

Active Decoupling and Optimal Control for Parallel NMR Spectroscopy

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

von der KIT-Fakultät für Elektrotechnik und Informationstechnik des
Karlsruher Instituts für Technologie (KIT)
angenommene

DISSERTATION

von

M.Sc. Mengjia He
geboren in Hubei, China

Tag der mündlichen Prüfung: July 13, 2026

Hauptreferent: Prof. Dr. sc. techn. ETH Jan Gerrit Korvink

Korreferent: Prof. Dr.-Ing. John Jelonnek



This document is licensed under a Creative Commons

Attribution-ShareAlike 4.0 International License (CC BY-SA 4.0):

<https://creativecommons.org/licenses/by-sa/4.0/deed.en>

Abstract

Parallel operation using a miniaturized detector array makes good use of the limited magnet bore volume and enhances nuclear magnetic resonance (NMR) measurement throughput. This thesis investigates active decoupling for parallel detection via optimal control pulses and introduces an accelerated optimal control framework based on the finite element method (FEM).

A multiphysics simulation workflow was developed to characterize radio-frequency (RF) coupling, quantify its impact on parallel excitation, and analyze the resulting mixture of parallel free induction decays (FIDs). In a parallel solenoid coil geometry, the circuit-level coupling was found to exceed direct field leakage by a factor of 10^3 . Microwave network modeling revealed that parallel excitation and acquisition share the same coupling matrix, with a modeling error on the order of 10^{-3} . These findings establish the feasibility of signal separation and active decoupling pulse design using the coupling matrix.

Building on the linear instantaneous mixture model of parallel FIDs, blind source separation was employed to extract coupling information and decompose composite signals. High-field (500 and 650 MHz) parallel spectra decomposition confirms that this method removes more than 90% of coupled components. The extracted coupling information was subsequently used to design cooperative optimal control pulses. Experimental validation demonstrated that these tailored pulses remain robust against RF coupling during parallel excitation.

For a parallel NMR probe equipped with parallel gradient coils, coherence-locking RF pulses designed by optimal control were implemented to mitigate gradient field spillover. Application of these pulses in pulsed gradient spin-echo (PGSE) and parallel heteronuclear single quantum coherence (HSQC) experiments verified their effectiveness in preserving the desired coherences.

Finally, single-spin optimal control was accelerated by over an order of magnitude by combining FEM with the method of moving asymptotes (MMA). By treating discretized time as spatial coordinates, the Liouville–von Neumann equation was reformulated as a linear system, efficiently yielding a joint solution of the spin trajectory and control gradient. The method of moving asymptotes, relying on the ensemble fidelities and gradients, achieves rapid convergence for a target fidelity of 0.995.

Zusammenfassung

Der parallele Betrieb mit einem Array miniaturisiert Detektoren nutzt das begrenzte Volumen der Magnetbohrung effizient aus und erhöht den Durchsatz bei kernmagnetischen Resonanzmessungen (NMR). Diese Arbeit untersucht die aktive Entkopplung für die Paralleldetektion mittels optimaler Kontrollpulse und stellt ein beschleunigtes Optimierungsverfahren auf Basis der Finite-Elemente-Methode (FEM) vor.

Ein Multiphysik-Simulationsworkflow wurde entwickelt, um die Hochfrequenz-(RF)-Kopplung zu charakterisieren, deren Einfluss auf parallele Anregung zu quantifizieren und die resultierende Mischung paralleler Free-Induction-Decays (FIDs) zu analysieren. In einer Geometrie mit parallelen Solenoidspulen wurde festgestellt, dass die schaltungstechnische Kopplung die direkte Feldüberlagerung um den Faktor 10^3 übersteigt. Die Modellierung mittels Mikrowellennetzwerken zeigte, dass parallele Anregung und Detektion dieselbe Kopplungsmatrix teilen, mit einem Modellierungsfehler in der Größenordnung von 10^{-3} . Diese Ergebnisse belegen die Machbarkeit der Signalseparation und der aktiven Entkopplungspuls-Entwicklung auf Basis der Kopplungsmatrix.

Aufbauend auf dem linearen Modell der momentanen Mischung paralleler FIDs wurde Blind-Source-Separation eingesetzt, um Kopplungsinformationen zu extrahieren und zusammengesetzte Signale zu zerlegen. Die Zerlegung paralleler Hochfeldspektren (500 und 650 MHz) bestätigt, dass diese Methode mehr als 90% der gekoppelten Komponenten entfernt. Die extrahierten Kopplungsinformationen wurden anschließend zur Entwicklung kooperativer optimaler Kontrollpulse genutzt. Experimentelle Validierung zeigte, dass diese maßgeschneiderten Pulse während der parallelen Anregung robust gegenüber RF-Kopplung bleiben.

Für eine parallele NMR-Sonde mit parallelen Gradientenfeldern wurden Kohärenzsichernde RF-Pulse, entworfen durch optimale Kontrolle, implementiert, um Gradienten-

feld Übersprechen zu kompensieren. Die Anwendung dieser Pulse in Pulsgradienten Spin Echo (PGSE) und parallelen Heteronuklear Single Quantum Coherence (HSQC) Experimenten bestätigte ihre Wirksamkeit beim Schutz der gewünschten Kohärenzen.

Schließlich wurde die Einzelspin-Optimalkontrolle durch die Kombination von FEM mit der Methode der beweglichen Asymptoten (MMA) um mehr als eine Größenordnung beschleunigt. Durch die Behandlung der diskretisierten Zeit als räumliche Koordinaten wurde die Liouville-von-Neumann-Gleichung zu einem linearen Gleichungssystem umformuliert, das effizient eine gemeinsame Lösung für die Spintrajektorie und den Steuergradienten liefert. Die Methode der beweglichen Asymptoten, die auf den Ensemblestreufidelitäten und deren Gradienten basiert, erreicht eine schnelle Konvergenz bis zu einer Zielfidelität von 0,995.

Nomenclature

Abbreviations

1D/2D/3D	One/two/three dimension
AC/DC	Alternating/Direct current
BSS	Blind source separation
CL	Coherence locking
DOF(s)	Degree(s) of freedom
EM	Electromagnetic
FD	Finite difference
FEM	Finite element method
FID	Free induction decay
GRAPE	Gradient ascent pulse engineering
HMQC	Heteronuclear multiple quantum coherence spectroscopy
HSQC	Heteronuclear single quantum coherence spectroscopy
LBFGS	Limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm
LvN	Liouville–von Neumann
MMA	Method of moving asymptotes
MRI	Magnetic resonance imaging
NMR	Nuclear magnetic resonance
OC	Optimal control
PGSE	Pulsed gradient spin echo
ppm	Parts per million
QOC	Quantum optimal control

RF	Radio frequency
SL	Spin locking
SNR	Signal to noise ratio
UR	Universal rotation

Symbols

I_c	Current on the coil
Δf	Spectral difference
δ	Compensation pulse term
η	Transfer fidelity of the pulse
$\hat{\mathbf{F}}$	Coupling matrix estimated from BSS method
$\hat{\mathbf{V}}_n$	Noise covariance
$\hat{\mathbf{V}}_s$	Signal covariance
$\mathbf{a}$	Incident wave in the microwave network
$\mathbf{b}$	Reflected wave in the microwave network
$\mathbf{E}$	Identity matrix
$\mathbf{F}$	Matrix transforming excitation wave to coil current
$\mathbf{G}$	Signal gain matrix in the reception stage
$\mathbf{p}_0$	Optimal control pulses
$\mathbf{p}$	Cooperative pulses
$\mathbf{S}_C$	S matrix of coil array
$\mathbf{S}_M$	S matrix of matching network
$\mathbf{S}_{CM}$	S matrix of combined coil and matching network
$\mathbf{S}$	Scattering matrix
ν_0	Resonance offset
ν_1	Nutation frequency or RF amplitude in Hz
B_0	Static magnetic field
B_1	Radio frequency magnetic field
C_A	Figure of merit of signal decomposition

Contents

Abstract	i
Zusammenfassung	iii
Nomenclature	v
1 Introduction	1
1.1 Motivation	1
1.2 Thesis outline	3
1.3 Main contributions	4
2 Fundamental theory	7
2.1 RF chain modeling	7
2.2 Nuclear spin Hamiltonian	11
2.2.1 Zeeman interactions	11
2.2.2 Dipole-dipole coupling	12
2.2.3 J -coupling	13
2.2.4 RF irradiation	13
2.2.5 Ensemble spin Hamiltonian	14
2.3 Average Hamiltonian theory	15
2.3.1 Magnus expansion	15
2.3.2 Zeeman rotating frame	16
2.3.3 Toggling frame by RF irradiation	18
2.4 Gradient ascent pulse engineering	19
2.5 Finite element method	22
2.5.1 Finite elements	23
2.5.2 Basis functions	24
2.5.3 Formulation of linear equations	26
2.5.4 Adjoint analysis	27

3	Multiphysics modeling for parallel NMR spectroscopy	29
3.1	Introduction	29
3.2	Multiphysics simulation workflow	30
3.3	EM simulation	32
3.3.1	B_0 field simulation	32
3.3.2	Tuning and matching	33
3.3.3	B_1 field calculation	38
3.4	Spin dynamics simulation	41
3.5	Wall clock time	44
3.6	RF coupling in parallel excitation	44
3.7	RF coupling in parallel acquisition	47
3.7.1	Modeling of FID mixing	47
3.7.2	Four-detector NMR simulation	49
3.7.3	Gain-matrix decomposition error	51
3.8	Conclusion	53
4	Radio frequency decoupling for parallel NMR spectroscopy	55
4.1	Introduction	55
4.2	Decoupling for parallel excitation	57
4.3	Signal decomposition with blind source separation	60
4.3.1	Blind source separation	60
4.3.2	Parallel acquisition experiment	63
4.3.3	Spectral resolution	64
4.3.4	Estimation error	67
4.4	Parallel excitation experiment	69
4.4.1	Experimental setup	69
4.4.2	Calibration test	70
4.4.3	Parallel excitation	72
4.5	End-to-end test	74
4.6	Conclusion	76
5	Gradient decoupling for parallel NMR spectroscopy	77
5.1	Introduction	77

5.2	Quantify the gradient spillover	78
5.2.1	Gradient field simulation	78
5.2.2	Measuring the spillover ratio	80
5.2.3	Field gradient dephasing	81
5.3	Coherence locking pulse	85
5.3.1	Optimal control formulation	85
5.3.2	Coherence-locking efficiency	89
5.3.3	Gradient imperfections	90
5.3.4	Heteronuclear decoupling	91
5.3.5	Homonuclear decoupling	94
5.4	Coherence locking in PGSE	95
5.5	Coherence locking in parallel HSQC	96
5.5.1	Simulation	96
5.5.2	Experiment	99
5.6	Simulation of coherence locking in parallel HMQC	101
5.7	Conclusion	104
6	Accelerated quantum optimal control for magnetic resonance pulse design	107
6.1	Introduction	107
6.2	Methodology	109
6.2.1	FEM solution of the Liouville-von Neumann equation	110
6.2.2	Adjoint analysis	115
6.2.3	Regularization	116
6.3	Numerical implementation	118
6.3.1	Matrix assembly and solving	118
6.3.2	Optimization algorithm	120
6.4	Results and discussion	122
6.4.1	Accuracy with small step	122
6.4.2	Speed for 1 and 2 spin systems	124
6.4.3	State-to-state optimal control	126
6.4.4	Propagator optimization	128
6.5	Conclusion	129

Contents

7	Conclusions	131
	Bibliography	133
	Publications	147
	Acknowledgements	149

Chapter 1

Introduction

1.1 Motivation

It is meaningful for a PhD thesis to trace the origins of the theories and technologies on which it is built. To begin with, several foundational discoveries played crucial roles in the understanding of spin. In 1897, Thomas Preston reported the splitting patterns of various atomic spectral lines in a magnetic field, an observation now known as the anomalous Zeeman effect, which could not be explained by the electron orbital theory at that time [1]. In 1922, Stern and Gerlach observed the double splitting of evaporated silver atom beams passing through an inhomogeneous magnetic field [2], a phenomenon that Bohr's atom model could not describe. In 1925, Pauli proposed his explanation of the doublet structure of the alkali spectra [3] and subsequently introduced an additional two-valued quantum property in his formulation of the exclusion principle [4]. This additional degree of freedom corresponds to the quantized behavior of electron spin. A rigorous description of a spin-1/2 particle was later provided by Dirac in 1928 through his relativistic wave equation [5].

The discovery of magnetic resonance dates back to 1938, when Isidor I. Rabi and his colleagues first demonstrated the reorientation of nuclear spin magnetic moments by photon absorption in a molecular beam of LiCl [6]. In 1946, Felix Bloch and Edward Purcell independently extended this phenomenon to bulk matter and observed the nuclear induction signals [7, 8], thereby establishing the foundation of nuclear magnetic resonance (NMR). Since then, magnetic resonance has evolved into two major technologies: NMR

spectroscopy and magnetic resonance imaging (MRI). Both have revolutionized their respective fields: MRI in medical diagnostics and NMR in molecular analysis, and related contributions have been recognized by multiple Nobel Prizes in Chemistry [9, 10] and Medicine [11, 12].

Magnetic resonance imaging (MRI) has become a standard clinical technique, offering non-invasive, high-resolution visualization of anatomical structures and physiological processes. In contrast, nuclear magnetic resonance (NMR) spectroscopy provides molecular-level insights into structure and dynamics through Fourier-transformed spectra, playing a crucial role in chemical reaction monitoring, drug discovery, and food analysis.

Despite their wide range of applications, NMR techniques face inherent limitations. The energy separation between the two Zeeman levels induced by an external magnetic field is given by $\Delta E = \gamma \hbar B_0 / 2\pi$. The corresponding population difference between the two levels, governed by the Boltzmann distribution, is exceedingly small. Under the high-temperature approximation, it can be expressed as

$$\rho_\alpha - \rho_\beta = \frac{\Delta E}{2k_B T}, \quad (1.1)$$

which corresponds to only about 4×10^{-5} for ^1H nuclei at 25 °C and 11.74 T. These tiny population differences give rise to the net spin alignment that produces the detectable signal, resulting in intrinsically low measurement sensitivity. Hyperpolarization techniques such as dynamic nuclear polarization [13, 14, 15], parahydrogen-induced polarization [16], and signal amplification by reversible exchange [17] overcome the thermal equilibrium limitations, dramatically enhancing signal strength and becoming a major focus in the magnetic resonance community.

When a large number of samples need to be analyzed, as in process monitoring or drug screening, high-throughput detection becomes essential. Enabled by microcoil and micro-volume probe technologies, flow-through NMR systems were developed to support automatic sample exchange and sequential data acquisition [18, 19, 20]. Alternatively, composite detection schemes employing bundles of capillaries have achieved high throughput with simplified hardware [21, 22].

Meanwhile, parallel NMR spectroscopy has emerged as an alternative approach [23], enabling the simultaneous measurement of multiple samples. Early implementations demon-

strated four-channel acquisition at 300 MHz [24] and 600 MHz [25] using custom probes equipped with multiple solenoid coils. Signal separation was achieved through the application of pulsed gradient fields [26] or alternative detection schemes [27]. The parallel NMR technique has found valuable applications in combination with classical pulse sequences [25], kinetic studies [28], and dissolution dynamic nuclear polarization [29]. More recently, chip-scale implementations have demonstrated portable NMR with integrated coils and time-interleaved pulses for radiofrequency (RF) decoupling [30].

To achieve truly independent parallel operation, each detector ideally includes its own RF coil, gradient coil, and shimming unit [31, 32, 33, 34]. However, when such multi-detector probes are installed within the confined bore of a superconducting magnet, residual electromagnetic (EM) coupling, which remains unavoidable despite hardware shielding, leads to inter-channel crosstalk. This issue has received relatively limited attention in the literature.

Accordingly, this thesis focuses on modeling inter-detector coupling in parallel NMR spectroscopy and developing active decoupling schemes using optimal control methods, intending to enable scalable, high-throughput measurements.

1.2 Thesis outline

This thesis focuses on decoupling parallel detectors and optimizing pulse sequences for parallel NMR spectroscopy. The work is organized into five main chapters (Chapters 2–6).

Chapter 2 introduces the theoretical foundations. The microwave network method is presented as a tool for analyzing RF coupling in parallel transmit/receive channels. A brief introduction is given to the spin Hamiltonian and the Liouville–von Neumann equation. Average Hamiltonian theory is outlined for analyzing pulse sequences. Gradient ascent pulse engineering (GRAPE), a widely used optimal control method, is briefly reviewed. The finite element method (FEM), employed both in magnetic field simulations and in optimal control, is also introduced.

Chapter 3 describes the development of a digital twin for parallel NMR spectroscopy. By combining electromagnetic simulations with spin dynamics simulations, the RF coupling effects in both the excitation and reception stages of parallel detection are revealed.

Chapter 4 demonstrates that blind source separation can be used to split composite free induction decays (FIDs) and provide coupling matrices that characterize RF coupling. Decoupling schemes for parallel excitation are developed and verified in parallel NMR spectroscopy. An end-to-end simulation further illustrates the applicability of this approach to broader experimental scenarios.

Chapter 5 proposes a gradient compensation strategy to counteract field spillover in probes with parallel gradient coils. Its effectiveness is validated through PGSE and parallel HSQC experiments.

Chapter 6 shows that combining FEM with the method of moving asymptotes (MMA) accelerates pulse optimization for ensemble single-spin systems by more than an order of magnitude compared with the commonly used L-BFGS–GRAPE method. This improvement highlights the potential of FEM–MMA for efficient pulse optimization in both NMR and MRI.

1.3 Main contributions

The main contributions of this thesis to the field of magnetic resonance, documented in three journal publications, are summarized below:

Digital twin for parallel NMR spectroscopy [35]. A digital twin was developed to simulate liquid-state parallel NMR spectroscopy. This digital twin quantifies inter-channel coupling during both excitation and acquisition, enables precise characterization of parallel NMR probes, identifies coupling-induced artifacts in experiments, and supports the design of decoupling pulse sequences. These capabilities are critical for advancing high-throughput NMR.

Gradient decoupling for pulsed gradient experiments [36]. A gradient field compensation scheme was proposed to mitigate spillover effects in pulsed gradient NMR experiments. This method provides an effective solution for parallel probes with independently switchable gradient coils. Combined with the digital twin, it supports scalable and high-throughput experimental designs.

Quantum optimal control using finite elements and moving asymptotes [37]. An optimization framework that combines FEM and MMA was developed for efficient optimal

control of ensemble single-spin systems. Compared to the widely used GRAPE method, this method achieved more than a tenfold speedup while maintaining high accuracy.

All experimental data and code developed in this thesis are available on GitHub: <https://github.com/kikioh>.

Chapter 2

Fundamental theory

2.1 RF chain modeling

High-field NMR spectroscopy typically operates at RF bands in the hundreds of MHz. In this frequency regime, the behavior of the RF chain is conveniently described using microwave network theory [38], in which scattering parameters characterize the incident and reflected traveling waves at specified terminals. Figure 2.1 illustrates an N -port microwave network, where a_i and b_i denote the incident and reflected voltage waves at port i , respectively. These quantities are related through the scattering matrix:

$$\mathbf{b} = \mathbf{S}\mathbf{a}. \quad (2.1)$$

An element of the scattering matrix is defined as

$$S_{ij} = \left. \frac{b_i}{a_j} \right|_{a_k=0, k \neq j}, \quad (2.2)$$

that is, S_{ij} is measured by applying an excitation to port j while terminating all other ports with matched loads and recording the response at port i .

If each port has characteristic impedance Z_0 , the terminal voltages and currents are




Figure 2.1: An arbitrary N -port microwave network, the a_i and b_i denote the incident and reflected voltage waves at port i , respectively.

related to the wave amplitudes by

$$\begin{aligned} V_i &= \sqrt{Z_0} (a_i + b_i), \\ I_i &= \frac{1}{\sqrt{Z_0}} (a_i - b_i). \end{aligned} \quad (2.3)$$

The impedance matrix $\mathbf{Z}$ is defined by

$$\mathbf{V} = \mathbf{Z}\mathbf{I}. \quad (2.4)$$

Combining Eqs. 2.3 and 2.4 yields the scattering matrix in terms of the impedance matrix:

$$\mathbf{S} = (\mathbf{Z} + Z_0\mathbf{E})^{-1}(\mathbf{Z} - Z_0\mathbf{E}), \quad (2.5)$$

where $\mathbf{E}$ is an identity matrix, and conversely, the impedance matrix in terms of $\mathbf{S}$ is

$$\mathbf{Z} = Z_0(\mathbf{E} + \mathbf{S})(\mathbf{E} - \mathbf{S})^{-1}. \quad (2.6)$$

Now consider an N -port coil array combined with a tuning and matching (T&M) network (Fig. 2.2). Once the capacitance values are determined, by using the port permutation, the scattering matrix of a $2N$ -port T&M network can be written as a block matrix [39]:

$$\mathbf{S}_M = \begin{bmatrix} \mathbf{S}_{11} & \mathbf{S}_{12} \\ \mathbf{S}_{21} & \mathbf{S}_{22} \end{bmatrix}, \quad (2.7)$$




Figure 2.2: Combination of a coil array with a tuning and matching (T&M) network, described using S-parameters. (a) Block diagram of a parallel probe consisting of a mutually coupled coil array, a matching network, and matched loads. The FID denotes the free induction decay; S_C , S_M , and S_{CM} are the S-matrices of the coil array, the T&M network, and their combination, respectively. Forward waves are a_1, a_2 and reverse waves are b_1, b_2 . (b) Schematic diagram of a coil array connected to the T&M network. Each coil is modeled by a series L and R , with C_t and C_m representing tuning and matching capacitances.

So the T&M network can be treated as a 2-port network whose scattering relations are

$$\begin{bmatrix} b_1 \\ b_2 \end{bmatrix} = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix}. \quad (2.8)$$

where subscripts 1 and 2 denote the connecting ports and the free ports, respectively. Meanwhile, a_1 and b_1 denote the backward and forward waves, respectively, with respect to the coil array. Accordingly, a_1 and b_1 satisfy

$$a_1 = S_C b_1. \quad (2.9)$$

Combining Eqs. 2.8 and 2.9 yields the overall scattering matrix

$$S_{CM} = S_{22} + S_{21} (E - S_C S_{11})^{-1} S_C S_{12}. \quad (2.10)$$

Moreover, $\mathbf{a}_1$ and $\mathbf{b}_1$ can be expressed in terms of $\mathbf{a}_2$ as

$$\begin{aligned}\mathbf{a}_1 &= (\mathbf{E} - \mathbf{S}_C \mathbf{S}_{11})^{-1} \mathbf{S}_C \mathbf{S}_{12} \mathbf{a}_2, \\ \mathbf{b}_1 &= \mathbf{S}_C^{-1} (\mathbf{E} - \mathbf{S}_C \mathbf{S}_{11})^{-1} \mathbf{S}_C \mathbf{S}_{12} \mathbf{a}_2.\end{aligned}\tag{2.11}$$

Considering that the coil currents and voltages are

$$\begin{aligned}\mathbf{I}_C &= \frac{1}{\sqrt{Z_0}} (\mathbf{b}_1 - \mathbf{a}_1), \\ \mathbf{V}_C &= \sqrt{Z_0} (\mathbf{b}_1 + \mathbf{a}_1).\end{aligned}\tag{2.12}$$

Taking the RF pulse $\mathbf{p}$ as the excitation $\mathbf{a}_2$, the transfer matrix from $\mathbf{a}_2$ to coil current $\mathbf{I}_C$ is

$$\mathbf{F} = \frac{1}{\sqrt{Z_0}} (\mathbf{S}_C^{-1} - \mathbf{E}) (\mathbf{E} - \mathbf{S}_C \mathbf{S}_{11})^{-1} \mathbf{S}_C \mathbf{S}_{12}.\tag{2.13}$$

At the reception stage, the reverse wave at coil terminal $\mathbf{a}_1$ consists of scattering energy induced by $\mathbf{b}_1$ and an additional contribution $\mathbf{b}_{\text{FID}}$ generated by the FIDs, i.e.,

$$\mathbf{a}_1 = \mathbf{b}_{\text{FID}} + \mathbf{S}_C \mathbf{b}_1,\tag{2.14}$$

The FID calculated by the reciprocity principle corresponds to an open-circuit voltage $\mathbf{V}_{1\text{on}}$ at the coil terminal, and $\mathbf{V}_{1\text{on}}$ satisfies

$$\mathbf{V}_{1\text{on}} = \sqrt{Z_0} (\mathbf{a}_1 + \mathbf{b}_1).\tag{2.15}$$

where $\mathbf{a}_1 = \mathbf{b}_1$ under open-circuit condition. Hence, $\mathbf{b}_{\text{FID}}$ can be given by $\mathbf{V}_{1\text{on}}$ as

$$\mathbf{b}_{\text{FID}} = \frac{1}{2\sqrt{Z_0}} (\mathbf{E} - \mathbf{S}_C) \mathbf{V}_{1\text{on}}.\tag{2.16}$$

Considering the matched loads at the T&M output, $\mathbf{a}_2 = 0$. The voltage gain matrix $\mathbf{G}$, defined by the relation $\mathbf{V}_2 = \sqrt{Z_0} \mathbf{b}_2 = \mathbf{G} \mathbf{V}_{1\text{on}}$, is

$$\mathbf{G} = \frac{1}{2} \mathbf{S}_{21} (\mathbf{E} - \mathbf{S}_C \mathbf{S}_{11})^{-1} (\mathbf{E} - \mathbf{S}_C).\tag{2.17}$$

Thus, the excitation coupling is quantified by $\mathbf{F}$, and the reception coupling by $\mathbf{G}$.

Noise coupling follows the same path as FID coupling until the signals reach the LNA. Assuming the noise originates from the coil array, the noise voltage covariance is

$$\hat{\mathbf{R}}_n = 4k_B T B \cdot \text{Re}[\mathbf{Z}_C], \quad (2.18)$$

where k_B is Boltzmann's constant, T the coil temperature, $\mathbf{Z}_C$ the coil array impedance matrix, and B the receiver bandwidth. The noise covariance at the load is

$$\hat{\mathbf{V}}_n = \mathbf{G}^* \hat{\mathbf{R}}_n \mathbf{G}^\top. \quad (2.19)$$

The signal-to-noise ratio (SNR) in channel m is

$$\text{SNR}_m = \frac{\hat{V}_{s,mm}}{\hat{V}_{n,mm}}, \quad (2.20)$$

where the signal covariance is

$$\hat{\mathbf{V}}_s = \mathbf{G}^* (\mathbf{V}_{\text{1on}}^* \mathbf{V}_{\text{1on}}^\top) \mathbf{G}^\top. \quad (2.21)$$

2.2 Nuclear spin Hamiltonian

The Schrödinger equation, formulated by Erwin Schrödinger in 1926 [40], describes the quantum dynamics of nuclei and electrons. In the context of NMR, it takes the form

$$\frac{\partial}{\partial t} \psi(t) = -i\mathbf{H}\psi(t), \quad (2.22)$$

where $\psi(t)$ is the nuclear spin state and $\mathbf{H}$ is the nuclear spin Hamiltonian. The Hamiltonian depends on nuclear spin interactions, which are detailed below.

2.2.1 Zeeman interactions

The magnetic moment operator is related to angular momentum by

$$\boldsymbol{\mu} = \gamma \mathbf{I}, \quad (2.23)$$

where γ is the gyromagnetic ratio and $\mathbf{I}$ is the nuclear spin operator. The spin Hamiltonian is the negative dot product of the magnetic moment with the magnetic field:

$$\mathbf{H}_Z = -\boldsymbol{\mu} \cdot \mathbf{B} = -\gamma(B_x \mathbf{I}_x + B_y \mathbf{I}_y + B_z \mathbf{I}_z). \quad (2.24)$$

For a strong longitudinal field $\mathbf{B} = B_0 \hat{z}$, this reduces to

$$\mathbf{H}_Z = -\gamma B_0 \mathbf{I}_z, \quad (2.25)$$

the nuclear Zeeman interaction.

In reality, electronic shielding modifies the local field experienced by the nucleus. The applied field B_0 induces electronic currents that generate a perturbation described by the shielding tensor δ :

$$\mathbf{B}_{\text{loc}} = (1 + \delta) \cdot B_0 \hat{z} = \begin{bmatrix} 1 + \delta_{xx} & \delta_{xy} & \delta_{xz} \\ \delta_{yx} & 1 + \delta_{yy} & \delta_{yz} \\ \delta_{zx} & \delta_{zy} & 1 + \delta_{zz} \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ B_0 \end{bmatrix}. \quad (2.26)$$

So the Zeeman Hamiltonian becomes

$$\mathbf{H}_Z = -\gamma B_0 (\delta_{xz} \mathbf{I}_x + \delta_{yz} \mathbf{I}_y + (1 + \delta_{zz}) \mathbf{I}_z). \quad (2.27)$$

In isotropic liquids, rapid molecular tumbling averages out the anisotropic components, leaving

$$\mathbf{H}_Z = -\gamma B_0 (1 + \delta) \mathbf{I}_z, \quad (2.28)$$

where δ is the isotropic chemical shift. In solids, only the last term of Eq. 2.27 remains [41] after secular approximation.

2.2.2 Dipole-dipole coupling

The direct spin–spin interaction can be treated using the point-dipole approximation, as the nuclei are sufficiently far apart. This interaction, known as the dipole–dipole coupling, depends on the internuclear distance and the angle between the internuclear vector and the

external magnetic field, and hence exhibits axial symmetry. The corresponding dipole–dipole spin Hamiltonian is expressed as

$$\mathbf{H}_{\text{DD}} = -\frac{\mu_0 \gamma_I \gamma_S \hbar}{4\pi r^3} [3(\mathbf{I} \cdot \hat{r})(\hat{r} \cdot \mathbf{S}) - (\mathbf{I} \cdot \mathbf{S})], \quad (2.29)$$

where μ_0 is the vacuum magnetic permeability, $\hbar$ is the reduced Planck constant, γ_I and γ_S are the gyromagnetic ratios of the two spins, $\hat{r}$ is the unit vector connecting the nuclei, and r is the internuclear distance. It can be shown that $\mathbf{H}_{\text{DD}}$ is traceless [42] and therefore averages to zero in liquids, where rapid molecular tumbling occurs.

2.2.3 *J*-coupling

Indirect spin–spin coupling mediated by bonding electrons is known as *J*-coupling. The *J*-coupling is typically treated as isotropic, although small anisotropic components may arise in anisotropic liquids and solids. However, these anisotropic components are indistinguishable from dipole–dipole couplings in measurement and are therefore usually neglected [41]. The isotropic *J*-coupling Hamiltonian, which is independent of the applied magnetic field, can be written as a bilinear interaction between spins *I* and *S* as follows:

$$\mathbf{H}_J = 2\pi J(\mathbf{I}_x \mathbf{S}_x + \mathbf{I}_y \mathbf{S}_y + \mathbf{I}_z \mathbf{S}_z). \quad (2.30)$$

For heteronuclear spin pairs, the secular approximation reduces this to

$$\mathbf{H}_J = 2\pi J \mathbf{I}_z \mathbf{S}_z. \quad (2.31)$$

2.2.4 RF irradiation

The RF field is applied to manipulate nuclear spin states. The field generated by an RF coil can be expressed as

$$\mathbf{B}_{\text{rf}} = B_1 \cos(\omega_{\text{rf}} t + \phi) \hat{x}, \quad (2.32)$$

where $\hat{x}$ defines the transverse reference axis, ϕ is the phase, and ω_{rf} is the oscillation frequency, typically close to the Larmor frequency ($\omega_{\text{rf}} \approx -\gamma B_0$). The corresponding RF

Hamiltonian is given by

$$\mathbf{H}_{\text{rf}} = -\gamma B_1 \cos(\omega_{\text{rf}} t + \phi) \mathbf{I}_x, \quad (2.33)$$

Modern NMR spectrometers can modulate B_1 and ϕ with microsecond precision, enabling high-fidelity spin manipulation via shaped pulses.

2.2.5 Ensemble spin Hamiltonian

For an ensemble of spins, such as the Avogadro number of proton spins in a water sample, it becomes impossible to treat each spin individually. An alternative approach, known as the density operator formalism, describes the quantum state of the entire ensemble. Considering the expectation value of an observable Q with operator $\mathbf{Q}$ is

$$\overline{Q} = \overline{\langle \psi | \mathbf{Q} | \psi \rangle} = \text{Tr}(\overline{|\psi\rangle\langle\psi|} \mathbf{Q}). \quad (2.34)$$

This motivates the density operator $\rho = |\psi\rangle\langle\psi|$, which carries the ensemble average information and is invariant under global phase shifts [42].

Differentiating ρ using Eq. 2.22 yields

$$\frac{\partial \rho}{\partial t} = -i\mathbf{H}|\psi\rangle\langle\psi| + i|\psi\rangle\langle\psi|\mathbf{H} = -i[\mathbf{H}, \rho], \quad (2.35)$$

the Liouville–von Neumann (LvN) equation.

For numerical computations, it is often recast in Liouville space as

$$\frac{\partial \boldsymbol{\rho}}{\partial t} = -i\mathcal{H}\boldsymbol{\rho}, \quad (2.36)$$

where $\boldsymbol{\rho}$ is column-wisely vectorized from the density matrix, and the commutation superoperator is

$$\mathcal{H} = \mathbf{E} \otimes \mathbf{H} - \mathbf{H} \otimes \mathbf{E}, \quad (2.37)$$

where $\mathbf{E}$ is an identity matrix with the same dimension as $\mathbf{H}$, the $n \times n$ Hamiltonian is then extended to a $n^2 \times n^2$ superoperator $\mathcal{H}$. Transforming to Liouville space facilitates the construction of a composite system comprising multiple non-interacting subsystems, which

account for the multiple samples in parallel NMR. Specifically, the sub-superoperators are arranged as a block-diagonal matrix, and the corresponding spin vectors are stacked column-wise.

2.3 Average Hamiltonian theory

2.3.1 Magnus expansion

The evolution of a spin system under a time-dependent Hamiltonian $\mathbf{H}(t)$ can be described by an effective average Hamiltonian over the time interval $[0, T]$, defined by

$$\exp \left[-i\overline{\mathbf{H}}T \right] = \exp_o \left[-i \int_0^T \mathbf{H}(t) dt \right], \quad (2.38)$$

where $\exp_o$ denotes the time-ordered exponential.

The average Hamiltonian $\overline{\mathbf{H}}$ can be expressed as a Magnus expansion [43]:

$$\overline{\mathbf{H}} = \overline{\mathbf{H}}^{(1)} + \overline{\mathbf{H}}^{(2)} + \overline{\mathbf{H}}^{(3)} + \dots, \quad (2.39)$$

with the first three terms given by

$$\begin{aligned} \overline{\mathbf{H}}^{(1)} &= \frac{1}{T} \int_0^T \mathbf{H}(t_1) dt_1, \\ \overline{\mathbf{H}}^{(2)} &= \frac{-i}{2T} \int_0^T dt_1 \int_0^{t_1} dt_2 [\mathbf{H}(t_1), \mathbf{H}(t_2)], \\ \overline{\mathbf{H}}^{(3)} &= \frac{-1}{6T} \int_0^T dt_1 \int_0^{t_1} dt_2 \int_0^{t_2} dt_3 \left([\mathbf{H}(t_1), [\mathbf{H}(t_2), \mathbf{H}(t_3)]] + [[\mathbf{H}(t_1), \mathbf{H}(t_2)], \mathbf{H}(t_3)] \right). \end{aligned} \quad (2.40)$$

The series converges rapidly if $\|\mathbf{H}(t)\|T \ll 1$, in which case the first-order term $\overline{\mathbf{H}}^{(1)}$ already provides a good approximation.

