, shown in Figure 4.5, one can draw comparisons to the standard approach of using 50Ω RF LNAs. These typically have noise factors down to 1.12 (NF=0.5 dB), start at frequencies in the tens of MHz, and require a 2D adjustment of a matching network, whereas the presented approach is seen to span all frequencies of interest, from zero-field to strong permanent magnet systems, and is useful with any coil that can be tuned to 50Ω ($F=1.97$) or above, improving its performance to $F=1.32$ at 290Ω .

The thermal noise's voltage spectral density expected for the tuned coil, after Q amplification, was $2.17 \text{ nVHz}^{-0.5}$ due to the 290Ω resistance. The measured noise value at the output was of $43.8 \text{ nVHz}^{-0.5}$ which, after computing the input-referred value, gives a total noise factor of 1.28 without any external shielding used, and which reasonably matches the simulated value. This result was possible due to the built-in 27 dB shielding on the stripline, and the precautions taken to keep out in-band noise, allowing for a near-ideal-performance front-end without an otherwise common heavy and bothersome shielding layer. For experimental confirmation, a copper/Metglas shielding box was used to re-characterise the noise, and no further reduction was found.

Having started the design discussion with the optimisation of signal reception, one must now consider the other requirements of an NMR experiment, namely that of power transmission for spin excitation. While transmission/excitation can be achieved with a monolithic transceiving coil, there are several benefits to using a secondary and decoupled transmission coil, beyond the homogeneity and sensitivity trade-off arguments already shown in Figure 4.1. These include the ease of simultaneous transmission and detection techniques, critical in samples with low T_2 values, the zero dead-time and ringing that is achieved, as the Tx and Rx coils are never simultaneously tuned, and the ability to achieve broadband matching of the Tx coil. The latter is possible due to its low Q , which would result in a low sensitivity if also used as a detector, but only means here a lower excitation efficiency, an acceptable trade-off. This approach is critical for permanent magnet systems, as frequency drifts will be measured during the T_1 relaxation time of the main ^1H sample, by electrically retuning the Rx coil, and probing an FC43 sample at the ^{19}F frequency. This allows for drift compensation using only one measurement channel, without losing detection time.

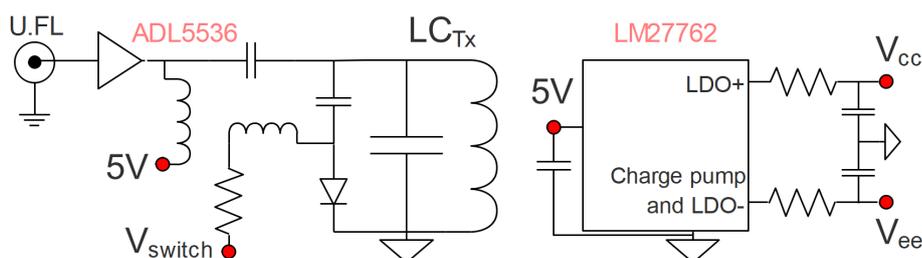


Figure 4.6: Schematic of the circuit used for switching and matching of the Tx coil, and for power amplification. The DC chip includes a positive low-dropout (LDO) regulator, a charge pump, and a negative LDO, which create the regulated voltages to feed the OPAMP. All support components are listed on the manufacturer’s datasheet and hence omitted here.

An independent Tx coil was thus devised to have its B_1 -field geometrically decoupled from the Rx coil, while being connected to the Tx channel through an amplifier with a datasheet-specified gain of 20.1 dB (the measured value was 18.35 dB) and P_{1dB} of 91 mW. The $50\ \Omega$ output leads to the matching network of the Tx coil, which was matched over the relevant bandwidth. The switch was designed so that the coil would only be correctly tuned when the pin diode was conducting, otherwise the coil was tuned to a shifted frequency, as shown in Fig. 4.7.

4 Inductive signal transception

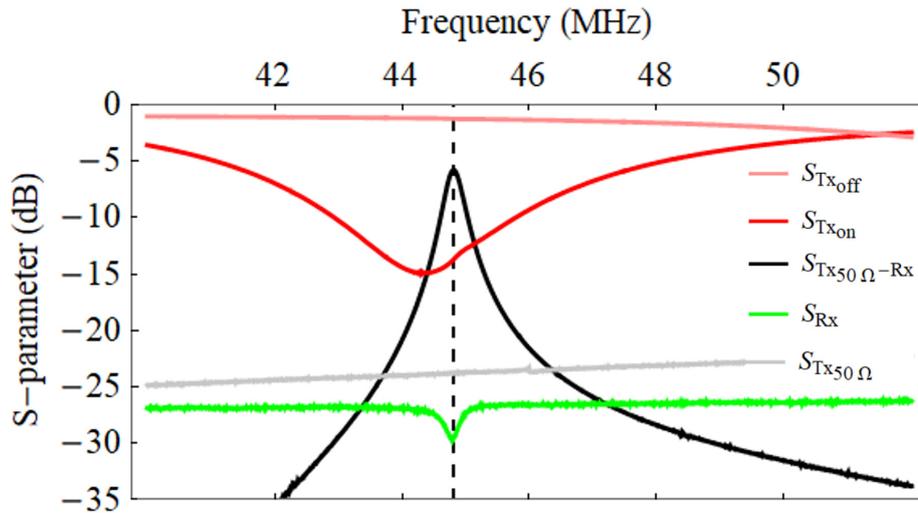


Figure 4.7: Scattering parameters of the Tx and Rx channels for different operational conditions, as measured on a network analyser (E5071C, Agilent Technologies). The legend showcases, from top to bottom, the insertion loss from the tuned stripline to the Tx coil with a 50Ω resistance, and the return loss of the amplified signal, the Tx coil with the switch on/off and of the Tx coil with resistive matching.

As seen on the $S_{12} : \text{Tx}_{50} - \text{Rx}$ curve, it is possible to tune the coil by maximising the transmission between the two accessible RF ports, Rx and Tx. The value shown is obtained using resistive matching of the excitation coil, and is thus a result of a path that includes the Rx – Tx coupling loss, the Q -amplification, the LNA gain, as well as the transmission losses. Removing the effects mentioned, and quantified before, one obtains a geometrical decoupling between Tx and Rx of 73 dB. This remarkable result is easily understood when considering that all field lines created by the stripline close on themselves inside its shield, meaning that they create no flux linkage across the solenoid, and thus induce no voltage.

Another useful characteristic of the design proposed, in addition to the high level of decoupling, is the ability to easily switch between Tx and Rx modes. This is done with a TTL-controlled line which biases two pin diodes, through an LC notch filter at the NMR frequency, so as to block incoming noise/outgoing power while allowing for fast switching of the DC level of the line. At 0 V (Rx mode, " S_{TxOFF} " in Fig. 4.7), the diode on the Rx coil does not affect measurement, and the excitation coil is tuned to a high frequency. At 5 V (Tx mode, " S_{TxON} "), the stripline is shorted, whereas the excitation coil has an increased capacitance to ground, thus adjusting its tuning to the NMR frequency.

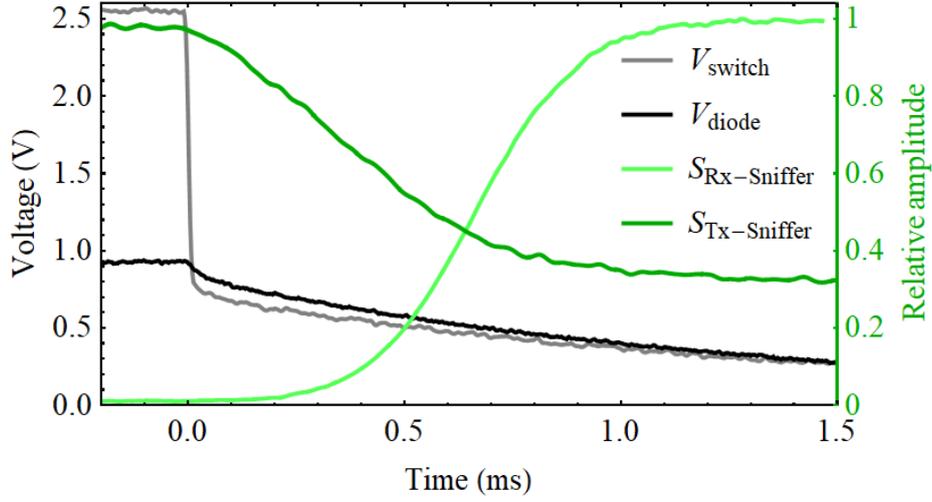


Figure 4.8: Evolution of the RF envelope in both coils, as excited with a loosely coupled sniffer coil, during a Tx-to-Rx switch transition.

The Tx coil was designed as an array of parallel straight tracks, perpendicular to the axis of the stripline. All lines were connected electrically in parallel and placed on the outermost layers of the PCB assembly. Considering that an infinite one-dimensional current sheet creates a uniform magnetic field, this arrangement constitutes the best approximation to a design that combines Tx-Rx decoupling, directional precision of the magnetic field and Tx field homogeneity. The connection of the two track planes resembles a long flattened solenoid, but without the chirality of solenoidal windings. For complex pulse sequences, Tx uniformity is of paramount importance, and thus the homogeneity of the excitation solenoid was computed. Using the dataset corresponding to the experiments shown on Figure 4.9, one can extract the sensitivity profile of the stripline and obtain the excitation angle, and thus field intensity at each pixel. This dataset could be used to compute the effective excitation angle distribution, leading to transverse magnetisation profiles. These, when computed for powers corresponding to two subsequent maxima, result in $A_{450}/A_{90} = 98.4\%$, an exceptionally high value.

4.1.4 Differential NMR

A novel feature of the detector arrangement was possible by leveraging the reciprocity principle in NMR. Consider a planar, circular NMR coil, which features an equal signal contribution of all spins around its azimuth, despite its excitation B_1 magnetic field being non-uniform in direction. This is a consequence of the introduction of a phase on the voltage generated by each spin, cancelling out the different excitation phases by reciprocity, which leads to all signals summing constructively. In the developed prototype, excitation and detection are performed by separate coils, and thus the effect remains present.

Due to its B_1 homogeneity, the Tx solenoid coil imparts a uniform spin phase across the sample volume ($\vec{e}_{Tx} \rightarrow \vec{x}\vec{x}$), whereas the Rx stripline coil, despite its local homogeneity, breaks symmetry across its thickness, so that the spins above ($\vec{e}_{Rx} \rightarrow \vec{y}\vec{y}$) and below ($\vec{e}_{Rx} \rightarrow -\vec{y}\vec{y}$) will be exactly 180° out of phase. Thus, the measured result will be due to the difference of the signal contributions from each sample, that is, one will obtain a differential measurement. This differential operation, formally explicit in the integral of $M(B_{Tx}) \cdot B_{Rx}$ in Equation 4.1, has a multitude of benefits in NMR, which are discussed below. Here, only the results from an imaging experiment are presented, since the post-shimming field homogeneity of the permanent magnet system used is not sufficient for spectroscopy. Nevertheless, Figure 4.9 clearly reveals the differential behaviour with a 180° phase shift between the two sample volumes.

Due to its usefulness in measurement theory, differential measurements have been attempted before [75] in NMR, with limited success, using two separate coils as signal sources, and an electrical signal subtractor. The correct differential subtraction of two separate NMR responses intrinsically requires an identical behaviour and thus the same precession magnetic field at both sample sites, for which proximity is critical. Using two separate coils a short distance apart, however, introduces cross-responses, in which the response of a sample is observed by the two coils. This led to the imperfect signal subtraction reported, which was achieved down to only 10%.

Geometrically-based differential NMR, on the other hand, intrinsically benefits from proximity of the two samples to the coil and thus overcomes this limitation. Furthermore, as any combination of common and differential-mode excitation can be made, the excitation angle, and thus the signal intensity, can be independently set for each sample, and one can therefore precisely subtract the responses of the two samples.

The benefits of introducing differential NMR as a standard technique are summarised in the list below, each illustrated through a common limitation in present setups, and can be extended to offset the limitations of many other approaches:

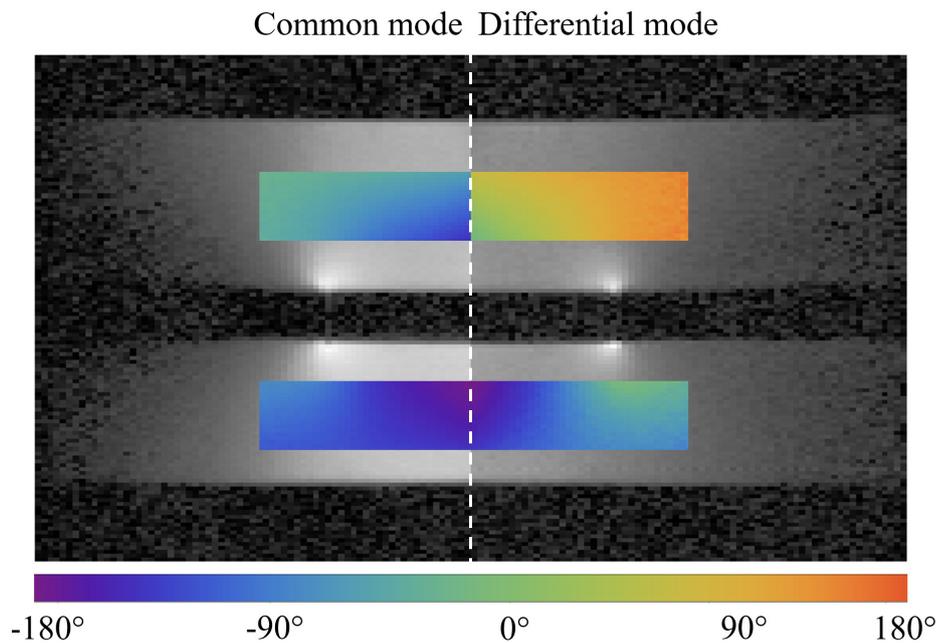


Figure 4.9: Results obtained from two completely identical FLASH MRI images of water samples with a maximal filling factor. The common-mode signal obtained from exciting and detecting with the stripline is compared to the differential-mode obtained by exciting with the solenoid instead. The area showing the signal's phase corresponds to the target volume. The only post-processing applied was a baseline phase shift.

- *Minimised influence of the measurement act;* in inductive NMR, for example, this is materialised in the cancellation of large, unwanted signals (e.g. solvents) and thus their corresponding induced currents. This becomes particularly beneficial in high-Q coils, such as striplines, which are known to cause radiation damping [76] when using high amplification.
- *Reduction of the information received in the signal;* these advantages show up in both the measurement procedure and the data handling. The cancellation of large peaks, for example, can drastically improve the resulting dynamic range and thus the necessary bit-depth/frequency of the ADC which, if insufficient, can significantly degrade the overall SNR [28]. Other NMR readout techniques, like NV-centre and electro-mechano-optical NMR [65], rely on small perturbations for which the simplification/reduction of the signal allows for a higher sensitivity and avoids response saturation.
- *Intrinsic removal of common-mode effects;* beyond the sample itself, common external effects and drifts can be removed. This can reduce other operational problems, such as cross-sensor interactions or coupling to other EM sources.

4.1.5 Conclusion

The presented detector showcases the first implementation of a stripline in low-field NMR (≤ 1 T), optimised for detection sensitivity on narrow flat sample spaces, performing near its theoretically best noise figure without the use of additional shielding, and with adequate operational times for liquid-state NMR. The design proposed has an active volume that can be reduced from the one shown in Figure 4.2 to the millimetre range, whilst improving its LOD, and it can measure a spectrum on sub-nL volumes, making it well-suited for microfluidic samples. While this would carry a decrease in the amplified resistance at resonance, and therefore a sharp increase of the effect of the OPAMP's noise, this effect can be easily offset by using multiple OPAMPs in parallel .

Simultaneously, the independent development of the detection and excitation coils led to several benefits, among which is the first demonstration of a geometry-based differential-mode measurement in NMR, shown in Figure 4.9. The geometrical phase shift technique can be readily adapted to any NMR or MRI application, allowing for multi-faceted simplifications in measurements involving *any* choice of two samples. As an outlook, further experiments in a high-field magnet of spectroscopic quality have confirmed the results and are being extended to best explore the new applications allowed by the differential technique. An expected expansion of the technique will be the use of other differential coil pairs, a planar package coil and a butterfly coil or even a coaxial coil and any cylindrical resonator, for example.

At a system level, the integration of the proposed detector into any complete system would only requires the implementation of the permanent magnet system, as seen previously in Section 3. Its thermally induced magnetisation drifts, usually a major concern with permanent magnets, can be rendered minimal, as the approach chosen allows for broadband excitation, continuous drift observation of a frequency-lock sample, and rapid electrical retuning of the resonators. The former is easily achievable due to the large tuning range of the resonator, which can be readjusted down to the ^{19}F frequency within a us time period.

4.2 Configurable, roll-up double-helix

The contents of this section are based on the author's contribution to the material published in "A multi-purpose, rolled-up, double-helix resonator" [25] with the permission of Elsevier.

4.2.1 Introduction

Departing from the use of gap magnets, and the applications they best serve, one finds that higher-end applications mostly employ cylindrical bores and thus cylindrical transceivers and samples. This approach, far more developed, is used with everything from Halbach permanent magnets to high-field, high-homogeneity superconducting magnets. However, the specifications of each application often require a redesign of the inductive front-end, be it through a change in the frequency of operation, thermal behaviour or field direction, and thus needs significant technical expertise and development effort. Simultaneously, whereas in high-field spectrometers the designs are fairly constrained by a need for a large self-resonance frequency, low-field front-end development is far less constricted and thus more complicated to optimise for each application.

