



**Industrial Silicon
Metal-Oxide-Semiconductor Spin Qubits as
Quantum Sensors for Single-Molecule
Magnet Qudits**

Zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften
(Dr. rer. nat.)

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

angenommene

Dissertation

von

M.Sc. Daniel Philipp Schroller

Tag der mündlichen Prüfung: 20. Februar 2026

Referent: Prof. Dr. Wolfgang Wernsdorfer

Korreferent: Prof. Dr. David Hunger

Prüfer: Prof. Dr. I. Pop

Prüfer: Prof. Dr. A. Mirlin

Prüferin: Prof. Dr. M. M. Mühlleitner



This document is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0): <https://creativecommons.org/licenses/by/4.0/deed.en>

Abstract

Single-molecule magnets (SMMs) provide chemically engineered, long-lived spin states that are attractive as qubits and quantum memories, but fast, local readout at the single-molecule level remains experimentally challenging. This thesis examines whether industrial silicon metal-oxide-semiconductor (SiMOS) spin qubits can serve as quantum sensors for SMMs, and which advances are required to approach single-molecule sensitivity.

Industrial SiMOS quantum dots in enriched ^{28}Si are operated as spin qubits with reliable spin-selective readout and coherence times of $T_2^* = (3.1 \pm 0.1) \mu\text{s}$ (Ramsey) and $T_2^{\text{HE}} = (92 \pm 10) \mu\text{s}$ (Hahn echo), establishing them as a suitable platform for quantum magnetometry. Quantitative estimates of the dipolar coupling identify the qubit-molecule separation imposed by the industrial gate stack as the main limitation for single-molecule readout. To amplify the molecular signal into the detectable range for current SiMOS devices, an ensemble of the SMM terbium bis(phthalocyanine), TbPc_2 , heavily diluted in a YPc_2 matrix, is deposited on the chip surface. In this hybrid architecture, TbPc_2 serves as a prototypical molecular qubit with long-lived spin states, while the SiMOS qubit provides fast, electrically controlled nanoscale readout. Using a compact rapid adiabatic passage during spin-selective tunnelling protocol, the qubit resonance can be followed in a robust and time-efficient manner across a 100 MHz window, enabling measurements of magnetic hysteresis, angular anisotropy, and slow relaxation of the TbPc_2 ensemble, with relaxation times of $(107 \pm 2) \text{ min}$ at 48 mK and $(0.8 \pm 0.3) \text{ min}$ at 140 mK. These results provide the first demonstration of a hybrid SiMOS-SMM sensing architecture.

Beyond single-dot sensing, double quantum dot schemes based on singlet-triplet qubits and Pauli spin blockade are developed, enabling magnetic-field-gradient sensing and low-field measurements of the absolute Zeeman splitting. Finally, self-consistent Schrödinger-Poisson simulations of the industrial gate stack show that the qubit-surface separation can be reduced to well below 10 nm with realistic design modifications, and a first generation of such sensing-optimised devices is fabricated in the 300 mm industrial cleanrooms of imec, outlining a concrete pathway towards single-molecule sensitivity.

Zusammenfassung

Einzelmolekülmagnete (Single-Molecule Magnets, SMMs) besitzen chemisch gezielt abstimmbare, langlebige Spinzustände und sind damit attraktive Kandidaten für Qudits und Quantenspeicher. Eine schnelle, lokale Auslese auf Einzelmolekülebene ist jedoch experimentell weiterhin anspruchsvoll. Diese Arbeit untersucht, ob industrielle Silizium-Metall-Oxid-Halbleiter-(SiMOS)-Spinqubits als Quantensensoren für SMMs eingesetzt werden können und welche Weiterentwicklungen erforderlich sind, um die Empfindlichkeit bis zur Einzelmolekülgröße zu steigern.

Industrielle SiMOS-Quantenpunkte in angereichertem ^{28}Si werden als Spinqubits mit zuverlässiger spinselektiver Auslese betrieben und zeigen Kohärenzzeiten von $T_2^* = (3.1 \pm 0.1) \mu\text{s}$ (Ramsey) und $T_2^{\text{HE}} = (92 \pm 10) \mu\text{s}$ (Hahn-Echo), was sie als geeignete Plattform für die Quantummagnetometrie etabliert. Quantitative Abschätzungen der dipolaren Kopplung identifizieren den durch den industriellen Gate-Stack vorgegebenen Qubit-Molekül-Abstand als zentrale Einschränkung für die Einzelmolekül-Auslese. Um das molekulare Signal für aktuelle SiMOS-Bauelemente in einen detektierbaren Bereich zu verstärken, wird ein Ensemble des SMM Terbium-bis(phthalocyanin), TbPc_2 , stark verdünnt in einer YPc_2 -Matrix, auf der Chipoberfläche deponiert. In dieser hybriden Architektur dient TbPc_2 als prototypisches molekulares Qudit mit langlebigen Spinzuständen, während das SiMOS-Qubit eine schnelle, elektrisch kontrollierte Auslese mit nanoskaliger Ortsauflösung bereitstellt. Mithilfe eines kompakten Protokolls der schnellen adiabatischen Passage während spinselektiven Tunnelns kann die Qubit-Resonanz robust und zeiteffizient über ein 100 MHz-Fenster verfolgt werden. Dies ermöglicht Messungen von magnetischer Hysterese, Winkelanisotropie und langsamer Relaxation des TbPc_2 -Ensembles mit Relaxationszeiten von $(107 \pm 2) \text{ min}$ bei 48 mK und $(0.8 \pm 0.3) \text{ min}$ bei 140 mK. Diese Ergebnisse stellen die erste Demonstration einer hybriden SiMOS-SMM-Architektur dar.

Über die Einzelquantenpunkt-Sensorik hinaus werden Doppel-Quantenpunkt-Konzepte auf Basis von Singulett-Triplet-Qubits und Pauli-Spin-Blockade entwickelt, die eine Messung von Magnetfeldgradienten sowie Niederfeldmessungen der absoluten Zeeman-Aufspaltung ermöglichen. Schließlich zeigen selbstkonsistente Schrödinger-Poisson-Simulationen des industriellen Gate-Stacks, dass der Qubit-Oberflächen-Abstand durch realistische Designanpassungen auf deutlich unter 10 nm reduziert werden kann. Eine erste Generation derart sensoroptimierter Bauelemente wird in den 300 mm-Industrie-Reinräumen von imec gefertigt und skizziert einen konkreten Weg hin zur Einzelmolekül-Empfindlichkeit.

Contents

Abstract	i
Zusammenfassung	iii
List of Figures	ix
List of Tables	xiii
1 Introduction	1
2 Silicon Spin Qubits Coupled to Single-Molecule Magnets	5
2.1 The Silicon MOS Spin Qubit	5
2.1.1 Loss-DiVincenzo Qubit	6
2.1.2 Electron Confinement	8
2.1.3 Silicon Band Structure and Valleys	9
2.1.4 Coherent Spin Manipulation	10
2.2 Single-Molecule Magnets	13
2.2.1 Fundamental Principles and Properties of TbPc ₂	14
2.3 Coupling SMMs to Silicon Spin Qubits	17
3 Experimental Methods and Instrumentation	25
3.1 Dilution Cryostat (Sionludi L)	25
3.2 Printed Circuit Board for Sample Interfacing	29
3.3 Cryogenic Cabling and Filtering	30
3.4 Integration of Measurement Electronics with the Cryostat	32
3.5 Single-Sideband Mixing	33
3.6 Pulse Shaping and Bias-Tee Compensation	35
3.7 FPGA-Based Control and Data Acquisition (ADwin Pro II)	37
3.8 Novel Printed Circuit Board: The Beast	42
4 The SiMOS Spin Qubit: Definition and Characterisation	49
4.1 Device Layout	49
4.2 The Single-Electron Transistor	51
4.2.1 Coulomb Blockade	52
4.2.2 Charge Sensing	54
4.2.3 Charge Noise Spectroscopy	55
4.3 Plunger Gates: Lever Arm Extraction	57
4.4 Electron Temperature Determination	59

4.5	Spin-Selective Readout	60
4.6	Spin Relaxation Time	62
4.7	Rapid Adiabatic Passage	66
4.8	Rabi Oscillations	69
4.9	Ramsey Experiment	73
4.10	Hahn-Echo	75
4.11	Discussion of Qubit Characterisation Results	77
5	Quantum Magnetometry with Single-Molecule Magnets Coupled to Spin Qubits	81
5.1	Rapid Adiabatic Passage during Spin-Selective Tunnelling (RPSS)	82
5.2	Calibration of the External Vector Magnet	91
5.3	Quantum Magnetometry Experiments on TbPc ₂	92
5.4	Outlook: Single-Molecule Sensing	97
6	Quantum Sensing Using Double Quantum Dot States	101
6.1	Introduction to Double Dot Systems	101
6.2	The Singlet-Triplet Qubit	103
6.3	Pauli Spin Blockade Readout	107
6.4	Singlet-Triplet Oscillations for Sensing	110
6.5	Exchange Oscillations	116
6.6	Towards ESR-Based Sensing with Pauli Spin Blockade	118
7	Simulation-Driven Optimisation of Gate Stack Architectures for Spin Qubit Sensors	121
7.1	Computational Methodology and Simulation Framework	122
7.2	Baseline Device: Simulation of the Original Device	125
7.3	Gate-Layout Optimisation: Single- vs. Split-Plunger Designs	126
7.4	Post-Processing: Buffered Oxide Etch Passivation Removal	128
7.5	Fabricated Device: Initial Tests	131
8	Conclusion & Outlook	135
A	Custom DC Line Fabrication	139
A.1	Fabrication of DC Lines	139
B	Design and Fabrication: The Beast Sample PCB and Connector Board	143
B.1	Layer Stack of The Beast	143
B.2	Accessories for The Beast	144
B.3	Assembly & Soldering Reference	145
C	The SiMOS Spin Qubit	155
C.1	Fabrication Details of Device A and B.	155
C.2	Device A Charge Noise Spectroscopy	155
C.3	Device Working Voltages	158
D	Calculation Tools	159
D.1	Impedance Determination of a CPW Sandwich	159

E Semiconductor Device Modelling with nextnano 161

List of Publications 163

Bibliography 165

Acknowledgments 183

List of Figures

2.1	Schematic illustration of the Bloch sphere and Loss-DiVincenzo model	7
2.2	Electron accumulation in a SiMOS qubit	8
2.3	Band structure and valley splitting of Si	10
2.4	Spin dynamics under resonant driving	11
2.5	Structure of $\text{Mn}_{12}\text{-ac}$ and Fe_8	14
2.6	Molecular structure of $[\text{TbPc}_2]^-$ and crystal packing	15
2.7	Energy levels of TbPc_2	16
2.9	Material stacks of gate-defined spin-qubit platforms	20
2.10	Scalable device architectures for on-chip quantum sensing	21
2.11	Distance dependence of dipolar coupling in the hybrid device.	23
3.1	Experimental volume above the mK stage	27
3.2	Schematic representation of the Sionludi L dilution refrigerator	28
3.3	Standard PCB	29
3.4	Low-pass filtering of the standard PCB	31
3.5	Measurement electronics with cryostat	32
3.6	Single-sideband mixing schematic and spectra	34
3.7	Bias-tee of standard PCB and compensation	36
3.8	Experimental bias-tee compensation calibration	37
3.9	ADwin readout implementation	38
3.10	Two-DAC voltage divider	40
3.11	Oversampling implementation	41
3.12	The Beast: Connector Board.	42
3.13	The Beast: SMD parts and CPWs	43
3.14	The Beast: Chip placement	44
3.15	The Beast: Schematic of the CPW geometry	44
3.16	The Beast: Tank circuit	46
3.17	The Beast: Thermometers	46
4.1	Device layout and heterostructure of the SiMOS qubit	50
4.2	Electrochemical potential diagram of an SET	52
4.3	SET Coulomb oscillations and barrier-gate map	53
4.4	Coulomb diamonds of SET	54
4.5	Charge sensing of an adjacent dot with the SET	55
4.6	Spectral noise density (SD) of the SET of device B	56
4.7	Lever-arm extraction of P1 and P2	58
4.8	Electron temperature measurement	60

4.9	Spin-to-charge conversion and single-shot readout	61
4.10	Single-shot Elzerman readout and blip analysis	63
4.11	Elzerman pulse sequence for T_1 measurements	64
4.12	Implementation and calibration of the Elzerman sequence	65
4.13	Spin relaxation time T_1	66
4.14	Working principle of rapid adiabatic passage	67
4.15	Microwave-power induced heating	68
4.16	Experimental implementation of RAP	69
4.17	Single-tone spectroscopy	71
4.18	Resonance tracking	71
4.19	Gauss fit of the resonance	72
4.20	Coherent control via Rabi oscillations	73
4.21	Rabi oscillations versus detuning	74
4.22	Ramsey pulse-sequence measurement	75
4.23	Hahn-echo measurement	76
5.1	SMM crystal placement	83
5.2	Magnetic-field alignment and lateral displacement	83
5.3	Simultaneous spin excitation and readout	84
5.4	Pulseforms for RPSS and ESR tracking	85
5.5	Spin-resonance detection using RPSS	86
5.6	Cumulative blip count analysis	88
5.7	Analysis of the RPSS visibility	89
5.8	Discrete-frequency excitation during spin-selective tunnelling	90
5.9	Tracking of the resonance frequency with the vector magnet	91
5.10	Silicon spin-qubit magnetometry of a $[\text{TbPc}_2]^-$ ensemble	93
5.11	Background measurement without the $[\text{TbPc}_2]^-$ crystal	94
5.12	μSQUID measurements of the $[\text{TbPc}_2]^-$ crystal	95
5.13	Polar plot of the resonance frequency	95
5.14	Time-resolved relaxation of the $[\text{TbPc}_2]^-$ ensemble	96
5.15	Sensitivity of pulse protocols for Larmor frequency determination	99
6.1	Double-dot device and charge stability map	102
6.2	Stability diagram of the $(2, 0) \leftrightarrow (1, 1)$ charge configuration	103
6.3	ST_0 qubit Bloch sphere and energy diagram	105
6.4	Detuning sweep over the $(1, 1) \leftrightarrow (2, 0)$ transition	109
6.5	Determination of PSB readout point	110
6.6	ST_0 oscillations	111
6.7	ST_0 oscillations as a function of magnetic field	112
6.8	ST_0 oscillations as a function of magnetic-field angle	113
6.9	ST_0 oscillations measured over extended times	114
6.10	Sensing test with ST_0 oscillations	115
6.11	Exchange oscillations	117
7.1	Gate-stack optimisation principle	122

7.2	Gate layout and dimensions of Device A	125
7.3	Electron density for a simulation bias derived from experiment	126
7.4	Reduced single-plunger layout and resulting electron densities	127
7.5	Split double-plunger layout and resulting electron densities	129
7.6	Asymmetric tuning with a split double-plunger	130
7.7	SiMOS device with locally etched passivation window	132
7.8	Custom-made split double-plunger device	133
A.1	DC lines above mK stage	142
B.1	The Beast: layer stack	143
B.2	Low-pass crosstalk reduction	144
B.3	The Beast: storage and ESD protection	145
B.4	The Beast: breakout PCB	146
B.5	The Beast: electrical circuit diagram	147
B.6	The Beast: SMD pin assignment	147
B.7	Connector board: SMD pin assignment	148
B.8	Copper layers of The Beast.	149
B.9	Copper layers of The Beast.	150
B.10	Copper layers of The Beast.	151
B.11	Copper layers of The Beast.	152
B.12	Copper layers of the connector board.	153
C.2	Spectral noise density (SD) of the SET of device A	157

List of Tables

3.1	ADwin process overview.	39
4.1	Overview of devices.	51
4.2	SET lever arm	54
4.3	Charge noise spectroscopy of device B	57
4.4	Qubit parameters of device A.	78
C.1	Gate stack along the z -axis	155
C.2	Charge noise spectroscopy of device A.	156
C.3	Gate voltages at the working points used for devices A and B.	158
E.1	Gate voltages used in simulations	161

1 Introduction

Quantum technologies have moved from theoretical proposals to laboratory devices that can be engineered, operated, and benchmarked [1–3]. This “second quantum revolution” is defined by the controlled preparation, manipulation, and readout of individual quantum systems, using superposition and entanglement as explicit resources [4]. Within this landscape, quantum computing, quantum communication, and quantum sensing form three closely connected pillars: computing targets classically intractable problems [5], communication aims at secure transmission of information [6], and sensing pushes measurement sensitivities towards the limits set by quantum mechanics [7].