When the condition $\|\mathbf{H}(t)\|T \ll 1$ does not hold, it is often useful to split the Hamiltonian into two parts [44]:

$$\mathbf{H}(t) = \mathbf{H}_A(t) + \mathbf{H}_B(t), \quad (2.41)$$

where $\mathbf{H}_A(t)$ contains the dominant interaction and $\mathbf{H}_B(t)$ the perturbation. The propagator for $\mathbf{H}_A(t)$ is $\mathbf{U}_A(t)$, which defines the toggling frame.

In this frame, the Hamiltonian is transformed to

$$\mathbf{H}^{\text{tog}}(t) = \mathbf{U}_A^\dagger(t) \mathbf{H}_B(t) \mathbf{U}_A(t), \quad (2.42)$$

and the propagator over $[0, T]$ is

$$\mathbf{U}^{\text{tog}}(T) = \exp_0 \left[-i \int_0^T \mathbf{H}^{\text{tog}}(t_1) dt_1 \right] = \exp[-i \bar{\mathbf{H}}_B T], \quad (2.43)$$

with $\bar{\mathbf{H}}_B$ expanded again via the Magnus series:

$$\bar{\mathbf{H}}_B = \bar{\mathbf{H}}_B^{(1)} + \bar{\mathbf{H}}_B^{(2)} + \bar{\mathbf{H}}_B^{(3)} + \dots \quad (2.44)$$

Each term can be computed using Eq. 2.40, in which $\mathbf{H}(t)$ is replaced by $\mathbf{H}^{\text{tog}}(t)$. The full propagator is then

$$\mathbf{U}(T) = \mathbf{U}_A(T) \mathbf{U}^{\text{tog}}(T) = \mathbf{U}_A(T) \exp[-i \bar{\mathbf{H}}_B T] \quad (2.45)$$

2.3.2 Zeeman rotating frame

In high-field NMR spectroscopy, the dominant term in the spin Hamiltonian is the Zeeman interaction between the nuclear spin and the strong static magnetic field. The reference frame that rotates at the Zeeman frequency is therefore referred to as the Zeeman rotating frame.

For a single spin subject to RF irradiation, the Hamiltonian consists of the Zeeman term and the RF term,

$$\mathbf{H} = \omega_0 \mathbf{I}_z + A \cos(\omega_{\text{rf}} t + \phi) \mathbf{I}_x, \quad (2.46)$$

where $\omega_0 = -\gamma(1 + \delta)B_0$ is the Zeeman frequency of the spin, ω_{rf} is the RF carrier frequency (close to ω_0), and $A = -\gamma B_1$ is the RF amplitude. The Hamiltonian is split into a dominant part and a perturbation,

$$\begin{aligned} \mathbf{H}_A &= \omega_{\text{rf}} \mathbf{I}_z, \\ \mathbf{H}_B &= (\omega_0 - \omega_{\text{rf}}) \mathbf{I}_z + A \cos(\omega_{\text{rf}} t + \phi) \mathbf{I}_x \end{aligned} \quad (2.47)$$

According to Eq. 2.42, the toggling-frame Hamiltonian is given by

$$\begin{aligned}
 \mathbf{H}^{\text{tog}}(t) &= e^{i\omega_{\text{rf}}\mathbf{I}_z t} \mathbf{H}_B(t) e^{-i\omega_{\text{rf}}\mathbf{I}_z t} \\
 &= (\omega_0 - \omega_{\text{rf}})\mathbf{I}_z + A \cos(\omega_{\text{rf}}t + \phi) e^{i\omega_{\text{rf}}\mathbf{I}_z t} \mathbf{I}_x e^{-i\omega_{\text{rf}}\mathbf{I}_z t}, \\
 &= (\omega_0 - \omega_{\text{rf}})\mathbf{I}_z + A \cos(\omega_{\text{rf}}t + \phi) (\mathbf{I}_x \cos(\omega_{\text{rf}}t) - \mathbf{I}_y \sin(\omega_{\text{rf}}t)) \\
 &= (\omega_0 - \omega_{\text{rf}})\mathbf{I}_z + \frac{A}{2} (\mathbf{I}_x \cos \phi + \mathbf{I}_y \sin \phi) + \frac{A}{2} (\mathbf{I}_x \cos(2\omega_{\text{rf}}t + \phi) - \mathbf{I}_y \sin(2\omega_{\text{rf}}t + \phi))
 \end{aligned} \tag{2.48}$$

Applying the first-order approximation, the effective Hamiltonian in the Zeeman rotating frame becomes

$$\overline{\mathbf{H}}_B^{(1)} = \frac{1}{T} \int_0^T \mathbf{H}^{\text{tog}}(t_1) dt_1 = (\omega_0 - \omega_{\text{rf}})\mathbf{I}_z + \frac{A}{2} (\mathbf{I}_x \cos \phi + \mathbf{I}_y \sin \phi) \tag{2.49}$$

where T can take several periods of the RF frequency ω_{rf} . Equation 2.49 shows that the effective RF amplitude is $A/2 = -\gamma B_1/2$, which quantifies the strength of the resonant RF field experienced by a given spin. By choosing an appropriate initial phase for the rotating frame,

$$\begin{aligned}
 \mathbf{U}_A(t) &= \exp [-i\mathbf{I}_z(\omega_{\text{rf}}t + \phi_{\text{ref}})] \\
 \phi_{\text{ref}} &= \pi, \quad \gamma > 0 \\
 \phi_{\text{ref}} &= 0, \quad \gamma < 0
 \end{aligned} \tag{2.50}$$

the effective amplitude can be written as $|\gamma B_1|/2$. In experiments, this quantity is referred to as the nutation frequency and is typically expressed in (angular) frequency units.

For a heteronuclear two-spin system, denoted as IS , the double-rotating frame is applied. The total Hamiltonian can be decomposed into two parts as follows:

$$\begin{aligned}
 \mathbf{H}_A &= \omega_{\text{rf}}^I \mathbf{I}_z + \omega_{\text{rf}}^S \mathbf{S}_z, \\
 \mathbf{H}_B &= (\omega_0^I - \omega_{\text{rf}}^I) \mathbf{I}_z + (\omega_0^S - \omega_{\text{rf}}^S) \mathbf{S}_z + 2\pi J (\mathbf{I}_x \mathbf{S}_x + \mathbf{I}_y \mathbf{S}_y + \mathbf{I}_z \mathbf{S}_z) \\
 &\quad + A^I \cos(\omega_{\text{rf}}^I t + \phi^I) \mathbf{I}_x + A^S \cos(\omega_{\text{rf}}^S t + \phi^S) \mathbf{S}_x.
 \end{aligned} \tag{2.51}$$

Applying the first-order approximation, the effective Hamiltonian is given by

$$\begin{aligned}
\overline{\mathbf{H}}_B &= \overline{\mathbf{H}}_B^{(1)} \\
&= \frac{1}{T} \int_0^T e^{i(\omega_{\text{rf}}^I \mathbf{I}_z t + \omega_{\text{rf}}^S \mathbf{S}_z t)} \mathbf{H}_B(t_1) e^{-i(\omega_{\text{rf}}^I \mathbf{I}_z t + \omega_{\text{rf}}^S \mathbf{S}_z t)} dt_1 \\
&= (\omega_0^I - \omega_{\text{rf}}^I) \mathbf{I}_z + (\omega_0^S - \omega_{\text{rf}}^S) \mathbf{S}_z + 2\pi J \mathbf{I}_z \mathbf{S}_z \\
&\quad + \frac{A^I}{2} (\mathbf{I}_x \cos \phi^I + \mathbf{I}_y \sin \phi^I) + \frac{A^S}{2} (\mathbf{S}_x \cos \phi^S + \mathbf{S}_y \sin \phi^S)
\end{aligned} \tag{2.52}$$

where T may span several periods of both ω_{rf}^I and ω_{rf}^S . The above computation considers the commutation of spin operators, i.e., $[\mathbf{I}_{x,y,z}, \mathbf{S}_{x,y,z}] = 0$. The bilinear interactions in the J -coupling terms $\mathbf{I}_x \mathbf{S}_x$ and $\mathbf{I}_y \mathbf{S}_y$ that oscillate rapidly in the Zeeman rotating frame are averaged out. This approximation is commonly referred to as the secular approximation [45] and has its conceptual origins in classical mechanics [46].

2.3.3 Toggling frame by RF irradiation

Under RF irradiation, when $||\mathbf{H}_{\text{rf}}|| \gg ||\mathbf{H}_z||$ and $||\mathbf{H}_{\text{rf}}|| \gg ||\mathbf{H}_J||$, the Zeeman rotating frame can be further transformed into the toggling frame of the RF field. For a spin system involving J -coupling, the Hamiltonian is decomposed as:

$$\begin{aligned}
\mathbf{H}_A &= \mathbf{H}_{\text{rf}}(t), \\
\mathbf{H}_B &= \mathbf{H}_z + \mathbf{H}_J,
\end{aligned} \tag{2.53}$$

where $\mathbf{H}_z$ is the chemical-shift (offset) Hamiltonian and $\mathbf{H}_J$ is the J -coupling Hamiltonian.

For broadband applications, where frequency offsets are comparable to RF amplitudes, the motion of the toggling frame can be defined by combining the RF term and the offset term. The Hamiltonian is decomposed as [47]:

$$\begin{aligned}
\mathbf{H}_A &= \mathbf{H}_z + \mathbf{H}_{\text{rf}}(t), \\
\mathbf{H}_B &= \mathbf{H}_J.
\end{aligned} \tag{2.54}$$

Here, $\mathbf{H}_A(t)$ can be decomposed into single-spin Hamiltonians:

$$\mathbf{H}_A(t) = \sum_i \mathbf{H}_{Ai}(t) = \sum_i (\omega_{iz} \mathbf{I}_{iz} + \omega_{ix} \mathbf{I}_{ix} + \omega_{iy} \mathbf{I}_{iy}), \tag{2.55}$$

where the subscript i refers to the i -th spin. Since the single-spin Hamiltonians mutually commute, the propagator factorizes as

$$\mathbf{U}_A(t) = \prod_i \mathbf{U}_{Ai}(t) = \prod_i \exp_o \left[-i \int_0^t \mathbf{H}_{Ai}(t_1) dt_1 \right]. \quad (2.56)$$

Average Hamiltonian theory provides a framework to approximate complex time-dependent spin dynamics [48]. It is especially powerful in analyzing RF-driven spin manipulations, composite pulses, and decoupling sequences.

2.4 Gradient ascent pulse engineering

Quantum optimal control (QOC) techniques have found widespread applications in quantum computing [49], quantum sensing [50], NMR [51], and MRI [52, 53]. The growing importance of QOC is reflected in the increasing number of publications in this field [54]. In general, the optimal control problem can be formulated as finding control parameters that steer a physical system from an initial state $|\rho_0\rangle$ to a target state $|\sigma\rangle$ under hardware and time constraints.

Among many approaches, Gradient ascent pulse engineering (GRAPE) [55], has proven to be a highly efficient algorithm over the last 20 years.

Consider an NMR quantum system described by a time-dependent Hamiltonian within the Liouville space:

$$\mathcal{H}(t) = \mathcal{H}_0 + \sum_k \omega_k(t) C_k, \quad (2.57)$$

where $\mathcal{H}_0$ is the internal Hamiltonian, C_k are control operators, and $\omega_k(t)$ are their time-dependent amplitudes.

The solution of LvN equation gives the final state as

$$\rho(T) = \exp_o \left[-i \int_0^T (\mathcal{H}_0 + \sum_k \omega_k(t) C_k) dt \right] \rho_0, \quad (2.58)$$

where T is the pulse duration and $\exp_o$ denotes a time-ordered exponential.

In practice, the control fields are discretized into N piecewise-constant segments of duration Δt_n . The propagation then becomes

$$\rho(T) = \mathcal{U}_N \cdots \mathcal{U}_2 \mathcal{U}_1 \rho_0, \quad (2.59)$$

where the n -th propagator is

$$\mathcal{U}_n = \exp \left[-i \left(\mathcal{H}_0 + \sum_k \omega_k(n) C_k \right) \Delta t_n \right]. \quad (2.60)$$

The quality of the control is measured by the overlap between the target state and the final state:

$$\eta = \langle \sigma | \rho(T) \rangle = \langle \sigma | \mathcal{U}_N \cdots \mathcal{U}_1 \rho_0 \rangle. \quad (2.61)$$

The fidelity gradient with respect to control amplitude $\omega_{k,n}$ is

$$\frac{\partial \eta}{\partial \omega_{k,n}} = \frac{\partial}{\partial \omega_{k,n}} \langle \sigma | \mathcal{U}_N \cdots \mathcal{U}_{n+1} \mathcal{U}_n \mathcal{U}_{n-1} \cdots \mathcal{U}_1 | \rho_0 \rangle. \quad (2.62)$$

Since only $\mathcal{U}_n$ depends on $\omega_{k,n}$, the derivative reduces to

$$\frac{\partial \eta}{\partial \omega_{k,n}} = \underbrace{\langle \sigma | \mathcal{U}_N \cdots \mathcal{U}_{n+1} \frac{\partial \mathcal{U}_n}{\partial \omega_{k,n}} \mathcal{U}_{n-1} \cdots \mathcal{U}_1 | \rho_0 \rangle}_{\text{backward propagation}} \quad (2.63)$$

Thus, evaluation requires one forward propagation from ρ_0 and one backward propagation from σ . The computational cost is $2N$ exponential-times-vector operations plus NK calculations of inner product, where N is the number of time steps and K the number of controls [42].

In the original GRAPE, $\partial \mathcal{U}_n / \partial \omega_{k,n}$ is approximated to first order for sufficiently small Δt_n ($\|\mathcal{H}_n\| \Delta t_n \ll 1$):

$$\frac{\partial \mathcal{U}_n}{\partial \omega_{k,n}} \approx (-i C_k \Delta t_n) \cdot \mathcal{U}_n. \quad (2.64)$$

Gradient-based optimization methods often require highly accurate derivatives. Higher-order finite difference (FD) could be applied to improve accuracy [56]. For example, the

second-order FD method is written as

$$\frac{\partial \mathcal{U}_n}{\partial \omega_{k,n}} \approx \frac{\mathcal{U}_n(\omega_{k,n} + h) - \mathcal{U}_n(\omega_{k,n} - h)}{2h} \quad (2.65)$$

which requires one additional forward propagator and one additional backward propagator per time step, yielding a computational cost similar to that of the spin-trajectory computation.

The auxiliary matrix formalism [57] evaluates exact derivatives by propagating an extended generator. The computation of first-order derivatives can be expressed as follows:

$$\begin{bmatrix} \mathcal{U}_n & \frac{\partial \mathcal{U}_n}{\partial \omega_{k,n}} \\ \mathbf{0} & \mathcal{U}_n \end{bmatrix} = \exp \left[-i \begin{bmatrix} \mathcal{H}_n & C_k \\ \mathbf{0} & \mathcal{H}_n \end{bmatrix} \Delta t_n \right], \quad (2.66)$$

where the propagator can be obtained via a Taylor expansion, and the derivatives are extracted from the resulting block matrix. For both the finite-difference and the auxiliary-matrix formalisms, the propagator derivatives $\partial \mathcal{U}_n / \partial \omega_{k,n}$ are never formed explicitly; instead, they are applied directly to the corresponding spin vector, which improves computational efficiency [58].

The GRAPE has been enhanced by combining with several optimization schemes. The quasi-Newton method updates the controls as

$$\omega^{s+1} = \omega^s - \alpha_s \mathbf{H}_s^{-1} \nabla \eta, \quad (2.67)$$

where $\nabla \eta$ is the fidelity gradient, $\mathbf{H}_s^{-1}$ is the approximate inverse Hessian matrix, calculated from a stack of history controls and gradients, the so-called L-BFGS algorithm [59]. The α_s can be calculated from a line search which finds an optimization step in the specified direction [60].

The Newton method updates according to

$$\omega^{s+1} = \omega^s - [\nabla^2 \eta]^{-1} \nabla \eta, \quad (2.68)$$

which achieves faster convergence but requires explicit Hessian evaluation, computationally prohibitive for systems with hundreds of controls. In practice, L-BFGS offers an effective balance between convergence rate and computational cost.

The convergence rates of L-BFGS and Newton–Raphson have been compared with steepest descent [61]. The computational cost of optimal control depends on the number of control channels, the system size, and the ensemble size, which can account for resonance offsets, field inhomogeneities, orientation distributions in solid-state NMR, etc. Figure 2.3 illustrates the wall-clock time of the LBFGS-GRAPe algorithm, as implemented in SPINACH [58]¹, for optimizing a single-spin excitation pulse.




Figure 2.3: Wall-clock time of the LBFGS-GRAPe algorithm in SPINACH. (a) Dependence on ensemble size with 400 time steps. (b) Dependence on the number of time steps with 1000 ensemble members. The pulse duration was 1 ms, covering a 20 kHz bandwidth with a nominal RF amplitude of 50 kHz and $\pm 20\%$ inhomogeneity. Iterations were terminated once the fidelity reached 0.995. The optimizations were performed on a workstation equipped with an AMD Ryzen 5950X 16-core processor (3.40 GHz base frequency) and 64 GB of RAM.

2.5 Finite element method

Considering the solution of a differential equation,

$$\mathcal{L}\phi = f, \quad (2.69)$$

where $\mathcal{L}$ is a differential operator, f is a source term, and ϕ is the unknown variable. The Ritz method [62] formulates a functional that is minimized with respect to the coefficients

¹Spinach is an open-source spin dynamics simulation library, available at <https://spindynamics.org>.

defining the trial solution. The functional for a trial function $\hat{\phi}$ is

$$\mathcal{F}(\hat{\phi}) = \frac{1}{2} \langle \mathcal{L}\hat{\phi}, \hat{\phi} \rangle - \langle f, \hat{\phi} \rangle, \quad (2.70)$$

where the trial function is expanded in terms of basis functions,

$$\hat{\phi} = \sum_i c_i v_i. \quad (2.71)$$

Minimization requires that the derivatives of $\mathcal{F}(\hat{\phi})$ with respect to the coefficients c_i vanish, yielding a set of equations for c_i .

The Galerkin method [63] instead enforces that the residual of the governing equation vanishes in a weighted integral sense:

$$R_i = \int_{\Omega} \omega_i (\mathcal{L}\hat{\phi} - f) d\Omega = 0, \quad (2.72)$$

where $\hat{\phi}$ is the trial function and the weighting functions are chosen as the basis functions themselves, i.e., $\omega_i = v_i$.

The central challenge in both methods is constructing suitable trial functions over the full domain, which becomes difficult in higher dimensions. The finite element method (FEM) addresses this by dividing the domain into subdomains (elements) and defining trial functions locally. The main steps of FEM are summarized below.

2.5.1 Finite elements

The first step in FEM is discretization of the domain, as illustrated in Fig. 2.4. In one dimension (1D), the domain is divided into line elements. In two dimensions (2D), triangular and rectangular elements are commonly used. In three dimensions (3D), tetrahedral and hexahedral elements are most frequently applied, with pyramidal and prismatic elements also possible. For complex geometries, triangular elements for 2D and tetrahedral elements for 3D are often preferred due to their flexibility.




Figure 2.4: Basic discretization elements, the black dots indicate the nodes of each element and their associated degrees of freedom to be solved. (a) Line element in one dimension. (b) Triangular and rectangular elements in two dimensions; the example shows a spirally curved resonator discretized using triangles. (c) Tetrahedral and hexahedral elements in three dimensions; the example shows a solenoid coil discretized using tetrahedra.

2.5.2 Basis functions

Within each finite element, basis (or shape) functions approximate the solution. The most common basis functions are Lagrange polynomials. For a 1D line element, the linear basis functions (Fig. 2.5a) are

$$N_1(x) = \frac{1-x}{2}, \quad N_2(x) = \frac{1+x}{2}, \quad x \in [-1, 1]. \quad (2.73)$$

For a quadratic line element (Fig. 2.5b), the basis functions are

$$N_1(x) = \frac{1}{2}x(x-1), \quad N_2(x) = 1-x^2, \quad N_3(x) = \frac{1}{2}x(x+1), \quad x \in [-1, 1]. \quad (2.74)$$

For a 2D triangular element (Fig. 2.5d), the linear basis functions are

$$N_1(x, y) = 1 - x - y, \quad N_2(x, y) = x, \quad N_3(x, y) = y, \quad (x, y) \in \{x \geq 0, y \geq 0, x + y \leq 1\}. \quad (2.75)$$

Interpolation with the Lagrange shape functions ensures continuity of function values across element boundaries, known as C^0 continuity. However, in problems such as beam




Figure 2.5: Basis functions for FEM elements. (a) Linear Lagrange, (b) quadratic Lagrange, and (c) cubic Hermite basis functions for a 1D element. (d) Linear Lagrange basis functions for a 2D triangular element.

bending, which are governed by higher-order differential equations, the first derivative must also be continuous, a requirement referred to as C^1 continuity. For such problems, Hermitian shape functions are employed to guarantee C^1 continuity.

The cubic Hermitian element uses four basis functions to perform cubic Hermite interpolation of the solution. As shown in Fig. 2.5c, the four shape functions for the 1D case are given by

$$\begin{aligned} N_1(x) &= \frac{1}{4}(1-x)^2(2+x), & N_2(x) &= \frac{1}{4}(1-x)^2(1+x), \\ N_3(x) &= \frac{1}{4}(1+x)^2(2-x), & N_4(x) &= \frac{1}{4}(1+x)^2(x-1), \quad x \in [-1, 1]. \end{aligned} \quad (2.76)$$

Compared to Lagrange interpolation, the degrees of freedom (DOFs) at each node for Hermite elements are doubled, representing both the function value and its first derivative.

Considering the Lagrange interpolation, the solution within an element e is approximated by

$$\hat{\phi}^e = \sum_{i=1}^n N_i^e \alpha_i^e, \quad (2.77)$$

where n is the number of element nodes, N_i^e is the element basis function associated with

node i , and α_i^e is the nodal coefficient. Each N_i^e is unity at its own node and zero at all other nodes, and vanishes outside its element.

2.5.3 Formulation of linear equations

Using the Galerkin method, the weighted residual for element e is

$$R_i^e = \int_{\Omega^e} N_i^e (\mathcal{L}\hat{\phi}^e - f) d\Omega^e = 0, \quad i = 1, \dots, n. \quad (2.78)$$

Substituting the expansion of $\hat{\phi}^e$ yields

$$\int_{\Omega^e} N_i^e \mathcal{L} \mathbf{N}^e d\Omega^e \alpha^e = \int_{\Omega^e} N_i^e f d\Omega^e, \quad i = 1, \dots, n, \quad (2.79)$$

which can be written in matrix form:

$$\mathbf{K}^e \alpha^e = \mathbf{b}^e, \quad (2.80)$$

where $\mathbf{K}^e$ is an $n \times n$ element stiffness matrix and $\mathbf{b}^e$ an $n \times 1$ load vector, given by

$$\begin{aligned} K_{ij}^e &= \int_{\Omega^e} N_i^e \mathcal{L} N_j^e d\Omega^e, \\ b_i^e &= \int_{\Omega^e} N_i^e f d\Omega^e. \end{aligned} \quad (2.81)$$

The element matrices are assembled into global equations by summing contributions from all elements. Entries corresponding to the same global node are added together, yielding

$$\mathbf{K} \alpha = \mathbf{b}, \quad (2.82)$$

where $\mathbf{K}$ is the global stiffness matrix and $\mathbf{b}$ the global load vector.

Equivalently, the system can be formulated using global basis functions:

$$\hat{\phi} = \sum_{j=1}^N N_j^g \alpha_j, \quad (2.83)$$

where N_j^g is the global basis function associated with node j , nonzero only in elements

connected to that node. Substitution into Eq. 2.72 gives

$$\sum_{j=1}^N \alpha_j \int_{\Omega} N_i^g \mathcal{L} N_j^g d\Omega = \int_{\Omega} N_i^g f d\Omega, \quad i = 1, \dots, N, \quad (2.84)$$

which has the same matrix form as Eq. 2.82, and the elements of the matrix and the load vector are given by

$$\begin{aligned} K_{ij} &= \int_{\Omega} N_i^g \mathcal{L} N_j^g d\Omega, \\ b_i &= \int_{\Omega} N_i^g f d\Omega. \end{aligned} \quad (2.85)$$




Figure 2.6: Banded structure of stiffness matrices in a 1D FEM model using (a) linear and (b) quadratic shape functions. The mesh has 10 elements, and each node has one degree of freedom.

The global stiffness matrix $\mathbf{K}$ is sparse and banded, with bandwidth depending on the order of the shape functions and DOFs per node. Figure 2.6 illustrates banded structures for 1D linear and quadratic FEM discretizations.

2.5.4 Adjoint analysis

In optimization problems, the stiffness matrix $\mathbf{K}$ may depend on control variables. Gradient-based optimization requires derivatives of an objective with respect to these controls. In FEM, the control derivatives are efficiently computed using the adjoint method.

Let the objective be

$$\mathcal{J} = \mathcal{J}(\alpha, \mathbf{c}), \quad (2.86)$$

where α is the FEM solution from $\mathbf{K}\alpha = \mathbf{b}$ and $\mathbf{c}$ are the control variables. Then

$$\frac{d\mathcal{J}}{d\mathbf{c}} = \frac{\partial \mathcal{J}}{\partial \mathbf{c}} + \frac{\partial \mathcal{J}}{\partial \alpha} \frac{d\alpha}{d\mathbf{c}}. \quad (2.87)$$

While computing the first term is straightforward, the second term involves

$$\frac{\partial \mathcal{J}}{\partial \alpha} \frac{d\alpha}{d\mathbf{c}} = \frac{\partial \mathcal{J}}{\partial \alpha} \mathbf{K}^{-1} \left(\frac{d\mathbf{b}}{d\mathbf{c}} - \frac{d\mathbf{K}}{d\mathbf{c}} \alpha \right). \quad (2.88)$$

Defining the adjoint vector λ via

$$\mathbf{K}^T \lambda = \left(\frac{\partial \mathcal{J}}{\partial \alpha} \right)^T, \quad (2.89)$$

the derivative becomes

$$\frac{d\mathcal{J}}{d\mathbf{c}} = \frac{\partial \mathcal{J}}{\partial \mathbf{c}} + \lambda^T \left(\frac{d\mathbf{b}}{d\mathbf{c}} - \frac{d\mathbf{K}}{d\mathbf{c}} \alpha \right). \quad (2.90)$$

This is known as the adjoint method, which reduces the computational cost of derivative evaluation to the same order as that of solving the original linear system. In numerical implementation, Eqs. 2.82 and 2.89 can be solved jointly using a pre-decomposition [64]. For a sparse banded square matrix encountered in quantum optimal control, the UMFPACK package within the SuiteSparse library [65] performs the LU factorization efficiently.

The FEM is a powerful and versatile tool in solving partial differential equations in engineering and theoretical modeling. Only a brief overview is presented here, as FEM is employed in two specific contexts in this thesis. First, it is employed in chapter 3 using COMSOL Multiphysics to simulate the magnetic fields of coil arrays, from which circuit parameters and field patterns are extracted. Second, FEM is used to solve quantum optimal control problems where spin evolution is governed by the Liouville–von Neumann equation. The derivation of stiffness matrices and adjoint-based gradient calculations for such problems is detailed in chapter 6.

Chapter 3

Multiphysics modeling for parallel NMR spectroscopy

Based on M. He, D. Faderl, N. MacKinnon, Y.T. Cheng, D. Buyens, M. Jouda, B. Luy, and J. G. Korvink, Communications Engineering, 3(1): 90, 2024.

3.1 Introduction

A straightforward way to increase the throughput of NMR spectroscopy is to detect multiple samples simultaneously via multiple channels, broadly termed parallel NMR [24]. Several previous studies utilize multiple RF coils for parallel excitation and reception [24, 25, 27]. Two broad approaches have been explored with regard to hardware design: The first method combines all coils into a single circuit with a shared tuning and matching (T&M) circuit. It separates signals using a pulsed gradient field, which reduces electronic components but collects noise from the entire coil array. This configuration is limited to cases [24, 30] which have a relatively high SNR. The second design deploys each coil with its own T&M circuit and receiver chain [25, 28], where each sample can be excited and detected without the use of gradient fields as a means to achieve effective parallelization.

A key challenge in NMR parallelism is ensuring exceptional electrical isolation among parallel channels [23], meaning that inter-coil coupling must remain sufficiently weak.

The method of suppressing coupling was recently discussed in a sample-centered shimming experiment [31], where the coil geometry was optimized to minimize the coupling effect. However, the RF shielding design is practically limited by space and probe complexity, so the coupling cannot be completely eliminated. This issue can lead to unwanted coupling of inductive coils, resulting in excitation pulse interference and signal transfer among multiple detection sites.

To investigate the impact of inter-channel coupling on the magnetic field distribution, spin dynamics, and the resulting NMR signals, a multiphysics simulation framework is required. This chapter introduces a digital twin of parallel NMR spectroscopy that combines full-wave electromagnetic and spin-dynamics simulations to analyze RF coupling effects and guide decoupling scheme design.

3.2 Multiphysics simulation workflow

Inspired by a simulation methodology for predicting MR noise [66], we developed a workflow for simulating parallel NMR, shown in Fig. 3.1. The workflow comprises three stages: magnetic field prediction, spin dynamics simulation, and pulse sequence optimization. In the first stage, the static magnetic field (B_0) and the RF field (B_1) within the samples are simulated based on well-defined coil geometries and electrical properties (e.g., permittivity and permeability). The resulting field maps are imported into a spin-dynamics module for predicting the spin evolution of multiple samples in the presence of inhomogeneous magnetic fields and inter-detector coupling. The third stage focuses on pulse optimization, where the strategies are adapted to account for inter-channel coupling.

To model the parallel excitation scenario (Fig. 3.2), synchronous pulses are applied across all channels. These RF pulses are amplified to a predetermined power level and delivered to the terminals of the T&M network. The resulting B_1 field within the samples is obtained by considering the simultaneous excitation of the RF coils. A composite spin system is then constructed to simulate spin evolution in multiple samples. This is achieved by combining the Hamiltonians of individual subsystems into a block-diagonal matrix, assuming negligible inter-sample coupling.

For parallel reception (Fig. 3.3), it is assumed that each channel acquires the FID synchronously, i.e., all detectors share the same sweep frequency and acquisition period.




Figure 3.1: A multiphysics simulation workflow consisting of three stages: magnetic field acquisition, spin dynamics calculation, and pulse sequence optimization. Each stage includes several sub-steps that are summarized in a box below and highlighted in an image above. DC denotes direct current, RF denotes radio frequency, B_0 is the static magnetic field, and B_1 is the radio-frequency magnetic field.




Figure 3.2: Modeling of the parallel excitation stage. Parallel pulses are processed to compute the coupled B_1 field, and a composite spin system is built to enable spin-dynamics simulations of multiple samples. All channels are excited with synchronous pulse sequences.

During acquisition, spin evolution is simulated to obtain the trajectory of the spin state, which is projected onto the coil sensitivity vector to compute the normalized FID, i.e., $\mathbf{m} \cdot \mathbf{b}^- = \langle \text{coil} | \rho \rangle$ [67]. The actual FID is calculated via the reciprocity principle [68]:

$$f = \int i\omega_0 |M| |B_1^-| e^{i\phi} \mathbf{m} \cdot \mathbf{b}^- dV. \quad (3.1)$$

Relaxation effects are neglected, and the magnetization magnitude $|M|$ is fixed at unity. The receive field is given by $|B_1^-| e^{i\phi} = B_x - iB_y$, computed in the same manner as the trans-

mit field B_1^+ , assuming a unit current in the relevant coil. In the receiver chain, the parallel FIDs contribute to the open-circuit voltages at the coil terminals, and an overall gain factor is applied to map these voltages to the T&M network terminals. The mixed signals at the output of the T&M network are subsequently separated using a signal decomposition module.




Figure 3.3: Modeling of the parallel reception stage. The magnetization is recorded by the RF coils, affected by coupling, and delivered to a post-processing module for signal decomposition. FID denotes free induction decay, μ is the magnetic moment, and B_0 is the static magnetic field.

In heteronuclear simulations, the relevant B_1 fields are computed for each resonance frequency. RF coupling between heteronuclear channels is neglected, as each channel is tuned to a distinct and isolated frequency. Signal acquisition for each nucleus is handled by an independent RF chain, each with its own coupling characteristics.

3.3 EM simulation

To detail the EM simulation without loss of generality, the water samples and solenoid RF coil arrays were selected for the EM simulations, and the background static magnetic field was chosen to be 11.74 T.

3.3.1 B_0 field simulation

The static B_0 field was simulated using the COMSOL `mfnc` module. As shown in Fig. 3.4a, the coil array was modeled as solenoidal copper windings, and the samples as cylindrical water phantoms, with a susceptibility of -9.035×10^{-6} . The background magnetic field

was applied through a magnetic flux boundary condition on the surrounding air domain, and the static field was computed using Gauss' law, $\nabla \cdot \mathbf{B} = 0$.

As shown in Fig. 3.4b, both the coil and the sample were discretized using tetrahedral elements, a standard choice for complex geometries. The resulting B_0 field was imported into MATLAB for post-processing via COMSOL Multiphysics with MATLAB LiveLink. For example, one cylindrical sample was divided into voxels along the axial (z , 8 values), radial (r , 3 values), and azimuthal (θ , 3 values) directions. Field values at voxel centers were obtained via linear interpolation from the FEM nodes. Finally, a zero-order shim was applied manually to remove mean field drift.




Figure 3.4: B_0 field simulation in the COMSOL mfn module. (a) Geometry of a four-solenoid coil array modeled with copper, with the samples simplified as cylindrical water phantoms. The background magnetic field was applied via a magnetic flux boundary condition on the surrounding air domain. (b) Enlarged view of one coil and one sample, both discretized with tetrahedral elements. The water sample was modeled as a cylinder with a radius of $r_s = 0.1$ mm and a length of $l_s = 2.5$ mm. The solenoid coil had a radius of $r_c = 0.5$ mm and a length of $l_c = 2.2$ mm, consisting of 20 turns of wire with a radius of 0.03 mm. The four coils were arranged on a circle of radius $R = 3.1$ mm in the yz -plane.

3.3.2 Tuning and matching

The schematic of a single-channel RF probe is shown in Fig. 3.5(a). A parallel NMR probe consists of multiple repetitions of this basic unit. Each coil is connected to a T&M circuit, which interfaces with the spectrometer. In transmit mode, the RF pulse generated by the synthesizer is phase-modulated by a timing device (pulse programmer) [41], amplified by

a power amplifier, and delivered to the coil via the T&M circuit. In receive mode, the FID signal detected by the coil is routed through the same T&M circuit to a low-noise amplifier. The amplified signal is subsequently mixed with a local oscillator for frequency conversion.

Figure 3.5(b) illustrates the circuit model of the coil array and the T&M network. The amplifiers are replaced by matched loads, under the assumption that they are all tuned and matched to the cable's characteristic impedance. As introduced in Section 2.1, we employ S-parameters to describe the interaction between the coil array and the T&M network. The forward and reverse traveling waves are denoted by a and b , respectively, and can incorporate both externally applied RF pulses and detected FID signals.