For these reasons, one set out to develop an adjustable low field design, fitting cylindrical samples of various lengths and diameters, targeting a large frequency range, whilst maintaining optimal performance. The resonator itself should have a small incremental volume, for insertion in narrow bores, a high filling factor and transverse-field geometry, so as to be compatible with both longitudinal and transverse precession fields. In addition, it would be favourable to have high winding numbers, both radially and axially, a direct interface to the electronics, and a commercial, low-cost fabrication method, as further elaborated on in Section 4.2.2.

The previously reported double-helix design is particularly well-suited as a transverse field resonator with high winding numbers, being readily fabricated with flexible printed circuit board (flexPCB) technology. The design consists of two tilted solenoids, as introduced in the 1970's during the development of transverse superconducting magnets [77], which were then repurposed into NMR as a higher sensitivity radio frequency transverse resonator [78].

4 Inductive signal transception

The initial RF design used a single coil with a tilt angle, leaving part of the induction field unused by the NMR receiver, namely the component along the axial B_0 field. A second version [79] used the same tilted coil, with windings connected in parallel, but this time with a secondary coil composed of floating loops with zero tilt angle. These removed the axial component of the excitation field perpendicular to the loops, through eddy currents, resulting in a transverse resonator. The authors reported an increase in sensitivity of the overall coil, meaning the loops were inductively coupled with the tilted coil and thus redirected the axial component of the signal. A third version with a double-helix and a single winding pair [80] was implemented, analytically approximated using the approach of Hoult [33], and reported to have a performance above that of a saddle coil, with both designs in their optimised layout. Building upon the state-of-the-art, the design presented below uses a complete double-helix and explores the benefits of arraying it with an optimal electrical connection, up to the SRF limit.

The coil's implementation results, explained in-depth in Section 4.2.3 and presented in Section 4.2.4, report on the frequency-dependent ideal winding configuration. The direct connection to on-flexPCB electronics, and the ability to adjust winding topology to different experimental conditions, are also shown. In particular, the ability to reconfigure the coil introduces the flexibility to run different experiments, at different frequencies, in different zones of the sample, or even just use a smaller section of the coil to acquire a smaller volume, at a larger sensitivity. Section 4.2.5 discusses the relevance of the results obtained, and several follow-up applications are suggested.

4.2.2 Theory

Before any design optimisation of a coil can be made, one must first constrain the design to the experiment's limitations, namely the direction of the B_0 magnetic field of the spectrometer in use, which must be perpendicular to the B_1 sensitivity field generated by the measurement coil. Even though solenoidal coils are considered the gold-standard for NMR receivers, these are often incompatible with highly homogeneous B_0 fields, which are typically used to generate longitudinal B_0 fields in cylindrical bores. In contrast, permanent-magnet setups, as used in portable NMR, have a transverse magnetic field due to the use of Halbach arrays or gap magnets, allowing for the use of both transverse and longitudinal resonators.

Another consequence of a sub-optimal pairing between a coil geometry and the B_0 field is the introduction of susceptibility mismatches. These originate from the introduction of a para/diamagnetic material close to the sample. Common susceptibilities differ in the ppm range, inducing a geometry-dependent magnetic field inhomogeneity with a similar relative difference. This effect can be mitigated using a long sample so as to have the interface away from the measurable spins, an approach available only for longitudinally-built resonators, or through the use of materials susceptibility-matched to the sample/coil, as used in transversely placed resonators. The latter approach manages to mitigate the problem to the sub-ppm range [81]. Through these considerations, it becomes clear that a transverse resonator is the preferable choice for a better filling factor of cylindrical bores at a low susceptibility mismatch to the sample. However, current transverse designs are limited to saddle, Alderman-Grant, or birdcage designs [35], known to be sub-optimal due to a low sensitivity, a need for high B_0 fields, and a high complexity, respectively. Most importantly, these designs are difficult to array with multiple windings placed either along the axial or radial directions, meaning their homogeneity and sensitivity cannot be improved beyond a certain baseline. The double-helix design thus opens up the possibility to create a heavily arrayed resonator with all the benefits of a transverse field.

NMR coils can be implemented in various dimensions and geometries, under varying constraints, but always with the objective of maintaining a high intrinsic sensitivity. A figure of merit defining a coil's performance as an electromagnetic sensor could be defined through the magnetic field generated perpendicularly to the static B_0 field and the coil's resistance, in Equation 4.1, with other dependencies of the SNR being well known and well quantified but harder to influence (e.g. coil temperature). The sensor's sensitivity is then proportional to the signal-to-noise ratio entering the first stage of amplification, usually a low noise amplifier, after which the measurement's SNR is stepwise degraded. This effect is minimised through a high-gain and low-noise LNA which reduces the impact of subsequent stages. Similarly, it is also possible to reduce the LNA's noise contribution by placing the inductor into a resonant circuit with a capacitor. This passive pre-amplification leads to the final SNR given by the adaptation of Friis' equation introduced in Equation 2.11. These two necessary steps immediately set a limitation to the design of the resonator, which must enable an impedance transformation to a resistance above 50Ω . This requires either a matching network to a standard 50Ω RF system, at which power transfer and the LNA's noise performance are optimal, or a high-impedance amplification with an OPAMP, which also requires an equivalent noise resistance in the hundreds of Ohms, as seen in Figure 4.5.

4 Inductive signal transception

Aiming to optimise their performance, NMR coils are commonly modelled using the simplified approach of a current path at a near DC frequency, with the analytical Biot-Savart law, or more comprehensively through a full-wave FEM solver, as was done in Section 4.1, and then sweeping all parameters of the coil. The latter awards few design guidelines when the structure is not easily computable, such as the one shown here, which has led to intensive research on analytically solvable designs, namely with a single winding/structure.

Using the former approach, acceptably accurate for low RF frequencies, one can gauge the expected performance dependencies with different connections of the basic unit depicted in Figure 4.10. Consider a multiple winding inductor with perfect coupling between N_r windings, rolled-up radially, as $\Delta R \ll R_0$, which are connected N_p times in parallel and then sequentially N_s times in series. Mutual inductive coupling between windings decreases with a factor $1/k$ for increasing incremental distance (i.e. from neighbour i to $i + 1$), and only negligible or scalar AC effects exist (e.g. zero parasitic capacitance and scalar increase of the resistance due to the skin effect). One can trivially write the total resistance of the coil as $R = N_R N_s R_0 / N_p$ and its impedance as $L = N_R^2 N_p^{-1} [N_s L_0 + 2M(1 - k^{N_s N_p / 2}) / (1 - k)]$, using the sum of a geometrical series. By conservation of the magnetic energy in an inductor, it is known that $(B_1/i)^2$ is proportional to the inductance, and one can thus state the SNR_L , Q_L and SRF_L as proportional to $\sqrt{L/R}$, L/R and $L^{-0.5}$ for large $N_s N_p$, and thus as:

$$\text{SNR}_L \propto \sqrt{Q_L} \propto \sqrt{N_R [L_0 + 2MN_s^{-1} / (1 - k)]} \quad (4.3)$$

$$\text{SRF}_L \propto N_R^{-1} \sqrt{N_p / [N_s L_0 + 2M / (1 - k)]} \quad (4.4)$$

Despite the significant insight given by the dependencies above, several design adjustments must be made to correct for model simplifications neglecting high-frequency and geometrical effects, with the following design process being proposed. As a first step, the series increase in the coil's resistance is highly favourable, as the introduction of more components/interfaces (e.g. capacitor, switch, solder-joint) is made negligible and the amplified resistance at resonance reaches above the required 50Ω previously mentioned. Beyond this point, it is preferential to generate current in parallel to avoid reducing the overall SRF through an increased parasitic capacitance and a higher inductance. However, this effect is limited in NMR applications due to the proximity effect, as a difference in impedance between parallel windings will shift the current to the edges of the tracks in parallel, drastically increasing resistance. To avoid this, the current must be forced to repeat once again, through a series connection, preferentially achieved by increasing N_r to improve sensitivity.

This inductor can then be introduced into a resonant circuit, most commonly represented as a parallel LC circuit including the inductance of the coil, based its dominant resistance and using the equivalent capacitance of the matching circuit.

This creates a bandwidth reduction, about ω_0 , and voltage amplification approximately proportional to Q , defined for an inductor as $\Im(Z)/\Re(Z)$:

$$Q_{\text{RLC}} = \frac{\omega_0}{\Delta\omega} = \frac{\omega_0 L}{R} \quad \text{for} \quad \omega_0 \approx \sqrt{\frac{1}{LC} - \frac{R^2}{L^2}} \quad (4.5)$$

Introducing only the additional noise of the capacitive system, which can be made negligible, this technique is used as a low-noise preamplification of the signal by creating a resonant amplification of the NMR signal.

As one departs from the lumped component model for an inductor to include a more realistic wave-like behaviour, prevalent at higher frequencies, a spatially varying phase of the current on the metallic conductor and thus also in the generated magnetic field comes into play. This can lead to an inefficient magnetic field or even complete field cancellation, and the length of the coil is therefore usually limited to $\lambda/10$ for the relevant frequency, as a rule of thumb. This phenomena can nonetheless re-emerge through the parasitic self-capacitance that a design has among its tracks or with respect to ground, leading to the phenomena of self resonance. As the impedance behaviour of this effect is similar to that of a series tuned RLC system, below the first resonance, one can characterise it as an equivalent capacitance in parallel, despite this having no physical meaning. Even without self-resonance concerns, the introduction of capacitance will result in losses in SNR and Q , due to the introduction of inefficient current paths.

Given the possibility to reduce parasitic capacitances through a narrowing of the tracks, and its inverse effect on the total resistance, one must analyse two effects which will dominate the resistance of a track. These originate from the intra and inter-winding responses to the AC magnetic fields generated, and are respectively known as the skin and proximity effects [82]. The skin effect introduces an increase of the DC resistance of a single standalone track with increasing frequency. For an infinite planar geometry or very high frequencies, one obtains the well known skin depth $\delta = (\pi f \sigma \mu)^{-0.5}$, for wire conductivity σ and permeability μ . This allows one to gauge whether this approximation is valid for the dimensions of the application. Unfortunately, common NMR frequencies entail a copper skin depth in the range of tens to hundreds of micrometers, a magnitude similar to that of the track dimensions used. Simultaneously, the proximity effect cannot be approximated using common strategies due to the similarly sized spacing between tracks. Despite the difficulty in analysis, this regime is known to be the one that most benefits from the arraying techniques presented [83].

4 Inductive signal transception

Following to the work of Cockroft [84], one can nonetheless obtain some insights into the effects present. The skin effect in a rectangular conductor is well understood and can be analytically approximated, for high frequencies, with $R_{AC} = k\sqrt{4fR_{DC}}$, for a geometrical factor k . This dependence functions as a maximiser when compared to the expected constant at near-DC regimes. However, due to the distances and frequencies at play, the coupling of both AC effects is far more complex and is reported to have been experimentally verified to be *non-monotonous* with regards to the inter-winding distance. This means that no further information can be used to optimise track width/spacing. The thickness of the track, t , can nonetheless be freely increased, up to a fabrication limit, to leverage the resistance's $\approx t^{-0.5}$ dependence.

4.2.3 Resonator design

The repeating unit of the resonator, as detailed in [80] and depicted in Figure 4.10, consists of two current loops carrying the same forward and return current, patterned with the same tilt angle ϕ . These structures can be arrayed axially, one above another, as well as radially, rolled around the sample tube N_R times. For N_R windings wound radially, the resulting single-helices are connected in series, and then again in series with the other single-helix row. Such a pattern is connected in parallel N_p times and arrayed along the length N_s times. Such a combination may be optionally series-tuned by an on-PCB capacitor, as indicated.

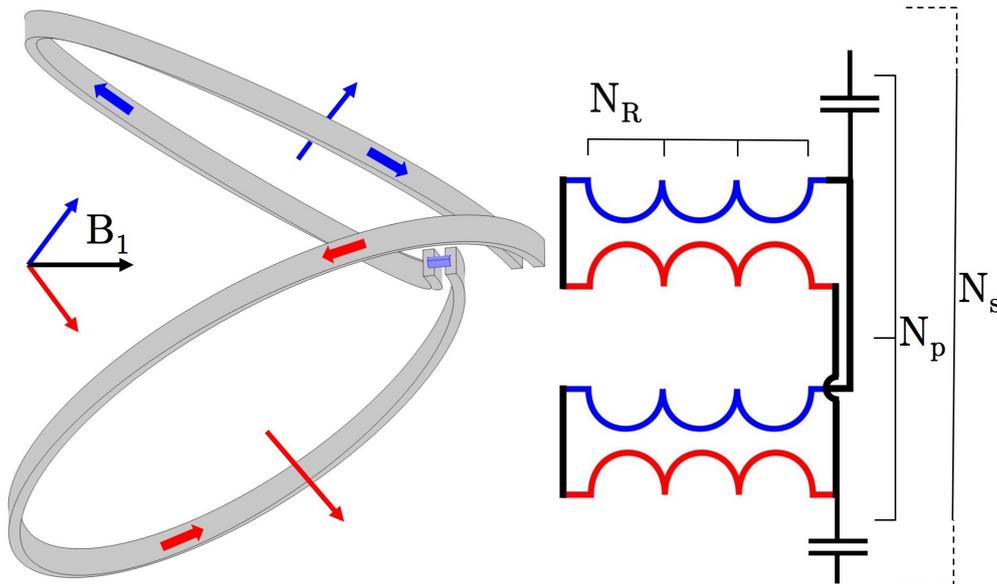


Figure 4.10: Single winding of a double-helix coil is depicted with its current directions and correspondent B_1 fields, and the nomenclature for the electrical connection in various schemes is indicated. Series-tuning capacitors shown are optional.

This basic unit was parameterised, so as to be implemented on the planar, Δ -thick flexPCB shown in Figure 4.11, by turning a cylindrical parameterisation of the coil, a function of r , θ and z , to a planar function of x and z only:

$$\left\{ \begin{array}{l} r = (R_0 + \Delta \lfloor \theta/2\pi \rfloor) \\ z = R_0 \tan(\phi) \cos(\theta) \\ 0 \leq \theta \leq 2\pi N_R \end{array} \right. \left\{ \begin{array}{l} x = \int_0^{2\pi n} \sqrt{r^2 + \left(\frac{dr}{d\theta}\right)^2} d\theta \\ = 2\pi R_0 n + \Delta \lfloor n \rfloor \left(n - \frac{\lfloor n \rfloor + 1}{2}\right) \\ z = R_0 \tan(\phi) \cos(2\pi n) \\ 0 \leq n \leq N_R \end{array} \right. \quad (4.6)$$

The arrayed version of the basic unit in Figure 4.10 constitutes a complex problem due to the multi-dimensionality of the unconstrained number of windings, geometrical dimensions, and possible electrical connections. When possible, the dimensions were set using heuristics based on considerations made in Section 4.2.2. The design-space was defined by the following dimensions:

- *The cylindrical sample-volume enclosed:* The inner diameter of the sample was defined as the normalising dimension and set to 5 mm, to match standard NMR tubes. The height of the sample, and thus the resonator, will significantly influence performance and, so as to minimise case-dependent end-effects, the total axial length of the coil was set at 50 mm to have a 1:10 ratio to the diameter. The net results will thus approximate that of an infinite coil with all the windings equally mutually-coupled.
- *The radial distribution of the design:* These dimensions correlate directly to the fabrication stack of the flexPCB, namely the thickness of the copper, dielectric core, and cover layers. The final stack consisted of two 35 μm copper layers, separated by a 50 μm adhesiveless polyimide layer and covered by two 25 μm cover layers, totalling $\Delta \approx 170 \mu\text{m}$. The inner radius of the coil, R_0 , was set as $2.55 \text{ mm} + \Delta/2$ to match the fabrication guidelines, which require a 10Δ to 20Δ semi-dynamic bending radius. This material stack was then itself stacked as the flexPCB was rolled-up $N_R = 5$ times, a constant across all the designs studied.

4 Inductive signal transception

- *The axial distribution of the design:* Given that the total patternable planar area shown in Figure 4.11B is constrained by the sample height and N_R , the copper tracks could be fully characterised by their width, spacing and tilt angle. The tilt angle of the windings to the horizontal was first set to $\phi = 51.43^\circ$, as analytically optimised in [80]. With any patterning being straightforward, track width and spacing were left unconstrained and part of the design-space on the fabrication set. Six combinations were thus fabricated above the lower bound of $175 \mu\text{m}$ arising from the lithography-based fabrication procedure.
- *The electrical connections between windings:* Due to the ability to access both ends of each individual winding, these could be connected in any series and/or parallel combination, or connected to lumped components, for series-tuning for example. This ability was also left unconstrained, for individual testing, and was enabled by the connection scheme explained and depicted in Figure 4.12.

The multi-winding double-helix coils were designed on Altium Designer (Altium Ltd., USA) and manufactured on 2-layer flexPCBs (PCBway, Shenzhen, China). Each flexPCB was rolled around a 5 mm NMR sample tube and constrained by soldering the locking pads above and below the windings so as to maintain structural cohesion and ensure the correct alignment of the various windings. To avoid delamination of the cover layer due to the stress on the pads, they were given elongated structures to better transfer the force to the substrate. Upon receiving the commercially prepared flexPCBs, these were assembled, geometrically constrained, and had the necessary SMD components soldered onto the coil. Assembly was a straightforward process as the locking pads were developed to be soldered sequentially and independently at low temperature, as the top layer pad on winding i only has a thermal connection to the overlapping bottom layer pad on layer $i + 1$, allowing for alternate roll-up and locking steps.