Realising these goals in hardware requires physical platforms that support large registers of well-controlled two-level systems (qubits) or higher-dimensional units (qudits) [8, 9]. Fault-tolerant quantum computing is expected to demand thousands to millions of physical qubits with high-fidelity gates and fast, efficient readout [10, 11]. No single implementation presently fulfils all of the DiVincenzo criteria [8] simultaneously: deterministic state preparation, long coherence, universal control, high-fidelity readout, and scalable fabrication. As a consequence, a variety of innovative approaches are being pursued to realise reliable and scalable quantum computing hardware, for example based on trapped ions [12, 13], photonic architectures [14, 15], superconducting [1, 16] and semiconducting electronics [2, 17], and magnetic molecules [18, 19]. Many of the DiVincenzo criteria can be met individually within each of these platforms, but no architecture yet combines them all in a single, mature technology.

In parallel with the drive towards universal quantum computing, quantum sensing has emerged as a powerful and more immediate application of small quantum systems [20]. Here, the key idea is to repurpose fragile quantum states as highly sensitive probes of external fields, forces, or temperatures [21, 22]. The same susceptibility that makes qubits hard to stabilise becomes an advantage when it is used to detect weak signals, and protocols originally developed for quantum information processing can be adapted to enhance sensitivity and spectral selectivity [20, 23]. Among the flagship quantum magnetometers are nitrogen-vacancy (NV) centres in diamond [24, 25], superconducting quantum interference devices (SQUIDs) [26, 27], and atomic vapour cells [28, 29], which together span operating regimes from room temperature to millikelvin and from macroscopic samples to the single-spin limit. Solid-state platforms are especially attractive for such nanoscale sensors, because they combine scalable lithographic control with on-chip integration, a feature that is central to the solid-state architecture explored in this work.

This thesis is positioned at the intersection of quantum computing and quantum sensing. It focuses on a hybrid solid-state architecture that aims to combine the advantages of

two distinct quantum systems while mitigating their respective drawbacks. In particular, industrial silicon metal-oxide-semiconductor (SiMOS) spin qubits are operated as quantum magnetometers and employed as local sensors for single-molecule magnet (SMM) qubits. The central theme of this work is to place these industrial SiMOS devices in a regime where they serve not only as building blocks for scalable quantum processors, but also as versatile, on-chip quantum sensors capable of probing individual molecular spin systems.

Silicon as a platform is the backbone of classical microelectronics and underpins virtually all modern information technology [30, 31]. This mature industrial ecosystem makes silicon an especially attractive host material for scalable qubits, because the same complementary metal-oxide-semiconductor (CMOS) technology that enabled the classical computing revolution can be leveraged for quantum devices [32, 33]. In this work, the spin qubits are realised in SiMOS quantum dots fabricated in the industrial 300 mm wafer line of the interuniversity microelectronics centre (imec) [34]. This foundry process provides precise control over gate-stack geometry and materials, high device yield, and direct compatibility with CMOS periphery. In the long term, such compatibility opens the prospect of integrating cryogenic classical control electronics and quantum devices on a single chip (“cryo-CMOS” [35, 36]), reducing wiring complexity and enabling large-scale quantum processor or sensor arrays.

In a SiMOS spin qubit, a single electron is confined in an electrostatically defined quantum dot at the Si/SiO₂ interface, and its spin- $\frac{1}{2}$ degree of freedom encodes the qubit. The Zeeman splitting of this spin in an applied magnetic field makes it an intrinsically sensitive probe of local magnetic fields: any additional field from nearby magnetic moments shifts the qubit resonance frequency. Silicon spin qubits are particularly well suited for such sensing tasks because of their resilience at high static magnetic fields and their small physical footprint ($\approx 80 \text{ nm} \times 80 \text{ nm}$ for SiMOS devices used here), which allows many qubits to be placed in close proximity to a target system. In principle, industrial fabrication can be used to realise multi-qubit sensor arrays [37, 38], where several SiMOS qubits simultaneously probe different regions of an engineered quantum device or molecular arrangement.

Despite these advantages, silicon spin qubits inherit important limitations from their materials stack. Charge noise arising from fluctuating charges in the oxide and the environment leads to fluctuations in the electrostatic confinement potential and thus in the qubit energy splitting [39, 40]. Interface disorder at the Si/SiO₂ boundary further affects the valley structure and orbital energies of the quantum dot, introducing device-to-device variability and additional decoherence channels [41, 42]. These imperfections currently limit both the coherence of individual SiMOS qubits and the reproducibility of large arrays.

Single-molecule magnets provide a complementary set of strengths to semiconductor spin qubits. They are molecular complexes in which a single magnetic ion or cluster is embedded in an organic ligand environment that can be designed and reproduced with chemical precision. As a result, SMMs are available as chemically identical copies, and their properties can be engineered at the level of the spin Hamiltonian by appropriate choice of metal ion, ligand field, and local symmetry [19, 43]. Many SMMs exhibit a large,

anisotropic spin ground state and a magnetic anisotropy barrier at the single-molecule level, leading to slow relaxation of the magnetisation and hysteresis reminiscent of classical nanomagnets, but in an object of molecular dimensions.

Among these systems, the terbium-phthalocyanine complex TbPc_2 has established itself as a workhorse platform [44]. Its electronic ground state realises a strongly anisotropic total angular momentum, coupled to the nuclear spin of the Tb ion, resulting in a compact multilevel structure that naturally implements a qudit rather than a simple qubit. The nuclear spin subspace, in particular, features remarkably long-lived dynamics, with spin-relaxation times on the order of $T_1 \approx 10\text{-}30$ s and inhomogeneous dephasing times $T_2^* \approx 200$ μs under suitable conditions [18, 45, 46]. Coherent quantum control at the single-molecule level has been demonstrated, including the implementation of multi-level algorithms such as Grover's search and the realisation of an iSWAP gate between selected levels of the TbPc_2 spin manifold [18, 47]. These results establish TbPc_2 as a compact, long-lived quantum memory and a prototypical molecular qudit.

At the same time, precisely those features that give rise to excellent coherence also make direct readout and circuit-level integration of SMMs challenging. The electronic states of TbPc_2 are well shielded from their environment, which suppresses decoherence but also weakens their coupling to external electrodes and fields. Previous experiments have achieved spin-state readout by incorporating TbPc_2 molecules into nanoelectronic transport structures, such as gold break junctions [18, 45, 48, 49] and carbon nanotubes [49, 50], where changes in magnetisation modulate the device conductance. While these approaches provide detailed insight into single-molecule spin physics, they suffer from significant drawbacks: large device-to-device variability, strong measurement back-action due to current transport through the molecule reported in break junctions [45], and poor compatibility with established semiconductor-based quantum device technologies. These limitations motivate the search for alternative readout strategies in which the SMM remains electrically insulating and is probed indirectly by a well-controlled quantum sensor, such as the industrial SiMOS spin qubits considered in this thesis. A key advantage of such non-transport-based readout is that it also enables the use of electrically insulating SMMs as quantum systems, thereby greatly increasing the flexibility in molecular design and choice of candidate species.

The central question of this work is whether industrial silicon spin qubits can act as viable quantum sensors for single-molecule magnets, and which device and protocol developments are required to approach single-molecule sensitivity. It will be shown that the key limitation of such a hybrid architecture is set by the dipolar coupling strength between the SiMOS qubit and the SMM, which in turn is determined by their physical separation within the SiMOS device material stack. As the current generation of industrial devices is not optimised for this sensing operation, the magnetic signal of a single TbPc_2 molecule is effectively amplified by employing an ensemble of crystallised, diluted TbPc_2 molecules on the chip surface.

Quantum magnetometry of this ensemble, performed with an industrial SiMOS qubit, serves as a proof of principle of the hybrid SiMOS-SMM platform. The experiment relies on a compact sensing protocol that combines spin-selective tunnelling and rapid

adiabatic passage of the SiMOS qubit to track its resonance and thereby infer the SMM magnetisation. In this way, canonical signatures of TbPc_2 are observed, including magnetic hysteresis, angular-dependent anisotropy, and thermally activated relaxation of the ensemble magnetisation.

Beyond single-dot operation, this thesis explores how a double quantum dot defined in the same SiMOS technology can be used as an complementary sensor for SMMs. In this regime, Pauli spin blockade provides a distinct readout mechanism at very low external magnetic fields, while singlet-triplet oscillations in the two-electron manifold render the device sensitive to field gradients across the double dot. A nearby TbPc_2 molecule can thus, in principle, be detected not only via its net dipolar field, but also through the spatial variation of this field on the scale of the interdot separation, with the singlet-triplet oscillations being largely insensitive to homogeneous changes of the absolute magnetic field.

Finally, the prospects for improving SMM-qubit coupling are addressed by means of Schrödinger-Poisson simulations of the industrial gate stack. These calculations demonstrate that the qubit-molecule separation can be substantially reduced by targeted gate-stack optimisation, without sacrificing few-electron operation of the SiMOS quantum dot. Guided by these simulations, a first generation of sensing-optimised SiMOS devices has been fabricated in the 300 mm industrial cleanroom of imec. Together, these experimental and simulation results outline a realistic pathway from ensemble TbPc_2 magnetometry towards single-molecule sensitivity in industrial SiMOS-SMM hybrid devices.

The thesis is organised as follows. Chapter 2 introduces the underlying concepts of the SiMOS qubit and SMMs, with particular emphasis on TbPc_2 , and develops the theoretical description of the dipolar coupling between the two. Chapter 3 summarises the experimental infrastructure required to operate the industrial SiMOS devices, including the dilution cryostat, wiring, printed circuit boards, microwave and pulse-generation chain. Chapter 4 presents the definition and characterisation of the industrial SiMOS spin qubit, covering charge sensing, noise spectroscopy, and coherent single-qubit control. Chapter 5 reports quantum magnetometry of a TbPc_2 ensemble using the SiMOS qubit as a nanoscale magnetic-field sensor. Chapter 6 explores double-quantum-dot devices for sensing, including singlet-triplet based schemes and Pauli spin blockade readout. Chapter 7 presents simulation-driven optimisations of gate-stack architectures aimed at enhancing the SMM-qubit coupling. The final chapter summarises the main results, discusses their implications for hybrid SMM-SiMOS platforms, and outlines future directions towards single-molecule sensing.

2 Silicon Spin Qubits Coupled to Single-Molecule Magnets

This chapter presents industrial SiMOS spin qubits as quantum sensors for single-molecule-magnet (SMM) qubits. It first summarises the SiMOS qubit, including its confinement and coherent control, before introducing key SMM concepts with emphasis on the TbPc₂ system used throughout this work. The coupling between the SiMOS qubit and TbPc₂ is quantified by translating device geometry into dipolar interaction strengths at relevant separations. The accessible qubit-SMM distances are then compared across different semiconductor spin qubit platforms.

2.1 The Silicon MOS Spin Qubit

Silicon is the foundation of modern information technology, with the semiconductor sector generating \$630.5 billion of global sales in 2024 [51]. In close analogy to conventional metal-oxide-semiconductor (MOS) transistors, silicon MOS spin qubits rely on the same production concepts: gate electrodes confine and control single electrons at the Si/SiO₂ interface, and the device stack is compatible with standard industrial flows. While the scaling of classical transistors approaches its physical limits, the quantum-computing approach in silicon is only just beginning. The unique advantage of industrial mass fabrication, including mature processes, reproducibility, and the prospect of large-scale integration, is therefore a strong motivation to investigate silicon qubits.

At the same time, several drawbacks are evident. Atomic-scale disorder at the Si/SiO₂ interface perturbs the electronic states and shifts the gate-voltage thresholds and operating points [52, 53], low-frequency charge noise leads to decoherence [39], and there is a wide device-to-device variation in valley splittings (Sec. 2.1.3) [54] that influences qubit stability and gate fidelity. These issues do not negate the promise of SiMOS, but they do limit present-day usability and motivate complementary hybrid strategies.

To dive into the topic, we begin by defining the qubit and briefly revisiting the Loss-DiVincenzo proposal for spin-based quantum computation in semiconductors [55].

Quantum Bit

The basic unit of quantum information is the quantum bit (qubit), a controllable two-level system. Unlike a classical bit, confined to the definite values 0 or 1, a qubit can be in coherent superposition of both basis states. A general pure state can be written as

$$|\Psi\rangle = a|0\rangle + b|1\rangle, \quad a, b \in \mathbb{C}, \quad |a|^2 + |b|^2 = 1, \quad (2.1)$$

with complex probability amplitudes a and b which are normalised. Superposition provides the foundation for quantum interference which is exploited by algorithms to steer probability amplitudes towards correct answers [56]. By choreographing unitary gate sequences that set relative phases, constructive interference amplifies the amplitudes of desired outcomes while destructive interference suppresses competing paths.

A second indispensable resource is entanglement: multi-qubit states that cannot be factored into products of single-qubit states exhibit correlations with no classical analogue. In such states, the subsystems are nonlocally correlated, and measurements on one immediately constrain the statistical predictions for the others irrespective of spatial separation. Superposition and entanglement together underwrite the quantum advantage.

It is convenient to visualise a single qubit on the Bloch sphere which is shown in the left panel of Fig. 2.1. Any pure state corresponds to a point on the unit sphere and can be parametrised by spherical angles (θ, ϕ) as

$$|\Psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle, \quad (2.2)$$

with the polar angle θ and the azimuthal angle ϕ . The states $|0\rangle$ and $|1\rangle$ are the north and south poles, respectively. Points on the equator represent equal-weight superpositions such as $(|0\rangle + |1\rangle)/\sqrt{2}$, while ϕ sets their relative phase.

The qubit state can be controlled by applying an oscillating magnetic field along the x-direction, which drives rotations of the qubit state around this axis. The amplitude, phase, and duration of the field control the rotation rate, axis, and angle of the qubit state. This control mechanism allows for precise manipulation of the qubit, enabling the implementation of various quantum operations. The specific pulse schemes used for qubit manipulation in this thesis are discussed in Sec. 2.1.4.

2.1.1 Loss-DiVincenzo Qubit

Daniel Loss and David P. DiVincenzo proposed encoding a qubit in the spin of a single, confined electron [55]. Motivated by rapid advances in nanofabrication, they identified gate-defined quantum dots as a practical way to realise this confinement. The two logical states are provided by the electron's spin- $\frac{1}{2}$ doublet, split by an external magnetic field \mathbf{B} via the Zeeman interaction

$$H_Z = \frac{g\mu_B}{\hbar} \mathbf{B} \cdot \mathbf{S}, \quad \mathbf{S} = \frac{\hbar}{2} \boldsymbol{\sigma}, \quad \boldsymbol{\sigma} = (\sigma_x, \sigma_y, \sigma_z), \quad (2.3)$$

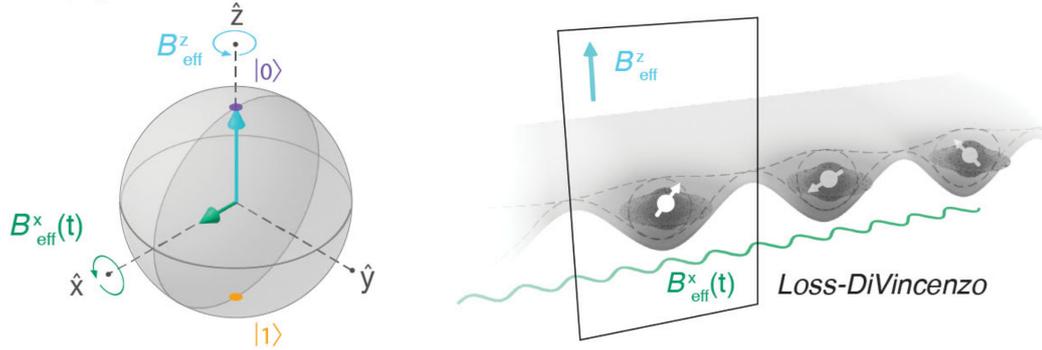


Figure 2.1: **Left** Geometric depiction of a qubit state on the Bloch sphere. The qubit state $|\Psi\rangle$ (black arrow) is mapped to a point on the unit sphere, defined by θ and ϕ . The quantisation axis is aligned with the static B_z^{eff} field (blue arrow), while the oscillating field B_x^{eff} (green arrow) induces coherent rotations of the state vector, enabling full qubit control. **Right** Schematic representation of spin qubits implemented in a linear array of quantum dots. Each dot contains a single confined electron, with its spin state controlled by an oscillating transverse magnetic field $B_x^{\text{eff}}(t)$, while a static magnetic field B_z^{eff} defines the quantisation axis. Taken from [54].

where g is the electron Landé g -factor and μ_B the Bohr magneton. Choosing $\mathbf{B} = B_z \hat{z}$ fixes the computational basis $\{|\uparrow\rangle, |\downarrow\rangle\}$ as the eigenstates of σ_z , to which one maps $|0\rangle$ and $|1\rangle$. The corresponding Zeeman energies are

$$E_Z = \pm \frac{1}{2} g \mu_B B_z, \quad (2.4)$$

so that the level splitting is $\Delta E_Z = g \mu_B B_z$ and the Larmor frequency is $\omega_0 = \Delta E_Z/\hbar$. Thus the static field sets both the qubit energy scale and its natural precession rate.