Figure 3.5: Schematic of the parallel detector. (a) One channel of the detector: the sample is placed inside a solenoid coil, which is tuned and matched, and connected to either the transmitter (Tx) or receiver (Rx) via a switch. PA denotes the power amplifier in the Tx path, and LNA denotes the low-noise amplifier in the Rx path. (b) Circuit model of the parallel detector using S-parameters. Amplifiers are substituted with matched loads, assuming all amplifiers are tuned and matched to the cable's characteristic impedance (50Ω). The symbols a and b indicate forward and reverse waves, respectively. The S_C , S_M , and S_{CM} denote the S-matrices of the coil array, the T&M network, and their combination.

The S-parameters of the four-solenoid coil array (S_C) were simulated using the COM-SOL emw module (Fig. 3.6). The RF simulation was performed with the frequency-domain solver, which computes the electric field by solving the wave equation with the finite ele-




Figure 3.6: B_1 -field simulation of a four-solenoid coil array in the COMSOL emw module. (a) Coil array geometry as in Fig. 3.4, with the air sphere surrounded by a perfectly matched layer (PML). (b) Enlarged view of one coil and one sample, both discretized with tetrahedral elements. Each solenoid coil was terminated and excited via a $50\ \Omega$ characteristic-impedance lumped port. PEC denotes a perfect electric conductor.

ment method. The operating frequency was set to 500 MHz, corresponding to the ^1H Larmor frequency at 11.74 T. The feed ports connected to the T&M network were modeled as $50\ \Omega$ lumped ports. To compute the full S-matrix, each port was sequentially excited while all others were terminated. For example, if port 1 was excited with the fundamental eigenmode E_1 , the reflection coefficient S_{11} is given by [69]:

$$S_{11} = \frac{\int_{\text{port1}} (E_1^c - E_1) \cdot E_1^* dS_1}{\int_{\text{port1}} E_1 \cdot E_1^* dS_1}, \quad (3.2)$$

and the transmission coefficients S_{j1} ($j \neq 1$) are given by

$$S_{j1} = \frac{\int_{\text{portj}} E_j^c \cdot E_j^* dS_j}{\int_{\text{portj}} E_j \cdot E_j^* dS_j}, \quad (3.3)$$

where E_j is the fundamental eigenmode on port j , and E_j^c is the computed field on port j .

Each coil can be independently tuned and matched using the diagonal elements of S_C , assuming sufficiently weak coupling. The $S_{C,21}$ in a typical solenoid array, as shown in Fig. 3.6, is on the order of $-60\ \text{dB}$; it decreases to below $-70\ \text{dB}$ for saddle-coil arrays at 45 MHz [31], and is expected to be even lower for stripline arrays of comparable geometric

size, as the latter confine the EM field more effectively within individual coils. However, if strong coupling degrades the match (e.g., $S_{CM,11} > -20$ dB), the capacitor values in the T&M network are optimized using the Nelder–Mead simplex algorithm [70], a derivative-free method for unconstrained optimization implemented in MATLAB as `fminsearch`.

After the T&M procedure, the coupling matrices introduced in Section 2.1 can be evaluated. The excitation coupling matrix $\mathbf{F}$ is defined as

$$\mathbf{F} = \frac{1}{\sqrt{Z_0}} (\mathbf{S}_C^{-1} - \mathbf{E}) (\mathbf{E} - \mathbf{S}_C \mathbf{S}_{11})^{-1} \mathbf{S}_C \mathbf{S}_{12}, \quad (3.4)$$

where Z_0 is the characteristic impedance and $\mathbf{E}$ is the identity matrix. For a two-coil array, the coupling strength can be quantified as

$$cs = \left| \frac{F_{21}}{F_{11}} \right|. \quad (3.5)$$

The reception coupling matrix $\mathbf{G}$ is given by

$$\mathbf{G} = \frac{1}{2} \mathbf{S}_{21} (\mathbf{E} - \mathbf{S}_C \mathbf{S}_{11})^{-1} (\mathbf{E} - \mathbf{S}_C). \quad (3.6)$$




Figure 3.7: The coupling strength (cs) as a function of the coil distance in a two-solenoid array. R is the distance between the centers of two solenoid coils positioned parallel to each other, and $r_c = 0.5$ mm is the coil radius, the same as in Fig. 3.4.

A two-solenoid coil array with the same dimensions as that shown in Fig. 3.4 was used to analyze inter-coil coupling. Figure 3.7 presents the coupling strength (cs) as a function

of the coil separation, showing a pronounced increase in coupling as the two coils approach each other, which requires optimization of tuning and matching.

Figure 3.8(a) shows the reflection at the two free ports as a function of coupling strength. In the case of individual T&M, both $S_{CM,11}$ and $S_{CM,22}$ increase with stronger coupling. Once the reflection exceeds -20 dB, optimization is applied to keep it below this threshold.

Figure 3.8(b) plots the Frobenius norm of the difference between the normalized matrices $\mathbf{F}_{\text{norm}}$ and $\mathbf{G}_{\text{norm}}$ as coupling increases. In both the individual and optimized T&M cases, $\|\mathbf{F}_{\text{norm}} - \mathbf{G}_{\text{norm}}\|$ remains on the order of 10^{-3} . Small fluctuations arise from the stochastic nature of the optimization. Under the assumption of identical coil geometries and identical tuning and matching (T&M) capacitance networks across all channels, it is verified that $\mathbf{G} = \sqrt{Z_0}/2 \cdot \mathbf{F}$. This implies that both matrices yield identical results once normalized [35].




Figure 3.8: Coupling analysis of a two-solenoid coil array. (a) Reflection coefficients at the two free ports, $S_{CM,11}$ (circles) and $S_{CM,22}$ (crosses), plotted as a function of coupling strength (cs). (b) Frobenius norm of the difference between $\mathbf{F}_{\text{norm}}$ and $\mathbf{G}_{\text{norm}}$ as a function of coupling strength. The dashed lines separate two T&M strategies: individual T&M (left) and optimization (right).

When the number of optimization variables becomes large, the Nelder–Mead simplex method may face convergence issues. To reduce dimensionality, the symmetric geometry of the coil array can be exploited by using identical capacitance sets for all channels. This significantly reduces the number of optimization parameters. Figure 3.9(a,b) shows results for a 16-solenoid array using identical capacitance sets: reflections at each port are suppressed below -30 dB. However, this comes at the cost of increased inter-coil coupling, which must be mitigated by active decoupling, as discussed in chapter 4. In contrast, using individual capacitance sets expands the parameter space and increases the risk of conver-

gence to local minima, as demonstrated in Fig. 3.9(c,d).




Figure 3.9: Optimization of the T&M network for a 16-solenoid coil array. (a,c) Convergence of the objective function $\text{obj} = \sum_n |S_{CM,nn}|^2$ using MATLAB's `fminsearch`, for identical (a) and individual (c) capacitance sets. (b,d) Reflection and total coupling at each port before (dashed) and after (solid) optimization, for identical (b) and individual (d) capacitance sets. The total coupling at port n is defined as $\sum_{m \neq n} |S_{CM,nm}|^2$.

3.3.3 B_1 field calculation

The B_1 field was calculated using a full-wave electromagnetic simulation, following the workflow in Fig. 3.10. The COMSOL simulation provides the port parameters, which are used to design the T&M network. Pulse parameters are then applied to compute the resulting coil currents. These currents are mapped onto the prototype B_1 field exported from the simulation to reconstruct the final field distribution. For a given frequency, full-wave simulation needs to be executed only once; subsequent evaluations for different pulses can be performed efficiently in MATLAB.

The B_1 field at a given position under parallel excitation can be expressed as the linear




Figure 3.10: Workflow for B_1 field simulation. S , I , and V denote S-parameters, currents, and voltages, respectively.

superposition of contributions from individual coils [71]:

$$B_1 = \sum_n B_{1n} = \sum_n c_n I_n, \quad (3.7)$$

where $B_{1n} = c_n I_n$ and I_n is the current in the n -th coil. Thus, the B_1 field under excitation $\mathbf{p}$ can be written as

$$\mathbf{B}_1 = \mathbf{T} \cdot \begin{pmatrix} I_1 \\ I_2 \\ \vdots \\ I_N \end{pmatrix} = \mathbf{T} \cdot \mathbf{F} \cdot \begin{pmatrix} p_1 \\ p_2 \\ \vdots \\ p_N \end{pmatrix}, \quad (3.8)$$

where $\mathbf{F}$ maps excitation $\mathbf{p}$ to coil currents, and $\mathbf{T}$ maps coil currents to B_1 field values. The number of rows of $\mathbf{T}$ depends on the number of voxels considered.

For M voxels across N samples, $\mathbf{F}$ is an $N \times N$ matrix (computed via Eq. 3.4), and $\mathbf{T}$ is an $M \times N$ matrix constructed from a complete port sweep. Considering one column of

$\mathbf{T}^\top$ corresponding to voxel m , we have

$$\begin{pmatrix} I_{11} & I_{12} & \dots & I_{1N} \\ I_{21} & I_{22} & \dots & I_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ I_{N1} & I_{N2} & \dots & I_{NN} \end{pmatrix} \cdot \mathbf{T}_m^\top = \begin{pmatrix} B_1 \\ B_2 \\ \vdots \\ B_N \end{pmatrix}, \quad (3.9)$$

so that $\mathbf{T}_m^\top = \mathbf{I}^{-1} \cdot \mathbf{B}$.

For a four-solenoid array (Fig. 3.6), the $\mathbf{F}$ matrix amplitude was computed as

$$|\mathbf{F}| = \begin{pmatrix} 0.4266 & 0.1084 & 0.0743 & 0.1089 \\ 0.1084 & 0.4265 & 0.1087 & 0.0743 \\ 0.0743 & 0.1088 & 0.4265 & 0.1088 \\ 0.1090 & 0.0744 & 0.1089 & 0.4266 \end{pmatrix}. \quad (3.10)$$

Selecting 72 voxels from each of the four samples yields $M = 288$. Solving Eq. 3.9 for each voxel produces the full $\mathbf{T}$ matrix. To obtain a representative 4×4 matrix, voxel rows belonging to the same sample are averaged, giving:

$$|\mathbf{T}| = 10^{-5} \cdot \begin{pmatrix} 1033 & 5.9902 & 2.0330 & 5.4047 \\ 5.9893 & 1033 & 5.4036 & 2.0342 \\ 2.0520 & 5.3936 & 1034 & 5.9281 \\ 5.3969 & 2.0536 & 5.9265 & 1033 \end{pmatrix}. \quad (3.11)$$

The ratio of off-diagonal to diagonal elements quantifies coupling. The circuit-level coupling matrix $\mathbf{F}$ exhibits coupling on the order of 10^{-1} , whereas the field spillover matrix $\mathbf{T}$ shows weaker coupling, on the order of 10^{-3} . This indicates that inductive coupling dominates over direct B_1 field spillover.

Figure 3.11 shows the B_1 field distribution when only the first channel is excited, corresponding to an input vector $\mathbf{p} = (1, 0, 0, 0)^\top$. Both amplitude and phase information are preserved from the full-wave simulation. The incident voltage vector $\mathbf{p}$ is applied at the T&M network terminal, assuming that the amplifier delivers the required power precisely to the RF probe. Hence, the amplifier must operate within its linear regime, and care should be taken when considering high-power pulses [72].




Figure 3.11: B_1 field of a four-solenoid coil array. (a) Amplitude and (b) phase of the transmit B_1 field ($B_1^+ = B_x + iB_y$) in the xoz cross-plane inside the sample. Only the first channel is excited from the T&M network terminal, i.e., $\mathbf{p} = (1, 0, 0, 0)^T$.

It should also be noted that the RF simulation, carried out in the frequency domain, provides steady-state results and does not account for transient responses of finite- Q circuits. This limitation has been addressed in previous studies by simulating the transient response [73] or by experimentally measuring the impulse response function [74]. To explicitly model coil currents and B_1 fields in the time domain, one can compute the transient effects using the impulse response function under the linear time-invariant (LTI) system approximation [74]. In practice, the transmission channel bandwidth of typical NMR hardware substantially exceeds the spectral width of the applied pulses. As a result, transient effects are generally negligible during pulse design, and it is customary to omit them.

3.4 Spin dynamics simulation

If inter-sample coupling is negligible, the evolution of multiple samples can be conveniently treated using isolated sub-spin systems accommodated by a block-diagonal Hamiltonian. The only mechanism by which samples can interact is through radiation damping, which we show to be sufficiently weak to neglect.

Radiation damping is typically interpreted as the process by which the FID signal in the coil induces an additional RF field B_{rad} , which drives the magnetization back toward equilibrium [75]. Because RF signals can transfer between channels, the FID generated

by one sample may induce a B_{rad} field in other samples through inter-coil coupling. If the radiation-damping field strength is directly proportional to the FID current in the coil, then the sample–sample coupling ratio is quantified by the ratio of induced currents.

Let I_1 and I_2 denote the FID currents in the primary coil and the coupled coil, respectively. The ratio I_2/I_1 can be calculated using the RF chain model introduced in Section 3.3.2, by applying an initial open-circuit voltage $V_{\text{1on}} = [1, 0]^T$ and computing the coil current I_C , assuming matched loads in the T&M network. The strength of B_{rad} depends on the coil filling factor η and quality factor Q , as well as on the sample concentration [75, 76]. It has been shown that a water sample produces a maximum B_{rad} of 16 Hz in a coil with $\eta Q = 10$ at 11.74 T [75]. Optimized microcoils, such as striplines, can achieve ηQ values on the order of several tens [77], thereby generating a higher B_{rad} . Assume the primary radiation-damping field is 50 Hz, the coupled radiation-damping field is given by

$$B_{\text{rad}} = 50 \text{ Hz} \cdot \frac{I_2}{I_1}, \quad (3.12)$$

as illustrated in Fig. 3.12. For a strong coupling strength of 10%, the coupled radiation-damping field is approximately 5 Hz, which can safely be neglected.




Figure 3.12: Quantification of sample–sample coupling via the radiation-damping effect.

Since each sample possesses its own spin state and Hamiltonian, and inter-sample coupling is neglected, a composite spin system is defined by stacking the individual Hamiltonians and spin state vectors. The total Hamiltonian is therefore a block-diagonal matrix,

where each block corresponds to a single sample. The spin state vectors are arranged consistently, as illustrated in Eq. 3.13. This representation allows each sample to evolve independently within a unified composite system.

$$\frac{d}{dt} \begin{pmatrix} \rho_1 \\ \rho_2 \\ \rho_3 \\ \vdots \end{pmatrix} = -i \begin{pmatrix} \mathbf{H}_1 & & & \\ & \mathbf{H}_2 & & \\ & & \mathbf{H}_3 & \\ & & & \ddots \end{pmatrix} \begin{pmatrix} \rho_1 \\ \rho_2 \\ \rho_3 \\ \vdots \end{pmatrix}. \quad (3.13)$$

Spin dynamics calculations were performed using **SPINACH**, a MATLAB-based library providing extensive functionality for NMR, optimal control, and related applications [58]. In the MATLAB implementation, variable arrays were defined for the following physical quantities:

- **Spin states and operators**

Unique indices were assigned to isotopes from different samples. Spin states and operators were generated using built-in **SPINACH** functions, for example:

```
rho = state(spinSystem, 'Lz', isoInd{m});
```

defines the spin state ρ as the $\mathbf{I}_z$ operator acting on the spins specified by `isoInd{m}`. Similarly,

```
Lx = state(spinSystem, 'Lx', isoInd{m});
```

defines the spin operator $\mathbf{I}_x$ for the same set of nuclei.

- **Hamiltonian**

The composite spin system was constructed, and the Zeeman Hamiltonian $\mathbf{H}_z$ was generated using

```
Hz = hamiltonian(assume(spinSystem, 'labframe', 'zeeman'));
```

where the static field was 1 T in `spinSystem` definition. The spin–spin coupling Hamiltonian $\mathbf{H}_c$ was generated using

```
Hc = hamiltonian(assume(spinSystem, 'labframe', 'coupling'));
```

In this syntax, **SPINACH** includes all spin–spin coupling terms. For liquid-state NMR simulations, however, only J-couplings contribute.

The voxel-wise internal Hamiltonian was then calculated as

$$\mathbf{H}_{0k} = \mathbf{H}_c + \mathbf{H}_z \cdot B_{0k},$$

where B_{0k} is the local static field at voxel k . All voxel-specific Hamiltonians were stored in a cell array.

- **Pulses**

For a given excitation vector $\mathbf{p}$ applied at the T&M network terminal, the local B_1 field was computed for each voxel and stored in a cell array. This local field determined the amplitude and phase of the voxel-wise RF pulse. The corresponding control Hamiltonian was added to the internal Hamiltonian. For voxel k of sample n , the total Hamiltonian was defined as:

$$H\{n, k\} = H_0\{n, k\} + C_x\{k\} * L_x\{n\} + C_y\{k\} * L_y\{n\};$$

where $L_x\{n\}$ and $L_y\{n\}$ are spin operators for sample n , and $C_x\{k\}$, $C_y\{k\}$ are real pulse amplitudes determined by the local B_1 field.

3.5 Wall clock time

The wall-clock time for the multiphysics simulations is summarized in Table 3.1. Simulations were performed on a workstation equipped with an AMD Ryzen 5950X 16-core processor (3.40 GHz base frequency) and 64 GB of RAM. Each sample was discretized into 72 voxels. For spin dynamics, each channel was modeled as a single ^1H spin subjected to a hard-pulse excitation and acquisition at 500 MHz.

As the number of channels increases, the RF simulation becomes the dominant computational cost, since each lumped port must be swept to compute the full $n \times n$ scattering matrix. For wall-clock times associated with large-scale spin system simulations, see the SPINACH literature [78].

3.6 RF coupling in parallel excitation

In parallel NMR experiments, multiple channels apply RF pulses simultaneously. Poor electrical isolation causes coupling, whereby the B_1 field generated by one channel excites

Channel number	DC simulation		RF simulation		Process B field	Spin dynamics	Total time
	DOF	Time	DOF	Time			
1	213497	6	969449	38	3.5	1.3	48.8
2	304195	13	2172051	158	21.5	1.6	194.1
4	587519	19	2830773	433	52.0	1.8	505.7
8	927966	24	5501691	2202	377.8	3.3	2607.1
16	2149472	62	7771127	10523	2151	10.4	12746 (3.5 h)

Table 3.1: Wall-clock time (seconds) for parallel NMR simulations. DOF denotes the number of degrees of freedom in the finite element simulation. DC denotes direct current, RF denotes radio frequency.

neighboring samples, leading to undesired manipulations. Consequently, conventional pulse calibration fails in parallel setups, since distortions arise not only from individual coil imperfections but also from leakage between channels.

The impact of inter-coil coupling on excitation efficiency was analyzed, in particular, the loss of fidelity as the coupling strength increases. In Fig. 3.13(a), by subjecting one ^1H spin to a single 90° pulse in both channels, a remarkable decrease in fidelity was observed as the coupling strength increased, indicated by the proximity of the two solenoid coils. The small discrepancy between the two lines arises from the distinct effects of the coupled pulse on each channel, as the proton chemical shifts in channel 1 and channel 2 were set to -2 ppm and 5 ppm, respectively.




Figure 3.13: Excitation fidelity $\langle \rho | I_x \rangle$ in parallel excitation. (a) Fidelity of parallel hard pulses as a function of coupling strength ($cs = |F_{12}/F_{11}|$). Each pulse has a 90° flip angle and a duration of $20 \mu\text{s}$. (b) Fidelity of optimized control pulses applied in parallel under selected coupling strengths. ν_0 and ν_1 denote resonance offsets and RF amplitudes, respectively. RF amplitudes were determined from the local B_1 field under fixed excitation power, resulting in slight variations with coupling strength. Channels are labeled chn. 1 and chn. 2.

Optimal control pulses have long been employed to compensate for inhomogeneities and offsets in single-coil experiments, yielding composite pulses with high transfer efficiency [55, 79, 80, 81] and even approaching physical excitation limits [82, 83, 84]. To investigate the effect of coupling on optimal control pulses, two broadband excitation pulses were generated using the GRAPE algorithm implemented in SPINACH (Section 2.4). The OC model considered an ensemble single-spin system incorporating distributed B_0 and B_1 field values as well as resonance offsets. Each cylindrical sample was divided into 72 voxels, and voxel-wise B_0 and B_1 fields extracted from the EM simulation were aligned to build Hamiltonians for the ensemble. The resonance offset was sampled over 5 kHz in 20 segments and combined with the field values.

For the shaped pulse, the duration was set to 2 ms and discretized into 100 equal time slices. Following a common experimental practice, the amplitude was fixed at unity, and only the phases were optimized, resulting in 100 control variables. The target operation was to transfer magnetization from $\mathbf{I}_z$ to $\mathbf{I}_x$. The ensemble-averaged fidelity was defined as

$$\bar{\eta} = \frac{1}{n_{\text{rf}} n_{\text{off}}} \sum_{n_{\text{rf}}} \sum_{n_{\text{off}}} \langle \rho | I_x \rangle, \quad (3.14)$$

and the LBFGS-GRAPE algorithm was employed, with optimization terminated once $\bar{\eta} > 0.995$.

Fig. 3.13(b) presents the excitation fidelity of optimal control pulses applied in parallel. In the absence of coupling, the optimized OC pulse achieved near-perfect fidelity. However, fidelity decreased significantly with increasing coupling strength, in a manner similar to the hard-pulse case. Furthermore, the two channels exhibited different sensitivity to coupling, which is attributed to the distinct pulse shapes of the two channels.

Although OC pulses offer precise spin manipulation, RF coupling distorts their waveforms by inducing additional currents in the coils. This reduces robustness to resonance offsets and field inhomogeneities, resulting in decreased transfer efficiency. The detrimental effect of coupling is heightened at higher pulse power, for instance, the functionality of the 10 kHz pulse is completely compromised at a coupling strength of 0.24, as shown in Fig. 3.14.




Figure 3.14: Transfer fidelity of optimized control pulses applied in parallel at different nominal pulse amplitudes, for a fixed coupling strength of $cs = 0.24$. Pulse duration was 2 ms, discretized into 100 time steps. Each pulse transfers a single spin from $\mathbf{I}_z$ to $\mathbf{I}_x$, accounting for resonance offsets and B_1 inhomogeneity. Channels are labeled chn. 1 and chn. 2. ν_0 and ν_1 denote resonance offset and RF amplitude, respectively.

3.7 RF coupling in parallel acquisition

3.7.1 Modeling of FID mixing

Coupling also affects the reception stage, causing FID cross-talk [27]. A coupling model is therefore required to describe how multiple FIDs combine during parallel acquisition. Consider N homonuclear channels operating in synchronous acquisition. Each coil records not only its own FID but also contributions from other samples, such that

$$s_i = f_{ii} + \sum_{j \neq i} f_{ij}, \quad (3.15)$$

where f_{ij} is the FID in channel i arising from sample j .

An overview of signal composition is shown in Fig. 3.15. The gain of the combined coil and T&M network, given in Eq. 3.6, relates the output signal to the open-circuit voltages




Figure 3.15: Signal transmission in the reception stage. The f_{ij} is the FID in channel i from sample j , s_i is the total FID in channel i , n_i is the input noise in channel i , $\mathbf{G}$ is the reception gain matrix, $\mathbf{x}$ are the detected signals, and $\mathbf{n}'$ are the modified noises.

induced by the FIDs. Neglecting noise for clarity, the detected signals can be written as

$$\begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{pmatrix} = \mathbf{G} \cdot \begin{pmatrix} f_{11} \\ f_{22} \\ \vdots \\ f_{NN} \end{pmatrix} + \mathbf{G} \cdot \begin{pmatrix} \sum_{j \neq 1} f_{1j} \\ \sum_{j \neq 2} f_{2j} \\ \vdots \\ \sum_{j \neq N} f_{Nj} \end{pmatrix}. \quad (3.16)$$

Expanding the first row,

$$x_1 = G_{11} \left(f_{11} + \sum_{j \neq 1} f_{1j} \right) + G_{12} \left(f_{22} + \sum_{j \neq 2} f_{2j} \right) + \cdots + G_{1N} \left(f_{NN} + \sum_{j \neq N} f_{Nj} \right). \quad (3.17)$$

In the weak-coupling regime, $G_{ij} \ll G_{ii}$ and the cross-coupling terms satisfy $f_{ij} \ll f_{ii}$; for example, $f_{ij} \approx 10^{-3} f_{ii}$ as given in Eq. 3.11 for the solenoid coil array. Hence, the contributions arising from their combinations, i.e., $G_{ij} f_{ij}$ ($i \neq j$), can be disregarded. This leads to

$$x_1 \approx G_{11} f_{11} + (G_{12} + R_{12} G_{11}) f_{22} + (G_{13} + R_{13} G_{11}) f_{33} + \cdots + (G_{1n} + R_{1N} G_{11}) f_{NN}, \quad (3.18)$$

where

$$R_{ij} = \frac{f_{ij}}{f_{jj}} = \frac{\int_{\text{sample } j} B_{ij} dV_j}{\int_{\text{sample } j} B_{jj} dV_j}, \quad f_{ij} = \int_{\text{sample } j} i\omega_0 M_j B_{ij} dV_j. \quad (3.19)$$

The Eq. 3.19 assumes uniform magnetization, e.g., $M_j = M_x$ for a given sample during acquisition. Here, R_{ij} denotes the ratio of the receiving B_1 fields between coils i and j when detecting the j -th sample, with the reception phase from the sample to the coil included in the receive B_1 field. Extending this derivation to all channels leads to

$$\begin{aligned} G_{11}f_{11} + (G_{12} + R_{12}G_{11})f_{22} + (G_{13} + R_{13}G_{11})f_{33} + \dots + (G_{1N} + R_{1N}G_{11})f_{NN} &= x_1, \\ (G_{21} + R_{21}G_{22})f_{11} + G_{22}f_{22} + (G_{23} + R_{23}G_{22})f_{33} + \dots + (G_{2N} + R_{2N}G_{22})f_{NN} &= x_2, \\ &\dots, \\ (G_{N1} + R_{N1}G_{NN})f_{11} + (G_{N2} + R_{N2}G_{NN})f_{22} + \dots + G_{NN}f_{NN} &= x_N. \end{aligned} \quad (3.20)$$

Marking $G'_{ii} = G_{ii}$, $G'_{ij} = G_{ij} + G_{ii}R_{ij}$ and picking up the noise gives

$$\mathbf{x} = \mathbf{G}' \cdot \mathbf{f} + \mathbf{n}', \quad (3.21)$$

where $\mathbf{f} = (f_{11}, f_{22}, \dots, f_{NN})^T$. Equation 3.21 indicates that the measured signals are the linear combination of the primary FIDs and the weak mutual correlations between the recorded signals in $\mathbf{x}$ can be disregarded. In the digital twin, the gain matrix can be computed, allowing $\mathbf{f}$ to be estimated as $\hat{\mathbf{f}} = \mathbf{G}'^{-1}\mathbf{x}$ under the assumption of a reasonable SNR.

3.7.2 Four-detector NMR simulation

Considering a four-detector NMR setup at 11.74 T magnet, each detector was double-resonant on ^1H and ^{13}C , with the $^1\text{H}/^{13}\text{C}$ channel pairs simultaneously excited and received under the assumption of perfect isolation between heteronuclear channels. The parallel spectra were simulated using the developed Multiphysics workflow. First, a direct-current FEM simulation of the four-solenoid array was performed to obtain the B_0 field, and an RF simulation of the same geometry was carried out to extract the S -parameters and prototype B_1 field. Subsequently, RF chain modeling was performed to construct the coupling




Figure 3.16: Signal decomposition of eight-channel NMR spectra using the gain matrix. (a) and (c) show the simulated coupled spectra, while (b) and (d) present the corresponding decomposed spectra. The “signal” denotes the primary component, and “interf.” indicates the interfering coupled component. An uncoupled ^1H – ^{13}C two-spin system was assigned for the sample in each detector. Synchronous 90° pulses of uniform duration ($20\ \mu\text{s}$) were applied to all channels, with the average RF amplitude set to 12.5 kHz.

matrices $\mathbf{F}$ and $\mathbf{G}$, and the resulting circuit parameters were combined with the prototype B_1 field to obtain the B_1^+ and B_1^- field patterns. The RF simulations and network modeling were performed independently for ^1H and ^{13}C .

The composite spin system was implemented in SPINACH, using the spherical tensor formalism as the basis operators and Liouville space for the algebra formalism. An uncoupled ^1H – ^{13}C two-spin system was assigned for the sample in each detector, and only

Zeeman interactions were considered. To account for spatial inhomogeneity, voxel-wise Hamiltonians were assembled in a cell array using the corresponding B_0 and B_1 field values.

During excitation, synchronous 90° pulses of uniform duration ($20 \mu\text{s}$) were applied to all channels, and spin evolution was performed voxel by voxel where the local B_1 was used for the single-pulse propagation. During reception, acquisition parameters were chosen according to the spectral bandwidth of each nucleus, i.e., 10 ppm for ^1H and 100 ppm for ^{13}C . A FID f_{ij} was obtained by summing voxel contributions for sample j according to the reciprocity principle (Eq. 3.1), and the detected signals $\mathbf{x}$ were computed by Eq. 3.16. Finally, the modified gain matrix $\mathbf{G}'$ was computed and applied to separate the primary and interfering components.

Fig. 3.16(a,c) shows the coupled signals containing interference. The interfering components are weaker in the ^{13}C channel, owing to the reduced coupling strength at lower frequency. Fig. 3.16(b,d) presents the decomposed signals, demonstrating a complete separation of the desired signals from cross-talk.

3.7.3 Gain-matrix decomposition error

Since the modified coupling matrix $\mathbf{G}'$ is derived from the reception gain $\mathbf{G}$ and the sensitivity ratio, the 2-channel decomposition error arising from random variations in B_1 and $\mathbf{G}$ was subsequently simulated.

B_1 field shift.

Since the B_1 field contributes to signal decomposition but cannot be predicted exactly, we evaluated the decomposition behavior under B_1 field perturbations. Specifically, FIDs were calculated using the correctly predicted B_1 field, and the B_1 values were then deliberately shifted during signal decomposition. The perturbed field was defined as

$$B'_1 = B_1 \cdot (1 + \alpha\chi), \quad (3.22)$$

where χ is a random variable drawn from a standard Gaussian distribution, $\chi \sim \mathcal{N}(0, 1)$, and α represents the shift amplitude. Figure 3.17 shows the decomposition results for two channels under different B_1 perturbations. Noticeable residual interference appears only




Figure 3.17: Two-channel signal decomposition results under different B_1 field shift amplitudes (α). The primary component (“signal”) and the interfering component (“interf.”) are shown together.

when $\alpha > 100$, corresponding to a completely erroneous prediction. This result indicates that the contribution of the B_1 field to signal decomposition is minimal.

Gain matrix perturbation

We next examined the decomposition efficiency under perturbations of the network gain matrix $\mathbf{G}$. The detected signals $\mathbf{x}$ were first calculated using the correctly predicted gain, and $\mathbf{G}$ was then deliberately perturbed during signal decomposition. Each element was modified according to

$$G'_{ij} = G_{ij} \cdot (1 + \alpha\chi), \quad (3.23)$$

where χ is a random variable following a standard Gaussian distribution, $\chi \sim \mathcal{N}(0, 1)$, and α denotes the perturbation amplitude.

Figure 3.18 shows the decomposition results for two channels under different perturbation levels. At $\alpha = 0.05$, no distinct interference peaks are visible in the decomposed




Figure 3.18: Two-channel signal decomposition results under different perturbation amplitudes (α) of the gain matrix $\mathbf{G}$. The primary component (“signal”) and the interfering component (“interf.”) are shown together.

signals. When α increases to 0.1, interference peaks become pronounced. These results indicate that decomposition is sensitive to the accuracy of $\mathbf{G}$, as the circuit-level combination contributes dominantly to signal mixing.

3.8 Conclusion

In this chapter, a multiphysics simulation workflow was developed to investigate inter-channel coupling in parallel NMR detectors. The simulations show that the robustness of optimal-control pulses degrades once the coupling strength reaches several tens of percent and the RF amplitude approaches 5 kHz. In addition, FID coupling can be modeled as a linear combination determined by a coupling matrix. The wall-clock time of the multiphysics simulations increases with the number of channels and is dominated by the B_1 field computation, reaching several hours for systems with 16 or more channels.

By integrating electromagnetic simulations of both static and RF fields with spin-dynamics calculations, the workflow enables a systematic characterization of RF coupling effects in

both the parallel excitation and reception. This capability is useful for assessing parallel probes and optimizing multi-detector geometries prior to physical construction.

The present model does not include the influence of cable connections between the probe and the spectrometer, which may affect the effective coupling strength. However, because parallel excitation and reception share the same RF chain, their coupling behavior can be approximately described by a common matrix, i.e., $\mathbf{G}_{\text{norm}} \approx \mathbf{F}_{\text{norm}}$. In practice, the exact coupling matrix can be determined by analyzing coupled signals in calibration experiments. The RF decoupling strategies in parallel NMR detection based on the coupling matrix are discussed in Chapter 4.

Chapter 4

Radio frequency decoupling for parallel NMR spectroscopy

Based on M. He, D. Faderl, N. MacKinnon, Y.T. Cheng, D. Buyens, M. Jouda, B. Luy, and J. G. Korvink, Communications Engineering, 3(1): 90, 2024.

The last chapter introduced a multiphysics simulation capable of modeling excitation distortions and signal compositions arising from inter-channel coupling. Building on this workflow, this chapter presents two decoupling strategies: one for the excitation stage and another for the reception stage. Section 4.2 describes a parallel pulse design method to mitigate inter-channel coupling during excitation. Section 4.3 proposes the use of blind source separation to extract individual FIDs from composite signals acquired with multiple detectors, thereby serving as a decoupling strategy at the reception stage. Finally, the excitation-stage decoupling approach is experimentally validated in Section 4.4 using a two-channel parallel spectroscopy experiment.

4.1 Introduction

In the realm of parallel detection, the integration of multiple receive coils represents a milestone within the MRI community [85, 86]. This approach, known as parallel imaging, has been widely adopted for its ability to enhance SNR [87] and reduce acquisition time

[88]. Its development also fostered the use of multiple transmit coils in so-called parallel transmission (pTx) [89], aimed at improving RF excitation performance. Applications include improving B_1 homogeneity [90], minimizing specific absorption rate [91], and RF shimming [92]. Spatially selective multidimensional excitation was further advanced by optimizing RF pulses to reduce duration while addressing the spatial sensitivity of transmit elements [93].

Because the transmit B_1 field strongly depends on the experimental platform, subject-specific optimization typically employs B_0 and B_1 maps as prior knowledge [94, 95, 96]. To avoid extensive field mapping, universal pulses have been developed, demonstrating robustness to subject-dependent inhomogeneities [97, 98]. These universal pulses, designed based on a database of B -field maps from representative subjects, have been further refined using the GRAPE algorithm. For instance, the simultaneous optimization of RF and gradient waveforms can address B_1 heterogeneity [99] and off-resonance effects [52].

In contrast, the parallel NMR concept is designed to increase throughput, with each coil exciting a distinct sample via its localized B_1 field. The central goal is independent operation of each channel, making coupling-induced parallel transmission a hardware imperfection that must be corrected. This is distinct from strongly coupled pTx in MRI, where pTx is intentionally employed to enhance excitation.