The different connection scenarios were possible due to the five type of connection pads on the coil, which are connectable by $0-\Omega$ resistors or simple solder bridges, as depicted schematically in Figure 4.11. In Figure 4.12, it becomes clear that the windings and their corresponding pads repeat along the height of the coil, which allows for an arbitrary starting and end point of the active section of the coil. This is possible as all other windings remain electrically floating and disconnected. Likewise, looking at how the windings are connected in parallel, it becomes clear that for the example shown, $N_p = 2$ could be extended to any value, and then put in series N_s times. It is also possible to use the centre pads to connect windings in series with capacitors instead, further improving performance with a series-tuned arrangement.

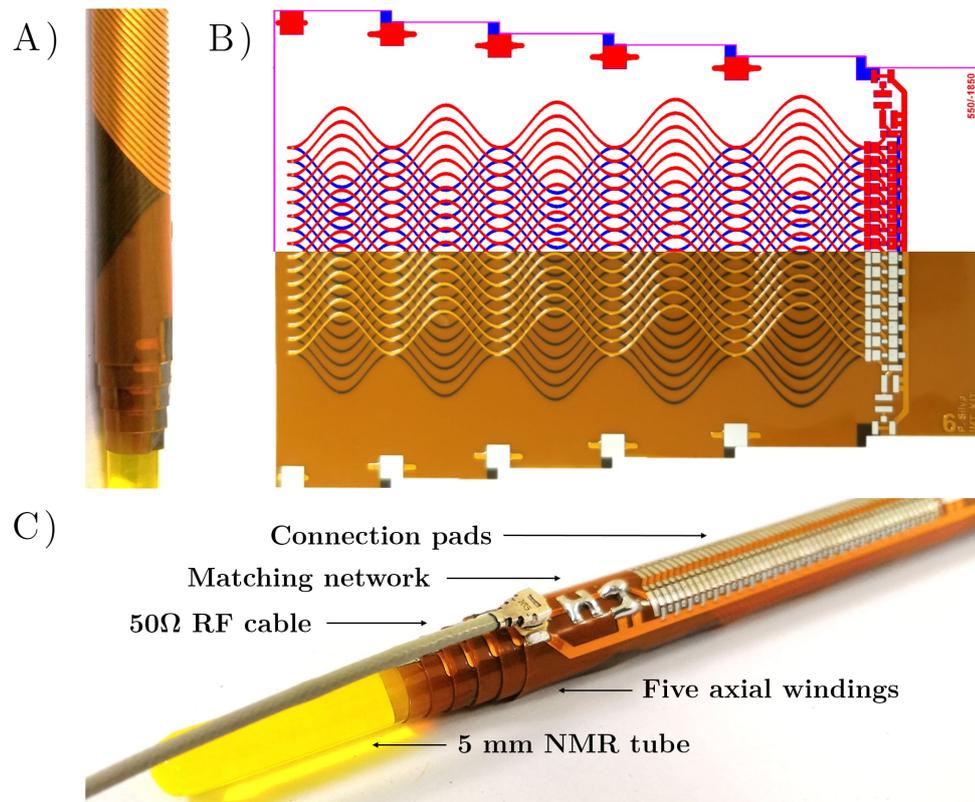


Figure 4.11: In B), the schematic (above) and photo (directly below) of the patterned traces of a flexPCB coil with five axial layers are shown. The configuration depicted has 550 μm wide tracks with a 1850 μm gap in-between, along a length of 50 mm. In A) and C), photographs of the assembled double-helix rolled around a 5 mm NMR sample tube are shown, depicting clearly the double-helix design and the interface with the on-coil pads/electronics.

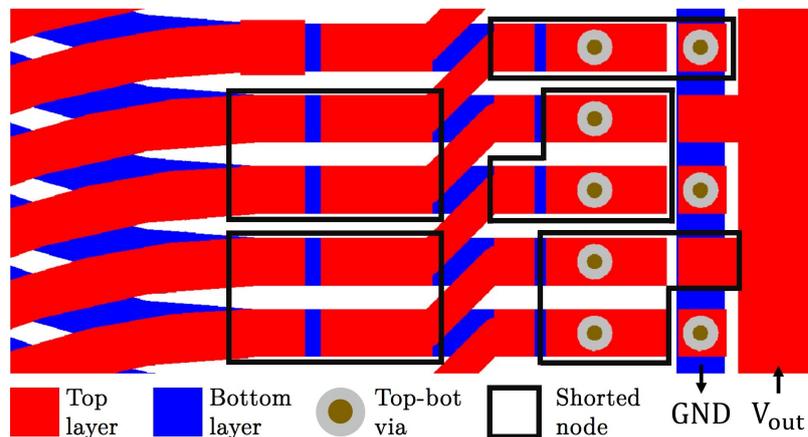


Figure 4.12: Schematic of the on-coil connection pads. The shorted node scheme indicated by the black frames using solder connections achieves a coil with $N_s = 2$ and $N_p = 2$.

4.2.4 Adjustable performance

To gauge the performance of the proposed coil, according to the insights of Section 4.2.2, the Q-factor of the inductor was measured and taken as a proxy for the NMR sensitivity when operating below its self-resonance. As some configurations had a low resistance (0.01-1 Ω), namely when most of the windings were connected in parallel, these values were measured using a 3-port measurement method [85] with a network analyser (DG8SAQ VNWA 3E, SDR-Kits). The S-parameter measurements were used to derive the Z-parameters, and subsequently the Q-factor and SRF, as shown in Figure 4.13 for all the configurations tested.

The remaining, unconstrained design-space of the resonator's was the axial density of copper, n_A , a metric for how close together the windings were patterned compared to their width, and the effective length of current in parallel, l_p , a metric for how the windings were electrically connected together. The latter value, when normalised to the inner diameter of the coil (d), ranges from a single winding "in parallel", which is then sequentially connected in series to complete the coil, to the full length of the coil connected in parallel, giving $l_p/d=36 \text{ mm}/5 \text{ mm}=7.2$.

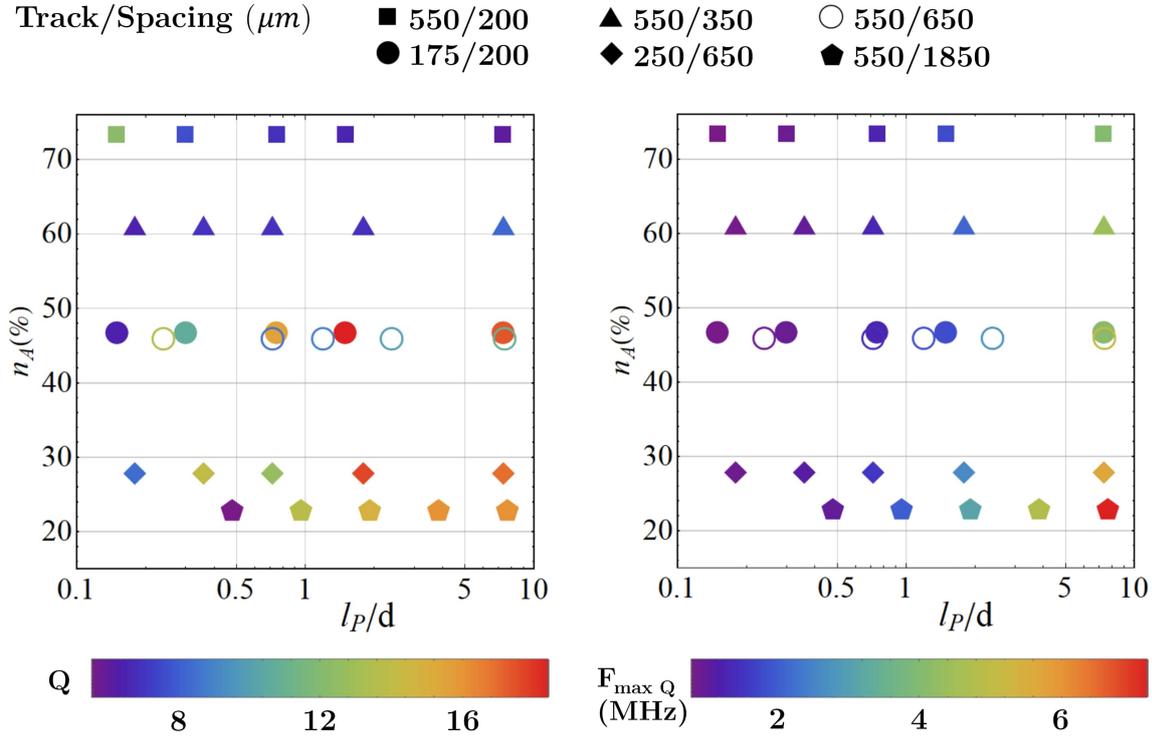


Figure 4.13: Frequency-dependent maximum Q for six track configurations and the frequency at which the maximum was found. Results are shown as a function of the axial copper density, n_A , and axial length connected in parallel, l_p , of each configuration.

Interpreting the results obtained, it is clear that the ideal operational frequency increases towards parallel configurations, regardless of n_A , and remains approximately constant when all the windings become fully mutually coupled at moderate n_A values, similarly to the SRF dependency shown in Section 4.2.2. The configuration showing the optimal quality factor of 17 at 1.9 MHz was the design with 175/200 μm track/spacing) with a third of its length connected in parallel, sequentially, which landed in the middle of the evaluated design-space. This position indicates a trade-off between increasing the axial copper density, for decreased resistance, and the increasing parasitic capacitance it entails.

A natural follow-up to these results is the comparison of the performance of the double-helix with that of a solenoid, while under a transverse B_0 field. Traditional solenoid implementations suffer from a difficulty in radial arraying, but using the roll-up method shown here, this can be overcome using a stack of planar spirals as the unitary winding. However, the need for a return current for the desired azimuthal current, going radially inwards, makes an high-performance implementation difficult to achieve. Furthermore, a stringent comparison of both designs would require a complete search of the design-space of both geometries and all their frequency-dependent optima, which is clearly not viable. For these reasons, and so as to still allow for a qualitative comparison to the design shown here, two flexPCBs with the same 550/1850 μm width/spacing dimensions were developed. Five radial windings were used in both designs, connected in parallel and put to the test.

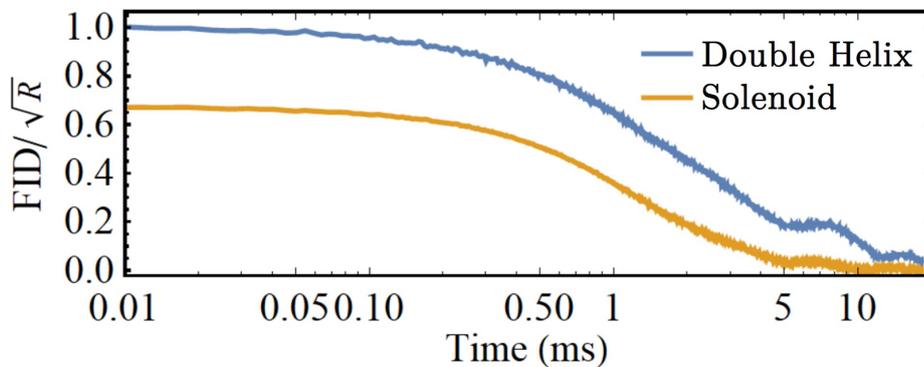


Figure 4.14: Envelope of the FID voltages, recorded under identical conditions and at a field of 1.05T, as obtained from equivalent solenoid and double-helix designs after tilt-angle optimisation.

To properly compare the performance of both designs, the functional shown in Equation 4.1 and used previously with the stripline design was computed through the procedure given by Hoult [33]. A 5 mm NMR tube with distilled water and copper sulphate was measured in a commercial 1.05 T system (ICON, Bruker Biospin) using both coils, and their respective input-referred voltages coming from the total water magnetisation of 3.5 mA m^{-1} were obtained.

4 Inductive signal transception

Measuring the impedance profile of the circuits, it was possible to fit them to a lumped model of the circuit and thus obtain the resonant voltage-amplification and resistance at the measuring frequency. These results are shown in Figure 4.14 and could be further developed to obtain the sensitivity of each of the coils, found to be $1.9 \times 10^{-6} \text{ TA}^{-1} \Omega^{-0.5}$ for the solenoid and $2.9 \times 10^{-6} \text{ TA}^{-1} \Omega^{-0.5}$ for the double-helix.

A particular advantage of a design composed of resolderable, floating windings is the ability to tune the electrical behaviour and geometrical profile to better suit, or even enable, a broad array of applications. The configurability showcased here can be leveraged in any other coil geometry, possibly with larger effects. As an example, given the ability to independently connect different sections of the coil, a dual-resonant matching network was implemented so as to have one section of the coil resonate at one frequency, assigned to ^1H , and another section resonant at a different frequency, ^2H .

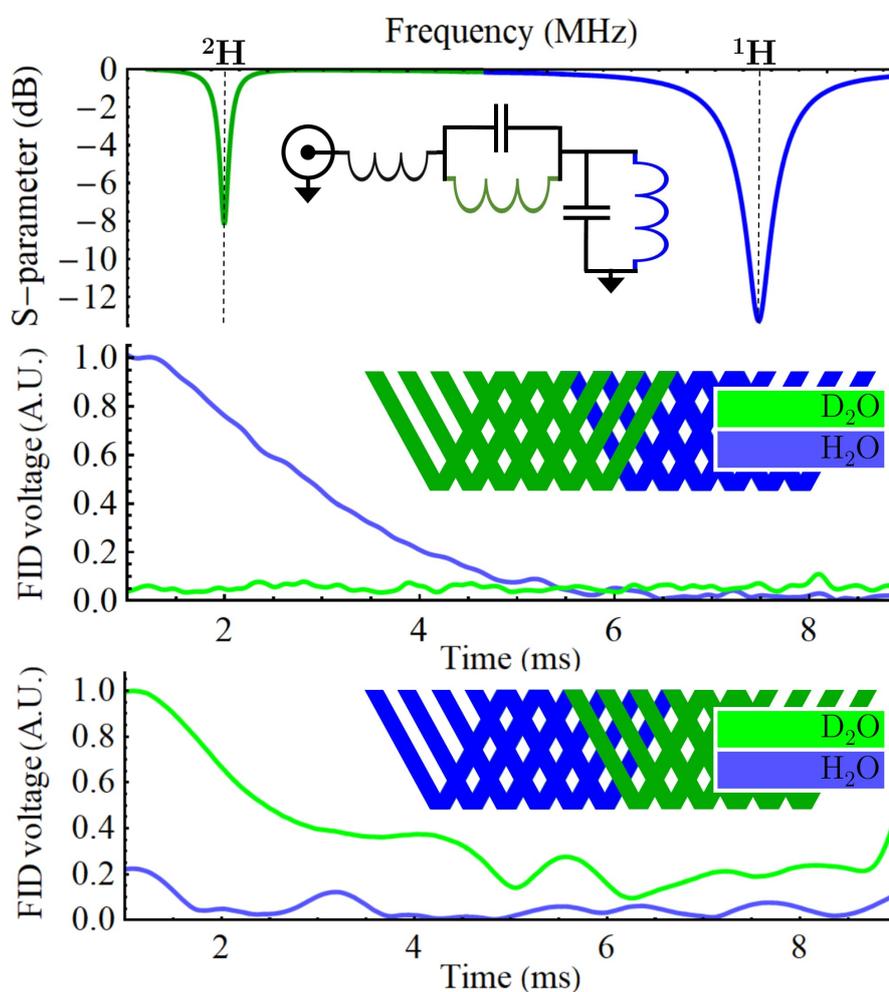


Figure 4.15: The S_{11} parameters of the single-port, double-resonant network represented are shown on top, with the resonance corresponding to each section of the coil colour-coded in blue (^1H) and green (^2H). FID envelope measurements of H_2O and D_2O samples in different sections of the coil, using the same pulse sequence, are shown below.

Two long capillaries filled with H₂O and D₂O, were measured sequentially at both frequencies, on the two sections of the coil, as shown in Figure 4.15. The corresponding nucleus' signals are seen to only be measurable in their respective coil section, demonstrating that the amplified current resonates only within the desired section, and that the coil can thus be geometrically split.

This approach could also be implemented with a section-specific electrical connection, which could perform optimally at different frequencies, due to the ability to independently establish the inter-winding electrical connections. The advantage becomes clear in Figure 4.16, where the same track configuration can perform with an optimal quality factor within an approximately 20-fold frequency range, a ratio encompassing the gyromagnetic ratios of all NMR-observable species.

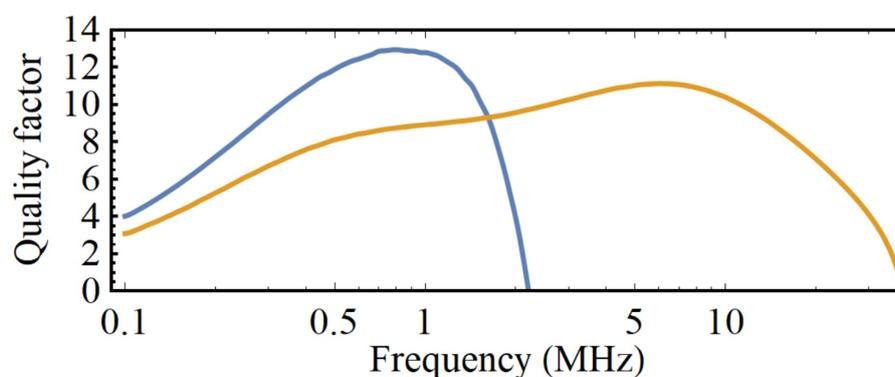


Figure 4.16: Frequency dependence of the quality factor for the 550/650 μm geometry connected in series (blue) and in parallel (yellow).

4.2.5 Conclusion

The work shown in this section demonstrates a new and successful implementation of a transverse resonator design with a more robust, repeatable, high-performing and fully commercial fabrication. Despite a severely unconstrained design-space, one could provide several insights into the design's performance and versatility, all understood through the appropriate RF theory.