While single-qubit control is achieved by applying a weak oscillating magnetic field transverse to \hat{z} and driving near ω_0 , scalable architectures also require controllable two-qubit interactions. In the proposal of Loss and DiVincenzo [55], these are implemented electrically via tunable exchange coupling between neighbouring quantum dots. As sketched in the right panel of Fig. 2.1, an array of quantum dots is formed, and barrier gates are used to control the wavefunction overlap between adjacent sites. By lowering a barrier for a controlled time interval, the Heisenberg exchange interaction is turned on and is described by [54]

$$H_H(t) = \frac{1}{\hbar^2} \sum_{\langle i,j \rangle} J_{ij}(t) \mathbf{S}_i \mathbf{S}_j = \frac{1}{4} \sum_{\langle i,j \rangle} J_{ij}(t) \sigma_i \sigma_j, \quad (2.5)$$

with $J_{ij}(t)$ the time-dependent exchange coupling between dots i and j . Properly timed exchange pulses implement $\sqrt{\text{SWAP}}$ operations. Combined with single-qubit rotations they yield standard entangling gates such as CNOT and controlled-phase, enabling multi-qubit computation. For more information on qubit gates refer to Ref. [54].

The physical origin of exchange is tied to the Pauli exclusion principle [57] and the antisymmetry of the total two-electron wavefunction. In the singlet configuration (antisymmetric in spin), a symmetric spatial state is allowed, which is bonding and lowers the energy

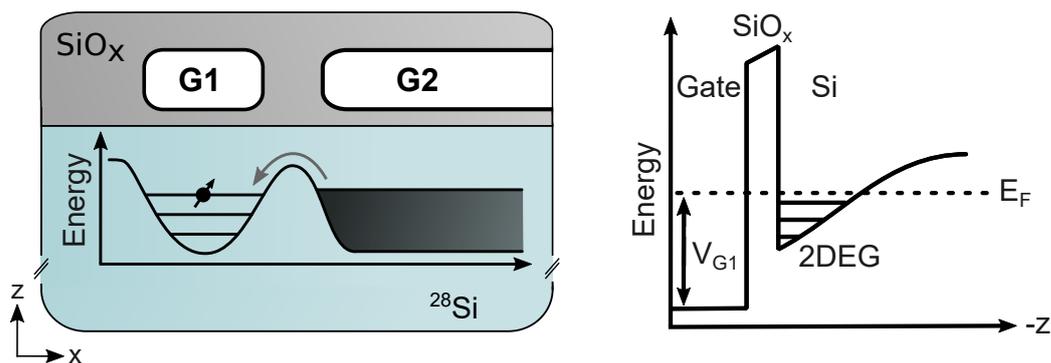


Figure 2.2: Electron accumulation in a SiMOS qubit. **Left** Cross-section sketch of the device, with ²⁸Si shown in blue, SiO_x in grey, and polycrystalline Si gates in white. A positive voltage on G1 and G2 lowers the conduction-band edge shown in a diagram (energy vs. x-direction). The reservoir has a continuum of states (dark shading), while the dot exhibits quantised levels. An electron can be loaded from the reservoir into the dot (grey arrow). Barrier gates are omitted in this simplified schematic. **Right** Vertical band diagram at the position of G1. A positive V_{G1} bends the conduction band below the Fermi level E_F in silicon, forming a triangular well and accumulating a two-dimensional electron gas (2DEG).

through delocalisation. In contrast, triplet states (symmetric in spin) require antisymmetric spatial states with a node between electrons, which are antibonding and therefore higher in energy. The resulting singlet-triplet energy splitting constitutes the exchange interaction and can be tuned electrically by gate voltages.

An experimental realisation of this scheme is presented in Ch. 6, where a readout based on Pauli spin blockade in a double dot [17] is implemented and coherent SWAP oscillations are demonstrated.

2.1.2 Electron Confinement

Silicon MOS spin qubits are built on a simple heterostructure: a gate electrode sits on a thin SiO_x layer that insulates it from the crystalline silicon substrate. A simplified schematic of this structure is shown in the left panel of Fig. 2.2. In practice, the “metal” gate is often realised as highly doped polycrystalline silicon rather than a true metal [58, 59].

A positive voltage on a surface gate bends the silicon conduction band downwards near the Si/SiO_x interface. Once the band edge falls below the Fermi level, electrons accumulate in a narrow triangular well beneath the interface as shown in the right panel of Fig. 2.2. Motion along z becomes quantised, while motion in the x - y plane remains free, thereby forming a two-dimensional electron gas (2DEG). The oxide ensures electrostatic control without any DC current from gate to 2DEG.

Lateral confinement is created electrostatically with patterned gates that sculpt the 2DEG into tunnel-coupled puddles. Gates biased less positively raise potential barriers, providing confinement along x and y . A plunger gate sets the dot potential minimum, while barrier gates tune the tunnel rates to a reservoir and to neighbouring dots. The specific gate geometry of the devices used in this work is introduced in Sec. 4.1.

Electron reservoirs in the 2DEG are tied to a well-defined ground reference by degenerately doped (n^{++}) implantation regions that form ohmic contacts to the metal interconnects and bond pads. Positively biased lead/accumulation gates then create a continuous low-resistance channel from the dot region to these implants, allowing controlled loading and unloading of single electrons.

Although the wafer starts from natural silicon, an epitaxial ~ 100 nm layer of isotopically enriched ^{28}Si (residual $^{29}\text{Si} \sim 400$ ppm) is grown before thermal oxidation to form the Si/SiO_x interface. Natural silicon contains 4.7% ^{29}Si nuclei ($I = \frac{1}{2}$), which generate slowly fluctuating hyperfine (Overhauser) fields and accelerated dephasing. Reducing the ^{29}Si fraction by over two orders of magnitude stabilises the Larmor frequency and extends coherence times [60, 61].

Silicon is an attractive host for gate-defined dots owing to abundant materials and mature CMOS-compatible processing [62]. The devices used in this thesis are fabricated in imec's state-of-the-art 300 mm pilot line [34], which brings modern process control and metrology to quantum-dot fabrication. It allows rapid iteration, statistical screening across many dies, and reproducible device recipes which are key advantages for scaling beyond single proof-of-concept qubits.

2.1.3 Silicon Band Structure and Valleys

To understand how electrostatic confinement at the Si/SiO_x interface shapes the spectrum of a SiMOS qubit, we first recall the key features of silicon's band structure. Crystalline silicon adopts the diamond-cubic lattice: a face-centred cubic lattice with a two-atom basis, in which each atom is tetrahedrally bonded to four neighbours [63]. As shown in the left panel of Fig. 2.3, silicon has an indirect band gap of $E_g \approx 1.17$ eV in the $T \rightarrow 0$ K limit [64]. The valence-band maximum lies at Γ , while the conduction-band minima are displaced from Γ along the $\langle 100 \rangle$ axes towards X . There are six such minima, or valleys, which are related by the cubic symmetry and are therefore degenerate in bulk, defining the valley degree of freedom in silicon [33].

In conventional electronic transport, the multiplicity of valleys typically does not change qualitative drift-diffusion behaviour [33]. In quantum-confined nanostructures, however, valley physics becomes central: intervalley coupling and valley-orbit interactions modify selection rules and energy spectra. A notable consequence is the erosion of Pauli spin blockade in double quantum dots when valley states are weakly split or mixed, which complicates spin initialisation, manipulation, and readout [65].

Vertical confinement together with the interfacial potential breaks cubic symmetry and lifts the sixfold degeneracy. In a (001) quantum well the four in-plane valleys ($\pm x, \pm y$) are raised relative to the two out-of-plane valleys ($\pm z$), which become the lowest in energy [33, 66]. The residual degeneracy within the $\pm z$ pair is further lifted by quantum confinement and by the vertical electric field, producing the valley splitting E_v . The magnitude of E_v depends sensitively on electric field and atomic-scale interface details such as interface abruptness, disorder, and local strain [67–69].

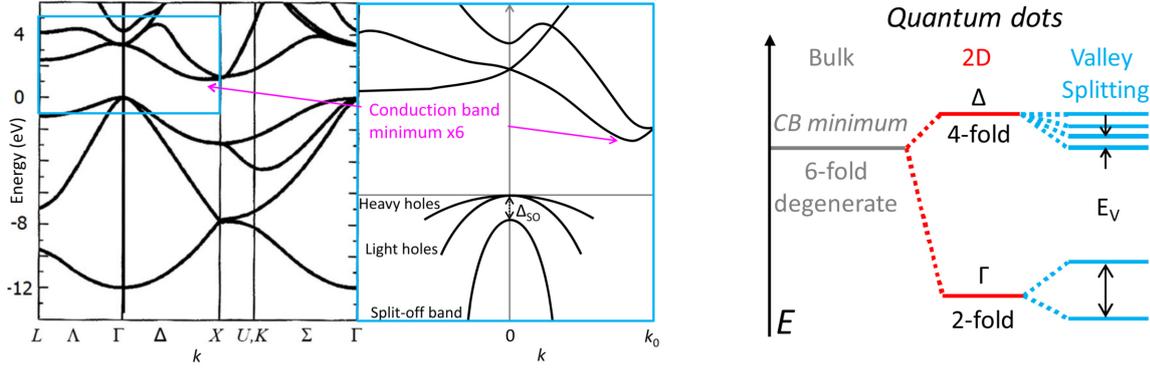


Figure 2.3: **Left** Electronic band structure of bulk silicon highlighting the indirect band gap ($E_g \approx 1.17$ eV at $T \rightarrow 0$ K) [64] and the six equivalent conduction-band minima along the Δ lines towards the X points. **Right** Schematic lifting of valley degeneracy in a confined silicon layer. Biaxial tensile strain raise the in-plane ($\pm x$, $\pm y$) valleys and lowers the out-of-plane ($\pm z$) valleys in energy (red). Confinement in the quantum well and electric fields further split the remaining pair by E_v (blue). Taken from [33].

Although it is generally desirable for E_v to exceed the Zeeman energy E_Z , this condition is not always satisfied. If $E_v < E_Z$, the excited spin state of the lower valley can approach the ground spin state of the upper valley, enabling spin-valley mixing via spin-orbit coupling or interface disorder. Practically, such mixing accelerates spin relaxation, with a pronounced “hot-spot” over applied magnetic field when the spin and valley levels become (near-) degenerate [41, 70]. Operation in this regime remains possible but introduces additional complexity that can degrade performance unless carefully managed.

Determining E_v was not an objective of this work. For SiMOS devices, reported valley splittings are typically ≥ 0.1 meV, i.e. well above the Zeeman energies used in this work [41, 71, 72]. The ESR frequencies employed were ≤ 18.5 GHz, corresponding to $E_Z = hf \approx 77$ μ eV, which lies below the cited SiMOS E_v range. In conclusion, valley effects did not play a limiting role in the measurements presented.

2.1.4 Coherent Spin Manipulation

Whilst the experimental details of state preparation and readout are presented in Ch. 4, here the basic control of a single electron spin is outlined. As discussed earlier, a static external magnetic field $\mathbf{B}_{\text{ext}} = B_z \hat{e}_z$ defines the quantisation axis and sets the Larmor precession frequency $\omega_0 = g\mu_B B_z / \hbar$.

Transitions between spin states $|\downarrow\rangle$ and $|\uparrow\rangle$ are driven by applying a magnetic-field perturbation perpendicular to the qubit’s quantisation axis, given by $\mathbf{B}_{\text{MW}}(t) = B_{\text{MW}} \sin(\omega_{\text{MW}} t + \delta) \hat{e}_x$. Here, B_{MW} denotes the field amplitude, ω_{MW} the drive frequency, and δ the initial phase. Since the drive frequencies in our experiments lie in the microwave range, the subscript “MW” is used.

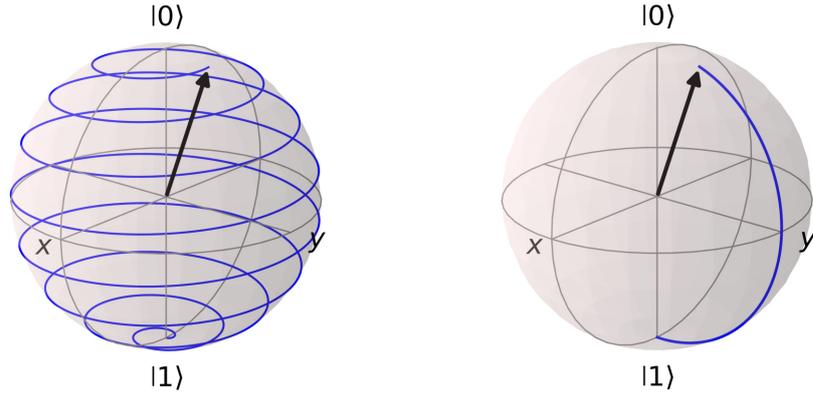


Figure 2.4: Schematic of spin dynamics under resonant transverse driving. The static magnetic field B_z defines the quantisation axis (z) and sets the Larmor frequency ω_0 . **Left** Lab-frame trajectory. The spin precesses about the effective field B_{eff} , producing a spiralling path on the Bloch sphere. **Right** Co-rotating frame (RWA). At resonance, B_{eff} points along x , yielding a uniform circular rotation about the x -axis at the Rabi frequency Ω_R .

Following the Zeeman interaction Hamiltonian in Eq. (2.3), the total Hamiltonian that includes both field contributions is

$$\begin{aligned}
 H(t) &= H_0 + H_1(t) \\
 &= \frac{g\mu_B}{\hbar} \mathbf{B} \cdot \mathbf{S} \\
 &= \frac{g\mu_B}{2} B_z \sigma_z + \frac{g\mu_B}{2} B_{\text{MW}} \sin(\omega_{\text{MW}} t + \delta) \sigma_x \\
 &= \frac{g\mu_B}{2} \begin{pmatrix} B_z & B_{\text{MW}} \sin(\omega_{\text{MW}} t + \delta) \\ B_{\text{MW}} \sin(\omega_{\text{MW}} t + \delta) & -B_z \end{pmatrix}.
 \end{aligned}$$

In this matrix form, the off-diagonal elements, proportional to the oscillating field B_{MW} , generate coherent coupling between $|\downarrow\rangle$ and $|\uparrow\rangle$. When the drive is tuned to the Larmor frequency, $\omega_{\text{MW}} = \omega_0$, the transverse excitation produces an effective magnetic field in the xy -plane that exerts a torque on the spin, rotating the Bloch vector away from the z -axis. In the laboratory frame this appears as precession about the composite effective field B_{eff} formed by B_z and $B_{\text{MW}}(t)$, yielding a spiral trajectory on the Bloch sphere, as sketched in the left panel of Fig. 2.4.

A more intuitive picture of Bloch-vector rotations is obtained in the co-rotating frame, which turns at the drive frequency ω_{MW} . The laboratory-frame state $|\Psi(t)\rangle$ is mapped to the rotating-frame state $|\tilde{\Psi}(t)\rangle$ by the time-dependent unitary operator

$$|\tilde{\Psi}(t)\rangle = U^\dagger(t) |\Psi(t)\rangle, \quad U(t) = \exp\left[-\frac{i}{2} \omega_{\text{MW}} t \sigma_z\right] = \begin{pmatrix} e^{-i\omega_{\text{MW}} t/2} & 0 \\ 0 & e^{i\omega_{\text{MW}} t/2} \end{pmatrix}.$$

This change of frame removes the rapid precession at ω_{MW} and leaves only the slow dynamics due to detuning and the transverse drive.