Mutual coupling within MRI transmit arrays has been extensively studied, yielding diverse decoupling strategies. These approaches including employing decoupling capacitances [100], passive networks [101, 102], or preamplifier decoupling via impedance mismatch [103]. A recent review summarized these methods [104]. While preamplifier decoupling requires substantial effort on system integration, active decoupling has attracted greater interest. One example entails controlling the currents in eight-channel transverse electromagnetic transmit coils to attain independent transmit sensitivities, with control coefficients designed based on the impedance matrix of the coil array [105, 106].

The active decoupling concept might be combined with the shaped RF pulses, which have become the principal approach for spin manipulation. Here, we explore parallel excitation by designing cooperative pulses that mitigate inter-channel coupling without requiring B -field mapping. The cooperative pulses produce the excitation, equivalent to that generated by each channel in the absence of coupling. This approach extends earlier concepts of non-identical, but cooperatively acting pulses introduced for heteronuclear Hartmann–

Hahn sequences [107], added-scan cooperativity [108], and consecutive scans [109].

4.2 Decoupling for parallel excitation

In parallel detector modeling, the B_1 field was correlated with the pulse sequence through two matrices,

$$\mathbf{B}_1 = \mathbf{T} \cdot \mathbf{F} \cdot \mathbf{p}, \quad (4.1)$$

where $\mathbf{p}$ is the excitation at the T&M network terminal, and $\mathbf{F}$ and $\mathbf{T}$ matrices indicate a two-step linear supercombination. Specifically, the $\mathbf{F}$ transforms the excitation pulse to coil currents, so it represents a circuit-level combination. This combination is position-independent and convenient to be canceled out by pulse compensation. The $\mathbf{T}$ matrix transforms the coil currents to the B_1 field, so it represents a field-level combination, which contributes to the direct field spill-over effect from one channel to another. In this way, the mutual coupling is separated into circuit-level combination and field-level combination. Generally, RF coils such as striplines and solenoids confine the B_1 field near their center, so the outgoing radiation is very weak and contributes only minimally. Consequently, under weak coupling conditions, the RF decoupling for parallel excitation is aimed at canceling out the circuit-level combination, while coupling induced by the field-level combination is ignored.

Figure 4.1a illustrates the principle of pulse compensation, where cooperative pulses are formed by augmenting parallel pulses with additional compensation terms. The cooperative pulses are designed to produce the same net effect as the parallel pulses would under individual execution. The procedure can be outlined in three steps:

1. Sequentially focus on each channel to generate the parallel pulses $\mathbf{p}_0$.
2. Introduce a compensation term δ_i for each channel, chosen to satisfy the elimination

condition, which leads to the following relation:

$$\mathbf{F} \cdot \begin{pmatrix} p_{01} + \delta_1 \\ p_{02} + \delta_2 \\ \vdots \\ p_{0n} + \delta_n \end{pmatrix} = \begin{pmatrix} F_{11} & & \\ & F_{22} & \\ & & \ddots \\ & & & F_{nn} \end{pmatrix} \begin{pmatrix} p_{01} \\ p_{02} \\ \vdots \\ p_{0n} \end{pmatrix}. \quad (4.2)$$

3. Obtain the cooperative pulses $\mathbf{p} = \mathbf{p}_0 + \delta$ by solving Eq. (4.2).




Figure 4.1: Parallel pulse compensation. (a) Schematic view of pulse compensation, where cooperative pulses are obtained by augmenting parallel pulses with additional compensation terms. (b) Waveforms at the input of the T&M network: cooperative pulses $\mathbf{p}$, parallel pulses $\mathbf{p}_0$, and compensation terms δ . The dark-colored traces (left axis) show amplitudes, with constant-amplitude $\mathbf{p}_0$ at a nominal RF amplitude of 2 kHz. The light-colored traces (right axis) show the corresponding phases. (c) Transfer fidelity obtained with $\mathbf{p}$ under coupling, compared with $\mathbf{p}_0$ in the absence and presence of coupling. The coil geometry, spin system, and transfer task are identical to those in Fig. 3.13, with coupling strength $cs = 0.24$. ν_0 and ν_1 denote the resonance offset and RF amplitude, respectively.

The first step considered multiple individual optimal control pulses as the parallel pulses, and each optimal control pulse behaves as a broadband excitation pulse, accounting for resonance offset (5 kHz bandwidth) and RF inhomogeneity determined by the imported B_1 field. The B_1 field, calculated at a reasonable excitation power, generated an average RF amplitude of 2 kHz.

In the second step, the compensation term δ was added to $\mathbf{p}_0$ to cancel the coupling effect, enabling the generation of coil currents consistent with individual excitation of the optimal control pulses. Figure 4.1b illustrates the normalized excitation waveform for two-channel pulse compensation, displaying pulse sequences in terms of amplitudes and phase with 100 time slices and a duration of 2 ms. Each element of $\mathbf{p}_0$ had a fixed amplitude of 1, whereas the amplitude and phase of $\mathbf{p}$ changed over time.




Figure 4.2: Waveforms of the optimal control pulse and the compensation term, and transfer fidelity of the cooperative pulse at varied RF coupling strength ($cs = |F_{12}/F_{11}|$). In the waveform plots, blue traces denote amplitude and orange traces denote phase (in units of π). The coil geometry, spin system, and optimization target are identical to those in Fig. 3.13. ν_0 and ν_1 denote the resonance offset and RF amplitude, respectively.

Figure 4.1c presents the simulated transfer fidelity under three excitation conditions: applying the optimal control pulses individually without coupling ($\mathbf{p}_0$ - no couple), the optimal control pulses in parallel with coupling ($\mathbf{p}_0$ - coupled), and the cooperative pulses

with coupling ($\mathbf{p}$ - coupled). When applying $\mathbf{p}_0$ in parallel, the coupling introduced B_1 field distortion and led to a drop in transfer fidelity. This drop varied depending on the resonance offset and local B_1 field. Due to the random initial guesses in the optimization of $\mathbf{p}_0$, the efficiency drop could differ between the two channels. The subsequent cooperative pulses $\mathbf{p}$ accounted for coupling and achieved the same efficiency as individual excitation of $\mathbf{p}_0$.

Under a wide range of coupling strengths, as long as Eq. 4.2 remains solvable, the cooperative pulses can achieve comparable transfer efficiency relative to that of the uncoupled case with $\mathbf{p}_0$ applied to isolated channels. Figure 4.2 shows that the cooperative pulses maintain over 96% mean fidelity with coupling strengths equivalent to 0.09, 0.24, and 0.36 of a 2-solenoid coil array.

When coupling information is characterized, simultaneous optimization is reformulated into channel-specific compensation, giving each channel greater freedom to address local inhomogeneity and interactions [110, 111, 112]. The drawback is that, as a portion of the cooperative pulse energy is devoted to counteracting coupling, its applicability may be limited by the available power, particularly under strong coupling conditions.

4.3 Signal decomposition with blind source separation

4.3.1 Blind source separation

The modeling of parallel acquisition concludes that the detected signals are the linear combination of the parallel FIDs, rewritten below:

$$\mathbf{x} = \mathbf{G}' \cdot \mathbf{f} + \mathbf{n}' . \quad (4.3)$$

In an experimental scenario, we noticed that splitting multiple FIDs falls within the realm of handling composite signals from multiple detectors. In this domain, blind source separation (BSS) [113, 114] stands out as an exceptional model. This method has numerous diverse applications, including speech recognition [115], image processing [116], and biomedical signal processing [117]. The BSS model can be described as in Fig. 4.3. The parallel signals $s(t)$ were imported into a mixing box, in which they are combined. A linear sensor system can modulate the incident signals in an instantaneous fashion [118]

or a convolutive fashion [119], and correspondingly, the solution for splitting varies on instantaneous or convolutive mixing assumption. The instantaneous linear mixtures can be modeled as

$$\mathbf{x}(t) = \mathbf{A}\mathbf{s}(t) + \mathbf{n}(t), \quad (4.4)$$

where $\mathbf{s}(t)$ vector contains the signals emitted by N sources, $\mathbf{x}(t)$ vector is the noisy instantaneous linear mixture of source signals, and $\mathbf{A}$ is an $M \times N$ mixing matrix. A separation solution estimates the mixing matrix and the source signals as

$$\mathbf{y}(t) = \mathbf{W}\mathbf{x}(t), \quad (4.5)$$

where $\mathbf{W}$ is a $N \times M$ splitting matrix, and $\mathbf{y}(t)$ is the estimated signal vector.




Figure 4.3: Blind source separation model. The vectors $\mathbf{s}(t)$, $\mathbf{x}(t)$, and $\mathbf{y}(t)$ denote the unknown signals, observed signals, and split (estimated) signals, respectively. N denotes the number of sources and M the number of sensors, with $M = N$ in the parallel NMR scenario.

Eq. 4.3 states that the received signals are instantaneous linear mixtures [114] of the FIDs, under the assumption that the reception gain is fixed, as the spectrometer's multi-receiver system is time-invariant. The source signals from different samples are considered statistically independent and are assumed to possess distinct spectral content. Here we employ the second-order blind identification (SOBI) algorithm [120], which recovers the sources by jointly diagonalizing a set of time-lagged correlation matrices. The SOBI relies solely on second-order statistics and requires only that the sources have distinct spectral content, a condition naturally satisfied by FIDs originating from different samples.




Figure 4.4: Flowchart of the SOBI algorithm. The superscript $+$ denotes the Moore–Penrose pseudoinverse, and $\dagger$ denotes the Hermitian transpose. The final step resolves the permutation and scaling ambiguities inherent to blind source separation.

As shown in Fig. 4.4, the algorithm first constructs a whitening matrix $\mathbf{Q}$ from the singular value decomposition of the data matrix $\mathbf{X}$, such that the whitened data $\mathbf{Z} = \mathbf{Q}\mathbf{X}$ satisfy $\mathbb{E}\{\mathbf{Z}\mathbf{Z}^\dagger\} = \mathbf{I}$. This reduces the remaining ambiguity of the mixing matrix to a unitary rotation. The rotation $\mathbf{U}$ is then fixed by approximate joint diagonalization of a set of time-lagged correlation matrices. The estimated mixing matrix and splitting matrix are given by

$$\begin{aligned}\hat{\mathbf{A}} &= \mathbf{Q}^+ \mathbf{U}, \\ \mathbf{W} &= \mathbf{U}^\dagger \mathbf{Q},\end{aligned}\tag{4.6}$$

This algorithm has been proven to be suitable to difficult contexts such as adverse SNR and sources with little spectral difference [120], although the scaling and permutation indeterminacies are present. The Kuhn–Munkres algorithm [121, 122] can be applied to align the estimated signals with the corresponding source signals, removing the permutation indeterminacy. And the scaling indeterminacy can be resolved by using the minimal

distortion principle [123], as given by

$$\mathbf{W}_s = \text{diag}(\mathbf{W}_p^{-1}) \cdot \mathbf{W}_p, \quad (4.7)$$

where $\mathbf{W}_p$ is the splitting matrix after permutation correction, and $\mathbf{W}_s$ is the splitting matrix we are looking for.

4.3.2 Parallel acquisition experiment

The parallel ^1H pulse-acquisition experiment was conducted to obtain data to test the SOBI method. The experiment was carried out using a 15.2 T Bruker NMR system and a custom-built two-stripline probe [124]. Each stripline coil was deployed with local shim coils powered by an external DC source to improve B_0 homogeneity. A global shim provided by the magnet was set to 1000 Hz/cm along the x-axis to increase the spectral separation between the two samples, as shown in Fig. 4.5. The ParaVision 360 was used to control the synchronous pulse and acquisition.




Figure 4.5: Global shimming setup for parallel acquisition experiment. Two striplines were inserted into the a 15.2 T magnet, and a global shim of 1000 Hz/cm was applied along the x-axis to increase the spectral separation between the two samples. Consequently, the static magnetic fields in the two samples were $B_0 \pm \Delta B_{0x}$. The samples were loaded into the detection zone through the PTFE tubes.

Fig. 4.6a gives the decomposition results for experimental data obtained with 2 water

samples. The average SNR of the two primary peaks improved to 45 dB by signal averaging. The primary peaks at 4 ppm and -1.8 ppm, were separated by applying a shimming gradient. The coupling strength extracted from signal decomposition is 13%. Phase correction and baseline correction were performed after signal decomposition. Compared to the original spectrum (grey lines), the split one (colored lines) removed over 90% of the coupled components, while maintaining the individual signals. Similarly, Fig. 4.6b shows that the composite signals from the acetone and the isopropanol were effectively split in the two channels. In the BSS for instantaneous mixtures, an estimated splitting matrix is obtained, and the recovered signals are generated by multiplying this matrix with the coupled signals, thereby preserving the spectral lineshape. The residual coupled components in the spectra indicate an estimation error of the splitting matrix, which is discussed below.




Figure 4.6: Two-channel ^1H signal decomposition using the SOBI method. (a) Synchronous excitation and acquisition using two water samples, NS=64. (b) Synchronous excitation and acquisition using acetone in channel 1 and isopropanol in channel 2, NS=256. Shared acquisition parameters in ParaVision included SW = 20 ppm, TD = 13157, TR = 1.5 s, FA = 20° , DR = 31.

4.3.3 Spectral resolution

Decomposition becomes more challenging for signals with low SNR and small spectral separations (Δf). To evaluate the resolution limit, each source was modeled with a single

frequency component. Figure 4.7 shows decomposition results for an eight-channel system using simulated FIDs under a homogeneous B_0 field, with a minimum $\Delta f = 3$ Hz and local SNR of 30 dB. By systematically varying the chemical shifts and SNR, the resolution limit as a function of SNR was determined, as summarized in Table 4.1. The results demonstrate that two signals separated by 3 Hz can be distinguished at SNR = 30 dB, while a resolution of 1.5 Hz is achievable at SNR = 40 dB. Higher SNR consistently leads to improved resolution.




Figure 4.7: Signal decomposition of an eight-channel system using simulated data, with each channel assigned a single spin (one chemical shift). The original signal is shown in grey, and the decomposed signal in blue. The zoomed-in view highlights two signals with a minimal spectral separation of $\Delta f = 0.006$ ppm (3 Hz) at a local SNR of 30 dB, which are successfully resolved.

SNR, dB	20	23	25	30	35	40	45
resolution, Hz	25	17.5	3.5	3.0	2.5	1.5	1.0

Table 4.1: Spectral resolution as a function of SNR under a homogeneous B_0 field. Higher SNR yields improved resolution.

To evaluate whether multiple spectral components can enhance the identification of




Figure 4.8: Signal decomposition of an 8-channel system using simulated FIDs under a homogeneous B_0 field. Each channel contains a random number of spins with chemical shifts randomly distributed in the range $[0, 10]$ ppm. The original signal is shown in grey, and the decomposed signal in blue. As highlighted in the zoomed-in view, two peaks in channel 2 and channel 4 are separated by only $\Delta f = 0.001$ ppm.

coupling and potentially exceed the resolution limit, signal decomposition is applied to complex spin systems. Table 4.2 lists the chemical shifts assigned to the 8 channels, and Figure 4.8 presents the decomposition results with an average SNR of 30 dB. Notably, the signals in channel 2 and channel 4 are extremely close in frequency (8.820 and 8.821 ppm). Since the SOBI method estimates the mixing matrix based on the overall spectral characteristics, the presence of a few closely spaced components does not significantly affect its performance.

channel	chemical shift, ppm
1	7.01, 2.97
2	0.87, 5.99, 2.54, 6.84, 2.22, 0.92, 8.821
3	8.98, 1.88, 6.56, 5.96, 8.87
4	5.04, 3.21, 5.47, 7.25, 5.74, 6.65, 8.820, 8.74
5	0.67, 8.56, 5.26, 2.53, 5.80, 2.30, 8.41
6	9.45, 5.54, 9.28, 6.20, 1.74, 5.10
7	0.86, 8.08, 1.36
8	5.70, 8.34, 3.28, 7.82, 2.62

Table 4.2: Randomly generated chemical shifts for 8 channels in Figure 4.8.

4.3.4 Estimation error

In highly parallel probes, the spectral differences between signals from different channels may be small. Combined with the complexity of spin systems, both SNR and linewidth can be compromised due to limited shimming quality. To assess these effects, we evaluate the performance of the SOBI method using a figure of merit in the context of a large number of mutually uncorrelated sources. As noted in the literature [125], when source signals are partially correlated, the separation capability degrades rapidly as the number of sources increases.

Each channel was modeled with a single harmonic signal

$$s_n(t) = \exp((i2\pi f_n - \alpha)t), \quad (4.8)$$

where $f_n = n$ and α is the decay rate, yielding a spectral linewidth of α/π in Hz. Each source signal contained 5000 sampling points over a 10-second duration with a fixed SNR of 30 dB. The coupling matrix was randomly generated as an $N \times N$ complex matrix for N channels.

The estimation error was quantified using the mixing matrix criterion [125]:

$$C_A = \frac{1}{N} \|\mathbf{E} - \hat{\mathbf{A}}^{-1} \mathbf{A}\|_1, \quad (4.9)$$

where $\mathbf{E}$ is the identity matrix, $\hat{\mathbf{A}}$ is the estimated coupling matrix in which the permutation and scale indeterminacies are removed, and $\mathbf{A}$ is the true coupling matrix. The calculated C_A values are presented in Fig. 4.9.




Figure 4.9: Mixing matrix criterion (C_A) as a function of channel number. For each channel number, 30 repetitions were performed and the mean C_A value is reported. The linewidth of each peak was α/π .

The results indicate that performance improves with an increasing number of sources, yielding smaller estimation errors and robust separation of uncorrelated signals. However, since the SOBI algorithm relies on distinct spectral features for each source, its separation efficiency deteriorates significantly when coupled components appear on the shoulders of primary peaks. Narrower linewidths α lead to reduced C_A values; for example, decreasing the linewidth from $0.4/\pi$ to $0.2/\pi$ improves the estimation efficiency by a factor of two (3 dB). As shimming quality remains the primary obstacle to implementing parallel detection, it may be worthwhile to test whether time-domain processing, such as apodization prior to decomposition, could improve performance, as this has not yet been investigated. And more robust decomposition techniques are required to handle densely packed peaks, broad spectral shoulders, and intrinsically correlated source signals. In this context, concepts developed for diffusion-ordered spectroscopy (DOSY) may provide useful strategies [126, 127, 128].

4.4 Parallel excitation experiment

4.4.1 Experimental setup

The parallel excitation experiments were conducted on a Bruker AVANCE NEO 11.7 T spectrometer equipped with a four-detector probe (Voxalytic GmbH), as shown in Fig. 4.10. The probe comprises four independent detectors, each containing a stripline, 6 local shim-coils (not shown), and a single-axis gradient coil. The striplines were individually tuned and matched using a T&M box supplied with the probe: two were double-resonant at $^1\text{H}/^{13}\text{C}$, and the other two were double-resonant at $^{19}\text{F}/^{13}\text{C}$. In addition to the global shimming provided by the magnet, the on-chip shimming coils were powered by a 24-channel direct-current source [124], with each channel having a maximum output current of ± 300 mA. Both shim currents and T&M adjustments were controlled via a laptop.




Figure 4.10: Experimental setup for the parallel NMR experiment. The parallel probe allows simultaneous operation of four detectors, each equipped with its own stripline, gradient coil, and local shim coils (not shown). Liquid samples were loaded and removed through the top of the probe. Each “U”-shaped PTFE tube, with one inlet and one outlet, was inserted into a stripline coil, which defines a detection volume of $3.9\ \mu\text{L}$. The local shim coils were powered by an external 24-channel current source, where each channel received a control voltage from a laptop and delivered a direct current of up to ± 300 mA. The T&M box provided voltage-controlled variable capacitances, enabling independent T&M of each detector.

For sample loading, each detector is equipped with an internal round tube, accessible through inlet and outlet connectors at the top of the probe. Samples were filled into syringes and manually pumped through these tubes into the individual detection zones of the

corresponding striplines.

The four gradient coils, each rated for a maximum current of 3 A, were powered by Bruker GREAT amplifiers: the micro-imaging amplifiers provided three split outputs via a custom cable, each driving one of three coils, while the fourth coil was powered by a separate amplifier. The gradient decoupling using these gradients is discussed in chapter 5. The parallel RF pulses were amplified by an eight-channel Bruker preamplifier and delivered to the probe through parallel RF cables. Synchronous excitation and acquisition were enabled by the spectrometer's multi-receiver functionality and controlled in TopSpin.

4.4.2 Calibration test

The workflow for generating cooperative pulses is illustrated in Fig. 4.11. First, the simulated B_1 fields were imported into the optimal control module implemented in SPINACH to design the optimal control pulses. Next, a calibration experiment was carried out to acquire coupled time-domain signals, which were analyzed with the BSS method to estimate the coupling matrix $\hat{\mathbf{F}}$. Finally, the cooperative pulses were computed according to

$$\mathbf{p} = \mathbf{F}^{-1} \mathbf{F}_{\text{diag}} \mathbf{p}_0, \quad (4.10)$$

where $\mathbf{F}_{\text{diag}}$ is the diagonal of $\mathbf{F}$, and $\mathbf{p}_0$ denotes the vector of individually optimized control pulses.

The optimal control pulses $\mathbf{p}_0$ had a bandwidth of 30 kHz and nominal RF amplitude of 30 kHz, with $\pm 20\%$ power scaling to account for B_1 inhomogeneities of a stripline. Both parameters were discretized into 100 ensemble segments. The pulse amplitude was fixed at unity, phases were optimized, and the 8 ms pulse was divided into 4000 bins.




Figure 4.11: Workflow for generating the cooperative pulses. The OC denotes optimal control, BSS denotes blind source separation, and $\mathbf{F}_{\text{diag}}$ is the diagonal of the coupling matrix $\mathbf{F}$.

Distilled water was used in the calibration experiment owing to its high signal amplitude. To facilitate separation of the two water peaks, the global y-shim value was set to a large value of 1.5×10^5 in TopSpin, producing a spectral difference of 5.3 ppm. The measured water peaks, referenced internally, were located at -0.7 ppm and 4.6 ppm for detectors 1 and 2, respectively.




Figure 4.12: Two-detector calibration with water samples. Grey traces show the original spectra, and colored traces show the split spectra. A large shimming gradient was applied along the y-axis to separate the two primary peaks. The TopSpin acquisition parameters included SW = 20 ppm, TD = 3072, TR = 5 s, RG = 10, NS = 64.

The detected FIDs were analyzed using the SOBI method for signal decomposition. The estimated coupling matrix was

$$\hat{\mathbf{F}}_{\text{norm}} = \begin{bmatrix} 0.9998 + 0.0004i & 0.0103 + 0.0201i \\ 0.0085 - 0.0081i & 0.9998 + 0.0004i \end{bmatrix}. \quad (4.11)$$

This matrix was then used for decomposition. In Fig. 4.12, the coupled component in detector 1 (D1) was successfully removed, whereas the one in detector 2 (D2) at -0.7 ppm persisted, indicating an error in coupling estimation. Two main factors contributed to this error: (i) the limited shimming quality in D1, where the peak showed a linewidth of 28 Hz together with a pronounced upfield shoulder; and (ii) strong impedance mismatches, with relative reflections of about 26% in D1 and 35% in D2, which led to estimation difference

between the amplitudes of F_{12} and F_{21} .

4.4.3 Parallel excitation

The cooperative pulses computed from Eq. 4.10 were tested in parallel excitation to validate the pulse compensation scheme. The power levels for the shaped pulses were determined by single-detector nutation experiments. Specifically, nutation experiments were performed on D1 and D2 separately to measure the RF amplitude at the maximum recommended power level, $P_{\max} = 20$ W. The power was scaled down to yield an RF amplitude of 30 kHz, corresponding to a $8.3 \mu\text{s}$ 90° pulse, assuming that the B_1 field scales with $\sqrt{P}$ in the linear regime of the amplifier.




Figure 4.13: ^1H spectra of the two water samples under three excitation conditions. Grey and colored dashed lines correspond to spectra obtained with the optimal control pulses applied in a single channel (p_0 , uncoupled) and in parallel excitation (p_0 , coupled), respectively. Colored solid lines depict spectra acquired with cooperative pulses during parallel excitation (p , coupled). Identical phase-correction parameters were applied to all three spectra in each subfigure. The TopSpin acquisition parameters included $\text{SW} = 20$ ppm, $\text{TD} = 3072$, $\text{TR} = 5$ s, $\text{RG} = 10$, $\text{NS} = 64$.

Fig. 4.13 shows the ^1H spectra of the two water samples, with the top and bottom plots corresponding to D1 and D2, respectively. Grey lines represent spectra obtained with the optimal control pulse in single-channel excitation (p_0 , uncoupled) as a reference. To enable comparison of excitation distortions, the three spectra in each plot share identical

phase-correction parameters, optimized according to the lineshape of the reference. Shimming with both global and local shim sets introduced a 0.9 ppm shift between the water peaks. In D1, spectra acquired with p_0 (coupled) exhibit amplitude and phase distortions due to coupling, whereas spectra with p (coupled) closely resemble the reference, indicating successful restoration by the cooperative pulse. In contrast, in D2 the spectra nearly overlap, consistent with weaker coupling of about 1%. In this case, the error in coupling estimation resulted in ineffective compensation by the cooperative pulse.




Figure 4.14: Two-detector pulse compensation with L-alanine and L-valine. Each subplot shows three ^1H spectra corresponding to different excitation conditions: p_0 (uncoupled), p_0 (coupled), and p (coupled). The left-side peaks arise from residual water, and the right-side peaks from TSP. Due to limited shimming quality, closely spaced peaks are not resolved. Identical phase-correction parameters were applied to all three spectra in each subplot. The TopSpin acquisition parameters included SW = 20 ppm, TD = 3072, TR = 10 s, RG = 18, NS = 64.

The same pulses were further tested with different samples, L-valine and L-alanine, to verify the efficiency of the cooperative pulses. Solutions of 250 mM L-valine and 250 mM L-alanine (Sigma-Aldrich) were prepared in D_2O (99.9%, Sigma-Aldrich) with 10 mM 3-(trimethylsilyl)propionic-2,2,3,3- d_4 acid sodium salt (TSP) as an internal reference. Figure 4.14 shows spectra of the L-alanine and L-valine samples excited with the same optimal

control pulses and cooperative pulses, exhibiting distortions similar to those in Fig. 4.13. Due to limited shimming quality, the closely spaced peaks from L-alanine and L-valine are not resolved; the downfield peaks arise from residual water in the D_2O solvent, while the peaks at 0 ppm correspond to TSP. Since the separation relies on the sources having distinct power spectra, coupling identification degrades when the parallel channels contain similar frequency components. The coupling is therefore calibrated beforehand with reference samples giving strong, well-separated spectra, and the result is then applied to the target measurement.

4.5 End-to-end test

Parallel excitation experiments demonstrated that cooperative pulses under coupling produced spectra comparable to those obtained with optimal control pulses in the absence of coupling. A more rigorous validation is comparing the excitation profiles under these two conditions. Such measurements typically require uniform power delivery and high shimming quality, which can be achieved using a Shigemi tube [129]. However, the present parallel probe does not permit tube replacement or excitation at a uniform power level.

As an alternative, the efficiency of pulse compensation at different SNR levels was simulated. The workflow is shown in Fig. 4.15, which follows the same procedure as Fig. 4.11, except that the parallel FIDs were generated by the multiphysics simulation, and noise of specified levels was added according to the desired SNR.




Figure 4.15: Workflow for generating cooperative pulses using multiphysics simulation. EM denotes electromagnetic, BSS blind source separation, $\mathbf{p}_0$ the optimal control pulses, $\mathbf{p}$ the cooperative pulses, and $\mathbf{F}_{\text{diag}}$ the diagonal of the coupling matrix $\mathbf{F}$.

Figure 4.16a illustrates a representative case with an average SNR of 10 dB, where the




Figure 4.16: Robustness of the RF decoupling to low SNR. (a) Example results for a two-channel simulation at $\text{SNR} = 10$ dB with a minimal separation of $\Delta f = 0.02$ ppm (10 Hz). The left, middle, and right panels show the original ^1H spectrum, the decomposed ^1H spectrum, and the fidelity (η) of the cooperative pulses, respectively. ν_0 is the resonance offset, ν_1 the RF amplitude, and $\bar{\eta}$ the averaged fidelity. (b) Mixing matrix criterion (C_A) as a function of SNR. (c) Mean fidelity of the cooperative pulses as a function of SNR. The mean fidelities of the optimal control pulses were $\bar{\eta}_1 = 0.9952$ for channel 1 and $\bar{\eta}_2 = 0.9951$ for channel 2.

spin system for each channel was randomly generated, including varying spin numbers and chemical shifts. At $\Delta f = 0.02$ ppm, the cooperative pulses achieved mean fidelities of 0.988 for channel 1 and 0.987 for channel 2, corresponding to 99.3% recovery relative to the optimal control pulse. Figure 4.16b demonstrates the robustness of C_A against SNR, with C_A converging to -17 dB (2%) as SNR increases to 50 dB. This error may be attributed to the weak correlations among the source signals, arising from cross-coupling between the coils and samples during reception. Nevertheless, a 2% decomposition error corresponds to the removal of over 95% of the coupled components, yielding cooperative-pulse fidelities comparable to those of the optimal control pulses, as shown in Fig. 4.16c.

4.6 Conclusion

This chapter proposed RF decoupling schemes for both the excitation and reception stages of parallel high-field NMR spectroscopy. In particular, in a linear time-invariant system, coupling information represented by a matrix can be extracted by separating the detected time-domain signals using the SOBI method. The estimated matrix was then employed to design cooperative pulses that cancel RF coupling under parallel excitation, a compensation scheme validated in a two-channel NMR spectroscopy experiment.

Excitation-side decoupling removes the need for time-interleaved pulse sequences and enables the direct application of powerful optimal control pulses in parallel NMR experiments. The reception decoupling based on blind source separation avoids the need to separate multiple signals by frequency division using additional field gradients. These two strategies increase the flexibility of parallel detection and facilitate the design of more sophisticated parallel pulse sequences.

The cooperative pulse design preserves single-channel independence. This allows each channel to allocate its degrees of freedom to address specific samples, making the scheme efficient for deployment in highly parallel detectors where the B_1 field is confined to individual coils. Looking ahead, one may envision NMR coil arrays in which field-level combinations become significant. In such cases, pulse compensation based solely on the coupling matrix may no longer suffice. Incorporating parallel transmit design strategies, as extensively developed in MRI [97, 99, 130], could then prove advantageous.

Chapter 5

Gradient decoupling for parallel NMR spectroscopy

Based on M. He, N. MacKinnon, D. Buyens, B. Luy, and J. G. Korvink, Magnetic Resonance, 6(2): 173–181, 2025.

5.1 Introduction

Linear magnetic field gradients are fundamental to magnetic resonance. In MRI, they encode spatial information into the signal, which is later recovered during image reconstruction. In NMR, pulsed field gradients are applied during evolution periods to select specific coherence transfer pathways, effectively replacing complex phase cycling. This reduces T_1 artifacts and shortens measurement time. In parallel NMR setups, multiple gradient coils can be integrated into individual detection channels. However, when parallel gradient pulses are applied, magnetic field leakage between the closely packed coils may occur. This effect, known as gradient field spillover, introduces B_0 inhomogeneity and causes spin dephasing.

A straightforward approach to mitigate spin dephasing to B_0 inhomogeneity is by transferring the spin state to longitudinal magnetization, e.g., I_z for single spins or $I_z S_z$ for coupled spins. However, this proves challenging when the initial state is unknown or time

constraints exist. Spin locking via RF pulses is an effective strategy for preserving spin magnetization, i.e., a long hard pulse with a defined phase is applied to protect a specific coherence, such as I_x or I_y . In heteronuclear experiments, continuous wave spin-locking fields have been utilized to render proton multiple quantum coherence immune to ^1H - ^1H J -coupling [131]. Spin locking induced crossing has a wide range of applications, including preparation of singlet states [132, 133, 134], excitation of long-lived states [135, 136, 137], and low field spectroscopy [138, 139]. Moreover, traditional spin-locking pulses have been adapted to address B_0 and B_1 inhomogeneity in quantifying the rotating frame relaxation time [140, 141].

Due to electromagnetic and temperature fluctuations and spin-spin interactions, developing coherence protection methods for manipulating quantum qubits is essential [142]. Examples include spin-locking of an electron spin for radiofrequency magnetometry [143], suppression of intrabath interactions and inhomogeneity using sequenced pulses [144], extension of coherence time via continuous microwave driving [145], and locking the dynamics of a single nitrogen vacancy (NV) spin qubit by microwave dressing [146].

Since gradient pulses are used to select coherence transfer pathways, only specific coherences are of interest. Similar to how spin locking or coherence protection in spin qubits is achieved using microwave irradiation, selected coherences such as I_- for a single spin or I_-S_+ for coupled spins may be preserved with specially designed RF pulses. In this chapter, we introduce coherence-locking pulses to protect the desired coherences against gradient field spillover and demonstrate their effectiveness in parallel NMR experiments.

5.2 Quantify the gradient spillover

5.2.1 Gradient field simulation

Consider a multi-detector geometry, as shown in Fig. 5.1a, each detector is equipped with a Helmholtz gradient coil and a stripline RF coil. The gradient current in detector 1 was represented by $i_1(t) = 3\sin(1000\pi t)$, as shown in Fig. 5.1b. This current generates a field gradient of 75 Gauss/cm in sample 1, although not explicitly shown. Two situations now arise at detector 2: first, the applied gradient at detector 1 can induce a current in the coil at detector 2, thereby creating an opposing field. Second, the gradient field generated at detector 1 is not fully shielded, so that this field spillover penetrates the sample at detector

2. The induced current i_2 in detector 2 due to the magnetic flux change was computed in Fig. 5.1d. Meanwhile, Fig. 5.1c portrays the magnetic field in sample 2 generated by i_1 and i_2 . Notably, the gradient amplitude at detector 2 from i_1 amounted to 0.5 Gauss/cm, resulting in a spillover ratio of 0.67%. This value is approximately 68 times higher than the maximum gradient generated from the induced current i_2 , echoing similarities with Lenz's law, where the induced current's magnetic flux opposes the change in flux but cannot fully counteract it. Consequently, the field generated by the induced current was disregarded in subsequent analyses.

The field spillover induces an additional gradient contribution alongside the primary gradient. When multiple gradient pulses are applied to select a coherence transfer pathway, the designed gradient ratio may undergo distortion due to field spillover.




Figure 5.1: Gradient simulation for a two detector array. (a) Geometry of the detector array. (b) The gradient coil in detector 1 was excited with a sin-shaped pulse. (c) The gradient field on sample 2 calculated with the excited current $i_1(t)$ and induced current $i_2(t)$, the ratio of maximum $B(i_1)$ and $B(i_2)$ is 68. (d) Induced current $i_2(t)$ in detector 2.