An initial, overarching conclusion arises from the *superior performance of the double-helix relative to an equivalent solenoid*. Literature establishes that an axially-wound solenoid performs a factor of 3.1 times better than a saddle-coil [33] and a single-winding double-helix performs 10.5% better than an equivalent saddle-coil [80]. This infers a top bound of a 2x better performance of a solenoid compared to a double-helix. Despite the use of a single radial layer and the simplified theoretical treatment, the results obtained qualitatively match the ones expected. Simultaneously, these relationships showcase that the fabrication method is sub-optimal for use with solenoid-like designs, as their performance is seen to be degraded below the level expected to come from their intrinsic advantage of generating a longitudinal field instead of a transverse one [35]. This suggests that a properly connected double-helix with *multiple radially wound layers* should become the preferred detector for intermediate-frequency applications, where the techniques presented can best be leveraged.

Simultaneously, one could also gauge the effect of *optimal inter-winding electrical connections* on performance. Comparing the standard rule-of-thumb assumptions for the best implementation, a large axial density of copper fully connected in parallel [79] or an $\approx 1:3$ width-to-spacing design connected in series [80, 78], with the best measured design at the optimal frequency, one obtains an improvement of up to 2.7 in the quality factor and consequently of $\approx 64\%$ in NMR detection sensitivity. This clearly establishes that a track layout must be tested/simulated in varying connection schemes to guarantee an optimal performance in any resonator design.

Beyond direct sensitivity improvements, the proposed design proves the ability and *sensitivity-driven need to reconfigure the system*. The Q-optimal track connection for the very same geometry, for example, can be reconnected to show a SRF between 1.1 and 30 MHz, which is a sufficient range to measure all NMR-active nuclei in an electrically-adjusted configuration. Beyond this, the ability to geometrically reconfigure the coil means that a single design/part can measure samples of variable size by shortening the coil, can measure multiple resonances by resonating differently in each section, or even include switching electronics for more complex behaviour.

Together, the results presented show a promising future for roll-up resonators in NMR, namely as fabrication and application-specific designs improve. With the thickness-to-bending-ratio limitation relaxed, better performing (i.e. thicker) double-helices on flexPCBs will be possible, leading to lower track resistances, diminishing intra-winding parasitic capacitance and copper-layer numbers above the two currently achievable. The latter will prove invaluable when designing decoupled double-helix pairs for longitudinal fields, a promising and easily achieved solution due to the design's patterning precision and ability to turn a planar offset into an angular one, when rolled-up. A promising research direction could rely on the axial/radial stacking of several resonators, to be rolled-up together with different N_r parameters and electrical connections, allowing for a large variation in self-resonant frequencies and thus opening the door to integrated EPR/DNP, for example. Fortunately, the magnetic field of a layer/coil-section far above its self-resonance is significantly reduced/shaped, which drastically reduces the mutual coupling to the other inductors rolled together. This can be used to stack multiple near-decoupled groups of resonances, in differing frequency ranges, in different layers/coil-sections, thus avoiding the existing limit of penta-resonant matching networks [86].

5 | Permanent magnet design

Magnetic field generation is an intrinsic requirement of MR equipment, providing both the source of spin polarisation and the precession field. Simultaneously, it poses the hardest engineering challenge among all the subsystems required to implement an NMR flowmeter/spectrometer as it requires a tight integration (i.e. high filling-factor) with front-end electronics and the sample itself, all while dominating the end-performance of the measurement. The following section explores several solutions aiming to simplify the design and fabrication of permanent magnet solutions, which have proven optimal for portable and low-cost NMR.

5.1 Magnetostatic reciprocity

The contents of this section are based on the author's contribution to the material in "Magnetostatic reciprocity for MR magnet design" published in *Magnetic Resonance in Medicine* [26].

5.1.1 Introduction

The vast majority of existing MR systems currently employ one of three magnetic field generation technologies: cryogenic superconducting magnets, dissipative electromagnets or rare-earth-based permanent magnet arrays. The preferred approach remains the use of superconducting magnets, due to their intensity's impact on SNR, but this effect is far from the only criteria, as seen from the upcoming use of Earth or even Zero magnetic field setups [55, 87]. Even despite a sharp dependence of the MR sensitivity on field intensity, the high cost of superconducting magnets has resulted in an increase in the use and development of compact permanent-magnet systems. These became an attractive solution, with an acceptable trade-off between cost and performance, starting in the 1980's [88] with the discovery of high-remanence and high-coercivity magnetic materials, such as Neodymium-based alloys.

5 Permanent magnet design

As better detection methods, hyperpolarisation techniques, and excitation schemes lower limits of detection, new applications based on these magnet systems have been made possible by offsetting the SNR loss stemming from the reduced magnetic field intensity. Beyond a cost reduction, permanent magnets have allowed for portable and non-standard MR applications and can be divided into narrow-bore spectroscopy magnets [89, 90], large-bore imaging magnets [91, 92, 93, 94], and other speciality magnets providing an external homogeneous volume [95, 96, 97, 98, 99], profiling gradients [100, 101, 102, 103, 104], or pre-polarisation fields [105, 106].

Despite the numerous published concepts, several key bottlenecks and cost factors remain, hindering the design and construction of suitably good quality magnet systems. Reported challenges include the complexity in using tens of source magnets [91, 92, 106], magnetisation strength error [91, 92, 106] of 1%, magnetisation direction error [91] of 2°, magnet size error [106] of 0.1 mm, as well as experimental difficulties regarding assembly-alignment errors and the large forces associated with strong permanent magnets. As target inhomogeneities in MR lie below the parts-per-million (ppm) range, *a priori* magnet deviations of 1% constitute a massive obstacle in achieving high homogeneity magnets. Current methodology requires, for example, the time-consuming analysis of each magnet in a large pool, so as to use only the most similar magnets [91], the individual characterisation and ordering of an array to reduce inhomogeneity [91, 92, 94], or the use of high-precision mechanical adjustment systems [89]. All of these approaches limit the number of magnets in a system, and as a result, the final field homogeneity. These efforts are often not possible or sufficient, and thus additional passive shimming techniques are required [91], in addition to active shimming.

So as to overcome these limitations, and building upon the self-assembly idea mentioned in [107], the long-known principle of magnetostatic reciprocity [108, 109] was considered for MR permanent magnet design. Geometrical reciprocity is ever-present in MR in the form of electromagnetic reciprocity [33], which creates a correspondence between a current distribution in the MR detection coil and the signal induced in it. Similarly, one can deduce a reciprocity between two regions in space and the magnets/fields contained therein. This allows for the quantitative *evaluation* of proposed magnet designs, as well as in planning the *assembly* of the geometry. An easier alignment/assembly allows for the use of a larger number of discrete magnets, which in turn allows for the discretisation of magnetisation distributions with better fidelity, and the use of smaller, cheaper and safer magnets [88].

The principle and its direct and indirect consequences are initially discussed in Section 5.1.2, subsequently enabling the computation of a material-independent figure of merit for magnetic systems in Section 5.1.2. In an equally analytical approach, the impact of an increasing number of magnets on the field quality of a Halbach array was researched, in Section 5.1.2, to showcase the benefits of being able to handle a larger number of magnets.

In an attempt to demonstrate the breadth of possibilities and insights allowed when designing with the principle, three different applications were proposed. The first leverages the minimum-energy state that comes with maximum magnetic coupling, for an easier auto-assembly of a Halbach magnet, and is shown in Section 5.1.4. A second application proposes an arbitrarily chosen development goal and shows how the reciprocity principle can elucidate the design process at each step, allowing for the complex behaviour shown in Section 5.1.5. The final application maximally leverages the principle, enabling a topological optimisation method to obtain the highest possible magnetic field in a single-sided magnet, which shows an unbound magnetic field strength in Section 5.1.6. The methods used for simulation, material modelling and result post-processing are shown separately in Section 5.1.3, for clarity, and the paper concludes with a discussion of the results and an outlook for the technique in Section 5.1.7.

5.1.2 Theory and corollaries

Consider an ensemble of magnetic sources M , a magnet being designed, and A , an "anchor" magnet serving as support tool, either virtual or real, for the design and assembly of M . Defining both through their magnetisation M , generating the magnetic field H_A and H_M in space, the energy reciprocity between the two can be stated as per [110]:

$$\int H_M \cdot M_A \, d\mathbb{R}^3 = \int H_A \cdot M_M \, d\mathbb{R}^3 \quad (5.1)$$

In other words, the magnetic vector potential's equivalent magnetic moment M_A has an energy in the magnetic moment's external field H_M that is equal to the energy of the original magnetic moment M_M in the external field of the vector potential H_A . If one assumes high-coercivity magnets with an effective tensorial permeability $\mu_0\mu_r$ and remanent field B_r , i.e., a model which approximates well the behaviour of common hard magnetic materials such as NdFeB compounds, it is possible to set $M = (\mu_r - 1)H + \mu_0^{-1}B_r$. Taking the integration volume \mathbb{R}^3 to volumes where M is non-zero, i.e., V_M and V_A , one

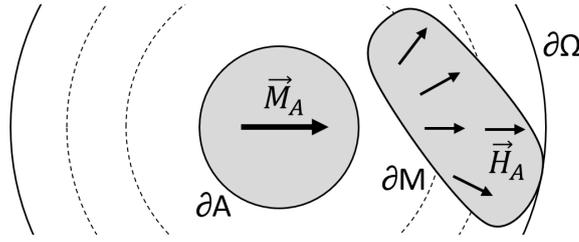


Figure 5.1: Representation of the domains involved in the reciprocity principle. As an example, the magnet M under development inside Ω should create a homogeneous field on the anchor volume A . Reciprocally, the field created by A on the volume M is shown and sets the direction that would better create a field in A with a uniform direction, i.e., a Halbach distribution.

can transform Eq. 5.1 to:

$$\begin{aligned} \int B_{r_A} \cdot H_M dV_A &= \int B_{r_M} \cdot H_A dV_M \\ + \left(\int [\mu_0(\mu_{r_M} - 1) \cdot (H_A + H_M)] \cdot H_A dV_M \right. & \quad (5.2) \\ \left. - \int [\mu_0(\mu_{r_A} - 1) \cdot (H_A + H_M)] \cdot H_M dV_A \right) \end{aligned}$$

For the purpose of interpretation, we first disregard the last two terms of Equation 5.2 and define A as an *a priori* design decision aiming to have an anchoring effect on M . Under these conditions, if the remanent magnetisation of the anchor, B_{r_A} , matches the field we wish to generate with the magnet, H_M , and thereby maximises the right side of the equation, then the total alignment of the remanent field being developed, B_{r_M} , with the field created by the anchor, H_A , is also maximising the alignment of the field created by that magnetisation to the target field.

Whereas this approach matches exactly the ideal coupling between two equal magnets, this remains an approximation when developing MR applications, as the permeability of samples and magnets differs. One can nonetheless estimate the relative inefficiency of a design in optimally placing magnetic energy in the anchor volume in which a sample will be placed. This inefficiency is dominated by the relative amplitude of the last two terms of Equation 5.2 w.r.t. the first one on the right. Beyond being a differential term with a similar integrand and integration volume, the last two terms are further multiplied by a reductive coefficient, $(\mu_r - 1)$, which takes small results in neodymium magnets as $\mu_{\parallel} = 1.03$ and $\mu_{\perp} = 1.12$ [111], for example. This clearly shows how the first term on the right remains dominant, energy-wise, and one can thus approximate the response through a complete linear superposition.

Considering that Equation 5.1 holds for an arbitrarily small M , one immediately derives that a higher anchor field in a location will also mean a larger contribution to the target field from magnetisation located there. This result is analogous to the principle of electromagnetic reciprocity in NMR, which tells us that the contribution of a spin in a specific location to the total signal, but not the design that creates a certain field shape. Similarly, the discussed method does not elucidate the boundary design which leads to an homogeneous target field.

Assuming $B_{target} = -B_{r_A}$ in Equation 5.1, one can see how maximum compliance with the target field is given by the minimum magnetostatic energy condition. This corollary indicates that a freely rotating magnet will tend to the direction that better generates the target field needed to minimise the energy with the anchor magnet, thus providing the force needed for *self-assembly*. Just as this corollary proves the hypothesis stated in [107], used to auto-assemble 8 magnets into a Halbach [71] array, it also elucidates an important limitation that arises, namely that of the self-interaction of the designed magnetisation distribution. The ideal magnetisation distribution has a minimum energy in its interaction with the anchor magnet, but not on its self-interaction energy, which has it evolving to a different, global energy minimum. For this reason, each individual magnet must be aligned individually, or the self-interaction must be countered with secondary magnetic fields, so as to have a high fidelity reproduction of the ideal distribution.

Knowing the remanent field distribution that maximises the magnetic field in one volume, one can thus compute the efficiency of an array in matching that distribution. In the case shown, the volume that can be filled with magnet material is contained between ∂A and $\partial \Omega$, corresponding to the domain Ω . The outer boundary is defined as the weakest contour line of H_A that intersects M , as indicated in Figure 5.1, but this definition can be adapted depending on the application and may be replaced by a symmetry line (e.g. profilometry applications can only populate a half-space). On this volume, one can compute the coupling parameter η , which computes how well the magnet under design, defined by its normalised remanent magnetisation $b_r = B_r / \max(|B_r|)$, is aligned with the field created by the anchor, as per:

$$\eta = \overline{H_A \cdot b_r} = \int (\vec{H}_A \cdot \vec{b}_r) dV_M \Bigg/ \int |\vec{H}_A| dV_\Omega \quad (5.3)$$

For example, in an ideal 2D Halbach magnet, the target field has perfectly circular contour lines, so that Ω is bounded by the least contributing magnetisation, the outer diameter of the magnet array, and the boundary of the circular sample, the inner diameter of the magnet array. These match exactly the placement and direction of magnets in an ideal 2D Halbach array and the coupling is thus optimal (i.e. $\eta = 1$). Most real implementations with sub-optimal filling factors, for example, will nonetheless have $\eta < 1$, while even $\eta > 1$ could be achieved if packing is optimal and the demagnetisation field contributes to a net-positive polarisation of the material. The coupling parameter compares designs quantitatively and provides a computable cost to inefficient but necessary design techniques, like introducing a dead-volume to increase homogeneity [91, 92], using an anti-aligned piece to cancel first order field gradients [99, 112], or overextending a magnet to emulate an infinite height [96].

High- n applications

A design choice that often directly impacts the performance and complexity of a magnetic arrangement is the number of individual parts involved. A large number of these will self-evidently allow for a better discretisation of a desired magnetisation profile. This introduces the possibility of a stronger magnetic field, by better matching the functional introduced in Equation 5.3, but the dependence on the experimental effects governing field inhomogeneity remains unclear.

For high homogeneity applications using permanent magnets, the centre of a 2D discrete Halbach field remains the ideal placement area for a test sample, due to a saddle point of the magnetic field intensity in the axis of the arrangement. A symmetrical placement of four magnets, or a combination thereof, will show a saddle point with zero first and second derivatives of the field intensity. This means that any deviation from an ideally-built discrete array can be modelled well, locally, by the total effect of each contribution on its directional, first-order derivative at the origin.

Consider the magnetic field generated [110] by one 2D cylindrical magnet of radius r_0 and uniform magnetisation M_0 , with an angle defined by a rotation matrix R_θ , and placed at $\{x_0, y_0\}$:

$$B(x, y, x_0, y_0, \theta) = \frac{M_0 r_0^2 \mu R_\theta \cdot \{(x - x_0)^2 - (y - y_0)^2, 2(x - x_0)(y - y_0)\}}{((x - x_0)^2 + (y - y_0)^2)^2} \quad (5.4)$$

Summing the effect of various cylinders allows for the analytical description of the field generated by an array of n magnets in a Halbach configuration, placed on the unitary

circle:

$$B_H(x, y) = \sum_{i=1}^n B(x, y, \cos(2\pi i/n), \sin(2\pi i/n), -4\pi i/n) \quad (5.5)$$

This equation enables the computation of $\nabla|B_H| \cdot \nu$ at the origin as a FOM, for any number of magnets. Assuming that the experimental uncertainty in magnetisation intensity, angular direction and piece dimensions have a small relative deviation, ϵ_i , one can use the small angle approximation and obtain the following figures of merit (FOM):

$$\text{FOM}_{\Delta Mag} \propto \text{FOM}_{\Delta r_0} \propto \nabla \left| \sum_{i=1}^n (1 + \epsilon_i) B_i \right| \cdot \nu = \sum_{i=1}^n c_i(\nu) \epsilon_i \quad (5.6)$$

$$\text{FOM}_{\Delta \theta} \approx \nabla \left| \sum_{i=1}^n \begin{pmatrix} 1 & -\epsilon_i \\ \epsilon_i & 1 \end{pmatrix} B_i \right| \cdot \nu \propto \sum_{i=1}^n c_i(\nu) \epsilon_i + n^{-1} \sum_{i,j=1}^n k_{ij}(\nu) \epsilon_i \epsilon_j \quad (5.7)$$

In Equation 5.7 one arrives at a total inhomogeneity that is dominated by a weighed linear sum of the various effects. Assuming that ϵ_i has a Gaussian distribution around a zero mean, it becomes clear that the field inhomogeneity will increase as \sqrt{n} as the the sum of statistically independent distributions of the magnet's behaviour.

Due to a linear superposition, the field scales with the number of magnets, n , while the effects dominating the inhomogeneity of the field, originating from fabrication errors, scale as \sqrt{n} , which means an overall improvement of the relative field inhomogeneity of $1/\sqrt{n}$. This shows a deep incentive in increasing the total amount of magnets in an assembly, which follows common fabrication approaches [91, 92, 106].