Inserting the rotating-frame transformation into the laboratory-frame Schrödinger equation $i\hbar \partial_t |\psi\rangle = H(t) |\psi\rangle$ gives [73]:

$$\begin{aligned}
i\hbar \partial_t |\tilde{\psi}(t)\rangle &= (U^\dagger(t)H(t)U(t) - i\hbar U^\dagger(t)\dot{U}(t)) |\tilde{\psi}(t)\rangle \\
&= \left[\frac{\hbar\Delta\omega}{2} \sigma_z + \frac{\hbar\Omega_R}{2} \cos(\delta) \sigma_x + \frac{\hbar\Omega_R}{2} \sin(\delta) \sigma_y + \frac{\hbar\Omega_R}{2} \begin{pmatrix} 0 & e^{-i2\omega_{\text{MW}}t-\delta} \\ e^{i2\omega_{\text{MW}}t+\delta} & 0 \end{pmatrix} \right] |\tilde{\psi}(t)\rangle \\
&\approx \left[\frac{\hbar\Delta\omega}{2} \sigma_z + \frac{\hbar\Omega_R}{2} \cos(\delta) \sigma_x + \frac{\hbar\Omega_R}{2} \sin(\delta) \sigma_y \right] |\tilde{\psi}(t)\rangle \\
&= \frac{\hbar}{2} \begin{pmatrix} \Delta\omega & \Omega_R e^{-i\delta} \\ \Omega_R e^{i\delta} & -\Delta\omega \end{pmatrix} |\tilde{\psi}(t)\rangle \\
&= H_{\text{RWA}} |\tilde{\psi}(t)\rangle
\end{aligned}$$

with the detuning $\Delta\omega = \omega_0 - \omega_{\text{MW}}$ and the Rabi frequency $\Omega_R = g\mu_B B_{\text{MW}}/\hbar$. The final line identifies the effective Hamiltonian H_{RWA} of the rotating-wave approximation (RWA). In making this approximation, the rapidly oscillating terms $\propto e^{\pm 2i\omega_{\text{MW}}t}$ are neglected because their frequency $2\omega_{\text{MW}}$ is much larger than the relevant rates $|\Delta\omega|$ and Ω_R .

The rotation axis in the xy plane can be tuned by the drive phase δ when the field is applied along \hat{x} , effectively rotating the control axis on the equator. In the co-rotating frame (right panel of Fig. 2.4) the motion reduces to a uniform circular trajectory about this axis. The rotation angle then depends on the field amplitude B_{MW} , the drive frequency ω_{MW} , and the chosen phase δ of the driving field.

The solution to the time-dependent Schrödinger equation in the co-rotating frame is $|\tilde{\psi}(t)\rangle = \exp(-\frac{i}{\hbar}H_{\text{RWA}}t) |\tilde{\psi}(0)\rangle$. The probability to find the spin in $|\uparrow\rangle$ after initialisation in $|\downarrow\rangle$ is [74]

$$P_{\uparrow}(t) = |\langle \uparrow | \tilde{\psi}(t) \rangle|^2 = \frac{\Omega_R^2}{\Omega_{\text{eff}}^2} \sin^2\left(\frac{\Omega_{\text{eff}}}{2} t\right)$$

after a manipulation of duration t , with the effective rotation frequency $\Omega_{\text{eff}} = \sqrt{\Delta\omega^2 + \Omega_R^2}$. At resonance ($\Delta\omega = 0$) this yields perfect spin inversion while for finite detuning the oscillation contrast is reduced by the factor $(\Omega_R/\Omega_{\text{eff}})^2$ and the oscillation frequency increases to Ω_{eff} . Experimentally, these oscillations, called Rabi oscillations, are presented in Sec. 4.8 alongside more advanced pulse schemes.

In the SiMOS devices used here, the oscillating magnetic field is generated by feeding a high-frequency current into a short-terminated coplanar waveguide (CPW) placed near the quantum dot. The short acts as a current-carrying conductor and produces a near-field magnetic field around the termination which is the standard electron spin resonance (ESR) implementation. According to the Biot-Savart law, the field magnitude from the conductor scales as $\propto 1/r$ with the distance r between the termination and the qubit, so an essential design goal is to minimise this distance to maximise B_{MW} and the achievable Rabi rate. In

the measurements presented, the on-chip antenna is positioned directly above the qubit, yielding an oscillating field predominantly in the plane of the chip and transverse to the static quantisation field.

The discussion so far has established the SiMOS spin qubit as a coherent, controllable two-level system. Beyond enabling conventional quantum-information protocols, this also turns the qubit into a sensitive, nanoscale probe of its magnetic environment. The next chapter introduces the small magnetic structures that are probed in this way.

2.2 Single-Molecule Magnets

In conventional magnets, properties such as remanence and hysteresis arise from the collective behaviour of many spins. The macroscopic magnetisation reflects long-range order established by exchange interactions between individual moments. By contrast, single-molecule magnets (SMMs) are isolated molecules that carry large moments and exhibit two stable magnetisation states despite the absence of intermolecular long-range order. This bistability is protected by an energy barrier that suppresses relaxation of the magnetisation. The discovery of magnetic bistability in Mn_{12} -acetate, $\text{Mn}_{12}\text{O}_{12}(\text{CH}_3\text{COO})_{16}(\text{H}_2\text{O})_4$ (Mn_{12} -ac), inaugurated the field in 1991 [75–78]. The molecular structure of Mn_{12} -ac is shown in Fig. 2.5, alongside another heavily investigated system, the Fe_8 cluster $[\text{Fe}_8\text{O}_2(\text{OH})_{12}(\text{tacn})_6]^{8+}$.

The technical interest in SMMs stems from their potential use in quantum information processing as long-lived quantum memories and qubits in which the logical state is encoded in the magnetisation of a single molecule [18]. A principal advantage of SMM-based qubits is that chemical synthesis delivers very large numbers of indistinguishable units. This striking reproducibility, combined with the chemical tunability of the spin Hamiltonian and environment [19, 79, 80], stands in sharp contrast to many present-day architectures (e.g. superconducting and semiconductor qubits), where scalability is hampered by device-to-device variability and the difficulty of reproducibly fabricating large arrays with uniform parameters. In short, producing many faithful copies, which is an essential prerequisite for a universal processor, is intrinsically more natural for molecules.

A further advantage is the inherent multilevel spin structure of many SMMs [19]. Consequently, usable quantum levels need not be restricted to a single-qubit subspace. TbPc_2 is the archetypal bis-phthalocyaninato terbium SMM used in this work. Its four hyperfine-resolved levels provide a naturally addressable four-level system (a qudit). This higher dimensionality can enable more efficient information encoding and manipulation [81]. These features motivate hybrid device concepts that allow one to read out and coherently manipulate SMMs without sacrificing their favourable properties. Previous approaches have achieved readout by incorporating TbPc_2 molecules into nanoelectronic transport structures, such as carbon nanotubes [49, 50] and gold break junctions [18, 46, 48, 82], where the magnetisation dynamics modulate the conductance. While such measurements have been highly informative, the through-current can introduce measurement back-action. In break-junction devices, signatures of current-induced perturbations were reported [45].

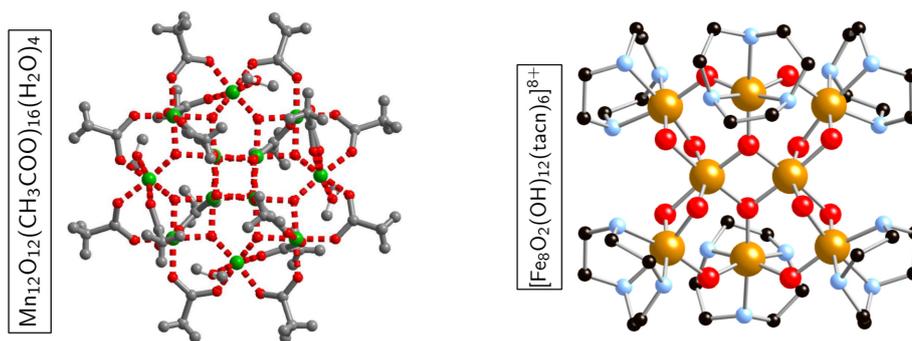


Figure 2.5: Molecular structures of $\text{Mn}_{12}\text{-ac}$ (**left**) and the Fe_8 cluster (**right**). Colour code: Mn green, Fe yellow, O red, N light blue, C grey/black; H atoms omitted for clarity. Adapted from [83].

To suppress back-action at the source, this thesis investigates a hybrid architecture in which a thin SiO_x layer separates a silicon spin qubit from the SMM, providing near-field coupling for control and readout whilst eliminating current transport through the molecule.

2.2.1 Fundamental Principles and Properties of TbPc_2

With the discovery that single lanthanide-phthalocyanine complexes of the form LnPc_2 ($\text{Ln} = \text{Tb}^{3+}, \text{Dy}^{3+}$) exhibit an energy barrier against magnetisation relaxation, a new era in molecular magnetism began in 2003 [44]. Owing to the strong spin-orbit coupling of lanthanide ions, the barrier in TbPc_2 is an order of magnitude larger than in $\text{Mn}_{12}\text{-ac}$ [19].

TbPc_2 hosts a Tb^{3+} ion with electron configuration $[\text{Xe}] 4f^8$. In Russell-Saunders coupling this yields $L = 3, S = 3$ and hence $J = L + S = 6$. The spin-orbit gap to the $J = 5$ multiplet is large (of order 2900 K in $k_B T$ units), so at cryogenic temperatures the dynamics are confined to the $J = 6$ manifold with $2J+1 = 13$ degenerate sublevels.

The phthalocyanine (Pc) ligands (approximately D_{4d} symmetry) define a molecular easy axis perpendicular to the Pc planes and split the $(2J+1)$ -fold degeneracy of the $J = 6$ multiplet. The ligand-field Hamiltonian is conveniently written in Stevens form [84]

$$H_{\text{lf}} = \sum_{n=1}^3 B_{2n}^0 O_{2n}^0 + B_4^4 O_4^4 + B_6^4 O_6^4 \quad (2.6)$$

where O_k^q are the Stevens operators and B_k^q the (experimentally determined) Stevens parameters [84, 85]. Strong axial terms of the ligand field select a ground-state doublet $|m_J = \pm 6\rangle$, separated from the first excited $|m_J = \pm 5\rangle$ doublet by approximately 600 K (~ 0.05 eV), rendering the molecule effectively Ising-like at low temperature.

Natural terbium is ^{159}Tb with nuclear spin $I = 3/2$. The hyperfine interaction between the electronic and nuclear moments $A_{\text{hyp}} \mathbf{I} \mathbf{J}$ splits each component of the ground electronic doublet into $2I+1 = 4$ sublevels with $m_I = \{-3/2, -1/2, +1/2, +3/2\}$. The nuclear electric

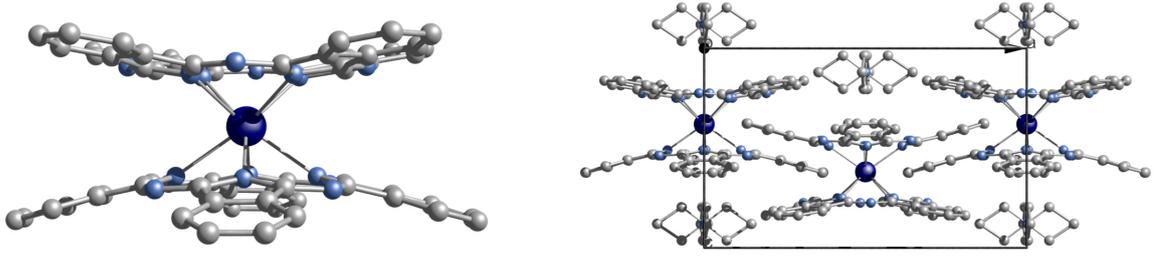


Figure 2.6: **Left** Side view of the molecular structure of $[\text{TbPc}_2]^-$. The central Tb^{3+} ion (dark blue) is coordinated by eight N atoms (light blue). C atoms are shown in grey; H atoms are omitted for clarity. **Right** Crystal packing of diluted $[\text{TbPc}_2]^-$ in the unit cell. Two $[\text{TbPc}_2]^-/[\text{YPC}_2]^-$ units are arranged with additional tetraethylammonium (TEA) ions to maintain charge neutrality.

quadrupole interaction $\mathbf{I}\hat{P}_{\text{quad}}\mathbf{I}$ produces unequal spacings between these levels. Including Zeeman coupling, a compact spin Hamiltonian for TbPc_2 is [18, 86]

$$H_{\text{TbPc}_2} = H_{\text{lf}} + \mu_B \mu_0 \mathbf{H} \hat{g} \mathbf{J} + A_{\text{hyp}} \mathbf{I} \mathbf{J} + \mathbf{I} \hat{P}_{\text{quad}} \mathbf{I} \quad (2.7)$$

where \mathbf{H} is the applied magnetic field and \hat{g} the g -tensor. The experimentally determined values of B_k^q , A_{hyp} , and \hat{P}_{quad} are given in Ref. [84]. The resulting electronic and hyperfine levels are illustrated in Fig. 2.7.

The energy splittings of the hyperfine-split levels are uneven and the transition frequencies are approximately 2.45 GHz, 3.13 GHz, 3.81 GHz between the individual $m_I = \{-3/2, -1/2, +1/2, +3/2\}$ state [87]. This is very convenient as the frequency range lies in the experimentally accessible microwave range and enables to chose which transition should be driven.

Magnetisation Reversal

Reversal of the $m_J = \pm 6$ magnetic moment can be accomplished by sweeping a magnetic field parallel to the molecular easy axis of TbPc_2 . At small fields the magnetisation switches in discrete steps at the hyperfine-resolved avoided crossings (labelled by m_I). These tunnel transitions through the anisotropy barrier are explained by quantum tunnelling of magnetisation (QTM): off-diagonal ligand-field terms, $B_4^4 O_4^4$ and $B_6^4 O_6^4$, admix $|+6\rangle$ and $|-6\rangle$ and open avoided crossings (see Fig. 2.7). When the longitudinal field is swept across an anticrossing with tunnel splitting $\Delta_{m,m'}$ (here m and m' denote the electronic quantum numbers m_J and m'_J), the probability to tunnel from $|J, m\rangle$ to $|J, m'\rangle$ is given by the Landau-Zener formula [88, 89]

$$P_{m,m'} = 1 - \exp \left[- \frac{\pi \Delta_{m,m'}^2}{2\hbar g \mu_B |m - m'| \mu_0 \frac{dH_{\parallel}}{dt}} \right], \quad (2.8)$$

where $\mu_0 dH_{\parallel}/dt$ is the sweep rate of the longitudinal field. Thus $P_{m,m'}$ increases with larger $\Delta_{m,m'}$ and slower sweeps. In TbPc_2 the four hyperfine-split anticrossings give

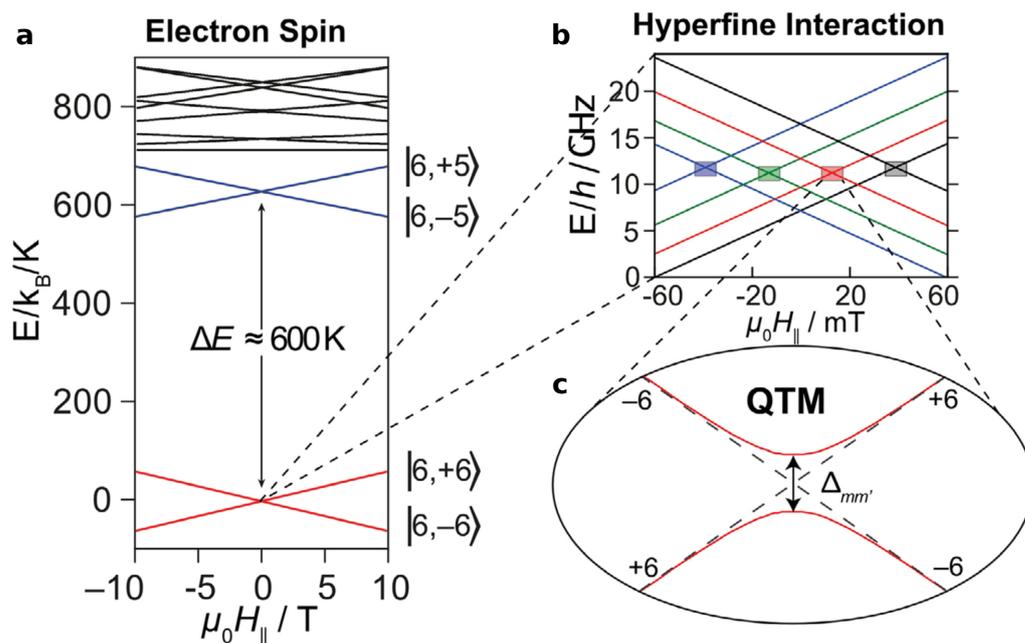


Figure 2.7: Energy spectrum of Eq. (2.7) for TbPc₂. **a** Electronic levels labelled by $|J, m_J\rangle$. The ≈ 600 K separation between $m_J = \pm 6$ and $m_J = \pm 5$ is annotated. **b** Zoom into the hyperfine-split ground doublet ($m_J = \pm 6$) showing $m_I = -3/2$ (black), $-1/2$ (red), $+1/2$ (green), $+3/2$ (blue). **c** Magnified view of the $m_I = -1/2$ avoided crossing illustrates quantum tunnelling of magnetisation (QTM). Taken from [19].

four characteristic magnetisation steps in the hysteresis (see μ SQUID data in Fig. 5.12), a hallmark signature of this SMM.