5.2.2 Measuring the spillover ratio

The gradient spillover ratio was measured using the four-detector probe introduced in chapter 4. As a first step, the maximum gradient strength of the parallel probe was determined with the pulsed gradient spin echo (PGSE) method [147], as shown in Fig. 5.2a. A 10% H₂O/90% D₂O solution with a known diffusion constant of $D = 1.9 \times 10^{-9} \text{ m}^2 \cdot \text{s}^{-1}$ [148] was used. The ratio of the signal to the maximum value is a function of the applied gradient amplitude [147],

$$\ln(I_g/I_0) = -[\gamma^2 \delta^2 G^2 (\Delta - \delta)] D, \quad (5.1)$$

where $\delta = 1 \text{ ms}$ is the gradient pulse length, $\Delta = 7 \text{ ms}$ is the interval between two gradient pulses, and $\gamma = 2.675 \times 10^8 \text{ rad} \cdot \text{s}^{-1} \cdot \text{T}^{-1}$ is the proton gyromagnetic ratio. A series of spectra were acquired by sweeping the gradient amplitude from 5% to 95%. Fitting the resulting signal amplitudes as a function of gradient amplitude using Eq. 5.1 yielded a maximum gradient of $G_{\text{max}} = 103.06 \text{ G/cm}$, see Fig. 5.2b. This measurement provided a calibrated reference for the subsequent quantification of gradient field spillover.




Figure 5.2: Measurement of the maximum gradient amplitude using a 10% H₂O/90% D₂O solution. (a) PGSE pulse sequence with a gradient pulse duration of $\delta = 1 \text{ ms}$ and a pulse separation of $\Delta = 7 \text{ ms}$. Black and white blocks denote 90° and 180° pulses, respectively. A "y" label below a pulse indicates a 90° phase; unlabelled pulses carry zero phase by default. These conventions hold throughout. Gradient pulses are drawn with their actual envelope; sine-shaped blocks represent sine-modulated gradient pulses. (b) Signal intensity as a function of gradient amplitude in the PGSE experiment. The maximum gradient amplitude was determined to be $G_{\text{max}} = 103.06 \text{ G/cm}$ at 100% gradient strength. Circles denote experimental data, and the solid line shows the fitted curve.

The gradient spillover ratio was then measured in a pulse-acquisition experiment, while a gradient pulse was applied to the second detector, as shown in Fig. 5.3a. The ratio of the signal to the maximum value is a function of the applied gradient amplitude,

$$\frac{I_g}{I_0} = \frac{\sin(kg)}{kg}, \quad (5.2)$$

where $k = l/2 \cdot \gamma \cdot R_G \cdot G_{\max} \cdot \int g(t) dt$, $l = 6.5$ mm is the detection zone length, and $\int g(t) dt = 0.9$ ms is the time integral of the trapezoidal gradient pulse. The gradient spillover ratio R_G is defined as $R_G = G_c/G_p$, where G_c is the maximum coupled gradient strength, and G_p is the maximum primary gradient strength. A series of spectra were acquired by sweeping the gradient amplitude from 5% to 95%. Fitting the resulting signal amplitudes as a function of gradient amplitude using Eq. 5.2 yielded the gradient spillover ratio of $R_G = 1.9 \times 10^{-3}$.




Figure 5.3: Measuring the gradient spillover ratio using a 10% H₂O/90% D₂O solution. (a) A gradient pulse was applied to detector 2 during the pulse-acquisition in detector 1; the trapezoidal-shaped block represents a trapezoidal-shaped gradient pulse. (b) Signal intensity as a function of coupling gradient amplitude in a pulse-acquisition experiment, the gradient spillover ratio was calculated to be $R_G = 1.9 \times 10^{-3}$. Circles denote experimental data, and the solid line shows the fitted curve.

5.2.3 Field gradient dephasing

For the evolution of p -quantum coherence under the influence of a field gradient pulse $G_z(t)$ of duration τ , the spin coherence after the gradient pulse is given by [41]

$$\rho(\tau) = \rho(0) \exp \left[-ip \int_0^\tau \gamma G_z(t) z dt \right], \quad (5.3)$$

where z is the spatial coordinate and γ the gyromagnetic ratio. The corresponding phase of the coherence is

$$\phi_p(z) = -p \int_0^\tau \gamma G_z(t) z dt, \quad (5.4)$$

which varies linearly with position, forming a helical modulation along the spatial coordinate. The wavelength of this phase helix is

$$\lambda = \frac{-2\pi}{p \int_0^\tau \gamma G_z(t) dt}, \quad (5.5)$$

so that the coherence phase can also be written as $\phi_p(z) = 2\pi z/\lambda$.

For a coherence transfer pathway involving N successive coherences $[p_1, p_2, \dots, p_N]$, Eq. 5.4 generalizes to

$$\phi_{p_1 \rightarrow p_N}(z) = -\sum_{i=1}^N p_i \int_{\tau_i} \gamma_i G_z(t) z dt, \quad (5.6)$$

with the corresponding helix wavelength

$$\lambda_{p_1 \rightarrow p_N} = \frac{-2\pi}{\sum_{i=1}^N p_i \int_{\tau_i} \gamma_i G_z(t) dt}. \quad (5.7)$$

To quantify the signal decay induced by gradient spillover, we consider a sample of length l . The observed signal is proportional to the spatial integration of the phase term over the sample volume:

$$s = \int_{-l/2}^{l/2} e^{i\phi(z)} dz = \int_{-l/2}^{l/2} e^{i2\pi z/\lambda} dz = \frac{\lambda}{\pi} \sin\left(\frac{\pi l}{\lambda}\right), \quad (5.8)$$

where λ is the phase helix wavelength of the pathway, cf. Eq. 5.7. In the limit $\lambda \rightarrow \infty$, the signal reaches its maximum value $s_{\max} = l$, indicating perfect refocusing of the selected pathway. For finite λ comparable to the sample dimension, the normalized signal intensity becomes

$$\frac{s}{s_{\max}} = \frac{\lambda}{\pi l} \sin\left(\frac{\pi l}{\lambda}\right), \quad (5.9)$$

which follows the well-known sinc dependence, as illustrated in Fig. 5.4.




Figure 5.4: Relative signal intensity as a function of helix wavelength, shown on a logarithmic scale.




Figure 5.5: Gradient-coupled HSQC pulse sequence. The black elements represent the standard HSQC sequence, where gradients with ratio 2 : 2 : 1 select the $S_- \rightarrow S_- \rightarrow I_+ \rightarrow I_-$ pathway. The orange blocks (G_c) indicate additional coupled gradient pulses from a second detector with ratio 2 : 2 : -1.

As an example, gradient coupling in a ^1H - ^{13}C heteronuclear single quantum coherence spectroscopy (HSQC) [149] experiment is illustrated in Fig. 5.5. The primary gradients,

with a ratio of 2 : 2 : 1, select the $S_- \rightarrow S_- \rightarrow I_+ \rightarrow I_-$ coherence transfer pathway. An additional set of coupled gradient pulses (orange) was applied from a parallel channel, temporally shifted relative to the primary gradients, thereby inducing spillover. Using Eq. 5.9, the HSQC signal intensity under such coupling was calculated and is plotted in Fig. 5.6. Each detector was subjected to sinusoidal gradient pulses of 1 ms duration, with spillover ratios varied between 0.1% and 1%. The primary gradient ratio was fixed at 2 : 2 : 1, while the coupled gradients were set to 2 : 2 : -1 (blue curve) and 2 : 2 : 1 (orange curve). Note that when the coupled gradient ratio matches the primary one, the net ratio remains unchanged, and no signal degradation occurs.




Figure 5.6: Relative intensity of the ^1H - ^{13}C HSQC signal as a function of gradient spillover ratio. The maximum gradient amplitude was 75 G/cm; all gradient pulses had a duration of 1 ms and a sinusoidal shape. The primary gradient ratio was $G_p = 1 : 1 : 0.5$. Coupled gradients were $G_c = 1 : 1 : -0.5$ (blue) and $G_c = 1 : 1 : 0.5$ (orange). Sample length was 8 mm.

In HSQC, a coherence pathway is fully refocused only when the combined effect of all gradient pulses cancels out. Any imbalance introduced by additional gradients leads to dephasing and signal attenuation. This situation arises when parallel gradient-enhanced experiments are performed with differing gradient ratios. Even small coupling ratios can have a substantial impact; for instance, a 0.2% spillover ratio resulted in a 20% HSQC signal loss (Fig. 5.6). Since strong gradients are critical in diffusion measurements, gradient spillover effects can become more pronounced in such experiments.

One possible solution is to preserve the target coherences during the coupled gradient

period using specially designed RF pulses, termed coherence locking. Fig. 5.7 illustrates the integration of coherence-locking pulses into an HSQC sequence. Three locking blocks were introduced alongside the coupled gradient pulses G_c . Two blocks (CL_1 , CL_2) were applied on ^{13}C to protect S_- coherence, while one block (CL_3) was applied on ^1H to protect I_- coherence. The design and optimization of these coherence-locking pulses, as well as their extension to parallel pulse sequences, will be addressed in the following section.




Figure 5.7: ^1H – ^{13}C HSQC pulse sequence with gradient coupling and coherence locking. Black blocks represent the primary gradient pulses ($G_p = 2 : 2 : 1$), orange blocks the coupled gradients ($G_c = 2 : 2 : -1$), and blue blocks the coherence-locking pulses (CL_1 , CL_2 on ^{13}C and CL_3 on ^1H).

5.3 Coherence locking pulse

5.3.1 Optimal control formulation

In this section, we design coherence-locking pulses (CLOC pulses) using optimal control theory. When a gradient pulse is applied, its shape and duration are fixed, while its amplitude may vary; the optimization of CLOC pulses follows this constraint. The gradient pulse contributes a time-dependent drift Hamiltonian,

$$\mathbf{H}_g(t) = A \sin\left(\frac{\pi t}{\tau}\right) \mathbf{H}_{z0}, \quad (5.10)$$

where $\mathbf{H}_{z0}$ is the Zeeman Hamiltonian assuming a 1 T magnetic field, and $\tau = 1$ ms is the gradient pulse duration. To accommodate the maximum B_0 drift of ± 0.06 G measured with the parallel probe, a wider range of ± 0.25 G was specified for the CLOC pulse optimization.

We begin with a spin dynamics description of the optimal control problem for coher-

ence locking. For an IS spin system, the rotating-frame Hamiltonian is written as

$$\mathbf{H}(t) = \mathbf{H}_{\text{int}} + \mathbf{H}_{\text{g}}(t) + \mathbf{H}_{\text{rf}}(t), \quad (5.11)$$

where the internal Hamiltonian, the gradient contribution, and the RF Hamiltonian are, respectively,

$$\begin{aligned} \mathbf{H}_{\text{int}} &= \omega_I \mathbf{I}_z + \omega_S \mathbf{S}_z + 2\pi J \mathbf{I}_z \mathbf{S}_z, \\ \mathbf{H}_{\text{g}}(t) &= B_{\text{g}}(t) \cdot \mathbf{H}_{z0} = B_{\text{g}}(t) \cdot (\Omega_I \mathbf{I}_z + \Omega_S \mathbf{S}_z), \\ \mathbf{H}_{\text{rf}}(t) &= \omega_{I,x}(t) \mathbf{I}_x + \omega_{I,y}(t) \mathbf{I}_y + \omega_{S,x}(t) \mathbf{S}_x + \omega_{S,y}(t) \mathbf{S}_y. \end{aligned} \quad (5.12)$$

Here, ω_I and ω_S denote resonance offsets, Ω_I and Ω_S are Larmor frequencies at 1 T, and $\omega_{I,x}$, $\omega_{I,y}$, $\omega_{S,x}$, and $\omega_{S,y}$ represent the four RF control amplitudes.

The $B_{\text{g}}(t)$ represents the time-dependent drift field introduced by the gradient, and $\mathbf{H}_{\text{int}} + \mathbf{H}_{\text{g}}(t)$ describes the drift Hamiltonian. To capture gradient spillover, an ensemble of drift Hamiltonians was considered. Since spillover introduces time- and space-dependent B_0 drifts, the sample was divided into n_{g} voxels along the z -axis, each voxel assigned its own drift Hamiltonian. The full Hamiltonian was then represented as a block-diagonal matrix comprising all voxel contributions.

Following the GRAPE method introduced in chapter 2, the fidelity function is defined as

$$\eta = \langle \sigma | \rho(\tau) \rangle = \text{trace} [\sigma^\dagger \rho(\tau)]. \quad (5.13)$$

For CLOC pulses, the target state σ coincides with the initial state. For non-Hermitian states, only the real part of Eq. 5.13 is considered. In this context, the target propagator $\mathbf{U}$ must commute with the initial state. For example, $[\mathbf{U}, \mathbf{I}_-] = 0$ ensures locking of the single-spin coherence $\mathbf{I}_-$, and $[\mathbf{U}, \mathbf{I}_- \mathbf{S}_-] = 0$ ensures locking of the double quantum coherence $\mathbf{I}_- \mathbf{S}_-$.

In practice, it is often desirable to protect several coherences within a single CLOC pulse, as multiple pathways may be relevant in a single experiment. For example, in gradient-enhanced HSQC, the gradient sign is alternated between scans to select either $\mathbf{S}_+$ or $\mathbf{S}_-$. Therefore, for a single-spin system, three coherences ($\mathbf{I}_z$, $\mathbf{I}_-$, $\mathbf{I}_+$) can be included

simultaneously, yielding an ensemble of source–target pairs and a cyclic optimal pulse. The fidelity in Eq. 5.13 generalizes to

$$\bar{\eta} = \frac{1}{n} \sum_{n_{\text{rf}}} \sum_{n_{\text{g}}} \sum_{n_{\text{off}}} \sum_{n_{\rho}} \eta, \quad (5.14)$$

where n_{rf} , n_{g} , n_{off} , and n_{ρ} denote the numbers of RF amplitudes, gradient strengths, resonance offsets, and source–target pairs, respectively.

The following parameters were used for single-spin optimization:

- Pulse duration: 1 ms, discretized into 1000 equivalent steps.
- Gradient pulse: sinusoidal shape, 1 ms duration, calculated over $n_{\text{g}} = 12$ voxels with $B_{\text{g,max}} = \pm 0.25$ G.
- ^1H resonance offset: 7 kHz bandwidth, $n_{\text{off}} = 51$.
- ^1H RF amplitude: $\nu_{1,\text{H}} = 6$ kHz with $\pm 20\%$ inhomogeneity, $n_{\text{rf}} = 15$.
- ^{13}C resonance offset: 6 kHz bandwidth, $n_{\text{off}} = 51$.
- ^{13}C RF amplitude: $\nu_{1,\text{C}} = 4$ kHz with $\pm 15\%$ inhomogeneity, $n_{\text{rf}} = 10$.

The optimization was performed using SPINACH v2.8 [58] with the LBFGS–GRAPE algorithm. For the single-spin case, calculations used 40 CPU cores at 2.1 GHz and completed in approximately one hour. The resulting CLOC pulse shapes for ^1H and ^{13}C are shown in Fig. 5.8.

For multi-spin cases, a more practical strategy is to establish robust coherence locking for individual spins and evaluate their decoupling performance within the designed bandwidth, thereby avoiding unnecessary model complexity.

The concurrent optimization becomes essential when J -coupling cannot be neglected or averaged out by single-spin pulses. The concurrent optimization for locking a double-quantum coherence in a ^1H - ^{13}C system was tested. The ^{13}C offsets ($n_{\text{off,C}} = 25$) and ^1H offsets ($n_{\text{off,H}} = 15$) were combined, and the RF inhomogeneity ($n_{\text{rf}} = 10$) across both channels were aligned. The average J -coupling was set to $J_{\text{HC}} = 145$ Hz. Optimization was carried out in two steps:

- Step 1: Bandwidths of ^1H and ^{13}C were each divided into 10 intervals, and RF




Figure 5.8: CLOC pulse shapes for ^1H (a) and ^{13}C (b). Both pulses are cyclic and universally preserve $\mathbf{I}_-$, $\mathbf{I}_+$, and $\mathbf{I}_z$.

inhomogeneity into 3 intervals. The optimization stopped after 500 iterations or once the fidelity $\eta > 0.9985$. The resulting pulse was used as the initial guess for Step 2.

- Step 2: All parameters were included. The optimization stopped after 50 iterations or once the fidelity $\eta > 0.995$.

The two-spin optimization used 64 CPU cores at 2.6 GHz and took about two days. If all relevant coherences are preserved, the optimization effectively targets a unitary propagator, i.e., $\mathbf{U} = \mathbf{I}$, uniquely defined within the Hilbert space [84]. The optimization for this target is more computationally expensive than the state-to-state transfer, especially in the two-spin case. The results from single-spin optimization provide a good initial guess.

5.3.2 Coherence-locking efficiency

The performance of coherence locking can be visualized through spin trajectory simulations, as illustrated in Fig. 5.9. The sample was divided into 12 voxels along the z -axis, and the spin trajectory for each voxel was computed while simultaneously applying a CLOC pulse and a coupled gradient pulse that induced a maximum B_0 drift of ± 0.25 G. Fig. 5.9a illustrates that the ensemble spin states start from $\mathbf{I}_-$ and oscillate between -1 and 1 while remaining confined within a narrow range, and go back to $\mathbf{I}_-$ at the end. Figs. 5.9b and 5.9c show similar results for spin states initially prepared in $\mathbf{S}_+$ and $\mathbf{I}_-\mathbf{S}_+$, respectively. For comparison, Figs. 5.9d–f display the corresponding trajectories without coherence locking, which clearly demonstrate the dephasing caused by the coupled gradient.




Figure 5.9: Ensemble spin trajectories under gradient spillover. (a–c) Trajectories starting from $\mathbf{I}_-$, $\mathbf{S}_+$, and $\mathbf{I}_-\mathbf{S}_+$ with CLOC pulses applied to ^1H , ^{13}C , and both ^1H and ^{13}C , respectively. (d–f) Corresponding trajectories without CLOC pulses. Real parts of the inner product are plotted. Both ^1H and ^{13}C were on resonance, the J -coupling constant was 145 Hz, and the RF amplitudes were 6 kHz for ^1H and 4 kHz for ^{13}C .

For a single-spin system, the coherence-locking efficiency was further analyzed as a function of resonance offset, RF inhomogeneity, and gradient field strength (Figs. 5.10 and 5.11). The efficiency was quantified using the fidelity function

$$\eta = \text{trace}(\mathbf{E}^\dagger \mathbf{U}), \quad (5.15)$$

where $\mathbf{E}$ is the identity operator (target propagator) and $\mathbf{U}$ is the actual propagator realized under the CLOC pulse.




Figure 5.10: Efficiency of the ^1H CLOC pulse. (a) Dependence on ν_0 and B_0 drift at $\nu_1 = 6$ kHz. (b) Dependence on ν_0 and ν_1 at B_0 drift equals to 0.25 G. (c) Dependence on ν_1 and B_0 drift at $\nu_0 = 0$.




Figure 5.11: Efficiency of the ^{13}C CLOC pulse. (a) Dependence on ν_0 and B_0 drift at $\nu_1 = 4$ kHz. (b) Dependence on ν_0 and ν_1 at B_0 drift equals to 0.25 G. (c) Dependence on ν_1 and B_0 drift at $\nu_0 = 0$.

5.3.3 Gradient imperfections

The gradient coil can be modeled as a resistor–inductor circuit, which behaves as a low-pass filter. Consequently, the applied gradient pulse undergoes smoothing and phase delay. As an example, the distorted gradient pulse shapes are shown in Fig. 5.12, where a clear phase delay is observed between the input and the output for both the trapezoidal and sine pulses.

These transient effects alter the temporal alignment between gradient and RF pulses, thereby degrading the efficiency of coherence locking. To evaluate this, we simulated coherence-locking performance under gradient imperfections.




Figure 5.12: Measured gradient pulse shapes without proper pre-emphasis. The black traces represent the input pulse detected from the DAC port of the gradient amplifier, and the red traces represent the output pulse detected from the I monitor of the gradient amplifier. (a) Trapezoidal pulse with a duration of 1 ms. (b) Sine pulse with a duration of 0.5 ms.

The quality factor of the gradient coil is given by

$$Q = \frac{\omega L}{R}, \quad (5.16)$$

where $\omega = \pi/\tau$ with τ the duration of a sinusoidal pulse. The corresponding phase delay of the circuit is

$$\phi = \tan^{-1}(Q). \quad (5.17)$$

The resulting coherence-locking efficiency as a function of delay phase and B_0 drift is shown in Fig. 5.13.

5.3.4 Heteronuclear decoupling

In multi-spin systems, J -coupling can interfere with coherence locking because the additional interaction terms cannot be averaged out during the locking period. Interestingly, CLOC pulses can also provide partial decoupling of heteronuclear interactions due to their rapid phase modulations, similar to noise decoupling techniques. To analyze this effect, we evaluated the effective coupling scale factor using average Hamiltonian theory, as outlined in chapter 2.

In the toggling frame defined by $\mathbf{H}_{\text{rf}} + \mathbf{H}_z$, the heteronuclear J -coupling Hamiltonian




Figure 5.13: Coherence-locking efficiency in the presence of gradient coil imperfections. (a) Transient response of the gradient coil at different delay phases. (b) Efficiency of locking the I -coherence as a function of delay phase and B_0 drift, calculated with the ^1H CLOC pulse.

takes the form [47]

$$\mathbf{H}_J^{\text{tog}} = 2\pi J \sum_{\beta, \gamma} a_{z\beta}^I(t) a_{z\gamma}^S(t) \mathbf{I}_\beta \mathbf{S}_\gamma = 2\pi \mathbf{I} \cdot \mathbf{J}(t) \cdot \mathbf{S}, \quad (5.18)$$

where $\mathbf{J}(t)$ is a 3×3 time-dependent coupling tensor with elements

$$c_{\beta, \gamma}(t) = a_{z\beta}^I(t) a_{z\gamma}^S(t), \quad (5.19)$$

and $\beta, \gamma \in \{x, y, z\}$. The rotation coefficients for I spin are

$$a_{z\beta}^I(t) = \langle \mathbf{U}_I^\dagger(t) \mathbf{I}_z | I_\beta \rangle, \quad (5.20)$$

where the single-spin propagator is

$$\mathbf{U}_I(t) = \exp_0 \left[-i \int_0^t \mathbf{H}_I(t') dt' \right], \quad (5.21)$$

with $\exp_0$ denotes the Dyson time-ordering exponential and

$$\mathbf{H}_I(t) = \omega_I \mathbf{I}_z + \omega_{I,x}(t) \mathbf{I}_x + \omega_{I,y}(t) \mathbf{I}_y. \quad (5.22)$$

The rotation coefficients $a_{z\gamma}^S(t)$ for S spin are obtained following a similar procedure.

The time-dependent tensor $\mathbf{c}(t)$ was averaged over the pulse duration τ :

$$\overline{c_{\beta,\gamma}} = \frac{1}{\tau} \int_0^\tau c_{\beta,\gamma}(t) dt. \quad (5.23)$$

The coupling scale factor is then defined as the Frobenius norm of the averaged tensor,

$$\chi = \text{norm}(\overline{\mathbf{c}}), \quad (5.24)$$

which is bounded between 0 and 1.

Fig. 5.14 shows the coupling scale factor χ for CLOC pulses designed to universally preserve $\mathbf{I}_+$, $\mathbf{I}_-$, and $\mathbf{I}_z$ of a single spin. A single CLOC pulse reduces the effective heteronuclear J -coupling by less than 10% across the designed resonance offset and B_1 inhomogeneity ranges (Figs. 5.14a,b). For pulse durations of 1 ms, heteronuclear couplings on the order of a few hundred Hz can be neglected since $\chi J \tau \ll 1$. This treatment also extends to couplings with a third spin in ^1H – ^{13}C systems, such as J_{CN} and J_{HN} in proteins [150, 151]. Fig. 5.14c illustrates that residual coupling remains when CLOC pulses are applied simultaneously to ^1H and ^{13}C . However, this does not imply a failure of decoupling, as χ represents a sum over all tensor components. The decoupling effect can also be assessed by directly simulating coherence locking of double-quantum coherences, as demonstrated for HMQC in Section 5.6.




Figure 5.14: Scale factor χ of the J_{HC} coupling as a function of resonance offset and B_1 amplitude. (a) CLOC pulse applied to ^1H with ^{13}C on resonance. (b) CLOC pulse applied to ^{13}C with ^1H on resonance. (c) CLOC pulses applied simultaneously to ^1H and ^{13}C , with matched ν_0 and ν_1 values.

5.3.5 Homonuclear decoupling

In principle, homonuclear J -coupling can also be incorporated into the optimal control model. For example, when optimizing a universal locking pulse for a single ^{13}C spin, the model may include two coupled ^{13}C spins, with a resonance offset applied to one spin while the other remains at zero offset. Denoting the two spins as I_1 and I_2 , the J -coupling Hamiltonian

$$\mathbf{H}_J = 2\pi J (\mathbf{I}_{1x}\mathbf{I}_{2x} + \mathbf{I}_{1y}\mathbf{I}_{2y} + \mathbf{I}_{1z}\mathbf{I}_{2z}), \quad (5.25)$$

is then added to the drift Hamiltonian, with J set to a representative value such as 40 Hz for J_{CC} . Including such terms increases the RF power required for the optimized pulse.




Figure 5.15: Simulated coherence-locking efficiency of a ^{13}C CLOC pulse in a two- ^{13}C spin system. Efficiency is defined as $\eta = \langle \rho_0 | \rho_T \rangle$, with ρ_0 the initial state and ρ_T the final state. $\delta\nu_0 = \omega_1 - \omega_2$, with $\omega_1 = 0$ and ω_2 swept from -3 to 3 kHz. $J = 50$ Hz, ν_1 is the RF amplitude. (a) $\rho_0 = \mathbf{I}_{1x} + \mathbf{I}_{2x}$, which commutes with $\mathbf{H}_J$. (b–c) $\rho_0 = \mathbf{I}_{1x}$, using a CLOC pulse without J -coupling compensation (b), and with compensation requiring increased RF amplitude of 6 kHz (c).

However, complete averaging of homonuclear J -coupling is generally not feasible when the chemical shift difference between the two spins is comparable to the coupling constant. This can be understood using the average Hamiltonian theory. The full Hamiltonian is

$$\mathbf{H} = \mathbf{H}_Z + \mathbf{H}_J + \mathbf{H}_{\text{rf}}, \quad (5.26)$$

where

$$\begin{aligned} \mathbf{H}_Z &= \omega_1 \mathbf{I}_{1z} + \omega_2 \mathbf{I}_{2z}, \\ \mathbf{H}_{\text{rf}} &= \omega_x (\mathbf{I}_{1x} + \mathbf{I}_{2x}) + \omega_y (\mathbf{I}_{1y} + \mathbf{I}_{2y}). \end{aligned} \quad (5.27)$$

Since $[\mathbf{H}_{\text{rf}}, \mathbf{H}_J] = 0$, the RF field alone cannot average out homonuclear couplings. When $|\omega_1 - \omega_2| \approx J$, the commutator $[\mathbf{H}_z, \mathbf{H}_J]$ is nearly zero, so the $\mathbf{H}_z$ term also fails to average out the coupling. Consequently, second-order spectra cannot be fully recovered by CLOC pulses in this regime. When $|\omega_1 - \omega_2| \gg J$, partial averaging of $\mathbf{H}_J$ by $\mathbf{H}_{\text{rf}} + \mathbf{H}_z$ becomes possible. Fig. 5.15 illustrates the effect of homonuclear coupling on the efficiency of a ^{13}C CLOC pulse.

5.4 Coherence locking in PGSE

The PGSE sequence introduces gradient pulses before and after the 180° pulse to measure the diffusion constant of liquids [147]. Here, coherence locking was evaluated with the PGSE. As shown in Fig. 5.16a, an additional gradient, G_C , was incorporated as the coupling component into a standard PGSE sequence, temporally shifted relative to the primary gradient G . When G_C disrupted the balance of gradient strengths around the inversion pulse, two CLOC pulses were applied, each aligned with one block of G_C , to counteract gradient spillover. These pulses should protect $\mathbf{I}_+$ during the first G_C period and $\mathbf{I}_-$ during the second. Alternatively, the same pulses can enforce effective cyclic propagation, which is a stricter condition and was chosen as the design target.




Figure 5.16: Coherence-locking test in PGSE using 0.6 M glycine-2- ^{13}C in D_2O . (a) Pulse sequence: the ^1H pulse and G were applied to detector 1, while G_C was applied to detector 2. (b) Normalized water peak intensity as a function of G_C amplitude, with G fixed at 10%. Acquisition parameters: TD = 1024, SW = 20 ppm, RG = 10, TR = 15 s, NS = 1.

In the PGSE experiment, the gradients G and G_C were driven by separate Bruker GREAT amplifiers, as described in chapter 4. A 0.6 M glycine-2- ^{13}C solution in D_2O

was used, loaded into detector 1 via a syringe. Only the global shim was applied using the Bruker shim system, since data were acquired solely from detector 1. Gradient amplitude sweeping was performed in a pseudo-2D manner: G was fixed at 10%, while the G_C amplitude was varied from 0 to 95% in 5% increments. The acquired data were Fourier transformed along the direct dimension and converted into 1D spectra corresponding to each G_C value.

Figure 5.16b presents the normalized water peak intensity as a function of G_C amplitude. Compared to the substantial signal loss in the uncoupled case, the inclusion of CLOC pulses effectively restored the signal intensity.

5.5 Coherence locking in parallel HSQC

5.5.1 Simulation

The pulse sequence for parallel HSQC is shown in Fig. 5.17, where standard HSQC schemes were implemented in detectors 1 and 2. In detector 1, the gradient ratio was set to 4 : 1 to select the $S_+ \rightarrow I_-$ pathway (Fig. 5.17a), while in detector 2 it was set to 4 : -1 to select the $S_- \rightarrow I_-$ pathway (Fig. 5.17b). Compensation CLOC pulses (blue) were inserted in detector 1 when gradient pulses were applied in detector 2, and vice versa. A slight temporal offset between gradients in the two detectors allowed insertion of the CLOC blocks. Each CLOC pulse was designed to lock the relevant coherence. For example, the first CLOC block in detector 1 was applied on ^{13}C to protect S_+ and its J -coupling product $I_z S_+$. The protected coherences are indicated below the CLOC blocks in Fig. 5.17.

If multiple coherences are required, e.g., in echo-antiecho acquisition, the CLOC pulse must preserve all relevant pathways. During the CLOC blocks, coherence evolution can be neglected and should be excluded from J -coupling delay calculations due to the heteronuclear decoupling effect.

Simulations were performed in SPINACH, assuming a homogeneous B_0 field (11.74 T) and homogeneous B_1 fields for the CLOC pulses. CLOC pulse shapes shown in Fig. 5.8 were used. Ideal $90^\circ/180^\circ$ propagators were assumed for excitation and inversion. Each sample was divided into 40 voxels along z , with maximum B_0 drifts of ± 25 G from the primary gradient and ± 0.2 G from the coupled gradient. Glycine and D-glucose were




Figure 5.17: Gradient-compensated parallel HSQC sequence. Blue blocks indicate CLOC pulses. In detector 1 (a), the gradient ratio was set to 4 : 1 to select the $S_+ \rightarrow I_-$ pathway; in detector 2 (b), it was set to 4 : -1 to select the $S_- \rightarrow I_-$ pathway. Black blocks represent $\pi/2$ pulses, white blocks π pulses, with phases set to 0 unless noted. $\Delta = 1/4J$. Protected coherences are labeled below the CLOC blocks.

assigned to detectors 1 and 2, respectively. The glucose spin system followed literature data for α -D-glucopyranose in D_2O (25 °C, pH 7.0) [152, 153]. To accelerate the simulation, product states up to the third-order spin correlations were included, and the spin system of glucose was split into $\binom{6}{2} = 15$ subsystems, with each subsystem containing two instances of ^{13}C spins, so that the homonuclear couplings, i.e., J_{CC} and J_{HH} , could be included. For each t_1 point, both S_+ and S_- were recorded by flipping the second gradient pulse, yielding quadrature detection. Data were apodized with a $\cos^2(x)$ half-bell function in both dimensions, Fourier transformed, and the real part was retained. Acquisition parameters are summarized in Table 5.1.

The simulated coherence-locking HSQC spectra are shown in Fig. 5.18. Noise was neglected, and no phase or baseline correction was required. Grey traces represent the reference without gradient coupling, orange traces the uncompensated case, and blue traces the CLOC-compensated case (offset for clarity). For glycine, the S_+ coherence contributes,

Parameters	Glycine (detector 1)		D-glucose (detector 2)	
	^1H	^{13}C	^1H	^{13}C
Gradient ratio	4:1		4:-1	
Sweep width, Hz	2000	5000	2000	6250
Transmit offset, Hz	2000	5000	2000	10625
Sampling points	512	1024	512	1024
Zero-fill points	1024	1024	1024	1024

Table 5.1: Acquisition parameters for simulating parallel HSQC spectra. Sweep widths were chosen to match the spectral ranges of glycine and D-glucose. A J value of 145 Hz was used for the evolution delay.

giving ^{13}C signals at $-\omega_C$ (negative chemical shifts). Compared with the reference, gradient coupling caused $\sim 50\%$ signal loss, whereas coherence locking preserved full signal amplitude without phase distortion.




Figure 5.18: Simulated parallel HSQC spectra of glycine (a) and glucose (b). Panels (c) and (d) show the corresponding 1D projections. Grey: reference (no coupling). Orange: with gradient spillover (no compensation). Blue: with CLOC compensation.

5.5.2 Experiment

The parallel HSQC experiment was conducted using the setup described in chapter 4. Two samples in D₂O (99.9%) were prepared: 0.6 M glycine-2-¹³C and 0.3 M D-glucose-¹³C₆ (Sigma-Aldrich). The solutions were loaded into syringes and manually pumped into the fluidic chambers of the respective detectors.

Detectors 1 (D1) and detector 2 (D2) were individually tuned and matched for ¹H/¹³C double resonance. The split X and Y micro-imaging gradients were connected to D1 and D2, respectively. Pulse calibration was performed separately: in D1, hard 90° pulses were 7.4 μs (¹H, 20 W) and 25 μs (¹³C, 25 W); in D2, 6.1 μs (¹H, 20 W) and 19.0 μs (¹³C, 25 W). Power levels were scaled down for the low-amplitude CLOC pulses. The parallel experiments under the three conditions were each repeated twice. Data from D1 were acquired with global shimming optimized for D1, and data from D2 were acquired with shimming optimized for D2.

The resulting spectra are shown in Fig. 5.19. Relative to the reference spectra (left column), gradient-coupled spectra (middle column) display severe artifacts from spin dephasing, whereas coherence-locking (right column) restores peak shape and SNR.

The 1D projections of glycine and glucose (Fig. 5.20) further highlight the effect. Reference spectra (grey) show intact peak intensities, whereas gradient coupling (orange) leads to signal suppression. With coherence-locking (blue), intensities are largely recovered. Ideally, broadband CLOC pulses should preserve ¹H and ¹³C coherences equally, yielding peak intensities identical to the reference. However, in Fig. 5.20b, the peaks at 3.8 ppm/72.8 ppm and 3.9 ppm/79.1 ppm display reduced intensity, reflecting degraded performance when homonuclear couplings are not compensated. Simulation results in Fig. 5.21 confirm this effect.

Residual discrepancies also stem from chemical shift dependence. Each CLOC pulse introduces minor phase errors: for example, 95% single-pulse fidelity compounds to 90% after two pulses. Furthermore, RF leakage between detectors contributes: when a CLOC pulse is applied to D1, weak RF signals (~ 1% at 500 MHz [35]) can reach D2, disturbing spin states. To avoid this, CLOC amplitudes were reduced to 6 kHz (¹H) and 4 kHz (¹³C), corresponding to 0.43 W and 2.3 W on detector 2. The higher RF power in the ¹³C channel illustrates the particular challenge of coherence locking for low-sensitivity nuclei when




Figure 5.19: Parallel HSQC spectra. Left: reference without gradient coupling. Middle: spectra under gradient spillover. Right: spectra with coherence-locking compensation. Spectra were exported from TopSpin and plotted in MATLAB. Acquisition parameters: TD1 = 256, TD2 = 1024, SW1 = 150 ppm, SW2 = 10 ppm, TR = 1 s, RG = 10, NS = 1, DS = 32.