5.1.3 Modelling and result standardisation

As one departs from analytical approaches to possible applications of the principle of reciprocity, numerical/computational solutions become key. However, while analytical formulas are intrinsically application agnostic, some work must be done so that simulations and other figures of merit can be applied across fields.

5 Permanent magnet design

Given the scale-independence of the laws governing magnetic fields, it becomes natural to normalise all the reported values so that their use is straightforward across materials and application sizes. For this reason, all values shown are presented as adimensional. Magnetic field flux intensity is normalised per unit remanent field B_r , meaning field intensity will scale directly with remanence improvements, provided permeability stays constant and coercivity scales accordingly. Dimensions are shown normalised by a constant length, critical to the application, and explicitly defined. The norm of field intensity gradients is normalised by B_r and multiplied by the characteristic length, to remove scale/material effects. Magnetic field inhomogeneity was defined as the relative standard deviation of the magnitude of the magnetic field, in a specific volume, and is thus shown in adimensional parts-per-million (ppm).

The simulations presented were performed in an FEM solver (COMSOL Multiphysics, COMSOL AB) by solving for the scalar magnetic potential, related to the magnetic field intensity as $H = -\nabla V_m$, with $\mu_0 \nabla \cdot (H + M) = 0$. No placement, magnetisation or angular errors were considered due to their demonstrably reducible impact. A linear constitutive relation was used, as explained in Section 5.1.2, and the validity of the model was checked after each simulation. Routines on the convergence of the mesh and numerical error were used to guarantee an error below 1% of all numerical values shown.

NdFeB magnets, due to their high coercive strength, are often modelled [71] with a constant marginal isotropic permeability ($\mu_r = 1.05$) and a constant remanent field for $M = (\mu_r - 1)H + \mu_0^{-1}B_r$. This is clear when observing simulated Halbach homogeneity profiles with radially-repeating patterns and is a crude approximation of the real non-linear anisotropic response. For this reason, and due to the difficulty in obtaining a complete model for NdFeB magnets, the values used were those reported in [111] due to their completeness: $B_r = 1.15$ T, $H_{c0^\circ} = 3.0$ T, $H_{c45^\circ} = 3.0$ T, $H_{c90^\circ} = 5.6$ T, $\mu_{\parallel} = 1.03$ and $\mu_{\perp} = 1.12$. These remanence and permeability values along with the directional coercivity were taken as a phenomenological model, shown in Figure 5.5D and are coherent with the results shown in [113]. Experimentally, one can nonetheless employ existing commercial grades of neodymium with better performance characteristics, for improved results.

Some of the presented applications require researching the self-alignment of arrays and thus the discovery of the equilibrium position. This was achieved in each case with a magnetostatic energy minimisation subroutine, which finds the local/global equilibrium position. This configuration is then used to compute the behaviour of the field, in each application. In Section 5.1.4, computing the non-corrected discrete Halbach array, a simple routine was used to iterate the angle α . The corrected alignment was then entered directly for comparison, but a check that a correction piece would be possible was performed beforehand.

When extending the technique to a continuous magnetisation distribution, this energy minimisation routine was performed on a scalar field, the angle of the magnetisation distribution on the available volume. The initial values were the Halbach distribution, as only a small perturbation is expected, allowing for a fast convergence. To guarantee this would be the global energy minimum, which is a non-trivial expectation because the magnet becomes much larger than the anchor, symmetry boundary conditions were removed, and the initial condition was set to be the one achieved after uniform magnetisation with a strong external field, as would be the case for a fabrication setting. The algorithm, despite not representing a physical evolution, returned a near-Halbach configuration even for the largest magnet, indicating that it would be the end configuration.

In Section 5.1.5, due to the axial sectional symmetry of the design, only one of the 20 sectors was simulated, thereby assuming the simultaneous and identical rotation of all magnets. An energy minimisation routine was used to find the equilibrium rotation angle of the cylinders when specifying the control-array at each control position.

In Section 5.1.6, density-based topology optimisation was employed through the maximisation of the objective function f_{OBJ} using the solid isotropic material with penalisation (SIMP) method [114].

The functional increased the magnetic field norm at height 1, while reducing the non-binary state of the density with the term $\rho(1 - \rho)$, ensuring a sharp transition using the term $|\nabla\rho|$, and guaranteeing the magnetic field remains within the constraints of the linear model through the 2D Heaviside function $\Theta(B - \mu_0 Hc)$, to avoid demagnetisation:

$$f_{\text{OBJ}} = |B| - k_1 |\nabla\rho| - k_2 \rho(1 - \rho) - k_3 \Theta(B - \mu_0 Hc) \rho \quad (5.8)$$

The density was linear in the remanent magnetisation and had an initial value of 1. Several steps were used, with various size-dependent k_i and mesh density, using both the globally convergent method of moving asymptotes (GCMMA) [115] and sparse nonlinear optimizer (SNOPT) [116] algorithms, to allow for a smooth convergence of the density.

5.1.4 Self-aligned magnets

An impactful application of the reciprocity principle is the automatic assembly targeting micro-scale applications, which, due to the small dimensions and the large relative forces involved, make assembly especially challenging. To best showcase the possible approaches, a micro-array and a powder magnet were simulated in their self-alignment and corrected alignment, attained through application-dependent correction pieces.

5 Permanent magnet design

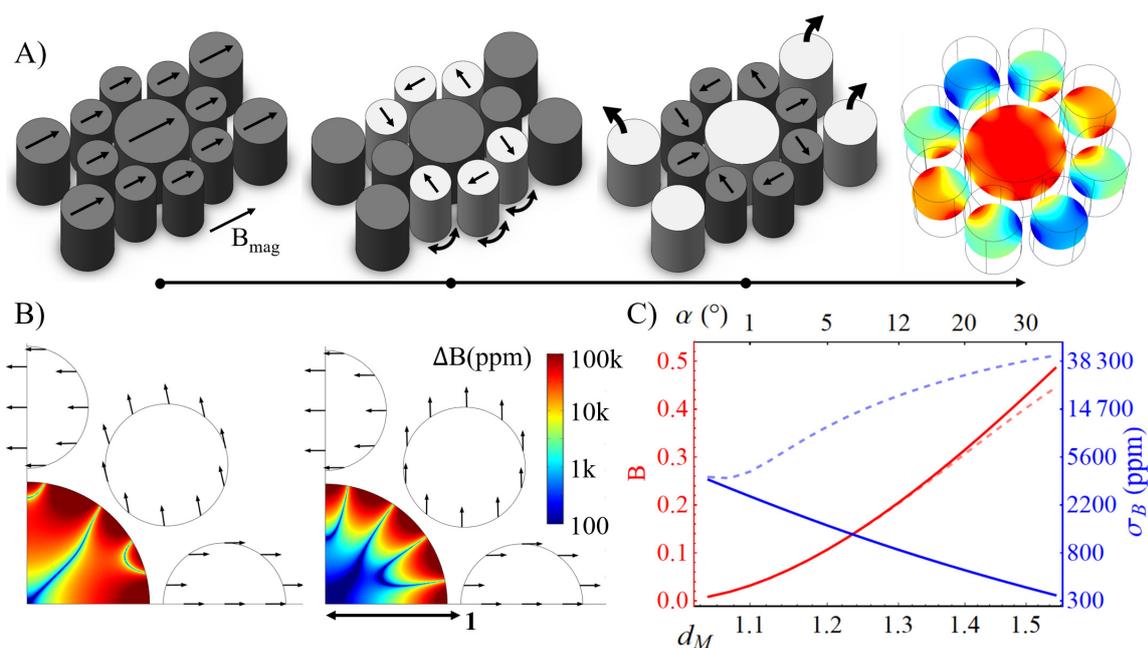


Figure 5.2: A) shows the alignment process: structures are deposited and magnetised uniformly, non-aligned structures are allowed to rotate and then fixed, correction structures are lifted off, and the resulting field is final. Structures active in a step are shown brighter, and the magnetisation direction is indicated as an arrow. B) and C) show the results in field intensity, homogeneity inside the half inner radius, and tilt angle, α , for the self-aligned (dashed lines) and corrected alignments (full lines).

A Halbach micro-array was initially conceived to be deposited as a single material layer on a substrate, uniformly magnetised, and allowed to self-align. This enables the use of a low fringe-field design, with a strong field intensity, in integrated applications, while using any desired magnetic material. A quarter of the array was simulated as a 2D model with sizes ranging up to the maximum packing of 8 circles in the available space Ω , between the normalised inner radius and d_M .

The results in Figure 5.2 clearly show the impact of outer radius d_M on field strength and homogeneity for one 8-magnet ring, which is further exacerbated as several magnet rings are employed. An improvement of two orders of magnitude is seen on the homogeneity, as first order field gradients cancel out in the correct alignment, thus showing that correction is a critical step. From Figure 5.2A, it is clear that the procedure uses only indirect handling of the magnets and requires only sequential liberation and fixation of magnets. The resulting ease of fabrication and the scale invariance of permanent magnets make these ideal for down-scaling magnetic arrays to a length-scale far below what is currently achievable for complex arrays.

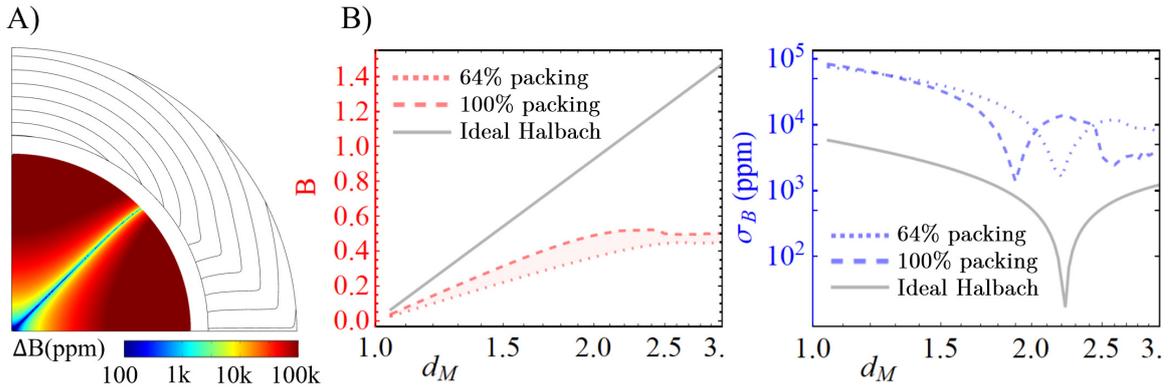


Figure 5.3: A) shows the powder implementation of the structures in Figure 5.2, in their self-aligned, minimum-energy condition. B) shows the dependence of field intensity and homogeneity, inside the inner half radius, on an increasing magnet size.

A natural follow-up to the results shown is the evolution from a finite number of well-defined magnetic structures to arbitrarily many, much like in a powder magnet. These allow for easy tooling, the drastic reduction of statistical variance errors as seen in Section 5.1.2 (by shifting these to fabrication precision), and large height micro-magnets, overcoming sputtering height limitations. A powder can reach a maximal theoretical sphere packing density of 74% in an FCC/HCP lattice but, without correct placement, will reach a random packing after vibration annealing, of 64% [117]. A mix of powders of staggered sizes can however reach arbitrarily close to a 100% packing factor, as smaller powders fill the gaps left by larger pieces, or even emulate a single phase with composite effects. Using two powder species with different temperature coefficients of remanence, α , one positive and one negative, it is possible to locally cancel temperature-induced drifts in field strength or even offset volume drifts by tuning the expansion of the binding agent. As an example, a thermally-compensated mix of NdFeB ($B_R = 1.45T$, $\alpha \approx -0.12\%/K$) and SmEr [118] ($B_R = 0.89T$, $\alpha \approx +0.11\%/K$), for example, shows an effective packing density of 75% when compared to NdFeB alone.

Figure 5.3 shows that the auto-alignment of a powder structure is possible and able to provide a significant field intensity, while leveraging all of the advantages mentioned above. Field intensity is seen to level-off as the self-interaction of the powder dominates over the alignment field of the anchor, which indicates a need for the use of correction structures. This is further emphasised by the large inhomogeneity gap between the self-aligned and the Halbach conditions.

5 Permanent magnet design

Due to the complexity of the endeavour and its application-specific nature, further investigation of this effect was deemed to be outside of the scope of the present manuscript. As a baseline, the ideal Halbach alignment was also simulated, which explicitly shows the critical effect of disregarding the susceptibility of hard magnetic materials. Whereas a full theoretical Halbach has a zero inhomogeneity, the real response of the magnetic material to the demagnetising field significantly alters this expectation to the behaviour shown in Figure 5.3B.

5.1.5 NMR profilometry

One of the advantages of an auto-assembled magnet is the ease at which it can be put together, usually a secondary concern as homogeneous MR magnets are only assembled once and then used in a static assembly. Similarly, MR profiling magnets are usually *set* for a specific working condition and thus field strength, discernible slice thickness, and penetration depth become magnet parameters that cannot be subsequently changed. However, with an auto-adjusting application one could dynamically tune the operational conditions as needed, making use of the coupling force introduced by a tuning piece.

By solving the inverse problem, the reciprocity principle readily facilitates application development and the interpretation of results. Setting out to generate an adjustable profiling magnet with a large penetration depth, an anchor was set at different heights above a plane below which magnets could be placed. The variation of the field direction on the design region, as the anchor is moved, reciprocally indicates the variation the magnets need to have to adjust to a target field at different penetration depths, as illustrated in Figure 5.4A with rotating cylinders. This rotation can then be implemented mechanically or with a magnetic control piece. Once more leveraging the principle, the contribution of each point in space to the field on the anchor space and on the cylinder space can be computed. With these, one can optimally find the region in space which strongly interacts with the cylindrical array, inducing torque, but has a minimum contribution to the target field, reducing any adverse effects. These pieces were implemented as cubes, for simplicity, as seen in Figure 5.4A.

The application requires a sectional symmetry to generate the nearly axially symmetric field which allows for a minimal discernible slice. As the alignment of all magnets at the same rotation angle is a meta-stable configuration of the cylinder-array, a scaffold is needed to maintain the relative angular positions. A stable configuration of the coherent cylinder-array is a parallel orientation pointing upwards/downwards and, due to the desired high packing density, the strong interaction between the cylinder-array elements makes it hard to shift the configuration away from its minimum energy alignment. For this reason, a set of control pieces with identical upward magnetisation is placed between the polarising magnets, creating a local energy minimum somewhat shallower than the absolute minimum, which allows for adjusting the strongly aligned magnet array with smaller pieces. In its final assembly, the freely-rotating array of cylinders is tunable by magnetic control pieces moving linearly, and their total field can now be simulated and analysed.

The figures of merit of the profiling magnet assembly are its field strength, its penetration depth, and the discernible slice thickness. The minimum thickness comes from the condition $T\gamma\nabla_z B > \Delta\omega_{\text{slice}}$ and was implemented, for the field in a cylindrical slice of radius R , as:

$$T_{\min}(R, z) = \frac{\text{Max}(\Delta|B|)|_{r<R}}{\text{Min}(\nabla_z|B|)|_{r<R}} \quad (5.9)$$

Due to the scale invariance of the field profile, all dimensions were normalised to the outer diameter of the available design volume (i.e. the magnet) and the geometry obtained was as follows:

From the results in Figure 5.4 and Table 5.1, the achieved results, beyond their novel continuous-tuning ability, outperform other reported designs/simulations in penetration depth, while having similar values of discernible thickness and field intensity/gradient, when compensating for the unavoidable distance decay. Further optimisation on the packing density was not attempted, as limits are fabrication dependent, and the centre was left unpopulated to allow for RF coils, meaning the performance could be increased further significantly.

Table 5.1: Normalised figures of merit, as described in Sect. 5.1.3, comparing the obtained results for the normalised penetration depth (d), magnetic field (B), normalised field gradient and minimum discernible thickness (T) to literature.

	d (%)	B	$\nabla_z B $	T (10^{-3})
Figure 5.4	25-68	0.008-0.07	0.2-0.3	<1
[100]	25	0.18	0.36	0.7
[101]	19	0.36	0.26	0.3
[102]	3	0.22	0.94	0.3
[103]	7	0.07	0.02	1.8

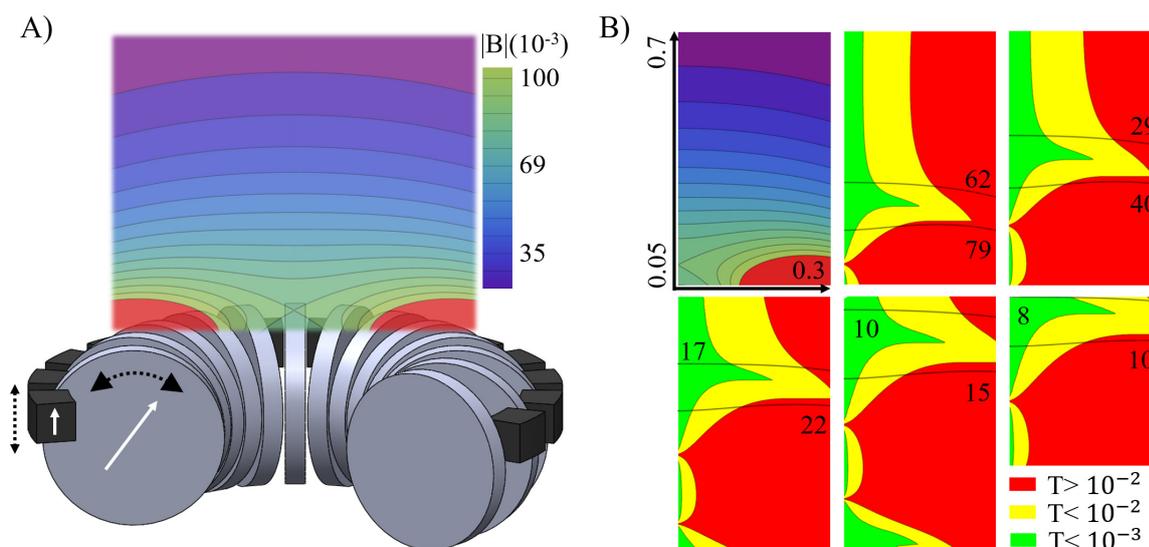


Figure 5.4: A) shows a section of the developed geometry, with the cylinder-array in grey, the control pieces in black and the magnetisation directions as white arrows. The resulting magnetic field intensity contours are shown at the minimum energy configuration and feature a uniform vertical gradient region above the magnet array. In B) the top left and centre plots show the magnetic field intensity and minimum discernible thickness for the configuration in A) as defined by Equation 5.9. The remaining figures show the profiles arising from actuation of the control pieces in $\{-0.06, 0.08\}$, rotating the cylinders by $\{-32^\circ, 27^\circ\}$. These are plotted on normalised dimensions going from 0.05 to 0.7 in relative penetration depth and 0 to 0.3 radially. The normalised magnetic field intensity, in normalised milliuunits, and its contour lines, are shown in black.