Additionally to QTM, phonon-assisted transitions can reverse the magnetic moment of TbPc₂ [90, 91]. This temperature-dependent relaxation is experimentally characterised in Sec. 5.3. In these inelastic processes, lattice phonons supply (or remove) the required energy and momentum—hence the term spin-lattice relaxation. Three types of processes can be distinguished. First, at low temperatures a *direct* one-phonon process will most probably reverse the moment by emission of a single phonon to the thermal bath. This process becomes more likely at higher magnetic fields and the relaxation time τ scales with $\tau \propto H^{-3}$ [92]. At higher temperatures, two-phonon processes become more likely: the molecular state is excited to an intermediate state $|e\rangle$ by absorbing a phonon of energy $\hbar\omega_1$ and relaxes by emitting a phonon of energy $\hbar\omega_2$ with $\omega_2 > \omega_1$. If $|e\rangle$ is a virtual state, the relaxation is a *Raman* process with temperature dependence $\tau \propto T^{-7}$. If $|e\rangle$ is a real crystal-field level, the relaxation is an *Orbach* process with activated behaviour $\tau \propto \exp(\hbar|\omega_2 - \omega_1|/(k_B T))$. Consequently, the exponential Orbach channel dominates at elevated temperatures.

Break-Junction Experiment

The most prominent implementation of single-molecule spin-state read-out is the electromigrated tunnel (“break”) junction [18, 45–48]. In this architecture a single TbPc₂

bridges a nanometre-scale gap between two gold electrodes. Charge transport proceeds through the delocalised, ligand-centred π -radical, which is exchange-coupled to the Tb^{3+} moment. Consequently, the device conductance depends on the electronic spin state. When the external magnetic field is swept across the hyperfine-resolved avoided crossings, the electronic moment flips while conserving the nuclear projection m_I , producing sharp conductance jumps that electrically reveal the nuclear state.

This platform enabled microwave manipulation of the nuclear spin manifold and yielded spin-relaxation times $T_1 \sim 10\text{-}30$ s [45, 46] and inhomogeneous dephasing times $T_2^* \approx 200$ μs [18]. Coherent quantum control at the single-molecule level has been demonstrated, including multilevel protocols such as Grover’s algorithm [18] and the realisation of an iSWAP gate [47].

At the same time, these experiments reveal measurement back-action: the tunnelling current through the SMM shortens the spin lifetime, with a clear linear dependence of T_1 on the inverse current [45]. The approach developed in this thesis outlines a roadmap towards read-out schemes that avoid direct transport through the molecule by using silicon qubits as proximal sensors, thereby mitigating current-induced perturbations of the SMM.

TbPc₂ Neutral and Anionic Form

Two charge states of TbPc_2 are commonly employed. The neutral (“radical”) species, $[\text{TbPc}_2]^0$, is charge-neutral and hosts a ligand-centred π -radical that enables transport measurements in the mentioned break junctions. For ensemble spectroscopy with minimal dipolar/exchange broadening, TbPc_2 is typically diluted into an isostructural diamagnetic YPc_2 matrix. However, ensemble measurements on neutral-radical crystals generally do not retain clean SMM behaviour: close stacking promotes strong intermolecular exchange between the ligand π -radicals, effectively coupling molecules and broadening or smearing the hyperfine-resolved signatures [93, 94]. Consequently, the crystalline ensemble studies presented in Ch. 5 and Ch. 6 use the anionic form $[\text{TbPc}_2]^-$, which avoids the ligand-centred radical and thereby suppresses intermolecular coupling. The unit cell is shown in the right panel of Fig. 2.6 and includes two $[\text{TbPc}_2]^-$ and two tetraethylammonium (TEA) molecules.

2.3 Coupling SMMs to Silicon Spin Qubits

Building on the introduction of the industrial SiMOS spin qubit and the essential properties of the TbPc_2 SMM qudit, this section develops the physical coupling mechanism and translates it into device-scale numbers. The SMM functions as the qudit resource, while the SiMOS spin qubit serves as a nanoscale, on-chip magnetometer that transduces the SMM state via magnetic dipolar coupling. Using the imec gate-stack geometry and layer

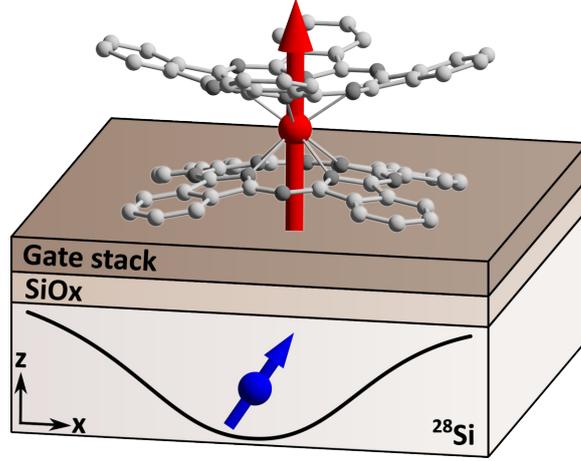


Figure 2.8: Hybrid system combining a single-molecule magnet and a silicon spin qubit. A single TbPc_2 molecule is placed above a silicon quantum-dot device and is magnetically coupled to the qubit via its dipolar field. The Tb^{3+} ion is depicted in red with its magnetic moment aligned with the molecular easy axis and indicated by a red arrow. The quantum-dot electron (blue) resides in isotopically enriched ^{28}Si within a lateral confinement potential (black) defined by the metallic gate stack (brown), which is separated from the silicon by an SiO_x layer. The gate stack and oxide together set the vertical separation between the molecular spin and the semiconductor qubit. Schematic not to scale; the quantum dot is an order of magnitude larger than the SMM.

dimensions, realistic coupling strengths are estimated at relevant separations. The section closes with a roadmap for hybrid-platform experiments.

Geometric Setup and Dipolar Coupling

The hybrid stack consists of a TbPc_2 SMM positioned above a SiMOS quantum-dot device. The separation between the two is set by the gate stack of the silicon qubit and the interfacial oxide (SiO_x). This is schematically shown in Fig. 2.8 with the large $J = 6$ moment of Tb^{3+} ($\mu \approx 9 \mu_B$ [95]) in red that couples to the $s = 1/2$ spin of the confined electron (blue) in the silicon quantum well.

The coupling is provided by the magnetic dipolar interaction, with the SMM approximated as a point dipole of moment $\boldsymbol{\mu}$ at displacement \mathbf{r} from the qubit. The resulting field is [96]

$$\mathbf{B}_{\text{SMM}}(\mathbf{r}) = \frac{\mu_0}{4\pi r^3} [3(\boldsymbol{\mu} \cdot \hat{\mathbf{r}}) \hat{\mathbf{r}} - \boldsymbol{\mu}] . \quad (2.9)$$

Two special orientations are

$$\mathbf{B}_{\text{SMM}}(\mathbf{r}) \stackrel{\mu \parallel \mathbf{r}}{=} 2 \frac{\mu_0 \mu}{4\pi r^3} \hat{\mathbf{r}}, \quad \mathbf{B}_{\text{SMM}}(\mathbf{r}) \stackrel{\mu \perp \mathbf{r}}{=} -\frac{\mu_0 \mu}{4\pi r^3} \hat{\boldsymbol{\mu}}, \quad (2.10)$$

showing that the field at the qubit location is twice as large on the dipole axis as in the equatorial plane. To maximise the contribution of $\boldsymbol{\mu}$ at the qubit, a parallel alignment is thus preferred, i.e. the SMM easy axis should be perpendicular to the chip surface, as shown in Fig. 2.8.

A switching event of the SMM magnetisation inverts the sign of $\mathbf{B}_{\text{SMM}}(\mathbf{r})$. The total magnetic field at the qubit is then

$$\mathbf{B}_{\text{tot}} = \mathbf{B}_{\text{ext}} \pm \mathbf{B}_{\text{SMM}}(\mathbf{r}), \quad (2.11)$$

which sets the silicon qubit's quantisation axis and resonance frequency.

The underlying detection of the electronic state of the TbPc_2 relies on changes in the qubit Larmor frequency induced by the added or subtracted field of the TbPc_2 . With the Zeeman Hamiltonian of Eq. (2.3), the Larmor frequency is

$$f_0 = \frac{g\mu_B}{h} |\mathbf{B}_{\text{ext}} \pm \mathbf{B}_{\text{SMM}}(\mathbf{r})|. \quad (2.12)$$

A magnetisation change of the SMM therefore leads to a small shift of the qubit resonance that can potentially be detected.

Readout of the TbPc_2 Nuclear States

A change in the electronic spin state of TbPc_2 produces a corresponding shift of the qubit resonance frequency. Following the break-junction experiment, the external field B_{ext} is swept between approximately ± 60 mT along the TbPc_2 easy axis. Four characteristic field values appear where the SMM reverses its magnetic moment via quantum tunnelling of magnetisation (QTM). These switching fields correspond to the hyperfine-resolved avoided crossings associated with the nuclear projections $m_I \in \{-3/2, -1/2, +1/2, +3/2\}$ (see Fig. 2.7b). Because each m_I branch shifts the position of the avoided crossing, the electronic moment flips at a different value of B_{ext} depending on the nuclear state. Identifying the switching field in a longitudinal sweep therefore provides a direct readout of m_I .

In contrast to the break-junction experiment, the SiMOS qubit requires a substantial static field for spin-to-charge readout. In the present devices, Elzerman readout (see Sec. 4.5) is implemented for $0.5 \text{ T} < B_{\text{ext}} < 0.7 \text{ T}$. This constant field cannot be applied along the TbPc_2 easy axis, as it would freeze the electronic state. Instead, it is applied transverse to the easy axis (along a hard axis) of the molecule. The transverse bias defines a stable quantisation axis for the qubit (GHz Larmor frequencies) and enhances the tunnel splitting $\Delta_{m,m'}$ of the ground doublet [97]. The corresponding increase in Landau-Zener tunnelling probability at each hyperfine avoided crossing facilitates state transfer without degrading nuclear-spin selectivity. Experimentally, controlled field components along the TbPc_2 easy axis are introduced by rotating the vector magnet field at fixed magnitude (see Ch. 5).

SiMOS as the Platform of Choice for Sensing

Beyond the SiMOS qubit structure discussed above, several semiconductor platforms have been engineered. Here, they are compared in terms of suitability for sensing, with emphasis on minimising the distance between a silicon qubit and the chip surface to enhance dipolar coupling to a SMM.

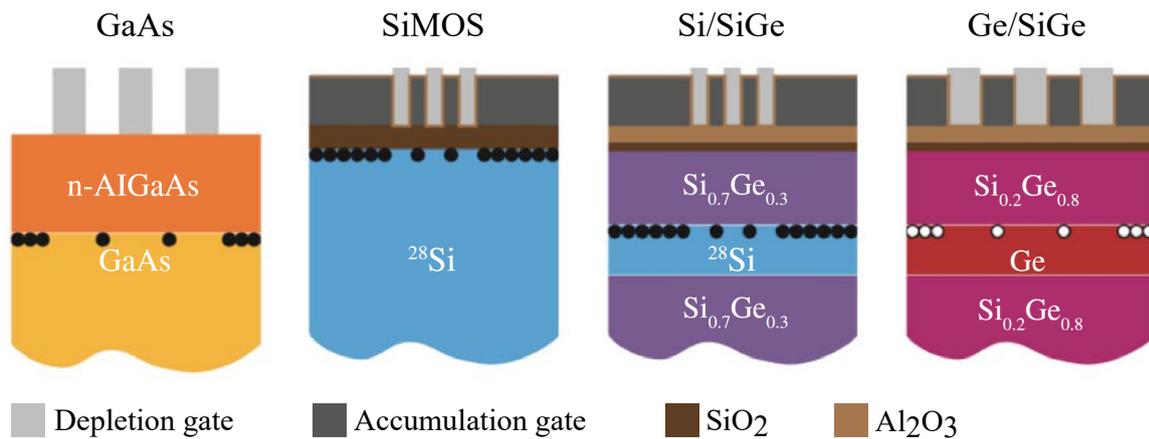


Figure 2.9: Schematic cross-sections of common gate-defined spin-qubit platforms: GaAs, SiMOS, Si/SiGe, and Ge/SiGe (left to right). Charge carriers are indicated by dots at the confinement plane (electrons: black; holes: white). In GaAs, a 2DEG is formed at the n-AlGaAs/GaAs interface and quantum dots are defined by depletion gates (light grey). In SiMOS, electrons are accumulated at the $^{28}\text{Si}/\text{SiO}_2$ interface using accumulation gates (dark grey) in combination with depletion gates for lateral confinement. In Si/SiGe, an electron quantum well in ^{28}Si is embedded between $\text{Si}_{0.7}\text{Ge}_{0.3}$ barriers, increasing the distance to the oxide. In Ge/SiGe, holes are confined in a Ge quantum well between $\text{Si}_{0.2}\text{Ge}_{0.8}$ barriers. Taken from [99].

Historically, coherent spin control was first demonstrated in modulation-doped AlGaAs/GaAs heterostructures [17]. In these devices a two-dimensional electron gas (2DEG) exists at zero gate bias due to remote donors in the AlGaAs layer. Surface gates biased negatively locally deplete the 2DEG to form quantum dots. For sensing, this depletion-mode geometry has the appealing feature that confinement can be achieved without an accumulation gate directly above the dot (Fig. 2.9), leaving local surface area for integration of an SMM. However, because all naturally abundant isotopes of Ga and As carry nuclear spin, hyperfine noise strongly limits coherence [17], and the AlGaAs spacer layer that separates the surface from the 2DEG is typically on the order of 50-100 nm [98], which is unfavourable for near-field sensing.

SiMOS emerged as a leading alternative thanks to nuclear-spin-free silicon isotopes ($^{28}\text{Si} \sim 92\%$, $^{30}\text{Si} \sim 3\%$) and the availability of isotopically enriched ^{28}Si with residual ^{29}Si concentrations as low as ~ 50 ppm [100]. The first single-qubit control in enriched ^{28}Si was reported in 2014 [101]. In SiMOS, the oxide thickness between the gate stack and the quantum dot is typically below 10 nm [99] (in device A of this thesis it is 8 nm; see Sec. 4.1), enabling very small spin-surface separations that are advantageous for sensing.

To mitigate trap-induced charge noise associated with the Si/SiO₂ interface [40], silicon quantum wells can be embedded between SiGe layers (Si/SiGe devices). This pushes the oxide further away, but introduces a SiGe spacer that generally exceeds ~ 40 nm [102–104], increasing the qubit-surface distance. The larger separation also weakens AC magnetic fields generated by conventional on-chip waveguides, so micromagnets together with electric-dipole spin resonance are commonly employed for fast control [54] which is a combination that is less favourable for magnetic-field sensing with an overlying SMM that is potentially influenced by the micromagnet.