Figure 5.20: 1D projections of parallel HSQC spectra for glycine-2- ^{13}C (a) and D-glucose- $^{13}\text{C}_6$ (b), using the sequence in Fig. 5.17. Grey: reference without gradient coupling. Orange: gradient spillover (no compensation). Blue: coherence-locking compensation.

wide bandwidths are required.




Figure 5.21: Simulated strong-coupling effects on D-glucose HSQC spectra. The relative intensity is the ratio of coherence-locking to reference peak heights. Simulations excluded (blue) and included (orange) homonuclear J couplings. Missing peaks reflect insufficient spectral resolution. Spin systems from [153, 154, 155].

5.6 Simulation of coherence locking in parallel HMQC

A compensation strategy similar to that used in the parallel HSQC was applied to the parallel HMQC pulse sequence [156], as illustrated in Fig. 5.22. In detector 1, the gradient pulse ratio was set to 2:2:1 to select the $\mathbf{I}_+\mathbf{S}_+ \rightarrow \mathbf{I}_-\mathbf{S}_+ \rightarrow \mathbf{I}_-$ coherence pathway. In contrast, in detector 2, the gradient pulse ratio was adjusted to 2:2:-1 to select the $\mathbf{I}_+\mathbf{S}_- \rightarrow \mathbf{I}_-\mathbf{S}_- \rightarrow \mathbf{I}_-$ pathway. The CLOC pulses were applied to both the ^1H and ^{13}C channels, depending on the double-quantum coherence to be protected.

The parallel HMQC sequence was simulated under homogeneous B_0 field conditions ($B_0 = 11.74$ T) and assuming ideal hard pulses, analogous to the HSQC simulation. The sample was segmented into 40 voxels along the z -direction. A maximum B_0 drift of ± 25 G was introduced by the primary gradient, and an additional drift of ± 0.2 G by the coupled




Figure 5.22: Gradient-compensated parallel HMQC pulse sequence with coherence locking indicated by the blue blocks. (a) Detector 1: gradient pulse ratio of 2:2:1 to select the $I_+S_+ \rightarrow I_-S_+ \rightarrow I_-$ pathway. (b) Detector 2: gradient pulse ratio of 2:2:-1 to select the $I_+S_- \rightarrow I_-S_- \rightarrow I_-$ pathway. Black blocks denote $\pi/2$ pulses and white blocks π pulses; all phases were set to zero unless specified. $\Delta = 1/4J$.

Parameter	Glycine (detector 1)		D-glucose (detector 2)	
	^1H	^{13}C	^1H	^{13}C
Gradient ratio	2:2:1		2:2:-1	
Spectral width (Hz)	2000	5000	2000	6250
Transmit offset (Hz)	2000	5000	2000	10625
Sampling points	512	1024	512	1024
Zero-filling points	1024	1024	1024	1024

Table 5.2: Acquisition parameters for simulating parallel HMQC spectra. The spectral widths were chosen to match those of the glycine and D-glucose samples. The evolution delay was set using $J = 145$ Hz.

gradient. The B_1 field of the CLOC pulse was linearly scaled with voxel position, being largest at the center and smallest at the edges, accounting for the B_1 pattern inside a stripline.

Glycine and D-glucose were used as representative samples in detectors 1 and 2, respectively. The spin systems were the same as in the HSQC simulation. For each detector, only

S_+ or S_- coherence was recorded in amplitude-mode acquisition. Data were apodized with a $\cos(x)$ half-bell function in both dimensions, Fourier transformed, and the absolute value was retained. Acquisition parameters are summarized in Table 5.2.




Figure 5.23: Simulated parallel HMQC spectra. Panels (a) and (b) show the 2D contours for glycine and glucose without gradient coupling. Panels (c) and (d) show the corresponding 1D projections. Grey: no gradient coupling; orange: with gradient spillover, no compensation; blue: with coherence-locking compensation (shifted for clarity).

Figure 5.23 shows the magnitude spectra. Grey lines represent the ideal case without gradient coupling. Orange 1D projections correspond to spectra with gradient spillover but without compensation; the resulting signal loss is much larger than in the HSQC case, because higher-order spin coherences are more sensitive to B_0 inhomogeneity. Blue projections correspond to spectra with coherence-locking compensation, indicating signal recovery in both dimensions.

As shown in Fig. 5.23(a,c), glycine has two magnetically equivalent ^1H spins, i.e., with the same chemical shift, giving rise to a singlet in the ^1H dimension. The ^{13}C - ^{13}C coupling, combined with the two equivalent ^1H - ^{13}C couplings, results in a triplet in the ^{13}C

dimension. In addition, the ^{13}C projection with coherence locking (blue trace) exhibits two symmetric spikes flanking the central peak, located just inside the two outer lines of the triplet, indicating residual ^{13}C - ^{13}C coupling that is only partially averaged out during the coherence-locking period. The same feature applies to the coherence-locked glucose spectra, although the tiny spikes are obscured in the more crowded spectra.

5.7 Conclusion

The coherence-locking (CLOC) pulses demonstrated robustness against a wide range of B_0 drifts, thereby mitigating gradient spillover in parallel gradient coils. This gradient-decoupling concept was validated experimentally in a PGSE experiment and in a parallel HSQC sequence.

These results establish CLOC pulses as an effective strategy for coherence protection in the presence of electromagnetic fluctuations and spin–spin interactions [142]. Such protection has enabled a wide range of applications in quantum physics, including electron spin locking for radiofrequency magnetometry [143], suppression of intrabath interactions and field inhomogeneity using sequenced pulses [144], and extension of coherence times via continuous microwave driving [145].

The limitations of this approach arise primarily from RF coupling between detectors and from homonuclear interactions in complex spin systems. As shown by the parallel HSQC experiments and HMQC simulations, single-spin coherences can be robustly locked, and double-quantum coherences (IS) can be preserved using two simultaneous locking pulses. However, extending this strategy to higher-order heteronuclear coherences is considerably more difficult, as multiple locking pulses introduce strong RF coupling. While optimal control can compensate for homonuclear coupling by increasing RF power, this strategy fails in the strong-coupling regime where second-order spectra arise. In such cases, the RF Hamiltonian commutes with the J -coupling term, while the Zeeman and J -coupling terms nearly commute with each other, preventing efficient decoupling. More advanced schemes, which jointly address RF coupling and gradient spillover, may be required for universal high-order coherence locking.

Extending this scheme to general parallel pulse sequences requires careful design. For instance, when executing HMQC on one detector and HSQC on another, the acquisitions

are no longer strictly synchronous, even though the relaxation delays on each detector can be adjusted to align the parallel sequences. Consequently, interactions between RF pulses and FID acquisition must be avoided, an issue that lies beyond the scope of this chapter.

Chapter 6

Accelerated quantum optimal control for magnetic resonance pulse design

Based on M. He, Y. Deng, B. Luy, J. G. Korvink, The Journal of Chemical Physics, 164 (6): 064114, 2026.

In designing RF and gradient decoupling schemes for parallel NMR spectroscopy, the optimal control method functions as the primary approach. The LBFGS-GRAPE method we currently use in the optimal control can require several hours on an HPC system to optimize a coherence-locking pulse. In this chapter, we present a strategy that improves computational efficiency by more than an order of magnitude. We investigate applying the finite element method, previously used in our electromagnetic simulations, to accelerate the solution of the Liouville–von Neumann equation, and employing the method of moving asymptotes (MMA) [157], a technique widely used in topology optimization, to speed up convergence.

6.1 Introduction

In magnetic resonance spectroscopy (MRS) and imaging (MRI), shaped radio frequency (RF) pulses are widely employed to drive the state of a spin system toward a desired target. The design of such pulses is typically framed as an optimal control problem [158].

Among gradient-based optimization techniques, the gradient ascent pulse engineering (GRAPE) method [55] has become one of the most widely used and efficient methods. The original GRAPE method discretizes the control time into piecewise-constant intervals and computes the spin trajectory via a bi-directional propagation, where gradients are obtained using a first-order finite-difference approximation under sufficiently fine time slicing.

A variety of extensions to GRAPE have since been developed, targeting more accurate derivative computation, advanced optimization strategies, and broader applicability to general quantum systems. The derivatives of the sub-propagator to control variables have been derived by differentiating the Taylor series [159]. A higher-order finite difference method was introduced to improve gradient accuracy [56]. The auxiliary matrix formalism (AUXMAT) was introduced to compute exact gradients and Hessians, leading to improved convergence [160]. Automatic differentiation, implemented on graphics processing units, was introduced to accelerate gradient computation [161, 162]. Beyond the gradient-descent update of the original formulation, several optimization algorithms have been incorporated into GRAPE frameworks, including Quasi-Newton approaches [56, 163, 164], such as the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm and its limited-memory version (L-BFGS), and Newton–Raphson methods [61] that provide quadratic convergence. The accelerated Newton–Raphson variant was developed to reduce the Hessian evaluation to $O(N)$ complexity, as demonstrated for a single-spin system [165]. Extension algorithms for open quantum systems demonstrated high-fidelity control when handling decoherence effects [162, 166].

As modern magnetic resonance systems operate at increasingly higher frequencies, various instrumental limitations have become more prominent, including RF power constraints [82, 83, 167] and hardware-induced distortions [74, 73, 168, 169]. Optimization problems also grow more complex when modeling multi-qubit systems exhibiting entanglement [170], heteronuclear spin systems with J -coupling [171], or crystalline orientation distributions in solid-state MRS [172]. In parallel transmit MRI, pulse design under safety constraints can involve thousands of control variables and hundreds of spatial voxels [91, 173, 174]. To address crosstalk in parallel MRS, compensation strategies incorporate an ensemble of B_0 distortions caused by gradient coils [36], along with multiple cooperative pulses tailored for parallel excitation [35].

These scenarios often require ensemble-based optimization, where system parameters

vary across ensemble members. In such cases, the gradient or Hessian must be computed for each ensemble element, and control updates are determined from the ensemble-averaged derivatives, as in the L-BFGS and Newton methods. This significantly increases the computational cost, potentially requiring several hours or even days of high-performance computing time.

In this work, we address the optimization of single-spin magnetic resonance pulses by combining the finite element method (FEM) [175, 176] with the method of moving asymptotes. Finite elements with linear, quadratic, and cubic Hermite basis functions provide an efficient and scalable solution for computing single-spin optimal-control gradients. By formulating the ensemble-averaged fidelity maximization as a constrained optimization, MMA converges more rapidly than L-BFGS and yields the lowest overall runtime among L-BFGS and Newton–Raphson methods.

6.2 Methodology

The Liouville-von Neumann (LvN) equation masters the evolution of a general spin system. The equation in Liouville space is given by

$$\begin{aligned}\frac{d}{dt}\boldsymbol{\rho}(t) + i\mathbf{L}\boldsymbol{\rho}(t) &= 0, \\ \boldsymbol{\rho}(0) &= \boldsymbol{\rho}_0,\end{aligned}\tag{6.1}$$

where $\boldsymbol{\rho}_0$ is the initial state. While ignoring the relaxation effect, the Liouvillian $\mathbf{L}$ can be decomposed into the internal part and the control part,

$$\mathbf{L}(t) = \mathbf{L}_{\text{int}}(t) + \sum_k x_k(t)\mathbf{L}_k,\tag{6.2}$$

where $\mathbf{L}_{\text{int}}$ could contain Zeeman interaction with the magnetic field and spin-spin couplings. The $\mathbf{L}_k$ represents a control operator and $x_k(t)$ is a time-dependent coefficient. A control sequence $\mathbf{x}(t)$ is applied to steer the spin system from an initial state $\boldsymbol{\rho}_0$ to a target state $\boldsymbol{\sigma}$ in a specified time duration T . A measurement of the control efficiency is the overlap between the target state and the actual final state $\boldsymbol{\rho}_T$, i.e.,

$$\eta = \langle \boldsymbol{\sigma} | \boldsymbol{\rho}_T \rangle,\tag{6.3}$$

which satisfies $-1 \leq \eta \leq 1$. We consider an ensemble spin system with N_{ens} members, for which the control amplitudes are limited, so that the optimal control problem can be defined as:

$$\begin{cases} \text{Find } \mathbf{x}(t), t \in [0, T], \\ \text{to maximize } f = \sum_{n=1}^{N_{\text{ens}}} \eta_n, \\ \text{constrained by } |x_k(t)| \leq x_k^{\text{max}}. \end{cases} \quad (6.4)$$

6.2.1 FEM solution of the Liouville-von Neumann equation

Linear elements

In the context of magnetic resonance, the FEM has previously been applied to solve the stochastic Liouville equation for chemically induced spin polarization problems [177, 178] and for electron spin resonance spectral simulations [179]. In Eq. 6.1, the spatial variables are absent relative to the stochastic Liouville equation. Within the Hamilton principle, the integration domain, time interval $[0, T]$, is discretized into N elements with $N + 1$ nodes. The control variables are represented as a piecewise-constant waveform $\mathbf{x}$, rendering the Liouvillian $\mathbf{L}$ constant within each element. The solution of Eq. 6.1 is approximated as a linear combination of shape functions,

$$\boldsymbol{\rho}(t) = \sum_{j=1}^{N+1} \alpha_j \phi_j(t), \quad (6.5)$$

where α_j has the same dimension as the spin state vector $\boldsymbol{\rho}(t)$, and $\phi_j(t)$ represents the j -th shape function. As shown in Fig. 6.1, by using the linear elements, the nodal shape functions globally defined at node j are expressed as

$$\phi_j^n(t) = \begin{cases} (t - t_{j-1})/\Delta t, & t_{j-1} \leq t \leq t_j, \\ (t_{j+1} - t)/\Delta t, & t_j \leq t \leq t_{j+1}, \\ 0, & \text{otherwise.} \end{cases} \quad (6.6)$$

Note that it's also possible to define the two linear shape functions locally at element j as

$$\phi_j^{e1}(t) = (t - t_j)/\Delta t, \quad \phi_j^{e2}(t) = (t_j + \Delta t - t)/\Delta t, \quad t_j \leq t \leq t_j + \Delta t. \quad (6.7)$$




Figure 6.1: View of the one-dimensional linear Lagrange global shape functions. The time interval between two nodes is uniformly set to Δt . The labels $\mathbf{L}_1, \mathbf{L}_2, \dots, \mathbf{L}_N$ represent the discretized Liouvillian for a piecewise-constant waveform.

In the following derivation, globally defined shape functions in Eq. 6.6 are used, and the superscript n is omitted.

The Galerkin's method [63] demonstrated that the integral of the weighted residual equals zero, i.e.,

$$\int_0^T \left(\frac{d\rho}{dt} + i\mathbf{L}\rho \right) \hat{\omega} dt = 0. \quad (6.8)$$

where ρ is the approximated solution expressed by Eq. 6.5, $\hat{\omega}$ is the chosen weighting function. Selecting each shape function as the weighting function, i.e., $\hat{\omega} = \phi_i$ ($i = 1, 2, \dots, N+1$), and substituting Eq. 6.5 into Eq. 6.8 gives the following linear equations,

$$\sum_{j=1}^{N+1} \alpha_j \int_0^T \left(\frac{d\phi_j}{dt} + i\mathbf{L}\phi_j \right) \phi_i dt = 0, \quad i = 1, 2, \dots, N+1. \quad (6.9)$$

In matrix form, this can be written compactly as:

$$\mathbf{K} \cdot \boldsymbol{\alpha} = \mathbf{f}, \quad (6.10)$$

where $\boldsymbol{\alpha}$ is the vectorized α_j ($j = 1, 2, \dots, N+1$), $\mathbf{f} = 0$ is a forcing function, often called the load vector, and $\mathbf{K}$ is an impedance function, often called the stiffness matrix. The terms load and stiffness appeared first in the finite element literature because of the

application to computational mechanics. An element of $\mathbf{K}$ is given by:

$$K_{ij} = \int_0^T \left(\frac{d\phi_j}{dt} + i\mathbf{L}\phi_j \right) \phi_i dt = \sum_{n=1}^N \int_{t_n}^{t_{n+1}} \left(\frac{d\phi_j}{dt} + i\mathbf{L}\phi_j \right) \phi_i dt. \quad (6.11)$$

Considering that ϕ_j is nonzero only within the elements connected to node j , the global stiffness matrix $\mathbf{K}$ can be assembled from the element stiffness matrices:

$$\mathbf{K} = \begin{bmatrix} K_{11}^{e,1} & K_{12}^{e,1} & & & & \\ K_{21}^{e,1} & K_{22}^{e,1} + K_{11}^{e,2} & K_{12}^{e,2} & & & \\ & K_{21}^{e,2} & K_{22}^{e,2} + K_{11}^{e,3} & K_{12}^{e,3} & & \\ & & K_{21}^{e,3} & \ddots & & \\ & & & & K_{22}^{e,N-1} + K_{11}^{e,N} & K_{12}^{e,N} \\ & & & & K_{21}^{e,N} & K_{22}^{e,N} \end{bmatrix}, \quad (6.12)$$

which reveals its banded structure, while the stiffness matrix of the j -th element is given by

$$K^{e,j} = \begin{bmatrix} \int_{t_j}^{t_{j+1}} \left(\frac{d\phi_j}{dt} + i\mathbf{L}_j\phi_j \right) \phi_j dt & \int_{t_j}^{t_{j+1}} \left(\frac{d\phi_{j+1}}{dt} + i\mathbf{L}_j\phi_{j+1} \right) \phi_j dt \\ \int_{t_j}^{t_{j+1}} \left(\frac{d\phi_j}{dt} + i\mathbf{L}_j\phi_j \right) \phi_{j+1} dt & \int_{t_j}^{t_{j+1}} \left(\frac{d\phi_{j+1}}{dt} + i\mathbf{L}_j\phi_{j+1} \right) \phi_{j+1} dt \end{bmatrix}. \quad (6.13)$$

Substituting the expressions for the linear shape functions yields:

$$K^{e,j} = \begin{bmatrix} -\frac{\mathbf{E}}{2} + \frac{i\mathbf{L}_j}{3}\Delta t & \frac{\mathbf{E}}{2} + \frac{i\mathbf{L}_j}{6}\Delta t \\ -\frac{\mathbf{E}}{2} + \frac{i\mathbf{L}_j}{6}\Delta t & \frac{\mathbf{E}}{2} + \frac{i\mathbf{L}_j}{3}\Delta t \end{bmatrix}. \quad (6.14)$$

Equation 6.14 defines the element stiffness matrix derived from a general Liouvillian. For a single-spin system, the local element spin vector is $\boldsymbol{\rho} \in \mathbb{C}^{4 \times 1}$, the local element Liouvillian is $\mathbf{L}_j \in \mathbb{C}^{4 \times 4}$, and $\mathbf{E}$ denotes the identity matrix. Each element stiffness matrix $K^{e,j} \in \mathbb{C}^{8 \times 8}$ couples two nodes, while the corresponding global stiffness matrix is $\mathbf{K} \in \mathbb{C}^{(4N+4) \times (4N+4)}$. The spin trajectory is obtained through four steps, presented in Algorithm 1.

Quadratic elements

To improve accuracy, quadratic shape functions were employed. The element shape functions are given by

$$\phi_1(\xi) = \frac{1}{2}\xi(\xi - 1), \quad \phi_2(\xi) = 1 - \xi^2, \quad \phi_3(\xi) = \frac{1}{2}\xi(\xi + 1), \quad \xi \in [-1, 1]. \quad (6.15)$$

Algorithm 1: Solve the linear system.

Input: Initial state ρ_0 , stiffness matrix $\mathbf{K}$ and load vector $\mathbf{f}$
Output: Solution vector α

Initialize $\alpha_{1:4} \leftarrow \rho_0$;

Update load vector: $\mathbf{f} \leftarrow \mathbf{f} - \mathbf{K}_{[:, 1:4]} \rho_0$;

Modify stiffness matrix: $\mathbf{K} \leftarrow \begin{bmatrix} \mathbb{E}_4 & \\ & \mathbf{K}_{[5:4N+4, 5:4N+4]} \end{bmatrix}$;

Solve linear system: $\mathbf{K}\alpha = \mathbf{f}$;

Assuming the Liouvillian is constant within an element of size $2\Delta t$, the element stiffness matrix can be written as a 3×3 block matrix:

$$K_{mn}^{e,j} = \int_{-1}^1 \left[\frac{d\phi_n}{d\xi} \frac{1}{\Delta t} + i\mathbf{L}_j \phi_n \right] \phi_m \cdot \Delta t d\xi, \quad (6.16)$$

where $\mathbf{L}_j$ denotes the Liouvillian on the j th element and $m, n \in \{1, 2, 3\}$. The element stiffness matrix is therefore

$$K^{e,j} = \begin{bmatrix} -\frac{1}{2} & \frac{2}{3} & -\frac{1}{6} \\ -\frac{2}{3} & 0 & \frac{2}{3} \\ \frac{1}{6} & -\frac{2}{3} & \frac{1}{2} \end{bmatrix} \otimes \mathbf{E} + \begin{bmatrix} \frac{4}{15} & \frac{2}{15} & -\frac{1}{15} \\ \frac{2}{15} & \frac{16}{15} & \frac{2}{15} \\ -\frac{1}{15} & \frac{2}{15} & \frac{4}{15} \end{bmatrix} \otimes i\mathbf{L}_j \Delta t \quad (6.17)$$

where $\mathbf{E}$ denotes the identity matrix. The global stiffness matrix is given by

$$\mathbf{K} = \begin{bmatrix} K_{11}^{e,1} & K_{12}^{e,1} & K_{13}^{e,1} & & & & \\ K_{21}^{e,1} & K_{22}^{e,1} & K_{23}^{e,1} & & & & \\ K_{31}^{e,1} & K_{32}^{e,1} & K_{33}^{e,1} + K_{11}^{e,2} & K_{12}^{e,2} & & & \\ & & K_{21}^{e,2} & \ddots & \ddots & & \\ & & & \ddots & K_{33}^{e,N-1} + K_{11}^{e,N} & K_{12}^{e,N} & K_{13}^{e,N} \\ & & & & K_{21}^{e,N} & K_{22}^{e,N} & K_{23}^{e,N} \\ & & & & K_{31}^{e,N} & K_{32}^{e,N} & K_{33}^{e,N} \end{bmatrix}. \quad (6.18)$$

For a single-spin system, the Liouvillian $\mathbf{L} \in \mathbb{C}^{4 \times 4}$, element stiffness matrix $K^{e,j} \in \mathbb{C}^{12 \times 12}$ connects three nodes, and the global stiffness matrix $\mathbf{K} \in \mathbb{C}^{(8N+4) \times (8N+4)}$.

Hermite elements

When a piecewise-linear waveform is adopted to mitigate instrumental distortions, the Hamiltonian should remain continuous across time intervals. Hermite shape functions enforce C^1 continuity between elements and are therefore well-suited. If one needs to




Figure 6.2: View of the piecewise-linear waveform, the Liouvillians are defined on the nodes.

optimize a piecewise-linear waveform in which the Liouvillian is continuous across time intervals [73], which is shown in Fig. 6.2. The cubic Hermite shape functions enforce C^1 continuity between elements and are therefore appropriate for this case. The element shape functions are given by

$$\begin{aligned}\phi_1(\xi) &= \frac{1}{4}(1 - \xi)^2(2 + \xi), & \phi_2(\xi) &= \frac{1}{4}(1 - \xi)^2(1 + \xi), \\ \phi_3(\xi) &= \frac{1}{4}(1 + \xi)^2(2 - \xi), & \phi_4(\xi) &= \frac{1}{4}(1 + \xi)^2(\xi - 1), & \xi \in [-1, 1].\end{aligned}\quad (6.19)$$

The element stiffness matrix can be expressed as a 4×4 block matrix, i.e.,

$$K_{mn}^{e,j} = \int_{-1}^1 \left[\frac{d\phi_n}{d\xi} \frac{2}{\Delta t} + i \left(\mathbf{L}_j + \frac{\xi + 1}{2} (\mathbf{L}_{j+1} - \mathbf{L}_j) \right) \phi_n \right] \phi_m \cdot \frac{\Delta t}{2} d\xi, \quad (6.20)$$

where $\mathbf{L}_j$ and $\mathbf{L}_{j+1}$ denote the Liouvillians at the left and right nodes of the j th element, respectively, and $m, n \in \{1, 2, 3, 4\}$. The element stiffness matrix is computed as

$$K^{e,j} = \begin{bmatrix} -\frac{1}{2} & \frac{1}{5} & \frac{1}{2} & -\frac{1}{5} \\ -\frac{1}{5} & 0 & \frac{1}{5} & -\frac{1}{15} \\ -\frac{1}{2} & -\frac{1}{5} & \frac{1}{2} & \frac{1}{5} \\ \frac{1}{5} & \frac{1}{15} & -\frac{1}{5} & 0 \end{bmatrix} \otimes \mathbf{E} + \begin{bmatrix} \frac{2}{7} & \frac{1}{14} & \frac{9}{140} & -\frac{1}{30} \\ \frac{1}{14} & \frac{1}{42} & \frac{1}{35} & -\frac{1}{70} \\ \frac{9}{140} & \frac{1}{35} & \frac{3}{35} & -\frac{1}{30} \\ -\frac{1}{30} & -\frac{1}{70} & -\frac{1}{30} & \frac{1}{70} \end{bmatrix} \otimes i\Delta t \mathbf{L}_j + \begin{bmatrix} \frac{3}{35} & \frac{1}{30} & \frac{9}{140} & -\frac{1}{35} \\ \frac{1}{30} & \frac{1}{70} & \frac{1}{30} & -\frac{1}{70} \\ \frac{9}{140} & \frac{1}{30} & \frac{2}{7} & -\frac{1}{14} \\ -\frac{1}{35} & -\frac{1}{70} & -\frac{1}{14} & \frac{1}{42} \end{bmatrix} \otimes i\Delta t \mathbf{L}_{j+1}. \quad (6.21)$$

when $j \geq 2$,

$$\frac{\partial \mathbf{f}}{\partial x_j} - \frac{\partial \mathbf{K}}{\partial x_j} \boldsymbol{\alpha} = \begin{bmatrix} \mathbf{0} & -\frac{\partial K^{e,j}}{\partial x_j} \\ & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{0} \\ \alpha_{4j-3:4j+4} \\ \mathbf{0} \end{bmatrix}. \quad (6.26)$$

The $\partial K^{e,j}/\partial x_j$ can be obtained by differentiating $K^{e,j}$ with respect to x_j in Eq. 6.14. Subsequently, multiplying $\mathbf{K}^{-1}$ by N vectors can be efficiently performed using the adjoint method [180], where an adjoint equation is defined as

$$\mathbf{K}^\top \boldsymbol{\lambda} = \left(\frac{d\eta}{d\boldsymbol{\alpha}} \right)^\top, \quad (6.27)$$

where $\boldsymbol{\lambda} \in \mathbb{C}^{(4N+4) \times 1}$ is the adjoint vector. After obtaining $\boldsymbol{\lambda}$, the gradient is computed as

$$\frac{d\eta}{d\mathbf{x}} = \boldsymbol{\lambda}^\top \left(\frac{d\mathbf{f}}{d\mathbf{x}} - \frac{d\mathbf{K}}{d\mathbf{x}} \boldsymbol{\alpha} \right). \quad (6.28)$$

With the adjoint method, the cost of gradient computation becomes comparable to that of solving the spin trajectory.

6.2.3 Regularization

To ensure smooth pulse shapes, we apply the Helmholtz filter [181], defined as

$$x^s(t) = R^2 \frac{d^2 x^s(t)}{dt^2} + x^c(t), \quad (6.29)$$

where R denotes the filter radius, set here to $R = T/130$. The term $x^c(t)$ corresponds to the original control variables, while $x^s(t)$ denotes the filtered (smoothed) variables. By approximating the solution with linear shape functions, Eq. 6.29 is solved through a separate linear system:

$$\mathbf{K}_h \mathbf{x}^s = \mathbf{f}_h, \quad (6.30)$$

where the stiffness matrix elements are

$$K_{ij} = \int_0^T R^2 \frac{d\phi_j}{dt} \frac{d\phi_i}{dt} + \phi_j \phi_i dt, \quad (6.31)$$

and the load vector elements are

$$f_i = \int_0^T x^c(t) \phi_i(t) dt. \quad (6.32)$$

The element stiffness matrix over interval $[t_j, t_{j+1}]$ is

$$K^{e,j} = \begin{bmatrix} \int_{t_j}^{t_{j+1}} R^2 \frac{d\phi_j}{dt} \frac{d\phi_j}{dt} + \phi_j \phi_j dt & \int_{t_j}^{t_{j+1}} R^2 \frac{d\phi_{j+1}}{dt} \frac{d\phi_j}{dt} + \phi_{j+1} \phi_j dt \\ \int_{t_j}^{t_{j+1}} R^2 \frac{d\phi_j}{dt} \frac{d\phi_{j+1}}{dt} + \phi_j \phi_{j+1} dt & \int_{t_j}^{t_{j+1}} R^2 \frac{d\phi_{j+1}}{dt} \frac{d\phi_{j+1}}{dt} + \phi_{j+1} \phi_{j+1} dt \end{bmatrix}. \quad (6.33)$$

By substituting linear shape functions, the equation simplifies to

$$K^{e,j} = \begin{bmatrix} \frac{\Delta t}{3} + \frac{R^2}{\Delta t} & \frac{\Delta t}{6} - \frac{R^2}{\Delta t} \\ \frac{\Delta t}{6} - \frac{R^2}{\Delta t} & \frac{\Delta t}{3} + \frac{R^2}{\Delta t} \end{bmatrix}. \quad (6.34)$$

Since the control variables $\mathbf{x}^c$ are piecewise constant, the load vector is computed as

$$f_i = \begin{cases} \frac{\Delta t}{2} x_1^c, & i = 1, \\ \frac{\Delta t}{2} (x_{i-1}^c + x_i^c), & 2 \leq i \leq N, \\ \frac{\Delta t}{2} x_N^c, & i = N + 1. \end{cases} \quad (6.35)$$

The Jacobian matrix that relates the smoothed variables to the control variables is

$$\frac{d\mathbf{x}^s}{d\mathbf{x}^c} = \mathbf{K}_h^{-1} \left(\frac{\partial \mathbf{f}_h}{\partial x_1^c}, \frac{\partial \mathbf{f}_h}{\partial x_2^c}, \dots, \frac{\partial \mathbf{f}_h}{\partial x_N^c} \right), \quad (6.36)$$

which yields

$$\frac{d\mathbf{x}^s}{d\mathbf{x}^c} = \frac{\Delta t}{2} [\mathbf{K}_h^{-1}]_{[1:N, 1:N+1]} \begin{bmatrix} 1 & & & & \\ 1 & 1 & & & \\ & 1 & \ddots & & \\ & & \ddots & 1 & \\ & & & 1 & 1 \\ & & & & 1 \end{bmatrix}_{(N+1) \times N}. \quad (6.37)$$

For pulses expressed in Cartesian coordinates, an additional hyperbolic tangent scaling function is applied to constrain the waveform amplitude within $[-1, 1]$:

$$x(t) = \frac{1 - e^{-\kappa x^s(t)}}{1 + e^{-\kappa x^s(t)}}. \quad (6.38)$$

Here, κ controls the steepness of the transition period; a typical choice is $\kappa = 10$, which allows the pulse to reach its maximum amplitude. The gradient of the scaled waveform $\mathbf{x}$ with respect to the smoothed variables $\mathbf{x}^s$ is a diagonal matrix:

$$\frac{d\mathbf{x}}{d\mathbf{x}^s} = \text{diag} \left(\frac{2\kappa e^{-\kappa x_i^s}}{(1 + e^{-\kappa x_i^s})^2} \right), \quad i = 1, 2, \dots, N. \quad (6.39)$$

By the chain rule, the gradient of the fidelity η with respect to the control variables $\mathbf{x}^c$ is given by

$$\frac{d\eta}{d\mathbf{x}^c} = \frac{d\eta}{d\mathbf{x}} \cdot \frac{d\mathbf{x}}{d\mathbf{x}^s} \cdot \frac{d\mathbf{x}^s}{d\mathbf{x}^c}. \quad (6.40)$$

6.3 Numerical implementation

6.3.1 Matrix assembly and solving

Depending on the interpolation element type, Eq. 6.14, 6.17, or 6.21 was used to generate the element stiffness matrices from the piecewise Liouvillians. As illustrated in Fig. 6.3, these element matrices were computed in a vectorized manner and stored in a three-dimensional array, where each slice corresponds to one time segment. This 3D array is then vectorized and assembled into a 2D sparse matrix using MATLAB's sparse function, with the index vectors generated separately.




Figure 6.3: Assembly of the stiffness matrix in MATLAB. The element stiffness matrices were stored as a 3D array, vectorized, and then transformed into a 2D sparse matrix using the sparse function.

The spin trajectory and the adjoint vector are jointly solved using an LU decomposition,

taking advantage of the banded sparse matrix structure:

```
dK = decomposition(K, 'banded');
alpha = dK \ f;
lambda = eta / dK;
```

Here, η denotes the extended objective defined in Eq. 6.23.

The gradient of the stiffness matrix in Eq. 6.28 is evaluated by multiplying a block-diagonal matrix with a vector containing the element-wise solutions, as illustrated in Fig. 6.4. Each block in the matrix represents the analytical gradient of a piecewise Liouvillian with respect to the local control parameters. The right-hand vector contains the state solutions obtained from Algorithm 1, and the corresponding degrees of freedom are arranged to align with the left-hand matrix according to Eqs. 6.25 and 6.26.




Figure 6.4: Computation of the gradient of the stiffness matrix with respect to the control variables. The left-hand gradient of $\mathbf{K}$ is constructed by stacking the analytic gradient matrices of K^e into a block-diagonal form, while the right-hand vector represents the element-wise spin states corresponding to each K^e on the left.

Two commonly used operator bases are available for constructing the Liouvillian: the Zeeman basis, which is built from Pauli matrices, and the irreducible spherical tensor (IST) basis [58]. These two bases yield stiffness matrices with different banded structures. To visualize the difference, we consider a single-spin system driven by two control channels (x and y) over 10 time steps. The resulting stiffness matrices are shown in Fig. 6.5. The matrix constructed with the IST basis is visibly sparser, containing fewer nonzero elements, which leads to improved computational efficiency.

The wall-clock time for gradient computation under the two bases was benchmarked for

both single-spin and two-spin systems, as shown in Fig. 6.6. Figures 6.6(a,b) correspond to linear elements, while Figs. 6.6(c,d) correspond to Hermite elements. In all cases, the IST basis outperforms the Zeeman basis, achieving approximately 20% higher computational efficiency. Therefore, the IST basis was adopted in the subsequent implementations.




Figure 6.5: Banded structure of the global stiffness matrix for a single-spin system with 10 time steps, using (a) the Zeeman basis and (b) the irreducible spherical tensor (IST) basis.