5.1.6 High-field designs

Beyond the applications shown, which emphasise fabrication and readjustment, the principle shown here is especially well suited for the design of strong magnets, as it near-optimally removes the degrees of freedom associated with angular alignment and mass-optimal magnet placement, allowing for development through shape/topology optimisation techniques.

As the field of spectroscopic NMR evolves, with experiments now being done in volumes below $(5 \text{ nm})^3$ using NV centre spectroscopy [119], superconducting magnets remain the most expensive and obtrusive element of the experimental setup. As an alternative, the author set out to create the strongest possible field using permanent magnets in a geometry that allows for easy access, such as a magnet integrated into a laboratory table, and with axial symmetry, to allow for easier fabrication and assembly. Such a setup would enable a passive, low-cost magnet with an encapsulation allowing for temperature control and multidirectional access to optical instrumentation or sample feed lines.

A critical limitation of projected/single-sided magnetic fields is that of strong gradients, as the field decays away from the magnet. Approaches based on creating a saddle point were researched and compared to achieved solutions and found to have a limiting performance. Given the targeted application of small-scale experiments, the use of electromagnetic coils to correct for the axial gradient are far better suited considering the l^{-2} scaling of gradient intensity per unit power into reduced scales.

As a starting point, magnetostatic reciprocity presets the optimal magnetisation direction on the half-space, a dipole field ([39]) centred at a normalised height of 1 above the plane and defined by its polar angle θ :

$$\vec{B}_d = \frac{\mu_0}{4\pi} \frac{|\vec{m}_z|}{(\rho^2 + (z-1)^2)^{3/2}} [1.5\sin(2\theta)\hat{e}_\rho + (3\cos^2(\theta) - 1)\hat{e}_z] \quad (5.10)$$

The contour lines for the intensity of the field determined the optimal placement of the magnets to be a half-ellipsoid with a smallest axis/radius of R_M . This result is reminiscent of a Halbach configuration which is known to have a logarithmic dependence on the outer radius. However, these must maintain a fully packed geometry to achieve homogeneity and thus quickly generate a demagnetised area with larger outer radii ([41]), which limits field intensity. To overcome this, a routine was established to optimally remove voxels which would become demagnetised. As the removal of a volume sharply impacts the demagnetising field on the nearest neighbours, an iterative algorithm was necessary and thus density-based topology optimisation techniques were used, obtaining the results in Fig. 5.5.

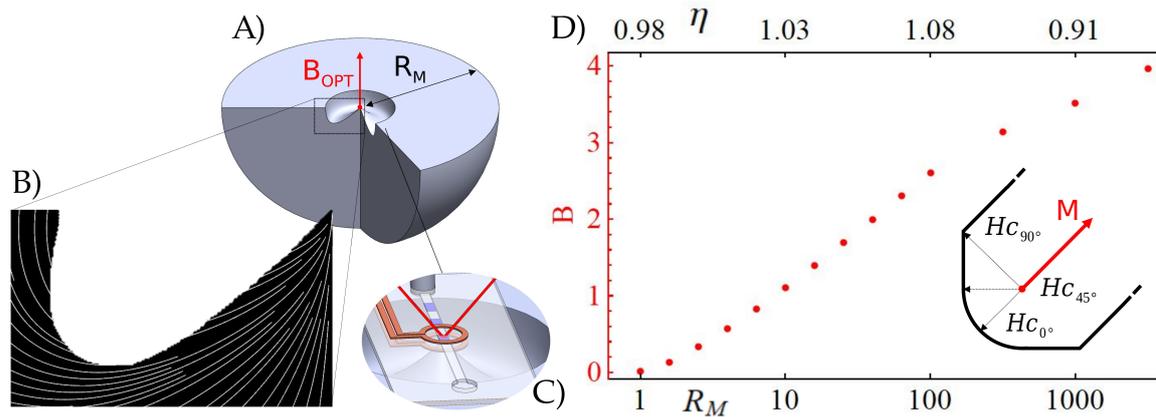


Figure 5.5: A) shows a section of the geometry (not drawn to scale) of a table-magnet that creates an optimal field at the tip of its on-axis protrusion, indicated by a red dot. B) Result of the topology optimisation for a magnet with diameter $R_M = 10^3$, with the magnetisation direction shown as white streamlines. C) An example application for microfluidic NMR. D) shows the dependence of the field intensity on the normalised magnet size, R_M . The inset right shows the phenomenological demagnetisation model for the magnetisation M used during optimisation.

The field intensity is seen to quickly attain a Halbach-like logarithmic scaling as the magnet becomes large ($R_M > 2$) compared to the normalised height, and maintains the logarithmic growth, albeit less pronounced, as demagnetised regions start being removed ($R_M > 100$). The coupling parameter explained through Eq. 5.3 is especially well-suited to quantitatively evaluate designs focused on intense magnetic fields, as intended here.

Before the necessary removal of magnetic material, the coupling parameter is close to unity after which they can be seen to subsequently diminish as volumes must be removed to avoid demagnetisation, showing the decreasing mass-efficiency of further increasing the amount of magnetic material. Values of η above one, while generally unexpected, are caused by the demagnetising field direction matching the alignment of the magnetisation at the centre/tip of the design, constructively magnetising that region. These results open a new avenue towards low-cost, high-performance spectroscopy and can easily be generalised, due to their scale independence, to any field requiring strong, localised magnetic fields, such as those needed for magnetic resonance force microscopy.

5.1.7 Conclusion

The results presented in this chapter offer several multidimensional insights into the development of permanent magnet systems for MR. By establishing a quantifiable metric for what constitutes a good magnet design, a clear beneficial dependence to the use of multiple magnets as well as a way to auto-align them, along with some exemplary applications of these methods, one hopes to facilitate the future development of high performance designs and the integration of advanced techniques, either numerical or experimental, into a now further constrained problem.

As discussed, the reciprocity principle clearly establishes a method for developing an intensity-optimal magnet, which should be thought of as the ultimate goal of the MR magnet designer, as the field's homogeneity can then be targeted through a wealth of other techniques. Despite its broad applicability and usefulness, namely in unintuitive applications, the principle carries an intrinsic limitation by relying on the superposition of linear responses, which breaks down when strong non-linear effects are present. Unfortunately, these are central to state of the art solutions, which have been able to achieve results beyond those possible solely with permanent magnets and up to 4 T [120], by leveraging the saturation magnetisation of soft-magnetic materials of up to 2.8 T [121]. On the other hand, this breakdown of linearity poses severe challenges in the development of shimming systems and design complexity trade-offs must thus be made.

The results shown further emphasise the well-known need for a large filling-factor of the magnet, similarly to what is attempted with resonators, as this has a large effect on the overall magnet's coupling parameter and thus volume/cost. This approach suggests a removal of in-bore passive shimming systems, which could be placed outside, benefiting from the negligible fringe field of Halbach magnets, for example, or even aid with the reduction of demagnetisation fields. As sample sizes become smaller, the use of active shimming becomes more favourable as the power efficiency of shimming coils increases and the absolute power dissipation decreases, with an approximate dependence of l^{-1} [122]. If the magnet size remains constant and the sample/shims are downscaled, a linear shimming profile generating the same field will span a larger gradient in the smaller enclosed volume, which introduces a further geometrical scaling factor of l^{-1} . At the same time, field inhomogeneity does not constitute a fundamental problem, as it only influences/removes the net measured signal but not the local contributions, which means it might be ultimately solved through non-magnetostatic methods. A successful example is that of "shimming" RF pulses [123], which periodically or continually compensate for this local deviation in phase and can therefore offer a versatile, RF-based solution to this limitation without the volume of passive solutions.

Simultaneously, another reciprocal concept was introduced in an attempt to reduce sources of inhomogeneity. While self-aligning arrays had long been an interesting phenomenon, they can now be brought into MR by leveraging their intrinsic advantages under the light of the analytical and numerical analysis presented. Despite the limitations of the technique, as magnetic self-interactions will always be present, a large part of the magnetostatic energy and thus the forces involved in the construction of a magnet assembly were shown to be self-correcting, solely through the use of an anchoring magnetic part. Additionally, the ability to auto-align a magnetic array appears to be a promising pathway towards breaking free from using a small number of magnets, limited by manufacturing complexity. Low numbers ultimately limit the overall homogeneity of designs, as shown by its scaling with magnet number, and by increasing them one can significantly limit manufacturing variations leaving performance susceptible only to assembly errors.

6 | Conclusion

Having reached the end of this manuscript with an in-depth view of the ideas explored, one can now review how they, together, fit into the purpose of the research work in this manuscript; advancing new component/system-level *concepts in low-cost and flow NMR*.

A first category of concepts explored were those targeting the performance per unit cost or the complexity of NMR flow/spectrometers, aiming to lower the bar for academic/commercial exploitation and therefore ease the dissemination of this important non-invasive tool. Roughly following the order of their introduction in this manuscript, the author targeted:

- *Fully commercial flow/spectrometers*: This initial target biased and drove all others, as achieving a low cost is key when considering the success of any NMR-based flowmeter. From the measured sensitivity of the system's stripline and its magnetic field intensity, one can now expect a water LOD of $\approx 2.2 \text{ nL Hz}^{-0.5}$ from low-field, compact spectrometer implementations, while using nothing but commercial parts. All the subcomponents of the proposed design can be purchased directly from assembled-PCB/SDR manufacturers and commercial workshops, and then assembled using nothing but a screwdriver, a wrench and a soldering iron.
- *Tunable electronics interfacing commercial SDRs*: While the raw sensitivity mentioned above is the core metric of any spectrometer, one cannot achieve this upper performance bound if the electronics cannot allow for a robustness to real world conditions, creating the need for in-built EMI shielding, on-line adjustments and high-quality amplification/read-out. With this motivation, a resonator was developed, showcasing an intrinsic 27 dB of EMI shielding, embedded electronics with an acceptably noisy amplification (i.e. 28 % of SNR loss), electrical frequency tuning and 68dB of amplification, which enable signal acquisition with any out-of-the-box commercial SDR board.

6 Conclusion

- *Flow volume selection and net-phase encoding:* Having shown a spectrometer concept that can measure spins at a high sensitivity, the author could further direct its sensitivity onto a specific volume inside a flowing chamber, where the flow was be shaped so as to be encoded with a near-uniform phase at each lamina. This opened the door to net-phase flow-encoding methods, from which a flow SNR (SNR_Q) of 1 to $5 \times 10^7 \text{ Hz}^{-0.5}$ ($SNR_Q/SNR_0 \approx 100$ to 500, without and with phase-wrapping, respectively) can be expected, when considering the LOD and 28% SNR loss mentioned above.
- *Universal, roll-up resonators:* While designing the implementation of the concepts introduced, the use of bore-based flowmetry systems was evaluated and found to be sub-optimal, however, a new SNR-improved solution to the implementation of cylindrical resonators could be developed. Employing a commercial fabrication and simple assembly, the new dual-helix parameterisation can now be repurposed to *any* magnet design, as it can be geometrically scaled, generates a universally-compatible transverse resonator and employs on-coil electrical tuning pads, for an adjustment to varying operating frequencies.
- *Magnet efficiency:* When attempting to optimise for the magnetic field in the prototype shown, it became clear that while one has a clear optimisation functional for the magnetic field, no metric to compare designs with vastly different magnet dimensions or weights had been proposed. In an attempt to normalise by the effect of increasing magnet cost/volume/weight, the author researched the corollaries of magnetostatic reciprocity and thus the *coupling parameter* between a design space and a target region, η , was introduced, offering a tool to quantifiably compare any two designs.
- *Self-alignment of magnets:* The above mentioned topic led directly to energetic corollaries and therefore an approximate energy reciprocity between an anchoring/assembly/virtual magnet and a structure being developed. Taking this, one could propose a simple self-assembly procedure, based on an idea in literature, whose extension into powder-based magnets or micro-manufacturing solutions was simulated. This could allow one to break the current field limit in micro-magnets that is the saturation/remanence field created in a gap, by using a more efficient Halbach array, or even implement self-assembled powder structures with magnetic or mechanical composite effects.

A second category of explored concepts were those whose natural implementation is better suited for high-performing applications rather than low-cost or compact approaches, despite those being the original motivation. As their implementation in higher performance systems drives further research, performance optimisation and reduced complexity, one can expect most of them to eventually be repurposed into low-end applications.

- *Ultra-high decoupling resonator pairs:* Motivated by a need for both high sensitivity reception and high homogeneity excitation solutions, two-coil approaches were researched. While allowing for the independent optimisation of both coils, a naturally arising concern is that of their decoupling. To solve this, the concept of reception fields that are fully-enclosed in (or anti-symmetrical to) the transmission coil was expanded on. These can easily be achieved with a spiral coil in a butterfly coil or a stripline in a solenoid, for example, with the latter being measured to achieve an unprecedented, purely-geometrical decoupling of 73 dB.
- *Differential NMR operation:* A feature of the magnetisation profiles generated by the pairs introduced above is their ability to generate a differential read-out of two samples, regardless of the pulse sequence employed. This purely geometrical effect was shown in two samples measured (i.e. imaged) to have either a 0° or a 180° phase relative to each other, thus constituting a common-mode or a differential-mode measurement, respectively. This feature is expected to be of significant importance in high-end systems where it can deliver both improvements in SNR and drastic simplifications when suppressing undesired signal backgrounds.
- *Electrically adaptable resonators:* Departing from high-frequency, low-winding resonators onto multi-dimensional applications, the author researched the ability to adaptively connect multiple windings within a resonator to allow for electrically-optimal responses and complex abilities. By connecting tens of windings all in parallel, all in series, in all combinations thereof and even splitting the resonator into two independently-acting sections, it could be shown that it is possible to change its SRF by $\approx 30x$, its optimal-Q frequency by $7x$ and its maximum-Q by over $2x$, all in a single design. The results derived from this concept should both influence multiple design decisions in future resonator implementations, as well as allow for the development of more electrically complex/varying resonators, all while avoiding new redesign or fabrication efforts.

6 Conclusion

- *High-field permanent magnets:* A final concept stemming from magnetostatic reciprocity was that of applications allowing for a formal solution to the inverse magnetostatic problem. As an example, the author started off by generating a magnetisation distribution in half-space that optimally generated a strong magnetic field in a point. With this optimal constraint, one could now avoid the natural upper bound of near-Halbach designs by removing magnetic material that would become demagnetised, by using a topology optimisation procedure. In an FEM simulation, it was shown that one could, in principle, generate an unbound magnetic field intensity at a point using standard, non-ideal permanent magnets, opening the door to low-cost, ultra-high fields in small-scale MR.

In the light of the the results obtained here, it is clear that low-cost, compact MR is a field ripe with potential for improvement and immediate use in various metrology applications, namely flowmetry. Its usual technical complexity can now be more easily bypassed with the use of smart system implementations and cheap, off-the-shelf solutions and one can therefore expect a blossoming of the field into both commercial applications and novel, more demanding research problems.

Bibliography

- [1] R. Mahadevia, A. J. Barker, S. Schnell, P. Entezari, P. Kansal, P. W. Fedak, S. C. Malaisrie, P. McCarthy, J. Collins, J. Carr, M. Markl, Bicuspid aortic cusp fusion morphology alters aortic three-dimensional outflow patterns, wall shear stress, and expression of aortopathy, *Circulation* 129 (6) (2014) 673–682. doi:10.1161/CIRCULATIONAHA.113.003026.
- [2] N. Spengler, J. Höfflin, A. Moazenzadeh, D. Mager, N. MacKinnon, V. Badilita, U. Wallrabe, J. Korvink, Heteronuclear micro-helmholtz coil facilitates μm -range spatial and sub-Hz spectral resolution NMR of nl-volume samples on customisable microfluidic chips, *PLoS One* 11 (5) (2016). doi:10.1371/journal.pone.0146384.
- [3] S. Dehkharghani, C. Fleischer, D. Qiu, M. Yepes, F. Tong, Cerebral temperature dysregulation: MR thermographic monitoring in a nonhuman primate study of acute ischemic stroke, *American Journal of Neuroradiology* 38 (4) (2017) 712–720. doi:10.3174/ajnr.A5059.
- [4] E. O. Fridjonsson, P. L. Stanwix, M. L. Johns, Earth's field NMR flow meter: Preliminary quantitative measurements, *Journal of magnetic resonance* 245 (2014) 110–115. doi:10.1016/j.jmr.2014.06.004.
- [5] T. M. Osan, J. M. Olle, M. Carpinella, L. M. C. Cerioni, D. J. Pusiol, M. Appel, J. Freeman, I. Espejo, Fast measurements of average flow velocity by low-field ^1H NMR, *Journal of magnetic resonance* 209 (2) (2011) 116–122. doi:10.1016/j.jmr.2010.07.011.
- [6] S. Richard, B. Newling, Measuring flow using a permanent magnet with a large constant gradient, *Applied Magnetic Resonance* 50 (2019) 5. doi:10.1007/s00723-018-1107-x.