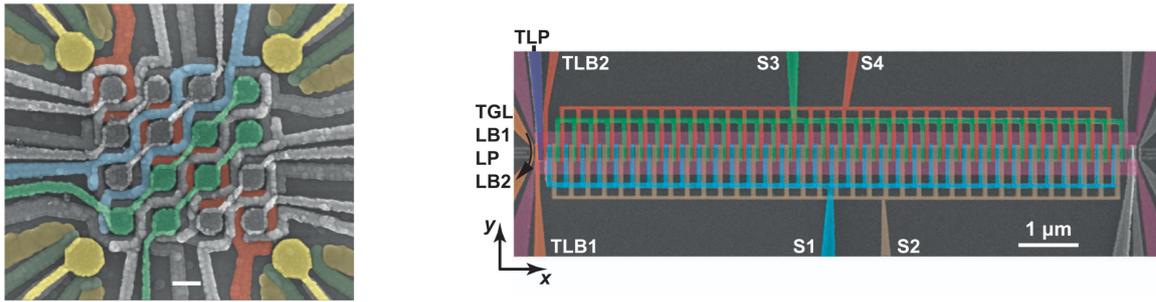


Figure 2.10: Scalable device architectures for on-chip quantum sensing. **Left** Quantum-dot crossbar array in Ge/SiGe. False-coloured scanning electron micrograph (top view) with individual gates coloured for clarity. Two staircase barrier-gate layers (red and blue) run orthogonally. A set of 16 quantum dots are defined beneath the plunger-gate layer (green) and four corner charge sensors (single-hole transistors, yellow) provide readout. Scale bar: 100 nm. The crossbar architecture controls 16 dots with only 23 control terminals. Taken from [37]. **Right** False-coloured scanning electron micrograph (top view) of a Si/SiGe shuttling (“QuBus”) device. Phase-shifted sinusoidal waveforms applied to the four shuttling gates (S1-S4) create a moving potential well that transports electrons along the x -direction. Taken from [108].

Ge/SiGe devices employ a Ge quantum well between SiGe barriers and accumulate holes rather than electrons. Single-hole manipulation has been demonstrated in 2020 [105]. The strong spin-orbit coupling of holes enables all-electric control without micromagnets, but typical spacers of ~ 20 -55 nm still separate the well from the surface [37, 106, 107], again limiting proximal coupling to external spins.

Gate-stack simulations (Ch. 7) indicate that accumulation layouts can be optimised so that a dot forms without placing a gate directly above it. With such sensing-oriented designs, the relevant separation is set by the oxide in SiMOS (only a few nanometres), whereas in Si/SiGe and Ge/SiGe it is the oxide *plus* the heterostructure spacer. This makes SiMOS the most favourable for strong dipolar coupling. Moreover, SiMOS devices typically incorporate an on-chip microwave antenna for ESR, which can potentially be repurposed to drive an SMM magnetically. In summary, the combination of ultra-shallow dots, the possibility of local oxide thinning or removal (Sec. 7.4), and native ESR hardware makes SiMOS the platform of choice for sensing applications.

Sensor Geometries

Recent progress in two-dimensional (2D) quantum-dot arrays and on-chip spin shuttling opens routes to spatially resolved sensor layouts (Fig. 2.10). Two-dimensional arrays of gate-defined quantum dots have been realised in semiconductor spin-qubit platforms [37, 38, 109, 110]. The largest implementation to date is the 4×4 Ge/SiGe crossbar array [37] shown in the left panel of Fig. 2.10, where interdot tunnel coupling in vertical and horizontal pairs is tunable over a ~ 20 GHz range. Although site-resolved single-qubit driving was not demonstrated, this architecture already illustrates the potential of semiconductor spin qubits to probe many lateral positions across a chip, enabling measurements of single molecules or microscopic particles placed above selected pixels.

Furthermore, lateral shuttling of single-electron qubits using conveyor-type gate sequences has been demonstrated [108, 111, 112]. For quantum computing, this enables coupling beyond nearest neighbours by translating an electron between dots over micrometre distances to interact with far-separated qubits. To date, charge shuttling has been shown over distances of order $10\ \mu\text{m}$, while preservation of spin coherence during motion has been demonstrated up to $\sim 0.56\ \mu\text{m}$ [112]. Beyond computation, such shuttling architectures offer movable sensor qubits: the spin can be positioned close to a target for enhanced coupling and then potentially displaced to a low-noise site for readout, effectively decoupling the optimal sensing and readout locations.

Coupling Estimations

The magnetic coupling between a TbPc_2 molecule and a SiMOS spin qubit is primarily set by their vertical separation. In the present (non-optimised) device, the gate stack and ESR antenna are encapsulated in a SiO_x passivation layer that fixes a minimum spacing of $r \approx 250\ \text{nm}$. A detailed layer stack is given in Tab. C.1. At this distance, the on-axis dipolar field of a Tb^{3+} moment $\mu \approx 9\ \mu_B$ [95] at the qubit position is

$$B(250\ \text{nm}) \stackrel{\mu \parallel r}{=} \frac{\mu_0 \mu_{\text{TbPc}_2}}{2\pi r^3} = \frac{\mu_0 9\ \mu_B}{2\pi (250\ \text{nm})^3} \approx 1\ \text{nT}. \quad (2.13)$$

A reversal of the TbPc_2 moment changes the local field by $\Delta B = 2B$, producing a Larmor shift of the silicon qubit

$$\Delta f = 2 \frac{g\ \mu_B}{h} B \approx 60\ \text{Hz}, \quad (2.14)$$

for $g \approx 2$ and $B \approx 1\ \text{nT}$. This is exceedingly small. For context, the qubit's homogeneous coherence time extracted from echo is $T_2 \approx 92\ \mu\text{s}$ (see Sec. 4.10), corresponding to a Lorentzian Fourier-limited linewidth of $1/(\pi T_2) \approx 3\ \text{kHz}$ [113]. Although this linewidth is only a rough indicator of detectability, it makes clear that the $\sim 60\ \text{Hz}$ single-molecule frequency step is far below the present SiMOS qubit's sensitivity at $r \approx 250\ \text{nm}$.

The r -dependence of Δf is sketched in Fig. 2.11. The red band marks the current separation ($r \approx 250\ \text{nm}$). The orange band indicates the approximate sensitivity window of the silicon qubit set by the Lorentzian linewidth. There is no overlap with the red band at present. To obtain a measurable response to switching events, a highly diluted ensemble is therefore employed in crystalline form: 5% $[\text{TbPc}_2]^-$ in a 95% diamagnetic $[\text{YPc}_2]^-$ matrix. In this configuration the molecular spins are sufficiently separated to preserve single-molecule behaviour (Sec. 2.2.1), yet the collective dipolar field at the qubit enhances the net shift into the tens-of-megahertz range (see measurements in Ch. 5), yielding a clear hysteresis signature.

Finally, while the orange region in Fig. 2.11 reflects the Fourier-limited linewidth of the qubit, the blue region highlights the sensitive operating range of the readout method developed in Ch. 5: a simultaneous rapid adiabatic passage and spin-selective tunnelling (RPSS) protocol tailored to detect frequency shifts in a $\sim 100\ \text{MHz}$ window from a TbPc_2 ensemble.

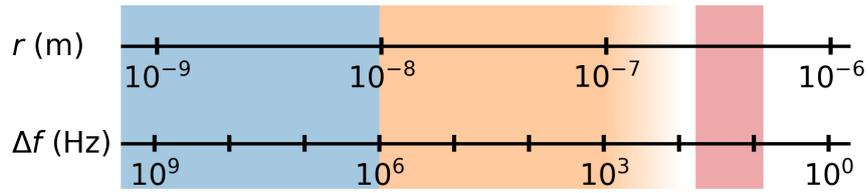


Figure 2.11: Distance dependence of dipolar coupling in the hybrid device. The two axes map the vertical separation r between a single $[\text{TbPc}_2]^-$ molecule and a silicon spin qubit to the corresponding qubit frequency shift Δf . In the current industrial device the gate stack and passivation set $r \approx 250$ nm (red), giving $\Delta f \approx 60$ Hz which is below the qubit's typical detection limit (orange, $T_2 \approx 92$ μs). Using a dilute ensemble amplifies the magnetic field by orders of magnitude and moves the response into the sensitive MHz range of the RPSS protocol (blue) described in Ch. 5.

Outlook: Single SMM Sensing

A central requirement for sensing a *single* SMM is to enhance the dipolar coupling between the molecular moment and the silicon spin qubit. Since the dipolar field decays as $B \propto 1/r^3$ with the separation r between SMM and qubit, the most effective lever is geometric: the physical distance between a surface-deposited TbPc_2 and the electron in the silicon quantum dot must be reduced, while the device remains tunable in the few-electron regime.

In the present industrial SiMOS device A (see Sec. 4.1 and Tab. C.1), the vertical separation between a TbPc_2 placed on the chip surface and the qubit is set by three stacked contributions above the channel: the gate stack, the on-chip ESR antenna and a SiO_x passivation layer. The gate stack and antenna are embedded within this passivation. This architecture is excellent for robustness and yield in standard qubit operation, but it pushes the molecule too far away from the active quantum region and thus strongly suppresses the dipolar field at the qubit.

Gate-stack simulations in Ch. 7 show that it is possible to alter the layout such that no gate needs to sit directly on top of the quantum dot used for sensing, while preserving device functionality. In parallel, Section 7.4 investigates post-processing routes in which the passivation layer is locally removed above the target dot by a selective hydrofluoric (HF) etch, opening a window to the surface without modifying the rest of the wafer. Together with an ESR antenna that is laterally displaced from the sensing site (at the cost of a modest reduction in Rabi rate), the minimum SMM-qubit distance is then set by the internal oxides of the gate stack.

For the purpose-designed layout of device A, this residual separation is approximately $r \approx 18$ nm. Using the dipolar field expression of Eq. (2.13), a TbPc_2 at this distance produces a field of $B(18 \text{ nm}) \approx 2.9$ μT at the qubit, corresponding to a frequency shift of $\Delta f \approx 160$ kHz.

Further thinning of the oxide by local HF etching can potentially reduce the separation towards $r \approx 5$ nm, where one expects

$$B(5 \text{ nm}) \approx 130 \mu\text{T} \quad \text{and} \quad \Delta f \approx 7.5 \text{ MHz}.$$

Both frequency shifts are orders of magnitude larger than the Fourier-limited linewidth of the present SiMOS qubit (~ 3 kHz). This indicates that, with suitably optimised device geometries and local post-processing, the field change associated with the magnetisation reversal of a single TbPc₂ should be resolvable, thereby paving the way towards single SMM sensing with SiMOS spin qubits.

3 Experimental Methods and Instrumentation

This chapter describes the experimental platform used to investigate hybrid quantum systems that couple single-molecule magnets to silicon spin qubits. Measurements were carried out in a custom $^3\text{He}/^4\text{He}$ dilution refrigerator with a base temperature below 50 mK, supporting stable device operation at millikelvin temperatures (Sec. 3.1).

We first cover the physical stack: the standard sample printed-circuit board (PCB) for device mounting and wiring (Sec. 3.2); the cryogenic cabling and filtering that link room temperature to the mK stage (Sec. 3.3); and the interfaces that integrate room-temperature instrumentation with the cryostat, including thermal anchoring and signal routing (Sec. 3.4).

We then describe the signal chain: microwave control via single-sideband IQ mixing (Sec. 3.5); high-speed gate pulses delivered through resistive bias-tee networks with pulse pre-emphasis and compensation (Sec. 3.6); and deterministic acquisition and device control using an ADwin Pro II (Sec. 3.7). We conclude with a newly developed high-density interface PCB (“The Beast”) that expands DC and RF I/O and provides on-board filtering and reflectometry (Sec. 3.8).

3.1 Dilution Cryostat (Sionludi L)

The cryogenic environment required for the experiments in this work was provided by a customised Sionludi L tabletop dilution refrigerator originally developed at the Néel Institute (CNRS, Grenoble) and optimised for compactness, rapid cooldown, and efficient helium use. It employs two independent helium cycles: a ^4He precooling loop and a closed ^3He - ^4He dilution circuit, which are thermally coupled via counter-flow heat exchangers. This configuration, combined with a direct fast-injection line for precooled mixture, enables a remarkably short cooldown time and reduces complexity in the precooling phase.

Operating at such low temperatures is crucial for suppressing thermal population of excited spin states. For an electron spin with $g \approx 2$ in a magnetic field of $B = 0.66$ T (used in all single-qubit experiments shown in Chs. 4 and 5), the Zeeman energy splitting is

$$\Delta E = g\mu_B B = 2 \times 5.788 \times 10^{-5} \text{ eV/T} \times 0.66 \text{ T} \approx 7.64 \times 10^{-5} \text{ eV}, \quad (3.1)$$

which corresponds to

$$\frac{\Delta E}{k_B} = \frac{7.64 \times 10^{-5} \text{ eV}}{8.617 \times 10^{-5} \text{ eV/K}} \approx 0.884 \text{ K}. \quad (3.2)$$

In thermal equilibrium, the excited-state occupation p_{exc} of a two-level system follows an Arrhenius-type dependence [63]:

$$p_{\text{exc}} = \exp\left(-\frac{\Delta E}{k_{\text{B}}T}\right). \quad (3.3)$$

At $T = 51$ mK, this yields

$$p_{\text{exc}} \approx \exp\left(-\frac{0.884}{0.051}\right) \approx 2.8 \times 10^{-8}. \quad (3.4)$$

However, as will be discussed in Sec. 4.4, the electron temperature T_{e} of the device is elevated due to imperfect thermalisation of the electrical lines, reaching $T_{\text{e}} = 130$ mK. This leads to an excited-state occupation of

$$p_{\text{exc}} \approx \exp\left(-\frac{0.884}{0.130}\right) \approx 1.1 \times 10^{-3}, \quad (3.5)$$

corresponding to approximately 0.1% of spins thermally populating the excited state even before any control pulses are applied. This finite thermal population imposes a fundamental limit on the qubit initialisation fidelity and underlines the critical importance of achieving and maintaining the lowest possible temperature in the experimental environment.

In the present configuration, the cryostat cools from room temperature to below 100 mK in approximately 3 h 20 min, without the use of any intermediate cryogens other than liquid helium. The base temperature achieved during stable operation is $T_{\text{base}} = 51$ mK, measured at the mixing-chamber plate using a calibrated ruthenium oxide thermometer. The mechanical layout follows a vertical, multi-stage arrangement reminiscent of a nested-doll geometry. Each thermal stage is mechanically supported by stainless-steel columns, enclosed by dedicated radiation shields, and isolated under high vacuum. The millikelvin (mK) stage is positioned at the top of the cryostat, directly above the mixing chamber, where the sample holder is mounted.

The cooling process begins in the ^4He cycle by filling the 4 K pot from an external liquid-helium storage vessel. The liquid helium is then routed through a counter-flow heat exchanger, where it pre-cools the ^3He - ^4He mixture of the secondary cycle. Once cooled to approximately 4 K, the mixture undergoes further cooling in a Joule-Thomson expansion stage before entering the dilution unit. In standard operation, the mixture is thermalised first to the still and subsequently flows into the mixing chamber via the normal injection line. A distinctive feature of the Sionludi design is the inclusion of a fast-injection line, which has a significantly larger cross-section than the normal injection path. This fast-injection line permits precooled mixture to circulate directly from the 1 K pot (also referred to as the turbo-condenser) into the mixing chamber during cooldown. By bypassing intermediate restrictions and delivering cooling power more directly, this arrangement shortens cooldown times by several hours compared to conventional dilution refrigerators.