6.3.2 Optimization algorithm

After obtaining the spin trajectory and the gradients, the control variables are updated using MMA, a widely used approach for large-scale, constrained nonlinear optimization problems, such as topology optimization [182, 183]. MMA is a gradient-based algorithm that leverages the values and gradients of the objective and constraint functions to iteratively construct and solve a sequence of convex subproblems. Each subproblem has a unique optimal solution that can be efficiently obtained via a dual approach [184]. To employ the MMA algorithm, the maximization of the ensemble fidelity is reformulated as a least-squares problem, i.e., by minimizing the following expression:

$$f(\mathbf{x}^c) = \sum_{n=1}^{N_{\text{ens}}} (1 - \eta_n)^2. \quad (6.41)$$




Figure 6.6: Comparison of wall-clock time for gradient computation using the Zeeman and IST bases in single-spin and two-spin systems. Linear elements are shown in (a,b), and Hermite elements in (c,d).

In the implementation, the least-squares objective was reformulated by converting the infidelities into $2N_{\text{ens}}$ linear constraints while setting the objective function to zero [185]:

$$\begin{aligned}
 f_0(\mathbf{x}^c) &= 0, \\
 f_n(\mathbf{x}^c) &= 1 - \eta_n, \quad n = 1, 2, \dots, N_{\text{ens}}, \\
 f_{N_{\text{ens}}+n}(\mathbf{x}^c) &= \eta_n - 1, \quad n = 1, 2, \dots, N_{\text{ens}}.
 \end{aligned} \tag{6.42}$$

Algorithm 2 outlines the pseudocode for solving the optimal control problem in Eq. 6.4. The control variables $\mathbf{x}^c$ consist of M channels, each with N discrete values. Iteration proceeds until the average fidelity $\bar{\eta}$ reaches the target (e.g., 0.995) or the maximum number of iterations is reached. For an ensemble of spin systems, the mesh is kept identical across all members, so the index matrix of the stiffness matrix remains unchanged. Each ensemble member has an individual Liouvillian, and the corresponding stiffness matrix is assem-

Algorithm 2: Find an optimal control pulse shape.

Input: Parameters of the spin system; Pulse parameters**Output:** Optimized control pulse $\mathbf{x}$ Initialize $\mathbf{x}^c \leftarrow \text{rand}(M, N)$; // Random initializationSolve η_n from Eq. 6.23 and $d\eta_n/d\mathbf{x}$ from Eq. 6.28;**for** $iter \leftarrow 2$ **to** $iter_{max}$ **do** **if** $\bar{\eta} \geq target$ **then** **break**; **else** Compute f_n from Eq. 6.42 and $df_n/d\mathbf{x}^c$ from Eq. 6.40; Update $\mathbf{x}^c$ using MMA; // Update control variables Solve $\mathbf{x}^s$ from Eq. 6.30 and $d\mathbf{x}^s/d\mathbf{x}^c$ from Eq. 6.37; // Smooth variables **if** $\mathbf{x}^c$ are Cartesian components **then** Solve $\mathbf{x}$ from Eq. 6.38 and $d\mathbf{x}/d\mathbf{x}^s$ from Eq. 6.39 ; // Scale variables **else** $\mathbf{x} \leftarrow \mathbf{x}^s$; Solve α from Algorithm 1; // Solve LvN equation Solve η_n from Eq. 6.23 and $d\eta_n/d\mathbf{x}$ from Eq. 6.28; // Adjoint analysis**return** $\mathbf{x}$;

bled to evaluate the spin trajectory and its gradient. Handling an ensemble system can be efficiently accelerated using MATLAB's parallel computing capabilities. In contrast, processing the variables involves solving a separate linear system, executed only once per iteration. For phase optimization, the smoothed variables $\mathbf{x}^s$ define the pulse shape; for Cartesian components (x and y) optimization, $\mathbf{x}^s$ is further scaled to the range $[-1, 1]$.

6.4 Results and discussion

6.4.1 Accuracy with small step

The accuracy and computational efficiency of the FEM approach for solving the single-spin system were evaluated. The step-by-step propagation method [55] is adopted as a standard reference, where the control gradient is computed using the AUXMAT. This referred method was executed using Spinach v2.8 [58], where step propagators are computed via the reordered Taylor expansion, summing low-order terms to machine precision

(2.22×10^{-16} on a 64-bit machine). All computations were performed in MATLAB 2025b on a PC equipped with an AMD Ryzen Threadripper PRO 3955WX 16-Cores processor (base frequency: 3.90 GHz) and 256 GB of RAM. For benchmarking the speed of the FEM method, MATLAB's parallel pool was disabled to remove the influence of parallel computing.

The relative error of the spin trajectory is defined as

$$\epsilon_\rho = \frac{1}{N+1} \sum_{i=1}^{N+1} \frac{|\rho_i^F - \rho_i^P|}{|\rho_i^P|}, \quad (6.43)$$

where N denotes the number of time steps, and ρ_i^P and ρ_i^F are the i -th spin vectors obtained from step-by-step propagation and FEM, respectively. Since the spin state vectors are normalized, $|\rho_i^P| = 1$. And the relative error of the gradient is computed as

$$\epsilon_{\text{grad}} = \frac{1}{N} \sum_{i=1}^N \left| \frac{\nabla_i^F - \nabla_i^{\text{auxmat}}}{\nabla_i^{\text{auxmat}}} \right|, \quad (6.44)$$

where ∇_i^{auxmat} and ∇_i^F are the partial derivatives of the objective with respect to the i -th control variable, computed via AUXMAT and FEM, respectively. The approximation error is mainly determined by the discrete step size $\|\mathbf{L}\|\Delta t$. In typical liquid-state NMR experiments, an RF amplitude of 10 kHz and a time step of 1 μs yield $\|\mathbf{L}\|\Delta t = 0.063$. In this analysis, $\|\mathbf{L}\|\Delta t$ was varied from 0.01 to 0.1, and for each value, 30 random pulse shapes were generated to compute the mean values of ϵ_ρ and ϵ_{grad} .

The performance of the linear and quadratic elements was compared with the AUXMAT and second-order finite difference (FD) methods, where the propagators were approximated by truncating the Taylor series. Figure 6.7 shows the benchmark results for linear and quadratic elements. Using the linear elements, the spin-trajectory error increases from 10^{-6} to 10^{-2} , while the gradient error rises from 10^{-4} to 10^{-1} . These accuracies are comparable to the AUXMAT/FD results obtained using second-order truncation. Employing quadratic shape functions significantly reduces these ranges to 10^{-7} – 10^{-4} for the trajectory and 10^{-6} – 10^{-3} for the gradient. Hence, quadratic elements show comparable accuracy to AUXMAT/FD using third-order truncation.

The difference between ϵ_ρ and ϵ_{grad} arises because in Eq. 6.43, $|\rho_i^P|$ is fixed at 1, whereas in Eq. 6.44, ∇_i^{auxmat} can become very small for some control variables and reaches the

approximation error of the linear or quadratic elements, resulting in an amplified relative error. An average ensemble fidelity of 0.995 can be achieved with a trajectory error below 10^{-3} , which for linear element approximations requires $\|\mathbf{L}\|\Delta t \leq 0.06$.




Figure 6.7: The relative spin trajectory (a) and gradient (b) errors for the three compared methods. These errors are calculated relative to a reference solution obtained using machine-precision step-by-step propagation combined with the auxiliary matrix formalism (AUXMAT) for gradient computation. The computation considered a single-spin system driven by a piecewise-constant waveform. In approximated AUXMAT and the second-order central finite difference (FD) method, the propagators were evaluated via Taylor-series expansions truncated at second and third order, as indicated in the legends. Note that with second or third order Taylor truncation, the errors of AUXMAT and FD nearly coincide because they are dominated by the truncation error in propagator computation. The finite difference used a differentiation step of $h = 0.1$. The pulse duration was set to $T = 0.5$ ms, the Liouvillian norm was $\|\mathbf{L}\| = 2 \times 10^4 \text{ rad} \cdot \text{s}^{-1}$, and the number of time slices N was swept from 100 to 1000.

The same procedure was applied to Hermite elements using a piecewise-linear waveform, and the relative error of the Hermite approximation is shown in Fig. 6.8. Unlike the case of piecewise-constant waveforms, the accuracy of the Hermite discretization depends on the smoothness of the linear waveform. Increasing the Helmholtz filter radius R enforces smoother pulse shapes, reducing the error and exceeding the quadratic-element accuracy when $R = T/42$. And cubic Hermite elements offer an accuracy intermediate to AUXMAT/FD methods using second- and third-order Taylor truncation.

6.4.2 Speed for 1 and 2 spin systems

The wall-clock time as a function of the number of time steps was evaluated for both a single-spin system and a 2-spin system. Figure 6.9 compares the runtimes of FEM,




Figure 6.8: The relative errors of the spin trajectory (a) and gradient (b) for three methods are shown, using a piecewise-linear waveform while keeping all other parameters identical to Fig. 6.7. Results from auxmat-2nd and auxmat-3rd are not shown because their errors are nearly identical to FD-2nd and FD-3rd, respectively. Note that the errors of the cubic Hermite elements depend on the pulse smoothness, where T is the pulse duration and R is the Helmholtz-filter radius controlling the smoothing.

AUXMAT, and second-order finite differences for the single-spin case. For a piecewise-constant waveform, when N ranges from 100 to 2000, the linear elements provide roughly a $7\times$ speedup over FD-2nd, while the quadratic elements yield about a $3.5\times$ speedup over FD-3rd. The wall-clock times for a piecewise-linear waveform are shown in Fig. 6.9b, where the Hermite elements achieve approximately a $10\times$ speedup over FD-2nd within the same range of N .

Corresponding results for the 2-spin system are presented in Fig. 6.10. In this case, both linear and quadratic elements require runtimes comparable to those of AUXMAT and finite-difference methods for similar accuracy. The Hermite elements achieve about a $5\times$ speedup over FD-2nd when N ranges from 100 to 2000.

It is worth noting that the FEM wall-clock time is piecewise linear; this behavior arises because stiffness matrix assembly and linear system solving are performed on large, sparse matrices. As the number of time steps increases, the working set eventually exceeds the CPU cache capacity. The computation thus transitions from a compute-bound regime, which is dominated by small dense matrix operations, to a memory-bound regime limited by RAM bandwidth. Consequently, a few jumps appear in the runtime curve, and its slope increases slightly in the large- N region.




Figure 6.9: Wall-clock time for computing spin trajectories and derivatives in a single-spin system. (a) Running time for the piecewise-constant waveform. (b) Running time for the piecewise-linear waveform. In both panels, $\|\mathbf{L}\|\Delta t = 0.063$ is held fixed, and the number of time steps N is swept from 10 to 2000. MATLAB’s parallel pool was disabled to eliminate the influence of parallel execution. Each benchmark was repeated 30 times for each method, and the average running time was recorded. Across the practically relevant regime of $N = 100$ to $N = 2000$, linear elements achieve an approximate $7\times$ acceleration relative to FD-2nd; quadratic elements achieve $3.5\times$ acceleration relative to FD-3rd; and cubic Hermite elements achieve about $10\times$ acceleration relative to FD-3rd.




Figure 6.10: Wall-clock time benchmarks for computing spin trajectories and gradients in a 2-spin system. (a) Running times for a piecewise-constant waveform. (b) Running times for a piecewise-linear waveform. Within the range of $100 \leq N \leq 2000$, linear elements show a $1.6\times$ reduction in running time relative to the FD-2nd, quadratic elements give comparable running times to FD-3rd, and cubic Hermite elements give an approximate $5\times$ speedup relative to FD-3rd.

6.4.3 State-to-state optimal control

The FEM combined with MMA was tested by optimizing a single-spin broadband excitation pulse, which transfers $\mathbf{I}_z$ to $\mathbf{I}_x$ over a 20 kHz offset range. The shaped pulse has

a duration of $500\ \mu\text{s}$, discretized into 500 time bins, and a nominal amplitude of 10 kHz with $\pm 20\%$ variations. Gradients were computed using three methods: FEM with linear elements, auxmat-2nd, and FD-2nd, which yield comparable numerical accuracy. Each method was repeated 20 times from different random initial guesses, and their convergence rates are shown in Fig. 6.11a. The similar convergence rates are primarily constrained by the MMA optimizer; its oscillatory convergence is analysed in the next section. In terms of runtime, the linear elements achieved an order-of-magnitude speedup over auxmat-2nd, while FD-2nd exhibits intermediate wall-clock time, as shown in Fig. 6.11b.




Figure 6.11: Performance of the excitation pulse optimization for an ensemble single-spin system. (a) Convergence rate comparison for gradients computed using FEM with linear elements, AUXMAT, and 2nd-order finite differences (FD). AUXMAT and FD used 2nd-order Taylor-truncated propagators. Each method was repeated 20 times from different random initial guesses, and the MMA algorithm was used for optimization. (b) Histogram showing the time consumption of the three methods. The shaped pulse steers $\mathbf{I}_z$ to $\mathbf{I}_x$ with RF amplitude 10 kHz and $\pm 20\%$ scaling ($n_{\text{rf}} = 5$). A 20 kHz bandwidth was discretized into $n_{\text{off}} = 51$ offsets, giving $N_{\text{ens}} = 255$. The $500\ \mu\text{s}$ pulse was piecewise constant with 500 segments, its amplitudes were fixed at 1, and its phases were optimized. The target ensemble fidelity was 0.995. To eliminate the influence of parallel computation, the parallel pool was disabled in MATLAB 2025b.

In the above test, a single-spin system was constructed for each ensemble member, and all fidelities and gradients obtained from either the FEM, AUXMAT, or FD methods were provided to the MMA algorithm, in which the ensemble constraints were built by Eq 6.42. When AUXMAT is combined with L-BFGS or Newton optimization, as implemented in Spinach, the Hamiltonians of the non-interacting subsystems can be assembled into a block-diagonal matrix, which allows the ensemble-averaged fidelity and gradient to be computed more efficiently. This block-diagonal acceleration was excluded from the comparisons.

6.4.4 Propagator optimization

The performance of MMA was tested against the L-BFGS and Newton-Raphson methods, with the latter two executed using Spinach v2.8. The test case involved optimizing a broadband universal rotation pulse, designed to realize the target propagator $\mathbf{U} = \exp(-i\pi\mathbf{I}_x/2)$ across a 20 kHz offset range. The shaped pulse has a duration of 500 μs , discretized into 500 time bins, with a nominal amplitude of 10 kHz and $\pm 10\%$ scaling variations. The linear algebra was implemented in Hilbert space, where a unique target propagator favored by the optimization algorithm can be defined [186]. Unlike the step-by-step propagation in the GRAPE method, the FEM model approximates spin evolution directly using a linear combination of basis functions. As a result, the effective propagator is not explicitly constructed, rendering propagator optimization currently infeasible within the FEM framework. AUXMAT was used to compute the gradient for both MMA and L-BFGS, and the Hessian matrix for the Newton-Raphson method; propagators were calculated using a Taylor series expansion truncated to machine precision. The convergence behavior shown in Fig. 6.12a indicates that L-BFGS achieves stable but slower convergence as the fidelity approaches the target. The Newton method converges more rapidly and stably, while MMA achieves a comparable convergence rate to the Newton-Raphson method, albeit with a non-monotonic curve. The wall-clock times, shown in Fig. 6.12b, indicate that MMA completes the optimization faster than both L-BFGS and the Newton-Raphson methods.

The oscillatory behavior of the MMA algorithm arises from the implicit Hessian (curvature) introduced by its moving-asymptote approximation. The empirical choices of the moving asymptotes implicitly generate an approximated, positive-definite diagonal Hessian, which may be inaccurate. For example, loose asymptotes yield smaller curvature and consequently an aggressive step that overshoots the local minimum. The algorithm then corrects this overshoot by moving back in the next iteration, leading to an oscillatory pattern. The modified variant, globally convergent MMA (GCMMA) [187], addresses this by introducing an additional inner iteration that plays a role analogous to a line search. This process forces more conservative approximations of both the objective and constraints, yielding slower but more stable convergence. However, each inner iteration requires re-computing the ensemble fidelities, whose cost is comparable to computing both fidelities and gradients together.




Figure 6.12: Performance of implementing a universal 90°_x rotation for an ensemble single-spin system. (a) Comparison of convergence rates for the MMA, L-BFGS, and Newton-Raphson methods, each method was repeated 40 times using different random initial guesses, AUXMAT was used to compute the gradient and Hessian. (b) Histogram of wall-clock time for three methods. The RF amplitude was 10 kHz with $\pm 10\%$ scaling ($n_{\text{rf}} = 5$). A 20 kHz bandwidth was sampled at $n_{\text{off}} = 51$ offsets, giving an ensemble size $N_{\text{ens}} = 255$. The $500\ \mu\text{s}$ pulse used 500 piecewise-constant slices with unit-fixed amplitude, and the phases were optimized. Optimization targeted an ensemble fidelity of 0.995 and terminated at 100 maximum iterations. Ensemble gradient and Hessian evaluations were accelerated using 16 workers in MATLAB R2025b.

6.5 Conclusion

In this work, we employed the FEM method to solve the Liouville–von Neumann equation for a single-spin system. Achieving gradient accuracy comparable to step-by-step propagation approach using second-order Taylor-truncated propagators, the use of linear elements provides an approximate order-of-magnitude speedup in gradient evaluation. Additionally, the propagator optimization benchmark confirms that the MMA optimizer converges faster than L-BFGS, highlighting the potential of this approach for time-constrained optimal control problems.

MRI pulse optimization inherently involves hardware constraints, including limits on RF amplitude and total energy, as well as optimization time as an additional constraint. Within this context, the FEM offers high computational efficiency for single-spin systems, while the MMA is particularly well suited for handling nonlinear constraints and achieving rapid early-stage convergence. Hence, their combination is a promising solution for accelerated MRI pulse optimization.

With the adjoint method, gradient computation is as efficient as trajectory evaluation, which is similar to the case in step-by-step propagation. Hence, the computational cost

is dominated by stiffness-matrix assembly and linear-system solving. While the 1-spin system attains the largest acceleration, this speedup drops by roughly a factor of 4 for 2 spins and is expected to diminish further as the spin system size increases, since memory usage then exceeds the CPU cache capacity and spills into slower RAM. When the parallelization is used to accelerate both the FEM [188, 189] and the step-by-step propagation method [78, 190, 191], it is also worth studying their relative efficiency as a function of the number of degrees of freedom.

Chapter 7

Conclusions

This thesis makes three main contributions to parallel high-field liquid-state NMR spectroscopy: RF decoupling and gradient decoupling, demonstrated in two-channel experiments at 500 MHz, and accelerated pulse optimization, demonstrated in a single-spin system.

By combining EM simulation with spin dynamics calculations, it was shown that RF coupling causes excitation distortions by the parallel pulses and mixing of the parallel FIDs. A linear-combination strategy was applied to design cooperative pulses that suppress RF coupling during parallel excitation, and the individual FIDs were recovered using blind source separation. The RF decoupling method assumes that coupling occurs as a linear mixture through a coupling matrix, since an RF pulse can be delivered precisely to the coil at a single RF frequency, for example 500 MHz. At GHz-range frequencies, however, time-domain distortions become significant and RF signals may couple in a convolutive manner. In such cases, BSS for convolutive mixture can be employed to identify the impulse response function [119], which could in turn guide the design of cooperative pulses.

In gradient decoupling, the coherence-locking pulse designed via optimal control was shown to be robust against gradient spillover-induced B_0 drift and variations in ensemble spin-system parameters. However, concurrent RF coupling, which is not addressed in the gradient decoupling scheme, necessitates lower pulse amplitudes and therefore restricts bandwidth. A more general strategy would address RF and gradient effects jointly. This approach, together with the execution of distinct pulse sequences in parallel channels, requires more advanced sequence design and greater hardware flexibility. For example, the

observe channel may be disabled while a homonuclear pulse is applied on another channel to prevent signal leakage between reception and excitation.

Pulse optimization using FEM with linear elements was more than an order of magnitude faster than step-by-step time propagation. The MMA algorithm with a least-squares objective converged faster than L-BFGS, although both methods rely on first-order gradient information; however, oscillatory convergence behavior limits their applicability in high-fidelity regimes, such as quantum gate optimization. A hybrid approach could use MMA to quickly reach an initial target and then switch to a more stable method for fine-tuning. Computation accuracy could be further improved by several FEM features, such as higher-order shape functions or adaptive moving meshes. Another direction is including spatial variables in the FEM model, either when diffusion becomes a degree of freedom [177, 192] or when optimal control of fluid samples [193] is concerned. Periodic sequences are used to drive a spin system to a steady state [194], with applications such as balanced steady-state free precession (bSSFP) [195] imaging and microwave-induced dynamic nuclear polarization [196]. Pulse optimization in this setting may be addressed by FEM, where a periodic boundary condition is imposed to enforce equality between the initial and final states.

Bibliography

- [1] Denis Weaire and Seamus O'Connor. Unfulfilled renown: Thomas preston (1860–1900) and the anomalous zeeman effect. *Annals of Science*, 44(6):617–644, 1987.
- [2] W. Gerlach and O. Stern. Der experimentelle Nachweis des magnetischen Moments des Silberatoms. *Zeitschrift für Physik*, 8(1):110–111, 1922.
- [3] W. Pauli. Über den Einfluß der Geschwindigkeitsabhängigkeit der Elektronenmasse auf den Zeemaneffekt. *Zeitschrift für Physik*, 31(1):373–385, 1925.
- [4] W. Pauli. Über den Zusammenhang des Abschlusses der Elektronengruppen im Atom mit der Komplexstruktur der Spektren. *Zeitschrift für Physik*, 31(1):765–783, 1925.
- [5] Paul Adrien Maurice Dirac. The quantum theory of the electron. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 117(778):610–624, 1928.
- [6] Isidor Isaac Rabi, Jerrold R Zacharias, Sidney Millman, and Polykarp Kusch. A new method of measuring nuclear magnetic moment. *Physical Review*, 53(4):318, 1938.
- [7] Felix Bloch. Nuclear induction. *Physical Review*, 70(7-8):460, 1946.
- [8] Edward M Purcell, Henry Cutler Torrey, and Robert V Pound. Resonance absorption by nuclear magnetic moments in a solid. *Physical Review*, 69(1-2):37, 1946.
- [9] J Jeneer, BH Meier, P Bachmann, and RR Ernst. Investigation of exchange processes by two-dimensional NMR spectroscopy. *The Journal of Chemical Physics*, 71:4546–4553, 1979.
- [10] Kurt Wüthrich. NMR with proteins and nucleic acids. *Europhysics News*, 17(1):11–13, 1986.
- [11] Paul C Lauterbur. Image formation by induced local interactions: examples employing nuclear magnetic resonance. *Nature*, 242(5394):190–191, 1973.
- [12] Peter Mansfield. Multi-planar image formation using NMR spin echoes. *Journal of Physics C: Solid State Physics*, 10(3):L55, 1977.
- [13] Albert W. Overhauser. Paramagnetic Relaxation in Metals. *Physical Review*, 89(4):689–700, 1953.
- [14] Albert W. Overhauser. Polarization of Nuclei in Metals. *Physical Review*, 92(2):411–415, 1953.
- [15] T. R. Carver and C. P. Slichter. Polarization of nuclear spins in metals. *Physical Review*, 92:212–213, 1953.
- [16] C Russell Bowers and Daniel P Weitekamp. Parahydrogen and synthesis allow dramatically enhanced nuclear alignment. *Journal of the American Chemical Society*,

- 109(18):5541–5542, 1987.
- [17] Ralph W Adams, Juan A Aguilar, Kevin D Atkinson, Michael J Cowley, Paul IP Elliott, Simon B Duckett, Gary GR Green, Iman G Khazal, Joaquín López-Serrano, and David C Williamson. Reversible interactions with para-hydrogen enhance NMR sensitivity by polarization transfer. *Science*, 323(5922):1708–1711, 2009.
- [18] Roger A. Kautz, Wolfgang K. Goetzinger, and Barry L. Karger. High-Throughput Microcoil NMR of Compound Libraries Using Zero-Dispersion Segmented Flow Analysis. *Journal of Combinatorial Chemistry*, 7(1):14–20, 2005.
- [19] Paul A. Keifer. High-resolution NMR techniques for solid-phase synthesis and combinatorial chemistry. *Drug Discovery Today*, 2(11):468–478, 1997.
- [20] Omar Nassar, Mazin Jouda, Michael Rapp, Dario Mager, Jan G. Korvink, and Neil MacKinnon. Integrated impedance sensing of liquid sample plug flow enables automated high throughput NMR spectroscopy. *Microsystems & Nanoengineering*, 7(1):30, 2021.
- [21] Alfred Ross, Götz Schlotterbeck, Hans Senn, and Markus von Kienlin. Application of chemical shift imaging for simultaneous and fast acquisition of NMR spectra on multiple samples. *Angewandte Chemie International Edition*, 40(17):3243–3245, 2001.
- [22] Dilara Faderl, Ajmal Chenakkara, Mazin Jouda, Neil MacKinnon, Alvar D. Gossert, and Jan G. Korvink. Accelerated screening of protein–ligand interactions via parallel T2-weighted ^{19}F -MRI. *Analytical Chemistry*, 96(24):9859–9865, 2024.
- [23] Ěriks Kupče, Lucio Frydman, Andrew G. Webb, Jonathan R. J. Yong, and Tim D. W. Claridge. Parallel nuclear magnetic resonance spectroscopy. *Nature Reviews Methods Primers*, 1(1):1–23, 2021.
- [24] Ernesto MacNamara, Ting Hou, Greg Fisher, Sheldon Williams, and Daniel Raftery. Multiplex sample NMR: An approach to high-throughput NMR using a parallel coil probe. *Analytica Chimica Acta*, 397(1):9–16, 1999.
- [25] H. Wang, L. Ciobanu, A. S. Edison, and A. G. Webb. An eight-coil high-frequency probehead design for high-throughput nuclear magnetic resonance spectroscopy. *Journal of Magnetic Resonance*, 170(2):206–12, 2004.
- [26] Ting Hou, Ernesto MacNamara, and Daniel Raftery. NMR analysis of multiple samples using parallel coils: Improved performance using reference deconvolution and multidimensional methods. *Analytica Chimica Acta*, 400(1):297–305, 1999.
- [27] Y. Li, A. M. Wolters, P. V. Malawey, J. V. Sweedler, and A. G. Webb. Multiple solenoidal microcoil probes for high-sensitivity, high-throughput nuclear magnetic resonance spectroscopy. *Analytical Chemistry*, 71(21):4815–4820, 1999.
- [28] Luisa Ciobanu, Dimuthu A. Jayawickrama, Xiaozhong Zhang, Andrew G. Webb, and Jonathan V. Sweedler. Measuring reaction kinetics by using multiple microcoil NMR spectroscopy. *Angewandte Chemie International Edition*, 42(38):4669–4672, 2003.
- [29] Yaewon Kim, Mengxiao Liu, and Christian Hilty. Parallelized ligand screening using dissolution dynamic nuclear polarization. *Analytical Chemistry*, 88(22):11178–11183, 2016.
- [30] Ka-Meng Lei, Dongwan Ha, Yi-Qiao Song, Robert M. Westervelt, Rui Martins, Pui-In Mak, and Donhee Ham. Portable NMR with parallelism. *Analytical Chemistry*, 92(2):2112–2120, 2020.

- [31] Yen-Tse Cheng, Mazin Jouada, and Jan Korvink. Sample-centred shimming enables independent parallel NMR detection. *Scientific Reports*, 12(1):14149, 2022.
- [32] Moritz Becker, Yen-Tse Cheng, Achim Voigt, Ajmal Chenakkara, Mengjia He, Sören Lehmkuhl, Mazin Jouada, and Jan G. Korvink. Artificial intelligence-driven shimming for parallel high field nuclear magnetic resonance. *Scientific Reports*, 13(1):17983, 2023.
- [33] Jing Yang, Yen-Tse Cheng, Sören Lehmkuhl, Juergen J. Brandner, Mazin Jouada, and Jan G. Korvink. Accelerating NMR spectroscopy via parallel SABRE in a continuous-flow system. *Review of Scientific Instruments*, 96(10):105203, 2025.
- [34] Hossein Esmaeilizadshali, Dominique Buyens, Jan Korvink, and Mazin Jouada. Miniaturization and Optimization of NMR Shimsets Allow Efficient Field Shimming. *Analytical Chemistry*, 98(20):14674–14682, 2026.
- [35] Mengjia He, Dilara Faderl, Neil MacKinnon, Yen-Tse Cheng, Dominique Buyens, Mazin Jouada, Burkhard Luy, and Jan G. Korvink. A digital twin for parallel liquid-state nuclear magnetic resonance spectroscopy. *Communications Engineering*, 3(1):1–13, 2024.
- [36] Mengjia He, Neil MacKinnon, Dominique Buyens, Burkhard Luy, and Jan G. Korvink. Coherence locking in a parallel nuclear magnetic resonance probe defends against gradient field spillover. *Magnetic Resonance*, 6(2):173–181, 2025.
- [37] Mengjia He, Yongbo Deng, Burkhard Luy, and Jan G. Korvink. Finite elements and moving asymptotes accelerate quantum optimal control—FEMMA. *The Journal of Chemical Physics*, 164(6):064114, 2026.
- [38] David M. Pozar. *Microwave Engineering, 4th Edition*. Wiley, Hoboken, NJ, 2011.
- [39] K.F. Warnick and M.A. Jensen. Effects of mutual coupling on interference mitigation with a focal plane array. *IEEE Transactions on Antennas and Propagation*, 53(8):2490–2498, 2005.
- [40] Erwin Schrödinger. Quantisierung als eigenwertproblem. *Annalen der Physik*, 79:361, 1926.
- [41] Malcolm H. Levitt. *Spin Dynamics: Basics of Nuclear Magnetic Resonance*. John Wiley & Sons, Chichester, 2008.
- [42] Ilya Kuprov. *Spin*. Springer, Cham, 2023.
- [43] Wilhelm Magnus. On the exponential solution of differential equations for a linear operator. *Communications on pure and applied mathematics*, 7(4):649–673, 1954.
- [44] Andreas Brinkmann. Introduction to average Hamiltonian theory. I. Basics. *Concepts in Magnetic Resonance Part A*, 45A(6):e21414, 2017.
- [45] Richard R. Ernst, Geoffrey Bodenhausen, and Alexander Wokaun. *Principles of Nuclear Magnetic Resonance in One and Two Dimensions*. Oxford University Press, Oxford, 1990.
- [46] Joseph Louis Lagrange. *Recherches sur les Équations Séculaires des Mouvements des Noeuds, et des Inclinaisons des Orbites des Planètes*. de l’Imprimerie Royale, Paris, 1778.
- [47] Steffen J. Glaser and Jens J. Quant. *Homonuclear and Heteronuclear Hartmann—Hahn Transfer in Isotropic Liquids*, volume 19, pages 59–254. Academic Press, 1996.
- [48] J. S Waugh. Theory of broadband spin decoupling. *Journal of Magnetic Resonance*, 50(1):30–49, 1982.

- [49] Florian Dolde, Ville Bergholm, Ya Wang, Ingmar Jakobi, Boris Naydenov, Sébastien Pezzagna, Jan Meijer, Fedor Jelezko, Philipp Neumann, Thomas Schulte-Herbrüggen, et al. High-fidelity spin entanglement using optimal control. *Nature Communications*, 5(1):3371, 2014.
- [50] Jack Saywell, Max Carey, Mohammad Belal, Ilya Kuprov, and Tim Freegarde. Optimal control of raman pulse sequences for atom interferometry. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 53(8):085006, 2020.
- [51] Navin Khaneja, Timo Reiss, Burkhard Luy, and Steffen J. Glaser. Optimal control of spin dynamics in the presence of relaxation. *Journal of Magnetic Resonance*, 162(2):311–319, 2003.
- [52] L. Van Damme, F. Mauconduit, T. Chambrion, N. Boulant, and V. Gras. Universal nonselective excitation and refocusing pulses with improved robustness to off-resonance for magnetic resonance imaging at 7 tesla with parallel transmission. *Magnetic Resonance in Medicine*, 85(2):678–693, 2021.
- [53] Eric Van Reeth, Hélène Ratiney, Kevin Tse Ve Koon, Michael Tesch, Denis Grenier, Olivier Beuf, Steffen J Glaser, and Dominique Sugny. A simplified framework to optimize mri contrast preparation. *Magnetic Resonance in Medicine*, 81(1):424–438, 2019.
- [54] Q Ansel, E Dionis, F Arrouas, B Peaudecerf, S Guérin, D Guéry-Odelin, and D Sugny. Introduction to theoretical and experimental aspects of quantum optimal control. *Journal of Physics B*, 57(13):133001, 2024.
- [55] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbruggen, and S. J. Glaser. Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance*, 172(2):296–305, 2005.
- [56] P. de Fouquieres, S. G. Schirmer, S. J. Glaser, and I. Kuprov. Second order gradient ascent pulse engineering. *Journal of Magnetic Resonance*, 212(2):412–7, 2011.
- [57] D. L. Goodwin and Ilya Kuprov. Auxiliary matrix formalism for interaction representation transformations, optimal control, and spin relaxation theories. *The Journal of Chemical Physics*, 143(8):084113, 2015.
- [58] H. J. Hogben, M. Krzystyniak, G. T. Charnock, P. J. Hore, and I. Kuprov. Spinach—a software library for simulation of spin dynamics in large spin systems. *Journal of Magnetic Resonance*, 208(2):179–194, 2011.
- [59] Richard H. Byrd, Jorge Nocedal, and Robert B. Schnabel. Representations of quasi-Newton matrices and their use in limited memory methods. *Mathematical Programming*, 63(1):129–156, 1994.
- [60] Dong C. Liu and Jorge Nocedal. FMINLBFGS: Fast limited memory optimizer. MATLAB Central File Exchange. <https://ww2.mathworks.cn/matlabcentral/fileexchange/23245-fminlbfgs-fast-limited-memory-optimizer>. Accessed: Jan. 15, 2024.
- [61] D. L. Goodwin and I. Kuprov. Modified newton-raphson GRAPE methods for optimal control of spin systems. *Journal of Chemical Physics*, 144(20):204107, 2016.
- [62] Walter Ritz. Über eine neue methode zur lösung gewisser variationsprobleme der mathematischen physik. *Journal für die reine und angewandte Mathematik*, 135:1–61, 1909.
- [63] B. G. Galerkin. *Rods and plates : series in some questions of elastic equilibrium of rods and plates*. National Technical Information Service, Springfield, VA, 1968.