Bibliography

- [7] E. Kirtil, S. Cikrikci, M. J. McCarthy, M. H. Oztop, Recent advances in time domain NMR and MRI sensors and their food applications, *Current Opinion in Food Science* 17 (2017) 9 – 15, *food Engineering and Processing*. doi:10.1016/j.cofs.2017.07.005.
- [8] F. Presciutti, J. Perlo, F. Casanova, S. Glöggler, C. Miliani, B. Blümich, B. G. Brunetti, A. Sgamellotti, Noninvasive nuclear magnetic resonance profiling of painting layers, *Applied Physics Letters* 93 (3) (2008) 033–505. doi:10.1063/1.2963026.
- [9] J. P. Marques, F. F. Simonis, A. G. Webb, Low-field MRI: An MR physics perspective, *Journal of Magnetic Resonance Imaging* 49 (6) (2019) 1528–1542. doi:10.1002/jmri.26637.
- [10] C. W. Windt, H. Soltner, D. van Dusschoten, P. Blümmler, A portable Halbach magnet that can be opened and closed without force: The NMR-CUFF, *Journal of magnetic resonance* 208 (2011). doi:10.1016/j.jmr.2010.09.020.
- [11] D. Ha, J. Paulsen, N. Sun, Y.-Q. Song, D. Ham, Scalable NMR spectroscopy with semiconductor chips, *Proceedings of the National Academy of Sciences* 111 (33) (2014) 11955–11960. doi:10.1073/pnas.1402015111.
- [12] C. H. Moon, J.-H. Kim, T. Zhao, K. T. Bae, Quantitative ^{23}Na MRI of human knee cartilage using dual-tuned $^1\text{H}/^{23}\text{Na}$ transceiver array radiofrequency coil at 7 Tesla, *Journal of Magnetic Resonance Imaging* 38 (5) (2013) 1063–1072. doi:10.1002/jmri.24030.
- [13] limemicro.com, [LimeSDR software-defined radio board](https://limemicro.com/products/boards/limesdr/), accessed Jul 2021.
URL limemicro.com/products/boards/limesdr/
- [14] R. C. Baker, *Flow Measurement Handbook: Industrial Designs, Operating Principles, Performance, and Applications*, 2nd Edition, Cambridge University Press, Cambridge, England, 2016. doi:10.1017/CBO9781107054141.
- [15] A. Bilgic, J. Kunze, V. Stegemann, Multiphase flow metering with nuclear magnetic resonance, *Technisches Messen* 82 (2015) 11. doi:doi:10.1515/teme-2015-0082.
- [16] C. Li, P. Wu, Z. Wu, K. Limnusun, R. Narayan, Highly accurate thermal flow microsensors for continuous and quantitative measurement of cerebral blood flow, *Biomed Microdevices* 17 (2015). doi:10.1007/s10544-015-9992-3.

- [17] C. Hoera, M. M. Skadell, S. A. Pfeiffer, M. Pahl, Z. Shu, E. Beckert, D. Belder, A chip-integrated highly variable thermal flow rate sensor, *Sensors and Actuators B: Chemical* 225 (2016) 42 – 49. doi:10.1016/j.snb.2015.11.009.
- [18] R. Vilares, C. Hunter, I. Ugarte, I. Aranburu, J. Berganzo, J. Elizalde, L. Fernandez, Fabrication and testing of a SU-8 thermal flow sensor, *Sensors and Actuators B: Chemical* 147 (2) (2010) 411 – 417. doi:doi.org/10.1016/j.snb.2010.03.054.
- [19] F. Del Cesta, P. Bruschi, M. Piotta, An integrated thermal sensor platform for multi-variable detection, in: 2015 XVIII AISEM Annual Conference, Trento, Italy, 2015, pp. 1–4.
- [20] M. Dijkstra, T. Lammerink, O. Pjetri, M. de Boer, J. Berenschot, R. Wiegerink, M. Elwenspoek, Thermal-wave balancing flow sensor with low-drift power feedback, *Journal of micromechanics and microengineering* 24 (2014) 5. doi:10.1088/0960-1317/24/5/055016.
- [21] W. Sparreboom, J. Van de Geest, M. Katerberg, F. Postma, J. Haneveld, J. Groenesteijn, T. Lammerink, R. Wiegerink, J. Lötters, Compact mass flow meter based on a micro coriolis flow sensor, *Micromachines* 4 (1) (2013) 22–33. doi:10.3390/mi4010022.
- [22] C. Shen, X. Lian, V. Kavungal, C. Zhong, D. Liu, Y. Semenova, G. Farrell, J. Albert, J. F. Donegan, Optical spectral sweep comb liquid flow rate sensor, *Opt. Lett.* 43 (4) (2018) 751–754. doi:10.1364/OL.43.000751.
- [23] C. Hoera, M. M. Skadell, S. A. Pfeiffer, M. Pahl, Z. Shu, E. Beckert, D. Belder, A chip-integrated highly variable thermal flow rate sensor, *Sensors and Actuators B: Chemical* 225 (2016) 42–49. doi:10.1016/j.snb.2015.11.009.
- [24] P. Silva, M. Jouda, J. Korvink, Geometrically-differential NMR in a stripline front-end, *Journal of Magnetic Resonance* 310 (2019) 106–659. doi:10.1016/j.jmr.2019.106659.
- [25] P. Silva, S. Torres, M. Jouda, D. Mager, J. Korvink, A multi-purpose, rolled-up, double-helix resonator, *Journal of Magnetic Resonance* 309 (2019) 106–599. doi:10.1016/j.jmr.2019.106599.
- [26] P. F. Silva, M. Jouda, J. G. Korvink, Magnetostatic reciprocity for MR magnet design, *Magnetic Resonance* 2 (2021) 607–617. doi:10.5194/mr-2-607-2021.

Bibliography

- [27] S. Nimbalkar, E. Fuhrer, P. Silva, T. Nguyen, M. Sereno, S. Kassegne, J. Korvink, Glassy carbon microelectrodes minimize induced voltages, mechanical vibrations, and artifacts in magnetic resonance imaging, *Microsystems and Nanoengineering* 5 (2019) 61. doi:10.1038/s41378-019-0106-x.
- [28] M. Jouda, E. Fuhrer, P. Silva, J. G. Korvink, N. MacKinnon, Automatic adaptive gain for magnetic resonance sensitivity enhancement, *Analytical Chemistry* 91 (2019) 2376–2383. doi:10.1021/ac100969k.
- [29] P. Hore, *Nuclear Magnetic Resonance*, Oxford chemistry primers, Oxford University Press, Oxford, England, 2015.
- [30] A. Abragam, *The Principles of Nuclear Magnetism*, Clarendon Press, Oxford, England, 1961.
- [31] M. H. Levitt, *Spin Dynamics: Basics of Nuclear Magnetic Resonance*, 2nd Edition, Wiley, New Jersey, United States, 2008.
- [32] C. Kittel, *Elementary statistical physics*, Wiley, New Jersey, United States, 1967.
- [33] D. Hoult, R. Richards, The signal-to-noise ratio of the nuclear magnetic resonance experiment, *Journal of Magnetic Resonance* 24 (1) (1976) 71–85. doi:10.1016/0022-2364(76)90233-X.
- [34] D. I. Hoult, The principle of reciprocity in signal strength calculations: A mathematical guide, *Concepts in Magnetic Resonance* 12 (4) (2000) 173–187. doi:10.1002/1099-0534.
- [35] J. Mispelter, M. Lupu, A. Briguet, *NMR probeheads for biophysical and biome: theoretical principles and practical guidelines*, 2nd Edition, Imperial College Press, London, England, 2006.
- [36] B. . Chung, C. Chien, H. Samuelli, R. Jain, Performance analysis of an all-digital BPSK direct-sequence spread-spectrum if receiver architecture, *IEEE Journal on Selected Areas in Communications* 11 (7) (1993) 1096–1107. doi:10.1109/49.233222.
- [37] T. J. Roupael, Chapter 5 - signal sampling and distortion, in: *Wireless Receiver Architectures and Design*, Academic Press, Boston, United States, 2014, pp. 263 – 303. doi:10.1016/B978-0-12-378640-1.00005-6.
- [38] R. R. Ernst, W. A. Anderson, Application of Fourier transform spectroscopy to magnetic resonance, *Review of Scientific Instruments* 37 (1) (1966) 93–102. doi:10.1063/1.1719961.

- [39] K. Seleznyova, M. Strugatsky, J. Kliava, Modelling the magnetic dipole, *European Journal of Physics* 37 (2) (2016) 025–203. doi:10.1088/0143-0807/37/2/025203.
- [40] W. F. Brown, *Magnetostatic principles in ferromagnetism*, North Holland Pub. Co., Amsterdam, Netherlands, 1962.
- [41] A. Insinga, C. Bahl, R. Bjørk, A. Smith, Performance of Halbach magnet arrays with finite coercivity, *Journal of Magnetism and Magnetic Materials* 407 (2016) 369–376. doi:10.1016/j.jmmm.2016.01.076.
- [42] I. K. Kang, H. H. Scholefield, A. P. Martin, Some concepts concerning the new high-permeability 50–50 nickel–iron alloys, *Journal of Applied Physics* 38 (3) (1967) 1178–1183. doi:10.1063/1.1709531.
- [43] E. Danieli, J. Perlo, B. Blümich, F. Casanova, Highly stable and finely tuned magnetic fields generated by permanent magnet assemblies, *Phys. Rev. Lett.* 110 (2013) 180801. doi:10.1103/PhysRevLett.110.180801.
- [44] R. Spencer, The time-domain matched filter and the spectral-domain matched filter in 1-dimensional NMR spectroscopy, *Concepts in magnetic resonance. Part A* 36A (2010) 255–264. doi:10.1002/cmr.a.20162.
- [45] T. E. Conturo, G. D. Smith, Signal-to-noise in phase angle reconstruction: Dynamic range extension using phase reference offsets, *Magnetic Resonance in Medicine* 15 (3) (1990) 420–437. doi:10.1002/mrm.1910150308.
- [46] A. Caprihan, E. Fukushima, Flow measurements by NMR, *Physics Reports* 198 (4) (1990) 195 – 235. doi:10.1016/0370-1573(90)90046-5.
- [47] O. Sigmund, K. Maute, Topology optimization approaches, *Structural and Multi-disciplinary Optimization* 48 (2013) 1031–1055. doi:10.1007/s00158-013-0978-6.
- [48] G. Suryan, Nuclear resonance in flowing liquids, *Proc. Indian Acad. Sci. (Math. Sci.)* 33 (1951) 107 – 111. doi:10.1007/BF03172192.
- [49] L. A. Colnago, F. D. Andrade, A. A. Souza, R. B. V. Azeredo, A. A. Lima, L. M. Cerioni, T. M. Osán, D. J. Pusiol, Why is inline NMR rarely used as industrial sensor? challenges and opportunities, *Chemical Engineering & Technology* 37 (2) (2014) 191–203. doi:10.1002/ceat.201300380.

Bibliography

- [50] E. Fukushima, Nuclear magnetic resonance as a tool to study flow, *Annual Review of Fluid Mechanics* 31 (1) (1999) 95–123. doi:10.1146/annurev.fluid.31.1.95.
- [51] C. J. Elkins, M. Markl, N. Pelc, J. K. Eaton, 4D magnetic resonance velocimetry for mean velocity measurements in complex turbulent flows, *Experiments in Fluids* 34 (4) (2003) 494–503. doi:10.1007/s00348-003-0587-z.
- [52] L. F. Gladden, A. J. Sederman, Recent advances in flow MRI, *Journal of Magnetic Resonance* 229 (2013) 2 – 11. doi:10.1016/j.jmr.2012.11.022.
- [53] S. J. Richard, B. Newling, Measuring flow using a permanent magnet with a large constant gradient, *Applied Magnetic Resonance* 50 (5) (2019) 627–635. doi:10.1007/s00723-018-1107-x.
- [54] T. Osán, J. Ollé, M. Carpinella, L. Cerioni, D. Pusiol, M. Appel, J. Freeman, I. Espejo, Fast measurements of average flow velocity by low-field ^1H nmr, *Journal of Magnetic Resonance* 209 (2) (2011) 116 – 122. doi:doi.org/10.1016/j.jmr.2010.07.011.
- [55] K. T. O’Neill, A. Klotz, P. L. Stanwix, E. O. Fridjonsson, M. L. Johns, Quantitative multiphase flow characterisation using an earth’s field NMR flow meter, *Flow Measurement and Instrumentation* 58 (2017) 104 – 111. doi:10.1016/j.flowmeasinst.2017.10.004.
- [56] E. Steinborn, K. Ruedenberg, Rotation and translation of regular and irregular solid spherical harmonics, in: *Advances in quantum chemistry*, Vol. 7, Elsevier, 1973, pp. 1–81. doi:10.1016/S0065-3276(08)60558-4.
- [57] R. Turner, A target field approach to optimal coil design, *Journal of physics D: Applied physics* 19 (8) (1986) L147. doi:10.1088/0022-3727/19/8/001.
- [58] W. A. Edelstein, F. Schenck, Current streamline method for coil construction, uS Patent 4,840,700 (1989).
- [59] L. K. Forbes, S. Crozier, Novel target-field method for designing shielded biplanar shim and gradient coils, *IEEE transactions on Magnetics* 40 (4) (2004) 1929–1938. doi:10.1109/TMAG.2004.828934.
- [60] R. Dengler, Self inductance of a wire loop as a curve integral, *Advanced Electromagnetics* 5 (2016) 1–8. doi:10.7716/aem.v5i1.331.

- [61] M. M. A. Condesso, [Implementation of biplanar coils for magnetic field generation](#), accessed Jul 2021 (2019).
URL fenix.tecnico.ulisboa.pt/downloadFile/281870113705157/MScThesis.pdf
- [62] I. M. Savukov, S. K. Lee, M. V. Romalis, Optical detection of liquid-state NMR, *Nature* 442 (2006) 1021–1024. doi:10.1038/nature05088.
- [63] R. J. Prance, A. Aydin, Acquisition of a nuclear magnetic resonance signal using an electric field detection technique, *Applied Physics Letters* 91 (4) (2007) 044–103. doi:10.1063/1.2762276.
- [64] I. M. Savukov, S. K. Lee, M. V. Romalis, Nanoscale imaging magnetometry with diamond spins under ambient conditions, *Nature* 455 (2008) 648–651. doi:10.1038/nature07278.
- [65] K. Takeda, K. Nagasaka, A. Noguchi, R. Yamazaki, Y. Nakamura, E. Iwase, J. M. Taylor, K. Usami, Electro-mechano-optical detection of nuclear magnetic resonance, *Optica* 5 (2) (2018) 152–158. doi:10.1364/OPTICA.5.000152.
- [66] N. Sun, T. Yoon, H. Lee, W. Andress, R. Weissleder, D. Ham, Palm NMR and 1-chip NMR, *IEEE Journal of Solid-State Circuits* 46 (1) (2011) 342–352. doi:10.1109/JSSC.2010.2074630.
- [67] A. L. R.E Ader, D. Songco, Utilization of a microprocessor in a pulsed NMR spectrometer, *Journal of Magnetic Resonance* (1969) 29 (1) (1978) 105 – 111. doi:10.1016/0022-2364(78)90035-5.
- [68] K. Takeda, Chapter 7 - Highly Customized NMR Systems Using an Open-Resource, Home-Built Spectrometer, Vol. 74 of *Annual Reports on NMR Spectroscopy*, Academic Press, 2011. doi:10.1016/B978-0-08-097072-1.00007-8.
- [69] C. J. Hasselwander, Z. Cao, W. A. Grissom, GR-MRI: A software package for magnetic resonance imaging using software defined radios, *Journal of Magnetic Resonance* 270 (Supplement C) (2016) 47 – 55. doi:10.1016/j.jmr.2016.06.023.
- [70] C. A. Michal, A low-cost multi-channel software-defined radio-based NMR spectrometer and ultra-affordable digital pulse programmer, *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering* 48B (2018) 3. doi:10.1002/cmr.b.21401.

Bibliography

- [71] K. Halbach, Design of permanent multipole magnets with oriented rare earth cobalt material, *Nuclear Instruments and Methods* 169 (1) (1980) 1 – 10. doi:10.1016/0029-554X(80)90094-4.
- [72] P. van Bentum, J. Janssen, A. Kentgens, J. Bart, J. Gardeniers, Stripline probes for nuclear magnetic resonance, *Journal of Magnetic Resonance* 189 (1) (2007) 104 – 113. doi:10.1016/j.jmr.2007.08.019.
- [73] V. Badilita, R. C. Meier, N. Spengler, U. Wallrabe, M. Utz, J. G. Korvink, Microscale nuclear magnetic resonance: a tool for soft matter research, *Soft Matter* 8 (2012) 10583–10597. doi:10.1039/C2SM26065D.
- [74] D. Hoult, The origins and present status of the radio wave controversy in NMR, *Concepts in Magnetic Resonance Part A* 34A (4) (2009) 193–216. doi:10.1002/cmra.20142.
- [75] M. A. Macnaughtan, A. P. Smith, P. B. Goldsbrough, R. E. Santini, D. Raftery, NMR difference spectroscopy with a dual saddle-coil difference probe, *Anal Bioanal Chem* 378 (2004). doi:10.1007/s00216-003-2374-3.
- [76] V. Krishnan, N. Murali, Radiation damping in modern NMR experiments: Progress and challenges, *Progress in Nuclear Magnetic Resonance Spectroscopy* 68 (2013) 41 – 57. doi:10.1016/j.pnmrs.2012.06.001.
- [77] A. V. Gavrilin, M. D. Bird, S. T. Bole, Y. M. Eyssa, Conceptual design of high transverse field magnets at the NHMFL, *IEEE Transactions on Applied Superconductivity* 12 (1) (2002) 465–469. doi:10.1109/TASC.2002.1018444.
- [78] Y. Sun, G. Maciel, The tilted coil for NMR experiments, *J. Magn. Reson. A* 105 (1993) 145–150. doi:10.1006/jmra.1993.1266.
- [79] E. Jeong, D. Kim, M. Kim, S. Lee, J. Suh, Y. Kwon, A solenoid-like coil producing transverse RF fields for MR imaging, *Journal of Magnetic Resonance* 127 (1997) 73–79. doi:10.1006/jmre.1997.1172.
- [80] J. Alonso, A. Soleilhavoup, A. Wong, A. Guiga, D. Sakellariou, Double helix dipole design applied to magnetic resonance: a novel NMR coil., *Journal of magnetic resonance (San Diego, Calif. : 1997)* 235 (2013) 32–41. doi:10.1016/j.jmr.2013.07.004.