When the mixing chamber reaches the phase-separation temperature of approximately 0.86 K, dilution cooling begins [114, 115]. In this regime, ^3He atoms dissolve into the

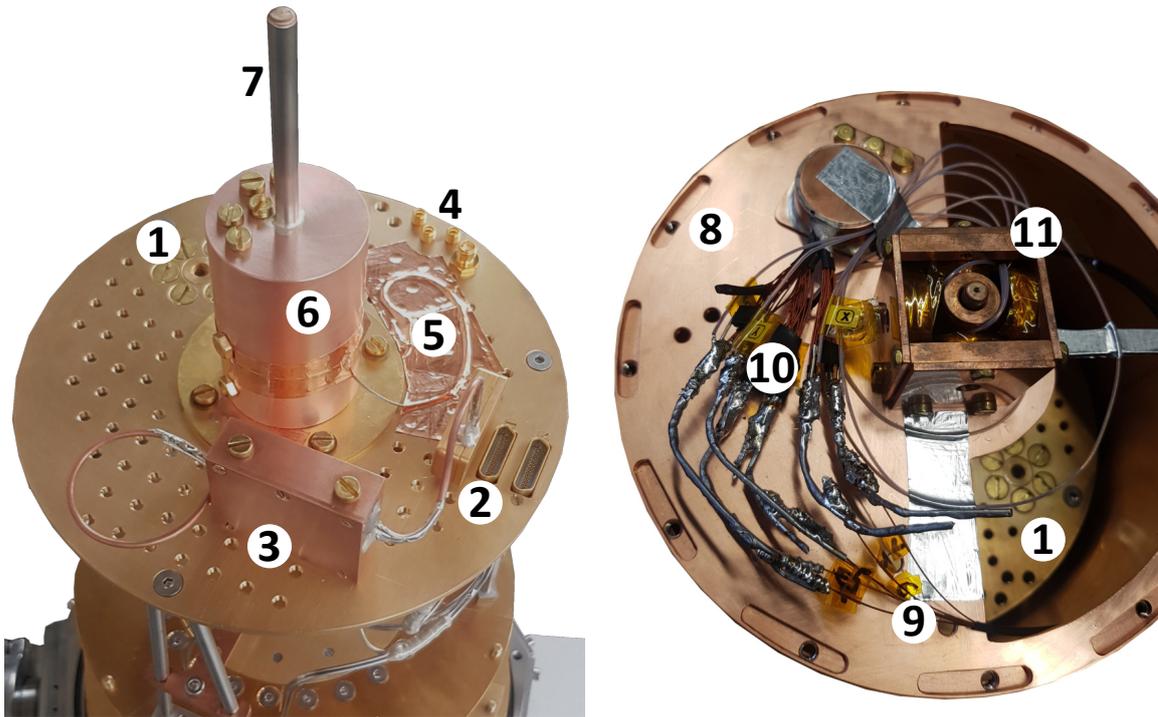


Figure 3.1: Experimental volume above the mK stage. **Left:** Side view of the Sionludi L showing: mK stage (diameter 17 cm) with connection to the dilution unit (1), DC line connectors (2), first low-pass filtering stage (3), RF connections (4), thermometer line (5), mK shield enclosing the sample PCB (6), and sample position (7). **Right:** Top view with the elevated 4 K plate (8) attached, showing the low- T_c superconductors (9), their thermal anchor (10), and 3D vector magnet (11).

^4He -rich phase, an endothermic process that absorbs heat from the surroundings due to the associated increase in entropy at constant total energy, in accordance with the second law of thermodynamics. The resulting difference in chemical potential between the concentrated and dilute phases generates an osmotic pressure gradient that drives ^3He atoms towards the still. At ~ 0.7 K, the saturation vapour pressure of ^3He is roughly an order of magnitude greater than that of ^4He , causing ^3He to evaporate preferentially. The evaporated ^3He is extracted from the still by a turbomolecular pump, recondensed via the counter-flow heat exchanger and Joule-Thomson expansion stage, and then reinjected into the mixing chamber, thereby completing the continuous cooling cycle.

A photograph of the upper stages of the cryostat used in this work is shown in Fig. 3.1, while a schematic view of its internal structure is presented in Fig. 3.2. Mounted on an elevated 4 K shield is a three-dimensional vector magnet system, developed by Julian Ferrero [116]. The magnet is wound from NbTi superconducting wire and provides a maximum field of 1 T along its primary x -axis, which coincides with the main rotational symmetry axis of the cryostat, and up to 0.35 T along each of the two orthogonal directions.

The compact tabletop geometry of the cryostat, measuring approximately 22 cm in diameter and 53 cm in height, allows convenient access to the experimental space and facilitates the integration of auxiliary instrumentation. All DC and RF wiring is routed through the

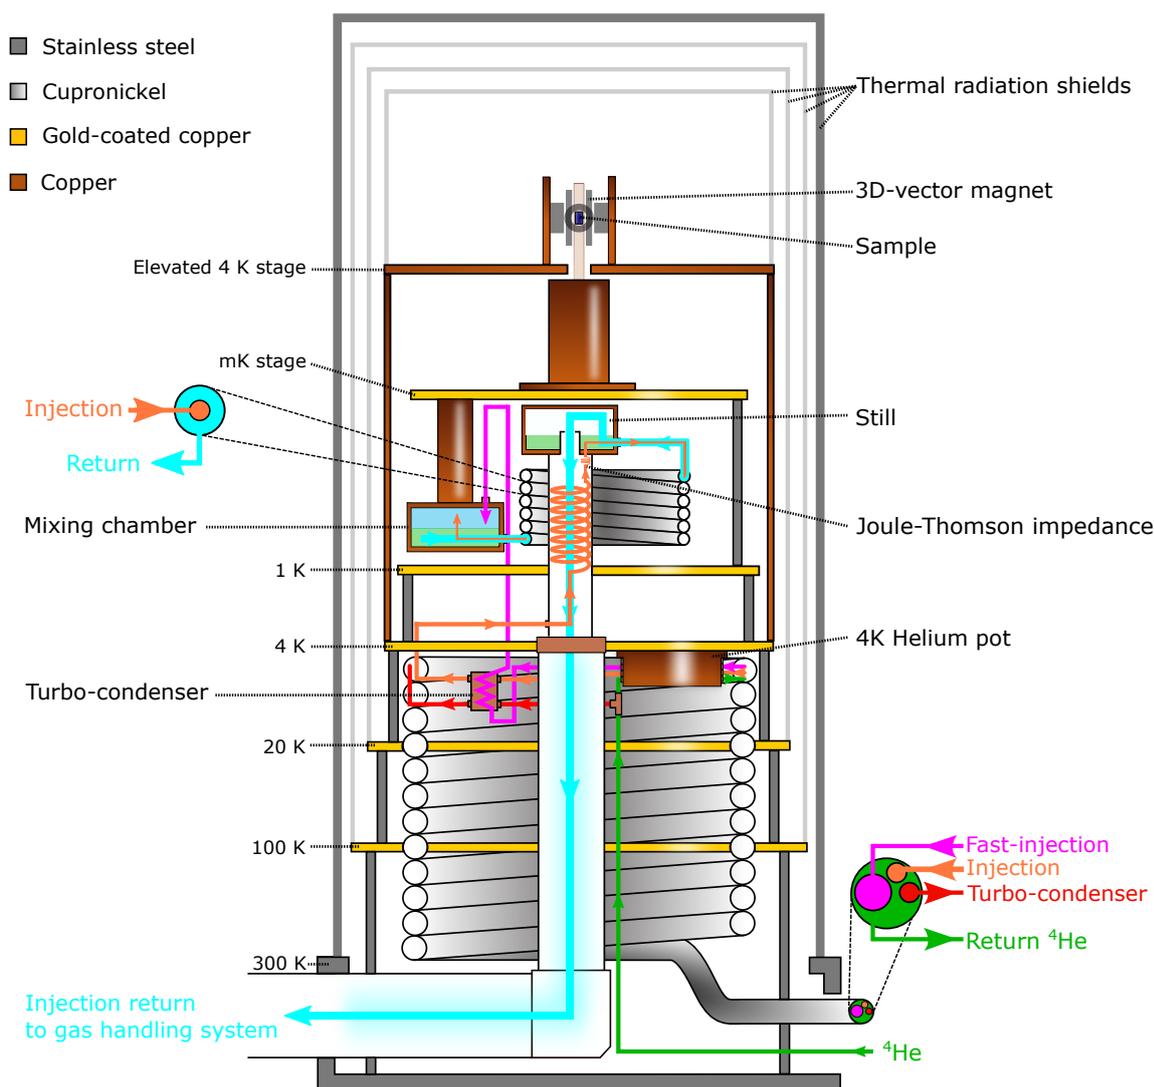


Figure 3.2: Schematic representation of the Sionludi L dilution refrigerator, nicknamed “Idefix”. The cryostat operates using only ^4He , eliminating the need for liquid nitrogen handling during regular operation. A superconducting 3D vector magnet is positioned around the sample region (blue) and thermally anchored to an elevated 4 K stage. The qubit sample, mounted on a printed circuit board (PCB), resides on the mixing chamber stage at millikelvin temperatures and is enclosed within a copper cylinder for thermal shielding.

nested thermal stages, with thermal anchoring implemented at each temperature stage to minimise conductive heat leaks to the millikelvin stage, as described in Sec. 3.3.

In summary, the Sionludi L refrigerator used in this work combines rapid cooldown capability, a stable base temperature of 51 mK, and a compact tabletop form factor, all without the need for a liquid-nitrogen pre-cooling stage. These characteristics make it particularly well suited for experiments requiring both fast turnaround and ultra-low thermal excitation of quantum devices.

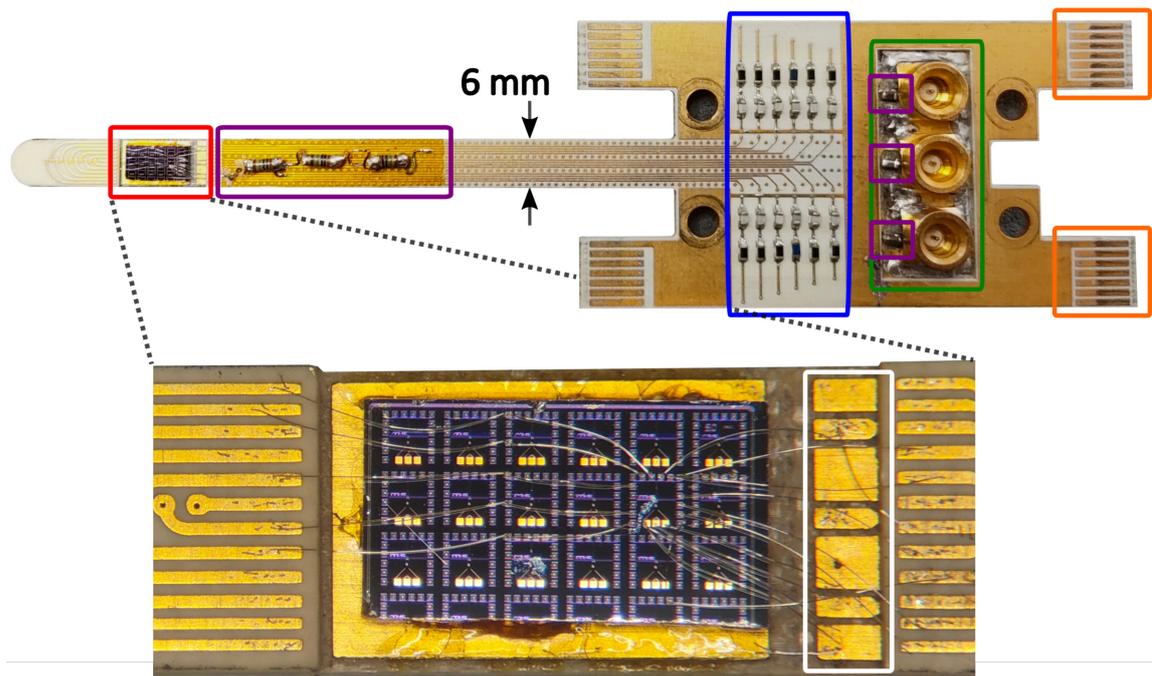


Figure 3.3: Standard PCB with individual components highlighted. **Upper picture:** The gold-coated copper surface is visible in yellow, with the chip position indicated in red. DC connections to the PCB are made via two PCIe connectors, whose pin outlines are highlighted in orange; additional PCIe connectors present on the board are not used during operation but can be used to prevent electrostatic discharge by application of a short-circuit adapter during transfer of the PCB. The second-stage DC filtering section (also present on the reverse side of the PCB) is marked in blue. High-frequency connections are provided by three SMP connectors shown in green. The bias-tee networks are indicated in purple and consist of a resistor (left box) and a capacitor (right box near the SMP connector). **Lower picture:** Zoom of the red-marked chip area. The chip is glued to the PCB using GE varnish, and electrical connections are made by wire bonding. The three coplanar waveguides approach from the right-hand side.

3.2 Printed Circuit Board for Sample Interfacing

All data presented in this thesis were acquired using one of the Wernsdorfer group's standard printed circuit boards (PCBs), designed by Tino Cubaynes. The PCB layout and functionality are described in detail below.

The standard PCB used to interface the sample with the cryogenic measurement setup is shown in Fig. 3.3. It provides 24 DC lines and 3 RF lines, with dedicated pads integrated into the layout to accommodate on-board RC filters for the DC connections. The second-stage DC filtering, implemented using surface-mount devices (SMDs), is visible in Fig. 3.3 and includes capacitances installed in series with the inner conductors of each RF line. The specific filter configurations and cut-off frequencies of the DC lines are discussed in Sec. 3.3.

Electrical connectivity for the DC lines is provided via two split-PCI connectors, one located on the top and one on the bottom of the PCB. This dual-connector arrangement enables galvanic grounding of the DC lines during sample exchange or storage, thereby reducing

the risk of electrostatic-discharge damage. The RF lines are connected through three SMP connectors and are routed to the sample region via individual coplanar waveguides designed to support frequencies up to 20 GHz.

The sample is mounted at the end of an extended cold finger, positioning it at the centre of the 3D vector magnet. It is fixed onto a 7 mm × 4 mm gold-coated copper plate, which serves as a back gate and is grounded in our experiments. Electrical contact between the sample and the PCB is established via wire bonds from the sample pads to the surrounding DC and RF bonding pads.

Resistive bias-tees can be implemented on the PCB by soldering a capacitor in series with the RF inner conductor and a resistor in series with the corresponding DC line. Near the sample, the RF and DC lines are joined by wire bonds to complete the bias-tee configuration. As shown in Fig. 3.3, bias-tees are installed on all three RF lines, consisting of 0402 SMD capacitors and metal-electrode leadless-face (MELF) resistors. Further details on the exact component values and electrical characteristics of the bias-tees are provided in Sec. 3.6.

3.3 Cryogenic Cabling and Filtering

As part of this thesis, 72 DC lines were installed in the cryostat in collaboration with Viktor Adam [119]. Each line consists of 70 cm of insulated Constantan wire (80 μm diameter) with a resistance of approximately 90 Ω. The lines are bundled in groups of 24 and housed within a FeCuNi capillary of 1.5 mm diameter. A thin coating of Apiezon grease is applied during assembly to reduce friction, thereby preventing insulation damage, and to improve thermal contact between the wires and the capillary. A detailed fabrication procedure is provided in Ch. A.

After assembly, the capillaries are galvanically connected to ground to provide electromagnetic shielding and mechanical stabilisation. To achieve this, silver-coated copper wire is affixed to the capillary using silver powder epoxy.

At millikelvin temperatures, the DC lines are filtered using a two-stage RC low-pass filter configuration, yielding combined cut-off frequencies of 250 Hz for static voltage gates and 59 kHz for DC lines used in current readout. Each DC line is also fitted with an LFCN-80+ surface-mount band-stop filter (Mini-Circuits), providing additional mid-frequency attenuation between 200 MHz and 4.5 GHz. A schematic of the filter configurations and the corresponding passive components is shown in Fig. 3.4. The LFCN-80+ filter and the first RC filter stage are soldered onto a dedicated filter PCB, enclosed in a copper housing, and thermally anchored to the mK stage. The second-stage RC filters are integrated directly onto the sample PCB, as illustrated in Fig. 3.3.