-
- [64] MathWorks. Decomposition – Matrix Decomposition for Solving Linear Systems. MATLAB Documentation. <https://www.mathworks.com/help/matlab/ref/decomposition.html>. Accessed: Feb. 15, 2025.
- [65] Timothy A. Davis. *Direct Methods for Sparse Linear Systems*, volume 2. SIAM, 2006.
- [66] Andreas Horneff, Michael Eder, Erich Hell, Johannes Ulrici, Jörg Felder, Volker Rasche, and Jens Anders. An EM simulation-based design flow for custom-built MR coils incorporating signal and noise. *IEEE Transactions on Medical Imaging*, 37(2):527–535, 2018.
- [67] Ilya Kuprov. Large-scale NMR simulations in liquid state: A tutorial. *Magnetic Resonance in Chemistry*, 56(6):415–437, 2018.
- [68] D. I. Hoult. The principle of reciprocity in signal strength calculations—a mathematical guide. *Concepts in Magnetic Resonance*, 12(4):173–187, 2000.
- [69] COMSOL AB. Theory for electromagnetic waves, frequency domain, 2019. https://doc.comsol.com/5.5/doc/com.comsol.help.rf/rf_ug_theory.06.12.html. Accessed: July 29, 2025.
- [70] Jeffrey C. Lagarias, James A. Reeds, Margaret H. Wright, and Paul E. Wright. Convergence properties of the nelder–mead simplex method in low dimensions. *SIAM Journal on Optimization*, 9(1):112–147, 1998.
- [71] Kurt Majewski. Simultaneous optimization of radio frequency and gradient waveforms with exact hessians and slew rate constraints applied to kT-point excitation. *Journal of Magnetic Resonance*, 326:106941, 2021.
- [72] Paul A. Keifer. 90° pulse width calibrations: How to read a pulse width array. *Concepts in Magnetic Resonance*, 11(3):165–180, 1999.
- [73] Uluk Rasulov, Anupama Acharya, Marina Carravetta, Guinevere Mathies, and Ilya Kuprov. Simulation and design of shaped pulses beyond the piecewise-constant approximation. *Journal of Magnetic Resonance*, 353:107478, 2023.
- [74] Zdeněk Tošner, Riddhiman Sarkar, Johanna Becker-Baldus, Clemens Glaubitz, Sebastian Wegner, Frank Engelke, Steffen J. Glaser, and Bernd Reif. Overcoming volume selectivity of dipolar recoupling in biological solid-state NMR spectroscopy. *Angewandte Chemie International Edition*, 57(44):14514–14518, 2018.
- [75] Xi-An Mao and Chao-Hui Ye. Understanding radiation damping in a simple way. *Concepts in Magnetic Resonance*, 9(3):173–187, 1997.
- [76] V. V. Krishnan. Radiation damping in microcoil NMR probes. *Journal of Magnetic Resonance*, 179(2):294–298, 2006.
- [77] P. J. van Bentum, J. W. Janssen, A. P. Kentgens, J. Bart, and J. G. Gardeniers. Stripline probes for nuclear magnetic resonance. *Journal of Magnetic Resonance*, 189(1):104–113, 2007.
- [78] Luke J. Edwards and Ilya Kuprov. Parallel density matrix propagation in spin dynamics simulations. *The Journal of Chemical Physics*, 136(4):044108, 2012.
- [79] Kyryl Kobzar, Burkhard Luy, Navin Khaneja, and Steffen J. Glaser. Pattern pulses: Design of arbitrary excitation profiles as a function of pulse amplitude and offset. *Journal of Magnetic Resonance*, 173(2):229–235, 2005.
- [80] Thomas E Skinner, Timo O Reiss, Burkhard Luy, Navin Khaneja, and Steffen J Glaser. Application of optimal control theory to the design of broadband excitation pulses for high-resolution NMR. *Journal of Magnetic Resonance*, 163(1):8–15,

- 2003.
- [81] Thomas E. Skinner, Timo O. Reiss, Burkhard Luy, Navin Khaneja, and Steffen J. Glaser. Reducing the duration of broadband excitation pulses using optimal control with limited RF amplitude. *Journal of Magnetic Resonance*, 167(1):68–74, 2004.
 - [82] Kyril Kobzar, Thomas E. Skinner, Navin Khaneja, Steffen J. Glaser, and Burkhard Luy. Exploring the limits of broadband excitation and inversion pulses. *Journal of Magnetic Resonance*, 170(2):236–243, 2004.
 - [83] Kyril Kobzar, Thomas E. Skinner, Navin Khaneja, Steffen J. Glaser, and Burkhard Luy. Exploring the limits of broadband excitation and inversion: II. Rf-power optimized pulses. *Journal of Magnetic Resonance*, 194(1):58–66, 2008.
 - [84] Kyril Kobzar, Sebastian Ehni, Thomas E. Skinner, Steffen J. Glaser, and Burkhard Luy. Exploring the limits of broadband 90° and 180° universal rotation pulses. *Journal of Magnetic Resonance*, 225:142–160, 2012.
 - [85] Ulrich Katscher and Peter Börnert. Parallel RF transmission in MRI. *NMR in Biomedicine*, 19(3):393–400, 2006.
 - [86] Klaas P. Pruessmann. Encoding and reconstruction in parallel MRI. *NMR in Biomedicine*, 19(3):288–299, 2006.
 - [87] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller. The NMR phased array. *Magnetic Resonance in Medicine*, 16(2):192–225, 1990.
 - [88] Michael Hutchinson and Ulrich Raff. Fast MRI data acquisition using multiple detectors. *Magnetic Resonance in Medicine*, 6(1):87–91, 1988.
 - [89] Antonello Sotgiu and James S. Hyde. High-order coils as transmitters for NMR imaging. *Magnetic Resonance in Medicine*, 3(1):55–62, 1986.
 - [90] Tamer S. Ibrahim, Robert Lee, Brian A. Baertlein, Allahyar Kangarlu, and Pierre-Marie L. Robitaille. Application of finite difference time domain method for the design of birdcage RF head coils using multi-port excitations. *Magnetic Resonance Imaging*, 18(6):733–742, 2000.
 - [91] Yudong Zhu. Parallel excitation with an array of transmit coils. *Magnetic Resonance in Medicine*, 51(4):775–784, 2004.
 - [92] Gregor Adriany, Pierre-Francois Van de Moortele, Florian Wiesinger, Steen Moeller, John P. Strupp, Peter Andersen, Carl Snyder, Xiaoliang Zhang, Wei Chen, Klaas P. Pruessmann, Peter Boesiger, Tommy Vaughan, and Kāmil Uğurbil. Transmit and receive transmission line arrays for 7 Tesla parallel imaging. *Magnetic Resonance in Medicine*, 53(2):434–445, 2005.
 - [93] Ulrich Katscher, Peter Börnert, Christoph Leussler, and Johan S. van den Brink. Transmit SENSE. *Magnetic Resonance in Medicine*, 49(1):144–150, 2003.
 - [94] M. A. Cloos, N. Boulant, M. Luong, G. Ferrand, E. Giacomini, M. F. Hang, C. J. Wiggins, D. Le Bihan, and A. Amadon. Parallel-transmission-enabled magnetization-prepared rapid gradient-echo T1-weighted imaging of the human brain at 7T. *NeuroImage*, 62(3):2140–2150, 2012.
 - [95] Cem M. Deniz, Leeor Alon, Ryan Brown, and Yudong Zhu. Subject- and resource-specific monitoring and proactive management of parallel radiofrequency transmission. *Magnetic Resonance in Medicine*, 76(1):20–31, 2016.
 - [96] Aurélien Massire, Alexandre Vignaud, Benjamin Robert, Denis Le Bihan, Nicolas Boulant, and Alexis Amadon. Parallel-transmission-enabled three-dimensional T2-weighted imaging of the human brain at 7 Tesla. *Magnetic Resonance in Medicine*,

- 73(6):2195–2203, 2015.
- [97] V. Gras, A. Vignaud, A. Amadon, D. Le Bihan, and N. Boulant. Universal pulses: A new concept for calibration-free parallel transmission. *Magnetic Resonance in Medicine*, 77(2):635–643, 2017.
 - [98] Vincent Gras, Franck Mauconduit, Alexandre Vignaud, Alexis Amadon, Denis Le Bihan, Tony Stöcker, and Nicolas Boulant. Design of universal parallel-transmit refocusing kT-point pulses and application to 3D T2-weighted imaging at 7T. *Magnetic Resonance in Medicine*, 80(1):53–65, 2018.
 - [99] C. Le Ster, F. Mauconduit, A. Massire, N. Boulant, and V. Gras. Standardized universal pulse: A fast RF calibration approach to improve flip angle accuracy in parallel transmission. *Magnetic Resonance in Medicine*, 87(6):2839–2850, 2022.
 - [100] Gregor Adriany, Edward J. Auerbach, Carl J. Snyder, Ark Gözübüyük, Steen Moeller, Johannes Ritter, Pierre-François Van de Moortele, Tommy Vaughan, and Kâmil Uğurbil. A 32-channel lattice transmission line array for parallel transmit and receive MRI at 7 tesla. *Magnetic Resonance in Medicine*, 63(6):1478–1485, 2010.
 - [101] P. Vernickel, P. Röschmann, C. Findekle, K.-M. Lüdeke, Ch. Leussler, J. Overweg, U. Katscher, I. Grässlin, and K. Schünemann. Eight-channel transmit/receive body MRI coil at 3T. *Magnetic Resonance in Medicine*, 58(2):381–389, 2007.
 - [102] Nikolai I. Avdievich, Jullie W. Pan, and Hoby P. Hetherington. Resonant inductive decoupling (RID) for transceiver arrays to compensate for both reactive and resistive components of the mutual impedance. *NMR in Biomedicine*, 26(11):1547–1554, 2013.
 - [103] Katherine Lynn Moody, Neal A. Hollingsworth, Feng Zhao, Jon-Fredrik Nielsen, Douglas C. Noll, Steven M. Wright, and Mary Preston McDougall. An eight-channel T/R head coil for parallel transmit MRI at 3T using ultra-low output impedance amplifiers. *Journal of Magnetic Resonance*, 246:62–68, 2014.
 - [104] Chang-Hoon Choi, Andrew Webb, Stephan Orzada, Mikheil Kelenjeridze, N. Jon Shah, and Jörg Felder. A review of parallel transmit arrays for ultra-high field mr imaging. *IEEE Reviews in Biomedical Engineering*, 17:351–368, 2024.
 - [105] Angelo Galante, Marco Fantasia, Piero Sebastiani, Antonello Sotgiu, and Marcello Alecci. Analytical RF coils decoupling: A theoretical approach for parallel transmission. In *Proceedings of the International Society for Magnetic Resonance in Medicine*, volume 25, 2017.
 - [106] P Vernickel, C Findekle, J Eichmann, and Ingmar Grässlin. Active digital decoupling for multi channel transmit MRI systems. *Proceedings of the International Society for Magnetic Resonance in Medicine*, 15, 2007.
 - [107] Teresa Carlomagno, Burkhard Luy, and Steffen J. Glaser. ”Kin” HEHAHA sequences, heteronuclear Hartmann–Hahn transfer with different bandwidths for spins I and S. *Journal of Magnetic Resonance*, 126(1):110–119, 1997.
 - [108] Michael Braun and Steffen J. Glaser. Cooperative pulses. *Journal of Magnetic Resonance*, 207(1):114–123, 2010.
 - [109] Michael Braun and Steffen J. Glaser. Concurrently optimized cooperative pulses in robust quantum control: Application to broadband ramsey-type pulse sequence elements. *New Journal of Physics*, 16(11):115002, 2014.
 - [110] Benjamin Görling, Stefan Bräse, and Burkhard Luy. HR-HSBC: Measuring heteronuclear one-bond couplings with enhanced resolution. *Magnetic Resonance in*

- Chemistry*, 50(S1):S58–S62, 2012.
- [111] Sebastian Ehni and Burkhard Luy. A systematic approach for optimizing the robustness of pulse sequence elements with respect to couplings, offsets, and B1-field inhomogeneities (COB). *Magnetic Resonance in Chemistry*, 50(S1):S63–S72, 2012.
 - [112] Sebastian Ehni and Burkhard Luy. BEBEtr and BUBI: J-compensated concurrent shaped pulses for 1H–13C experiments. *Journal of Magnetic Resonance*, 232:7–17, 2013.
 - [113] L. Tong, R.-w. Liu, V.C. Soon, and Y.-F. Huang. Indeterminacy and identifiability of blind identification. *IEEE Transactions on Circuits and Systems*, 38(5):499–509, 1991.
 - [114] Christian Jutten and Jeanny Herault. Blind separation of sources, part I: An adaptive algorithm based on neuromimetic architecture. *Signal Processing*, 24(1):1–10, 1991.
 - [115] Taihui Wang, Feiran Yang, and Jun Yang. Convolutional transfer function-based multichannel nonnegative matrix factorization for overdetermined blind source separation. *IEEE/ACM Transactions on Audio, Speech and Language Processing*, 30:802–815, 2022.
 - [116] Sheng Chang, Yunkai Deng, Yanyan Zhang, Qingchao Zhao, Robert Wang, and Ke Zhang. An advanced scheme for range ambiguity suppression of spaceborne sar based on blind source separation. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–12, 2022.
 - [117] Chen Ye, Kentaroh Toyoda, and Tomoaki Ohtsuki. Blind source separation on non-contact heartbeat detection by non-negative matrix factorization algorithms. *IEEE Transactions on Biomedical Engineering*, 67(2):482–494, 2020.
 - [118] K. J. Pope and R. E. Bogner. Blind signal separation i. linear, instantaneous combinations: I. linear, instantaneous combinations. *Digital Signal Processing*, 6(1):5–16, 1996.
 - [119] K. J. Pope and R. E. Bogner. Blind signal separation ii. linear, convolutional combinations: II. linear, convolutional combinations. *Digital Signal Processing*, 6(1):17–28, 1996.
 - [120] A. Belouchrani, K. Abed-Meraim, J.-F. Cardoso, and E. Moulines. A blind source separation technique using second-order statistics. *IEEE Transactions on Signal Processing*, 45(2):434–444, 1997.
 - [121] James Munkres. Algorithms for the assignment and transportation problems. *Journal of the society for industrial and applied mathematics*, 5(1):32–38, 1957.
 - [122] H. W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955.
 - [123] K. Matsuoka. Minimal distortion principle for blind source separation. In *Proceedings of the 41st SICE Annual Conference. SICE 2002.*, volume 4, pages 2138–2143, Osaka, Japan, 2002. Soc. Instrument & Control Eng. (SICE).
 - [124] Moritz Becker, Yen-Tse Cheng, Achim Voigt, Ajmal Chenakkara, Mengjia He, Sören Lehmkuhl, Mazin Jouda, and Jan G. Korvink. Artificial intelligence-driven shimming for parallel high field nuclear magnetic resonance. *Scientific Reports*, 13(1):17983, 2023.
 - [125] Jérôme Bobin, Jérémy Rapin, Anthony Larue, and Jean-Luc Starck. Sparsity and adaptivity for the blind separation of partially correlated sources. *IEEE Transactions on Signal Processing*, 63(5):1199–1213, 2015.

- [126] Mathias Nilsson and Gareth A. Morris. Improved DECRA processing of DOSY data: correcting for non-uniform field gradients. *Magnetic Resonance in Chemistry*, 45(8):656–660, 2007.
- [127] Mathias Nilsson and Gareth A. Morris. Speedy component resolution: An improved tool for processing diffusion-ordered spectroscopy data. *Analytical Chemistry*, 80(10):3777–3782, 2008.
- [128] Enping Lin, Yu Yang, Yuqing Huang, and Zhong Chen. High-resolution reconstruction for diffusion-ordered NMR spectroscopy. *Analytical Chemistry*, 92(1):634–639, 2020.
- [129] Stella Slad, Wolfgang Bermel, Rainer Kümmerle, Daniel Mathieu, and Burkhard Luy. Band-selective universal 90° and 180° rotation pulses covering the aliphatic carbon chemical shift range for triple resonance experiments on 1.2 ghz spectrometers. *Journal of Biomolecular NMR*, 76(5):185–195, 2022.
- [130] Raphaël Tomi-Tricot, Vincent Gras, Bertrand Thirion, Franck Mauconduit, Nicolas Boulant, Hamza Cherkaoui, Pierre Zerbib, Alexandre Vignaud, Alain Luciani, and Alexis Amadon. Smartpulse, a machine learning approach for calibration-free dynamic RF shimming: Preliminary study in a clinical environment. *Magnetic Resonance in Medicine*, 82(6):2016–2031, 2019.
- [131] Stephan Grzesiek and Ad Bax. Spin-locked multiple quantum coherence for signal enhancement in heteronuclear multidimensional NMR experiments. *Journal of Biomolecular NMR*, 6(3):335–339, 1995.
- [132] Stephen J. DeVience, Ronald L. Walsworth, and Matthew S. Rosen. Nuclear spin singlet states as a contrast mechanism for NMR spectroscopy. *NMR in Biomedicine*, 26(10):1204–1212, 2013.
- [133] Stephen J. DeVience, Ronald L. Walsworth, and Matthew S. Rosen. Preparation of nuclear spin singlet states using spin-lock induced crossing. *Physical Review Letters*, 111(17):173002, 2013.
- [134] Bogdan A. Rodin, Alexey S. Kiryutin, Alexandra V. Yurkovskaya, Konstantin L. Ivanov, Satoru Yamamoto, Kazunobu Sato, and Takeji Takui. Using optimal control methods with constraints to generate singlet states in NMR. *Journal of Magnetic Resonance*, 291:14–22, 2018.
- [135] Anna Sonnefeld, Aiky Razanahoera, Philippe Pelupessy, Geoffrey Bodenhausen, and Kirill Sheberstov. Long-lived states of methylene protons in achiral molecules. *Science Advances*, 8(48):eade2113, 2022.
- [136] Danila A. Barskiy, Oleg G. Salnikov, Alexey S. Romanov, Matthew A. Feldman, Aaron M. Coffey, Kirill V. Kovtunov, Igor V. Koptug, and Eduard Y. Chekmenev. NMR spin-lock induced crossing (SLIC) dispersion and long-lived spin states of gaseous propane at low magnetic field (0.05T). *Journal of Magnetic Resonance*, 276:78–85, 2017.
- [137] Stephen J. DeVience, Ronald L. Walsworth, and Matthew S. Rosen. Probing scalar coupling differences via long-lived singlet states. *Journal of Magnetic Resonance*, 262:5761, 2015.
- [138] Stephen J. DeVience, Mason Greer, Soumyajit Mandal, and Matthew S. Rosen. Homonuclear J-coupling spectroscopy at low magnetic fields using spin-lock induced crossing. *ChemPhysChem*, 22(20):2128–2137, 2021.
- [139] Kirill V. Kovtunov, Milton L. Truong, Danila A. Barskiy, Igor V. Koptug, Aaron M. Coffey, Kevin W. Waddell, and Eduard Y. Chekmenev. Long-lived

- spin states for low-field hyperpolarized gas MRI. *Chemistry: A European Journal*, 20(45):14629–14632, 2014.
- [140] Baiyan Jiang and Weitian Chen. On-resonance and off-resonance continuous wave constant amplitude spin-lock and $T_{1\rho}$ quantification in the presence of B_1 and B_0 inhomogeneities. *NMR in Biomedicine*, 31(7):e3928, 2018.
- [141] Maximilian Gram, Michael Seethaler, Daniel Gensler, Johannes Oberberger, Peter M. Jakob, and Peter Nordbeck. Balanced spin-lock preparation for B_1 -insensitive and B_0 -insensitive quantification of the rotating frame relaxation time $T_{1\rho}$. *Magnetic Resonance in Medicine*, 85(5):2771–2780, 2021.
- [142] Kevin C. Miao, Joseph P. Blanton, Christopher P. Anderson, Alexandre Bourassa, Alexander L. Crook, Gary Wolfowicz, Hiroshi Abe, Takeshi Ohshima, and David D. Awschalom. Universal coherence protection in a solid-state spin qubit. *Science*, 369(6510):1493–1497, 2020.
- [143] M. Loretz, T. Rosskopf, and C. L. Degen. Radio-frequency magnetometry using a single electron spin. *Physical Review Letters*, 110(1):017602, 2013.
- [144] A. M. Waeber, G. Gillard, G. Ragunathan, M. Hopkinson, P. Spencer, D. A. Ritchie, M. S. Skolnick, and E. A. Chekhovich. Pulse control protocols for preserving coherence in dipolar-coupled nuclear spin baths. *Nature Communications*, 10(1):3157, 2019.
- [145] Andrew J. Ramsay, Reza Hekmati, Charlie J. Patrickson, Simon Baber, David R. M. Arvidsson-Shukur, Anthony J. Bennett, and Isaac J. Luxmoore. Coherence protection of spin qubits in hexagonal boron nitride. *Nature Communications*, 14(1):461, 2023.
- [146] S. Rohr. Synchronizing the dynamics of a single nitrogen vacancy spin qubit on a parametrically coupled radio-frequency field through microwave dressing. *Physical Review Letters*, 112(1), 2014.
- [147] E. O. Stejskal and J. E. Tanner. Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics.*, 42(1):288–292, 1965.
- [148] M Holz and H Weingartner. Calibration in accurate spin-echo self-diffusion measurements using ^1H and less-common nuclei. *Journal of Magnetic Resonance (1969)*, 92(1):115–125, 1991.
- [149] Geoffrey Bodenhausen and David J. Ruben. Natural abundance nitrogen-15 NMR by enhanced heteronuclear spectroscopy. *Chemical Physics Letters*, 69(1):185–189, 1980.
- [150] David M LeMaster and Frederic M Richards. Proton-nitrogen-15 heteronuclear NMR studies of escherichia coli thioredoxin in samples isotopically labeled by residue type. *Biochemistry*, 24(25):7263–7268, 1985.
- [151] Yizhou Liu and James H Prestegard. Measurement of one and two bond N–C couplings in large proteins by TROSY-based J-modulation experiments. *Journal of Magnetic Resonance*, 200(1):109–118, 2009.
- [152] Elin Alexandersson and Gustav Nestor. Complete ^1H and ^{13}C NMR spectral assignment of d-glucofuranose. *Carbohydrate Research*, 511:108477, 2022.
- [153] K. Kamieńska-Trela. One-bond ^{13}C - ^{13}C spin-spin coupling constants. In G. A. Webb, editor, *Annual Reports on NMR Spectroscopy*, volume 30, pages 131–230. Academic Press, 1995.
- [154] Mattias U. Roslund, Petri Tähtinen, Matthias Niemitz, and Rainer Sjöholm. Com-

- plete assignments of the ^1H and ^{13}C chemical shifts and $J_{\text{H,H}}$ coupling constants in NMR spectra of D-glucopyranose and all D-glucopyranosyl-D-glucopyranosides. *Carbohydrate Research*, 343(1):101–112, 2008.
- [155] Bingwu Yu, Hugo van Ingen, Subramanian Vivekanandan, Christoph Rademacher, Scott E. Norris, and Darón I. Freedberg. More accurate $^1J_{\text{CH}}$ coupling measurement in the presence of $^3J_{\text{HH}}$ strong coupling in natural abundance. *Journal of Magnetic Resonance*, 215:10–22, 2012.
- [156] Jesus Ruiz-Cabello, Geerten W Vuister, Chrit T. W Moonen, Peter van Gelderen, Jack S Cohen, and Peter C. M van Zijl. Gradient-enhanced heteronuclear correlation spectroscopy. theory and experimental aspects. *Journal of Magnetic Resonance*, 100(2):282–302, 1992.
- [157] Krister Svanberg. The method of moving asymptotes—a new method for structural optimization. *International Journal for Numerical Methods in Engineering*, 24(2):359–373, 1987.
- [158] Christiane P. Koch, Ugo Boscain, Tommaso Calarco, Gunther Dirr, Stefan Filipp, Steffen J. Glaser, Ronnie Kosloff, Simone Montangero, Thomas Schulte-Herbrüggen, Dominique Sugny, and Frank K. Wilhelm. Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe. *EPJ Quantum Technology*, 9(1):1–60, 2022.
- [159] Ilya Kuprov and Christopher T. Rodgers. Derivatives of spin dynamics simulations. *The Journal of Chemical Physics*, 131(23):234108, 2009.
- [160] D. L. Goodwin and Ilya Kuprov. Auxiliary matrix formalism for interaction representation transformations, optimal control, and spin relaxation theories. *The Journal of Chemical Physics*, 143(8):084113, 2015.
- [161] Nelson Leung, Mohamed Abdelhafez, Jens Koch, and David Schuster. Speedup for quantum optimal control from automatic differentiation based on graphics processing units. *Physical Review A*, 95(4):042318, 2017.
- [162] Mohamed Abdelhafez, David I Schuster, and Jens Koch. Gradient-based optimal control of open quantum systems using quantum trajectories and automatic differentiation. *Physical Review A*, 99(5):052327, 2019.
- [163] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen. Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework. *Physical Review A*, 84(2):022305, 2011.
- [164] Jesper Hasseriis Mohr Jensen, Frederik Skovbo Møller, Jens Jakob Sørensen, and Jacob Friis Sherson. Approximate dynamics leading to more optimal control: Efficient exact derivatives. *Physical Review A*, 103(6):062612, 2021.
- [165] David L. Goodwin and Mads Sloth Vinding. Accelerated newton-raphson GRAPE methods for optimal control. *Physical Review Research*, 5(1):L012042, 2023.
- [166] Zi-Jie Chen, Hongwei Huang, Lida Sun, Qing-Xuan Jie, Jie Zhou, Ziyue Hua, Yifang Xu, Weiting Wang, Guang-Can Guo, Chang-Ling Zou, Luyan Sun, and Xu-Bo Zou. Robust and optimal control of open quantum systems. *Science Advances*, 11(9):eadr0875, 2025.
- [167] David Joseph and Christian Griesinger. Optimal control pulses for the 1.2-GHz (28.2-T) NMR spectrometers. *Science Advances*, 9(45):eadj1133, 2023.
- [168] Eric R. Lowe, Stefan Stoll, and J. P. Kestner. Optimizing EPR pulses for broadband excitation and refocusing. *Journal of Magnetic Resonance*, 369:107807, 2024.

- [169] Uluk Rasulov and Ilya Kuprov. Instrumental distortions in quantum optimal control. *The Journal of Chemical Physics*, 162(16):164107, 2025.
- [170] David L. Goodwin, Pranav Singh, and Mohammadali Foroozandeh. Adaptive optimal control of entangled qubits. *Science Advances*, 8(49):abq4244, 2022.
- [171] Sebastian Ehni and Burkhard Luy. BEBEtr and BUBI: J-compensated concurrent shaped pulses for 1H – 13C experiments. *Journal of Magnetic Resonance*, 232:7–17, 2013.
- [172] Z. Tosner, M. J. Brandl, J. Blahut, S. J. Glaser, and B. Reif. Maximizing efficiency of dipolar recoupling in solid-state NMR using optimal control sequences. *Science Advances*, 7(42):eabj5913, 2021.
- [173] Mads S. Vinding, Bastien Guérin, Thomas Vosegaard, and Niels Chr. Nielsen. Local sar, global sar, and power-constrained large-flip-angle pulses with optimal control and virtual observation points. *Magnetic Resonance in Medicine*, 77(1):374–384, 2017.
- [174] M. S. Vinding, B. Skyum, R. Sangill, and T. E. Lund. Ultrafast (milliseconds), multidimensional rf pulse design with deep learning. *Magnetic Resonance in Medicine*, 82(2):586–599, 2019.
- [175] Thomas J. R. Hughes. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Courier Corporation, Mineola, NY, 2012.
- [176] G. Stang and G. Fix. *An Analysis of the Finite Element Method*. Prentice Hall, Englewood Cliffs, NJ, 1973.
- [177] Gary P. Zientara and Jack H. Freed. The variational method and the stochastic-liouville equation. I. A finite element solution to the CIDN(E)P problema. *The Journal of Chemical Physics*, 70(6):2587–2598, 1979.
- [178] Gary P. Zientara and Jack H. Freed. The variational method and the stochastic-liouville equation. III. infinite elements for CIDN(E)P. *The Journal of Chemical Physics*, 71(2):744–749, 1979.
- [179] Arthur E. Stillman, Gary P. Zientara, and Jack H. Freed. The variational method and the stochastic-liouville equation. II. ESR spectral simulation via finite elements. *The Journal of Chemical Physics*, 71(1):113–118, 1979.
- [180] Yang Cao, Shengtai Li, Linda Petzold, and Radu Serban. Adjoint sensitivity analysis for differential-algebraic equations: The adjoint DAE system and its numerical solution. *SIAM Journal on Scientific Computing*, 24(3):1076–1089, 2003.
- [181] B. S. Lazarov and O. Sigmund. Filters in topology optimization based on Helmholtz-type differential equations. *International Journal for Numerical Methods in Engineering*, 86(6):765–781, 2011.
- [182] Ole Sigmund. Morphology-based black and white filters for topology optimization. *Structural and Multidisciplinary Optimization*, 33(4):401–424, 2007.
- [183] Ji-Hong Zhu, Wei-Hong Zhang, and Liang Xia. Topology optimization in aircraft and aerospace structures design. *Archives of Computational Methods in Engineering*, 23(4):595–622, 2016.
- [184] Krister Svanberg. An algorithm for optimum structural design using duality. *Mathematical Programming Study*, 20:161–177, 1982.
- [185] Krister Svanberg. Some modelling aspects for the matlab implementation of mma, 2004. Technical report, KTH Royal Institute of Technology.
- [186] Kyril Kobzar, Sebastian Ehni, Thomas E. Skinner, Steffen J. Glaser, and Burkhard

- Luy. Exploring the limits of broadband 90° and 180° universal rotation pulses. *Journal of Magnetic Resonance*, 225:142–160, 2012.
- [187] Krister Svanberg. A Class of Globally Convergent Optimization Methods Based on Conservative Convex Separable Approximations. *SIAM Journal on Optimization*, 12(2):555–573, 2002.
- [188] Wolfgang Bangerth, Carsten Burstedde, Timo Heister, and Martin Kronbichler. Algorithms and data structures for massively parallel generic adaptive finite element codes. *ACM Transactions on Mathematical Software*, 38(2):14:1–14:28, 2012.
- [189] Daniel Arndt, Wolfgang Bangerth, Denis Davydov, Timo Heister, Luca Heltai, Martin Kronbichler, Matthias Maier, Jean-Paul Pelteret, Bruno Turcksin, and David Wells. The deal.II finite element library: Design, features, and insights. *Computers & Mathematics with Applications*, 81:407–422, 2021.
- [190] José M. Rodríguez-Borbón, Xian Wang, Adrián P. Diéguez, Khaled Z. Ibrahim, and Bryan M. Wong. TRAVOLTA: GPU acceleration and algorithmic improvements for constructing quantum optimal control fields in photo-excited systems. *Computer Physics Communications*, 296:109017, 2024.
- [191] José M. Rodríguez-Borbón, Xian Wang, Adrián P. Diéguez, Khaled Z. Ibrahim, and Bryan M. Wong. VAN-DAMME: GPU-accelerated and symmetry-assisted quantum optimal control of multi-qubit systems. *Computer Physics Communications*, 307:109403, 2025.
- [192] Anupama Acharya, Madhukar Said, Sylwia J. Barker, Marcel Utz, Bruno Linclau, and Ilya Kuprov. Chemical hydrodynamics of nuclear spin states. *Science Advances*, 11(43):eady9103, 2025.
- [193] Mehrdad Alinaghian Jouzdani, Mazin Jouda, and Jan G. Korvink. Optimal control flow encoding for time-efficient magnetic resonance velocimetry. *Journal of Magnetic Resonance*, 352:107461, 2023.
- [194] N. Jbili, K. Hamraoui, S. J. Glaser, J. Salomon, and D. Sugny. Optimal periodic control of spin systems: Application to the maximization of the signal-to-noise ratio per unit time. *Physical Review A*, 99(5):053415, 2019.
- [195] A Oppelt, R Graumann, H Barfuss, H Fischer, W Hartl, W Schajor, et al. FISP—a new fast MRI sequence. *Electromedica*, 54(1):15–18, 1986.
- [196] Shebha Anandhi Jegadeesan, Yujie Zhao, Robert Hunter, Hassane El Mkami, Maximilian Keitel, Callum Musselwhite, Graham Smith, Guinevere Mathies, and Ilya Kuprov. Quantum optimal control of steady orbits. arXiv preprint arXiv:2606.15383, 2026.

Publications

Journal articles

1. **Mengjia He**, Dilara Faderl, Neil MacKinnon, Yen-Tse Cheng, Dominique Buyens, Mazin Jouda, Burkhard Luy, and Jan G. Korvink. A digital twin for parallel liquid-state nuclear magnetic resonance spectroscopy. *Communications Engineering*, 3(1): 90, 2024.
2. **Mengjia He**, Neil MacKinnon, Dominique Buyens, Burkhard Luy, and Jan G. Korvink. Coherence locking in a parallel NMR probe defends against gradient field spillover. *Magnetic Resonance*, 6(2): 173-181, 2025.
3. **Mengjia He**, Yongbo Deng, Burkhard Luy, and Jan G. Korvink. Finite elements and moving asymptotes accelerate quantum optimal control – FEMMA. *The Journal of Chemical Physics*, 164 (6): 064114, 2026.

Conference contributions

1. **Mengjia He**, Neil MacKinnon, Mazin Jouda, Yen-Tse Cheng, Burkhard Luy, Jan G. Korvink. Pulse compensation and signal decomposition for parallel NMR experiments. EuroMar, Glasgow, UK, 2023. Poster Presentation.
2. **Mengjia He**, Neil MacKinnon, Burkhard Luy, Jan G. Korvink. Scheme for defending gradient field spillover in multi-detector NMR by coherence locking. EuroMar, Bilbao, Spain, 2024. Poster Presentation.
3. **Mengjia He**, Neil MacKinnon, Yongbo Deng, Burkhard Luy, and Jan G. Korvink. Optimization of magnetic resonance pulse using finite element method. EuroMar, Oulu, Finland, 2025. Poster Presentation.

Acknowledgements

I would like to express my deepest gratitude to my supervisor, Prof. Jan G. Korvink, for giving me the opportunity to pursue this doctoral research. I am sincerely thankful for his continuous support over the past five years, for his insightful supervision, and for his guidance in advancing this work. His encouragement and thoughtful advice have been essential throughout the course of this research.

I am grateful to Dr. Neil MacKinnon for introducing me to the field of NMR and for his generous support. I thank him for the fruitful and helpful discussions during both monthly and individual meetings, for his assistance with the spectrometer and experiments, and for his encouragement to participate in conferences. I also appreciate the unforgettable trip to the North Cape.

My special thanks go to Prof. John Jelonnek for serving as my second examiner and for the detailed and inspiring discussions. I am also grateful to Prof. Christian Koos and Prof. Tobias Huber-Loyola for encouraging me to reflect on my thesis from different perspectives, and to Prof. Bernhard Holzapfel for kindly chairing the examination.

I would like to thank Prof. Burkhard Luy for his continuous support, especially for his insightful comments on each sub-project, for the spin dynamics lecture, and for the joint experiments at IBG.

I would like to thank Yongbo Deng for many helpful discussions over the past two years, and especially for his valuable suggestions on numerical optimization.

I sincerely thank the Micro-NMR group for their support in both academic work and daily life. To Mazin Joude, Dominique Buyens, Yen-Tse Cheng, and Dilara Faderl for their assistance with the experiments and for their valuable contributions to this thesis. To Jianyi Liang and Jing Yang for helping me settle in and move house several times in Karlsruhe.

To Chun Him Lee for helpful discussions and assistance with gradient measurements.

I would like to thank IMT for providing a supportive working environment. I acknowledge the KIT SCC for providing computing time on the bwUniCluster and for organizing HPC training sessions.

I thank the Spinach developer team for developing this simulation library, which greatly supported my work.

I gratefully acknowledge financial support from the Helmholtz Association through the Joint Lab Virtual Materials (VirtMat) Design, from the European Research Council under ERC Synergy Grant HiSCORE (No. 951459), and from the Deutsche Forschungsgemeinschaft (DFG) through CRC 1527 HyPERiON.

I also thank the KIT VirtMat team, in particular Prof. Wolfgang Wenzel, Ilona Koroncz, and Haoran Duan, for the inspiring seminars, workshops, and retreats.

Finally, I would like to express my heartfelt gratitude to my family for their constant understanding, encouragement, and unconditional love throughout my doctoral journey.

Eggenstein-Leopoldshafen, July 2026