- [81] M. C. Wapler, J. Leupold, I. Dragonu, D. von Elverfeld, M. Zaitsev, U. Wallrabe, Magnetic properties of materials for MR engineering, micro-MR and beyond, *Journal of Magnetic Resonance* 242 (2014) 233–242. doi:10.1016/j.jmr.2014.02.005.
- [82] T. Peck, R. Magin, P. Lauterbur, Design and analysis of microcoils for NMR microscopy, *Journal of magnetic resonance. Series B* 108 (1995) 114–24. doi:10.1006/jmrb.1995.1112.
- [83] S. C. Grant, L. A. Murphy, R. L. Magin, G. Friedman, Analysis of multilayer radio frequency microcoils for nuclear magnetic resonance spectroscopy, *IEEE Transactions on Magnetics* 37 (4) (2001) 2989–2998. doi:10.1109/20.947051.
- [84] J. D. Cockcroft, Skin effect in rectangular conductors at high frequencies, *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 122 (790) (1929). doi:10.1098/rspa.1929.0038.
- [85] Keysight Technologies, Ultra-low impedance measurements using 2-port measurements, literature.cdn.keysight.com/litweb/pdf/5989-5935EN.pdf, online application note; accessed 23 May 2019.
- [86] H. Davoodi, M. Jouda, J. G. Korvink, N. MacKinnon, V. Badilita, Broadband and multi-resonant sensors for NMR, *Progress in Nuclear Magnetic Resonance Spectroscopy* 112-113 (2019) 34 – 54. doi:10.1016/j.pnmrs.2019.05.001.
- [87] M. C. D. Tayler, T. Theis, T. F. Sjolander, J. W. Blanchard, A. Kentner, S. Pustelny, A. Pines, D. Budker, Instrumentation for nuclear magnetic resonance in zero and ultralow magnetic field, *Review of Scientific Instruments* 88 (9) (2017) 091 – 101. doi:10.1063/1.5003347.
- [88] M. Zakotnik, C. Tudor, Commercial-scale recycling of NdFeB-type magnets with grain boundary modification yields products with designer properties that exceed those of starting materials, *Waste Management* 44 (2015) 48–54. doi:10.1016/j.wasman.2015.07.041.
- [89] E. Danieli, J. Perlo, B. Blümich, F. Casanova, Small magnets for portable NMR spectrometers, *Angewandte Chemie International Edition* 49 (24) (2010) 4133–4135.
- [90] K. Chonlathep, T. Sakamoto, K. Sugahara, Y. Kondo, A simple and low-cost permanent magnet system for NMR, *Journal of Magnetic Resonance* 275 (2017) 114–119. doi:10.1016/j.jmr.2016.12.010.

Bibliography

- [91] C. Hugon, F. D'Amico, G. Aubert, D. Sakellariou, Design of arbitrarily homogeneous permanent magnet systems for NMR and MRI: Theory and experimental developments of a simple portable magnet, *Journal of Magnetic Resonance* 205 (1) (2010) 75–85. doi:10.1016/j.jmr.2010.04.003.
- [92] H. Soltner, P. Blümmler, Dipolar Halbach magnet stacks made from identically shaped permanent magnets for magnetic resonance, *Concepts in Magnetic Resonance Part A* 36A (4) (2010) 211–222. doi:10.1002/cmr.a.20165.
- [93] C. W. Windt, H. Soltner, D. van Dusschoten, P. Blümmler, A portable Halbach magnet that can be opened and closed without force: The NMR-CUFF, *Journal of Magnetic Resonance* 208 (1) (2011) 27–33. doi:10.1016/j.jmr.2010.09.020.
- [94] C. Z. Cooley, M. W. Haskell, S. F. Cauley, C. Sappo, C. D. Lapierre, C. G. Ha, J. P. Stockmann, L. L. Wald, Design of sparse Halbach magnet arrays for portable MRI using a genetic algorithm, *IEEE Transactions on Magnetics* 54 (1) (2018). doi:10.1109/TMAG.2017.2751001.
- [95] S. Utsuzawa, E. Fukushima, Unilateral NMR with a barrel magnet, *Journal of Magnetic Resonance* 282 (2017) 104–113. doi:10.1016/j.jmr.2017.07.006.
- [96] J. L. Paulsen, L. S. Bouchard, D. Graziani, B. Blümich, A. Pines, Volume-selective magnetic resonance imaging using an adjustable, single-sided, portable sensor, *Proceedings of the National Academy of Sciences* 105 (52) (2008) 20601–20604. doi:10.1073/pnas.0811222106.
- [97] W.-H. Chang, C.-Y. Chung, J.-H. Chen, D. W. Hwang, C.-H. Hsu, C. Yao, L.-P. Hwang, Simple mobile single-sided NMR apparatus with a relatively homogeneous B_0 distribution, *Magnetic Resonance Imaging* 29 (6) (2011) 869–876. doi:10.1016/j.mri.2011.02.026.
- [98] A. E. Marble, I. V. Mastikhin, B. G. Colpitts, B. J. Balcom, A compact permanent magnet array with a remote homogeneous field, *Journal of Magnetic Resonance* 186 (1) (2007) 100–104. doi:10.1016/j.jmr.2007.01.020.
- [99] B. Manz, A. Coy, R. Dykstra, C. Eccles, M. Hunter, B. Parkinson, P. Callaghan, A mobile one-sided NMR sensor with a homogeneous magnetic field: The NMR-MOLE, *Journal of Magnetic Resonance* 183 (1) (2006) 25–31. doi:10.1016/j.jmr.2006.07.017.

- [100] S. Rahmatallah, Y. Li, H. Seton, I. Mackenzie, J. Gregory, R. Aspden, NMR detection and one-dimensional imaging using the inhomogeneous magnetic field of a portable single-sided magnet, *Journal of Magnetic Resonance* 173 (1) (2005) 23–28. doi:[10.1016/j.jmr.2004.11.014](https://doi.org/10.1016/j.jmr.2004.11.014).
- [101] M. V. Landeghem, E. Danieli, J. Perlo, B. Blümich, F. Casanova, Low-gradient single-sided NMR sensor for one-shot profiling of human skin, *Journal of Magnetic Resonance* 215 (2012) 74–84. doi:[10.1016/j.jmr.2011.12.010](https://doi.org/10.1016/j.jmr.2011.12.010).
- [102] W.-H. Chang, J.-H. Chen, L.-P. Hwang, Single-sided mobile NMR with a Halbach magnet, *Magnetic Resonance Imaging* 24 (8) (2006) 1095–1102. doi:[10.1016/j.mri.2006.04.005](https://doi.org/10.1016/j.mri.2006.04.005).
- [103] A. E. Marble, I. V. Mastikhin, B. G. Colpitts, B. J. Balcom, A constant gradient unilateral magnet for near-surface MRI profiling, *Journal of Magnetic Resonance* 183 (2) (2006) 228–234. doi:[10.1016/j.jmr.2006.08.013](https://doi.org/10.1016/j.jmr.2006.08.013).
- [104] S. Luo, L. Xiao, X. Li, G. Liao, H. Liu, Z. Wang, Z. Sun, W. Liu, Y. Xu, New magnet array design for downhole NMR azimuthal measurement, *Magnetic Resonance Imaging* 56 (2018) 168–173. doi:[10.1016/j.mri.2018.09.027](https://doi.org/10.1016/j.mri.2018.09.027).
- [105] M. C. Tayler, D. Sakellariou, Low-cost, pseudo-Halbach dipole magnets for NMR, *Journal of Magnetic Resonance* 277 (2017) 143–148. doi:[10.1016/j.jmr.2017.03.001](https://doi.org/10.1016/j.jmr.2017.03.001).
- [106] H. Raich, P. Blümli, Design and construction of a dipolar Halbach array with a homogeneous field from identical bar magnets: NMR mandhalas, *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering* 23B (1) (2004) 16–25. doi:[10.1002/cmr.b.20018](https://doi.org/10.1002/cmr.b.20018).
- [107] C. Chandrana, J. Neal, D. Platts, B. Morgan, P. Nath, Automatic alignment of multiple magnets into Halbach cylinders, *Journal of Magnetism and Magnetic Materials* 381 (2015) 396–400. doi:[10.1016/j.jmmm.2015.01.011](https://doi.org/10.1016/j.jmmm.2015.01.011).
- [108] A. R. Insinga, R. Bjørk, A. Smith, C. R. H. Bahl, Globally optimal segmentation of permanent-magnet systems, *Phys. Rev. Applied* 5 (2016) 064014. doi:[10.1103/PhysRevApplied.5.064014](https://doi.org/10.1103/PhysRevApplied.5.064014).
- [109] S. G. Mikhlin, *Variational Methods in Mathematical Physics*, Pergamon Press, Oxford, England, 1964. doi:[10.1002/zamm.19650450439](https://doi.org/10.1002/zamm.19650450439).

Bibliography

- [110] J. Jeans, *The Mathematical Theory of Electricity and Magnetism*, Cambridge U. Press, Cambridge, England, 1908. doi:10.1038/078537a0.
- [111] M. Katter, Angular dependence of the demagnetization stability of sintered Nd-Fe-B magnets, *IEEE Transactions on Magnetics* 41 (10) (2005) 3853–3855. doi:10.1109/TMAG.2005.854872.
- [112] J. Perlo, F. Casanova, B. Blümich, Single-sided sensor for high-resolution NMR spectroscopy, *Journal of Magnetic Resonance* 180 (2) (2006) 274–279. doi:10.1016/j.jmr.2006.03.004.
- [113] G. Martinek, H. Kronmüller, Influence of grain orientation of the coercive field in Fe-Nd-B permanent magnets, *Journal of Magnetism and Magnetic Materials* 86 (2) (1990) 177–183. doi:10.1016/0304-8853(90)90119-B.
- [114] M. P. Bendsoe, N. Kikuchi, Generating optimal topologies in structural design using a homogenization method, *Computer Methods in Applied Mechanics and Engineering* 20 (1988) 197–224. doi:10.1016/0045-7825(88)90086-2.
- [115] K. Svanberg, A class of globally convergent optimization methods based on conservative convex separable approximations, *SIAM Journal of Optimization* 12 (2002) 555–573. doi:10.1137/S1052623499362822.
- [116] P. E. Gill, W. Murray, M. A. Saunders, Snopt: An sqp algorithm for large-scale constrained optimization, *SIAM Journal of Optimization* 12 (4) (2002) 979–1006. doi:10.1137/S0036144504446096.
- [117] H. M. Jaeger, S. R. Nagel, Physics of the granular state, *Science* 255 (5051) (1992) 1523–1531. doi:10.1126/science.255.5051.1523.
- [118] C. H. Chen, W. Gong, M. H. Walmer, S. Liu, G. E. Kuhl, Behavior of some heavy and light rare earth-cobalt magnets at high temperature, *Journal of Applied Physics* 91 (10) (2002) 8483–8485. doi:10.1063/1.1453324.
- [119] T. Staudacher, F. Shi, S. Pezzagna, J. Meijer, J. Du, C. A. Meriles, F. Reinhard, J. Wrachtrup, Nuclear magnetic resonance spectroscopy on a (5-nm)³ sample volume, *Science* 339 (6119) (2013) 561–563. doi:10.1126/science.1231675.
- [120] M. Kumada, T. Fujisawa, Y. Hirao, M. Endo, M. Aoki, T. Kohda, I. Bolshakova, R. Holyaka, Development of 4 tesla permanent magnet, in: *Proceedings of the 2001 Particle Accelerator Conference*, Vol. 5, 2001, pp. 3221–3223. doi:10.1109/PAC.2001.988064.

- [121] M. Mehedi, Y. Jiang, P. K. Suri, D. J. Flannigan, J.-P. Wang, Minnealloy: a new magnetic material with high saturation flux density and low magnetic anisotropy, *Journal of Physics D: Applied Physics* 50 (37) (2017). doi:10.1088/1361-6463/aa8130.
- [122] J. G. Korvink, N. MacKinnon, V. Badilita, M. Jouda, “Small is beautiful” in NMR, *Journal of Magnetic Resonance* 306 (2019) 112–117. doi:10.1016/j.jmr.2019.07.012.
- [123] D. Topgaard, R. W. Martin, D. Sakellariou, C. A. Meriles, A. Pines, “Shim pulses” for NMR spectroscopy and imaging, *Proceedings of the National Academy of Sciences* 101 (51) (2004) 17576–17581. doi:10.1073/pnas.0408296102.

Acknowledgments

Just as you made it past all the sciencey bits, it appears that so have I. For someone who always claimed to have no interest in doing a PhD, having had such a successful and enjoyable journey is still somewhat of a mystery. I can only conclude it to be owed to all the people involved in this undertaking, to whom I am deeply grateful, namely:

- all the people who read, advised and reviewed this thesis. I am sure it wasn't your first, but I very much appreciated the help, motivation, feedback and constructive criticism.
- all the people who guided, managed and supported me during my PhD. Your availability, soft-touch management and creative freedom made all the difference in what the process ended up being.
- all the people who helped me navigate the dark depths of RF and the mysteries of spectroscopy. No number of acknowledgement sections is enough to thank you for all the pain you spared me.
- all the people who made lunch and afternoon snacks, milestones of the day, all the more enjoyable and stimulating. Whatever problems we did not solve, we'll get to someday.
- all the people who, in my personal life, made this journey and the stay in my #1 city, far more enjoyable than they would have been otherwise. All these moments, and those everywhere else, will forever be cherished and remembered.
- the reviewers of this thesis manuscript: Prof. Jan Korvink and Prof. Arno Kentgens.
- the funding sources that made this work possible: Bürkert GmbH, E.G.O AG, and the Institute of Microstructure Technology of the KIT.

List of Figures

1.1	Examples of measurements of fluid properties using magnetic resonance.	1
1.2	State of the art systems using NMR to measure total flow.	3
1.3	Hardware techniques typically employed in low-cost NMR systems. . . .	4
1.4	Different transception techniques applied to flow quantification.	6
2.1	Qualitative response of magnetic materials to an external magnetic field.	16
2.2	Homogeneity improvement when using to field-homogenizing materials.	18
2.3	Responses of platinum and carbon electrodes in high-field MR applications.	22
3.1	Photograph and breakdown of the NMR-based flowmeter developed. . .	28
3.2	Simulation of the magnetic field of a C-shaped magnet.	28
3.3	Sample holder used for NMR flowmetry.	30
3.4	Optimised sensitivity profile of an RF stripline.	30
3.5	Effect of the developed shim set on NMR linewidth.	32
3.6	Disassembled view of the RF front-end developed for flowmetry.	33
3.7	Pulse sequence block used to encode flow.	34
3.8	Signal amplitude and phase responses under increasing flow encodings. .	35
3.9	Flow measurement SNR under increasing encoding strengths.	36
4.1	Sensitivity map of a stripline optimised for SNR or excitation homogeneity.	42
4.2	Disassembled view of a modular, stripline-based NMR front-end.	43
4.3	Recorded FID of water using the NMR front-end.	44
4.4	Schematic of the front-end's reception chain.	45
4.5	Noise factor evolution with varying coil impedance.	46
4.6	Schematic of the front-end's transmission chain.	47
4.7	Scattering parameters of the dual-coil, decoupled front-end.	48
4.8	Time evolution of coil voltages during a switch event.	49
4.9	MRI images measured in common and differential mode configurations. .	51
4.10	Schematic of a rollable double-helix resonator.	58

List of Figures

4.11	Photographs of the implemented double-helix resonator.	61
4.12	On-PCB connection pads for coil adaption.	61
4.13	Q-factor and optimal RF frequency for various patterning configurations.	62
4.14	Sensitivity comparison between a double-helix and a solenoid.	63
4.15	Response of a resonator with differently-tuned coil sections.	64
4.16	Q-factor dependence on winding connections in series and in parallel.	65
5.1	Representation of the domains involved in the reciprocity principle.	72
5.2	Alignment process and field response of a self-aligned micromagnet.	78
5.3	Magnetic field intensity and homogeneity of a self-aligned powder magnet.	79
5.4	Field gradient homogeneity of a tunable profilometry magnet.	82
5.5	Schematic and performance of a topology-optimized table magnet	83

List of Tables

- 1.1 Comparative analysis of different flowmetry encoding and transception methods and their respective turndown ratio and accuracy. 5
- 3.1 Simulated performance of a seven coil shim set generating spherical harmonics. 31
- 5.1 Normalised figures of merit, as described in Sect. 5.1.3, comparing the obtained results for the normalised penetration depth (d), magnetic field (B), normalised field gradient and minimum discernible thickness (T) to literature. 81