The cryostat was further equipped with 12 RF lines to deliver microwave signals to the sample. Each RF line consists of four sections of CuNi coaxial cable: (1) from room temperature to the 4 K stage (thermally anchored to the 100 K and 20 K stage), (2) from 4 K to the 1 K stage, (3) from 1 K to the 50 mK stage, and (4) from 50 mK to the sample

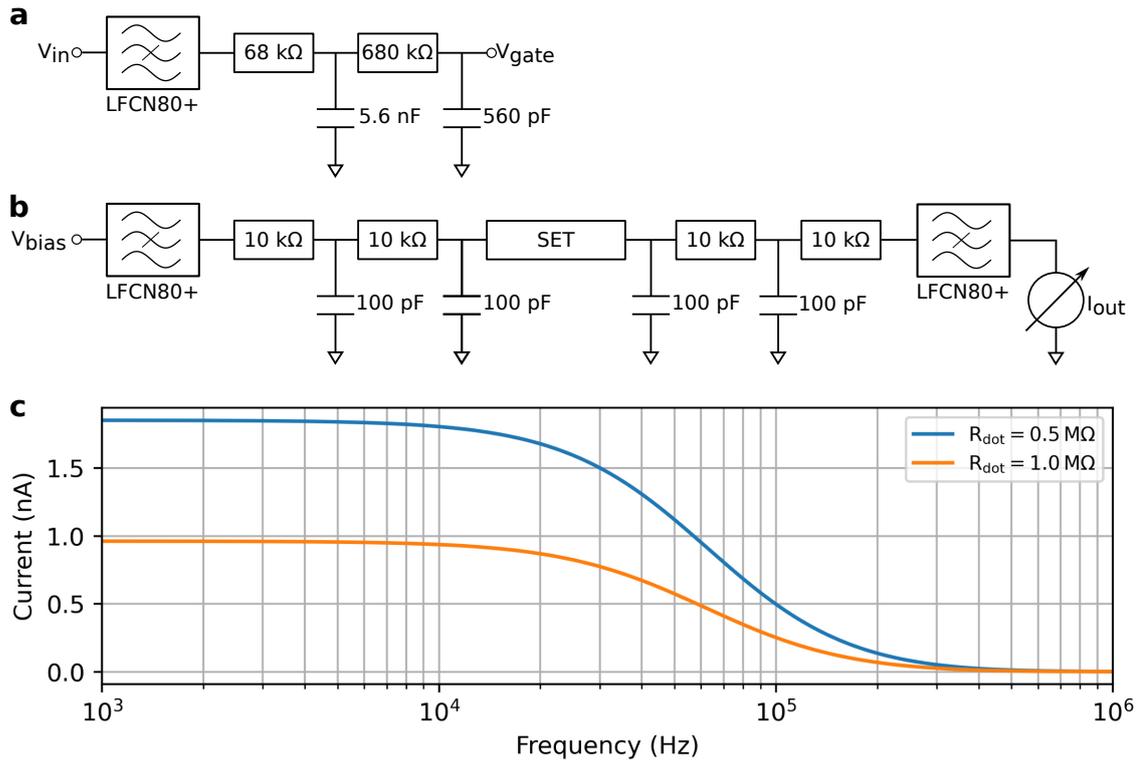


Figure 3.4: Low-pass filtering of the standard PCB at millikelvin temperatures. **a** Two-stage low-pass filter arrangement used for gate voltages, with a -3 dB point at 250 Hz. An additional Mini-Circuits LFCN-80+ band-stop filter is included, providing a stop band in the range of 200-4500 MHz. The first low-pass stage is located on the mixing chamber stage, while the second stage is implemented directly on the PCB. **b** Filter arrangement of the DC lines used for readout, designed with a higher cut-off frequency. A two-stage low-pass filter leads to the single-electron transistor (SET), with a -3 dB point at 59 kHz. The current then flows through the filter configuration in reverse until it reaches the readout electronics. The -3 dB point of the readout path is 40 kHz. **c** LTspice simulation of the signal path shown in **b**, excluding the LFCN-80+ filter, with an applied 1 mV amplitude. Two SET resistances (0.5 M Ω and 1 M Ω) illustrate the current-frequency response; these values are typical for SET charge sensors [117, 118].

PCB. Within the cryostat, SMP connectors are employed due to their compact form factor and frequency range up to 26.5 GHz, while SMA connectors are used for the hermetic RF feedthroughs. Additional cryogenic attenuators (Rosenberger 19AS101) are installed at the 4 K, 1 K, and 50 mK stages to improve the signal-to-noise ratio (SNR) at the sample.

In the measurements reported in this thesis, only 24 DC lines and three RF lines were used for qubit manipulation, as the standard sample PCB accommodates this number. With the novel PCB design presented in Sec. 3.8, however, all 72 DC lines and 12 RF lines installed in the cryostat can be fully utilised.

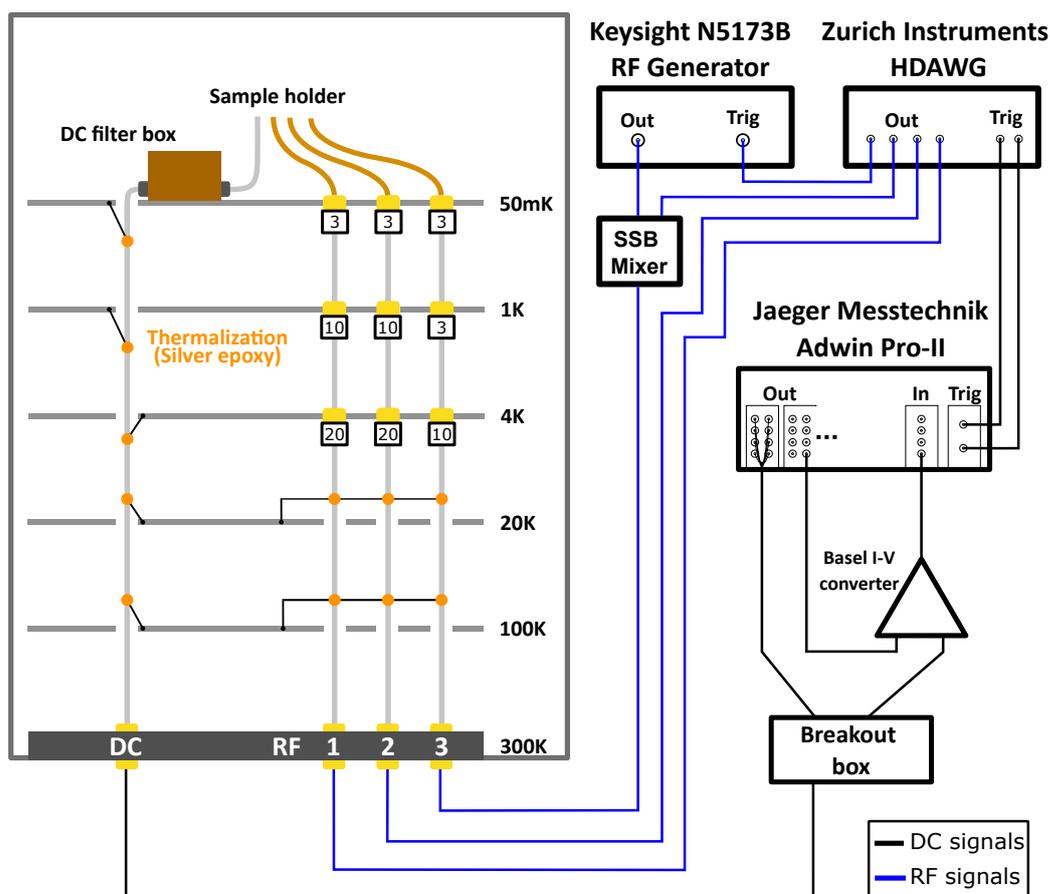


Figure 3.5: Integration of the measurement electronics with the cryostat. DC voltages (black lines) are routed from the ADwin system via a breakout box into the cryostat, while output DC signals are returned via the IV converter to the ADwin inputs. Time-domain control is achieved using the HDAWG, which triggers the microwave generator and mixes its MHz-frequency signal to the GHz range using the RF tone. Gate manipulation pulses are sent directly from the HDAWG to the cryostat. Attenuation values for the RF path are indicated.

3.4 Integration of Measurement Electronics with the Cryostat

The overall measurement architecture used for qubit experiments is illustrated in Fig. 3.5, which depicts the full signal chain from the room-temperature control electronics to the cryogenic wiring. All DC lines are controlled by an ADwin Pro II real-time control system (Jaeger Messtechnik), offering a voltage resolution of $305 \mu\text{V}$ from its 16-bit $\pm 10 \text{ V}$ outputs, and enter the cryostat through an external breakout box (a shielded fan-out enclosure that converts multiway connectors to individual connectors). The readout current is converted to voltage at room temperature by an SP983c current-voltage (IV) converter (Basel Precision Instruments) and subsequently digitised using a Pro II AIn-F-4/16 analogue input module, which supports sampling rates of up to 4 MHz. For the standard amplification setting of 10^8 V/A applied throughout this work, the IV converters have a specified bandwidth of 24 kHz.

Fast voltage pulses for qubit manipulation are generated by a HDAWG arbitrary waveform generator (Zurich Instruments), which provides four output channels operating at a sampling rate of 2.4 GHz. Microwave control is implemented using a Keysight N5173B microwave source, delivering frequencies up to 20 GHz. Both signals are combined using a Marki Microwave SSB-0618 single-sideband mixer to generate the required sideband tones for qubit control. The principle of IQ mixing and its role in these experiments is discussed in Sec. 3.5. All instruments are computer-controlled via the Python-based Spin-Suite, developed by Julian Ferrero [116] within the qkit measurement framework [120].

Inside the cryostat, DC wiring is carefully filtered and thermally anchored at multiple temperature stages. As indicated by the orange markers in Fig. 3.5, thermal anchoring is realised using silver epoxy. For clarity, only the three RF lines and one DC capillary containing 24 lines used in this thesis are shown in Fig. 3.5. Each RF line passes through staged attenuation at the 4 K, 1 K, and 50 mK plates to improve the signal-to-noise ratio and suppress high-frequency noise.

This integrated setup enables coherent qubit control while maintaining a low-noise, thermally optimised environment, ensuring stable operating conditions for the experiments presented in this work.

3.5 Single-Sideband Mixing

Microwave pulses for qubit control in this thesis are generated with a Marki Microwave SSB-0618 single-sideband mixer. In general, an SSB mixer allows one to manipulate a high-frequency carrier (here ~ 18.5 GHz) using low-frequency baseband signals (here < 100 MHz). This enables full amplitude, phase, and frequency control without retuning the microwave source using cost-effective baseband electronics.

An ordinary (double-sideband) mixer is a three-port device that multiplies a local-oscillator (LO) tone with a baseband or intermediate-frequency (IF) signal. Let

$$v_{\text{LO}}(t) = A_{\text{LO}} \cos(\omega_{\text{LO}} t), \quad v_{\text{IF}}(t) = A_{\text{IF}} \cos(\omega_{\text{IF}} t).$$

The RF output is proportional to the product

$$v_{\text{RF}}(t) \propto v_{\text{LO}}(t)v_{\text{IF}}(t) = \frac{A_{\text{LO}}A_{\text{IF}}}{2} \left[\cos((\omega_{\text{LO}} - \omega_{\text{IF}})t) + \cos((\omega_{\text{LO}} + \omega_{\text{IF}})t) \right], \quad (3.6)$$

using $\cos(x) \cos(y) = \frac{1}{2} [\cos(x - y) + \cos(x + y)]$. Thus, ordinary mixing creates two sidebands: the lower sideband (LSB) at $\omega_{\text{LO}} - \omega_{\text{IF}}$ and the upper sideband (USB) at $\omega_{\text{LO}} + \omega_{\text{IF}}$ (Fig. 3.6 left). In practice there is also finite LO and IF feedthrough. While small leakage may be tolerable, continuous RF carrier leakage heats the device and raises the electron temperature, reducing ground-state occupation and capping qubit initialisation fidelity. In multi-tone control, out-of-band leakage can also drive off-resonant rotations, thereby lowering gate fidelities.

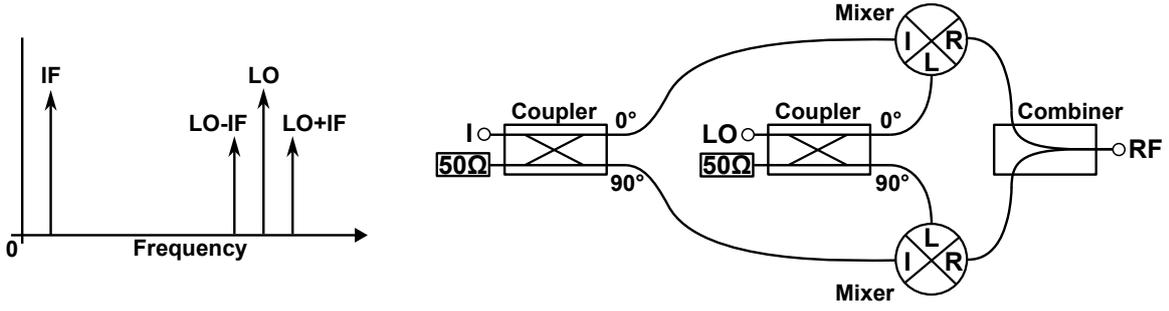


Figure 3.6: Single-sideband (SSB) mixing used for qubit control. **Left** Ordinary (double-sideband) mixing produces both the lower ($LO-IF$) and upper ($LO+IF$) sidebands around the LO. **Right** In an SSB architecture, the LO and IF are each split by 90° hybrids to drive two mixer cores in quadrature; their RF outputs are then recombined so that one sideband interferes destructively while the other adds constructively.

Principle of Single-Sideband (SSB) Mixing

Single-sideband (SSB) mixing suppresses one of the two image sidebands so that, ideally, only the upper sideband (USB) or the lower sideband (LSB) remains. In our setup the SSB mixer is configured to select the LSB. In practice, an SSB mixer comprises two mixer cores operating in parallel (Fig. 3.6 right). The local-oscillator (LO) signal is split into two copies with a 90° phase shift—one per mixer. Likewise, the baseband or intermediate-frequency (IF) signal is split into two copies that are also 90° out of phase, forming the in-phase (I) and quadrature (Q) channels. Each mixer multiplies its LO input with its respective IF input, and the two outputs are then recombined such that one sideband interferes destructively while the other adds constructively [121, 122].

Feed the LO into a 90° hybrid coupler to obtain

$$LO_I = \cos(\omega_{LO}t), \quad LO_Q = \sin(\omega_{LO}t),$$

and split the IF likewise into

$$\text{USB choice: } IF_I = \cos(\omega_{IF}t), \quad IF_Q = -\sin(\omega_{IF}t),$$

$$\text{LSB choice: } IF_I = \cos(\omega_{IF}t), \quad IF_Q = +\sin(\omega_{IF}t).$$

with the choice determining which sideband is selected. This depends on which 90° hybrid output port is fed into the I and Q mixers.

Each mixer core multiplies its LO and IF inputs:

$$\begin{aligned} V_I &\propto \cos(\omega_{LO}t) \cos(\omega_{IF}t) \\ &= \frac{1}{2} [\cos((\omega_{LO} - \omega_{IF})t) + \cos((\omega_{LO} + \omega_{IF})t)], \\ V_Q &\propto \sin(\omega_{LO}t) (\mp \sin(\omega_{IF}t)) \\ &= \mp \frac{1}{2} [\cos((\omega_{LO} - \omega_{IF})t) - \cos((\omega_{LO} + \omega_{IF})t)], \end{aligned}$$

where the upper/lower sign corresponds to the USB/LSB choice of IF phase above.

Finally, recombining the two mixer outputs yields

$$V_{\text{RF}} \propto V_I + V_Q \propto \cos((\omega_{\text{LO}} \pm \omega_{\text{IF}})t),$$

so that, in the ideal case, the desired sideband (USB or LSB) remains while the unwanted sideband is fully cancelled.

3.6 Pulse Shaping and Bias-Tee Compensation

In our setup, a resistive bias-tee allows fast AC manipulation to be superimposed on a static gate voltage. While bias-tees are commonly implemented with an inductor in the DC path, we employ a resistive feed instead. An inductor placed close to the sample would experience continuously changing high magnetic fields during sensing experiments and, by Faraday's law, develop unwanted induced voltages. In the implementation shown in the left panel of Fig. 3.7, the series coupling capacitor $C = 15 \text{ nF}$ and the DC feed resistor $R = 7 \text{ M}\Omega$ form a first-order high-pass as seen by the RF drive, with time constant

$$\tau = RC \approx 105 \text{ ms}.$$

An uncompensated step applied at the RF input produces an exponential response at the gate:

$$V_{\text{gate}}(t) = V_{\text{step}} e^{-t/\tau},$$

i.e., the bias-tee differentiates slow components and removes DC [123]. Consequently, a nominally flat (square) pulse at the RF port will relax towards baseline at the gate on the timescale τ . To restore a square gate waveform of duration T , we pre-emphasise the drive with a small linear term,

$$V_{\text{RF}}(t) = V_0 (1 + \alpha t),$$

where α is tuned experimentally. This pre-distortion flattens the top of $V_{\text{gate}}(t)$, but it also means that V_{RF} and V_{gate} no longer return to zero simultaneously. Simply forcing V_{RF} to 0 V after the voltage step would induce a negative transient at the gate due to the high-pass coupling.

To end the pulse with both V_{RF} and V_{gate} at 0 V, the gate waveform must satisfy a zero-area (zero-DC) condition,

$$\int V_{\text{gate}}(t) dt = 0.$$

A straightforward way to enforce this is to append a opposite-polarity (inverse) segment after the main pulse so that the net area cancels. In practice, using a briefer inverse segment with a proportionally larger amplitude resets the baseline faster, enabling back-to-back pulse sequences without cumulative offsets.

The effect is illustrated in the right panel of Fig. 3.7. The target is a 1 mV square gate pulse of 200 ms (blue). The required RF drive (orange) includes (i) a small linear pre-emphasis to flatten the top and (ii) a inverse section to satisfy the zero-area condition. Both traces

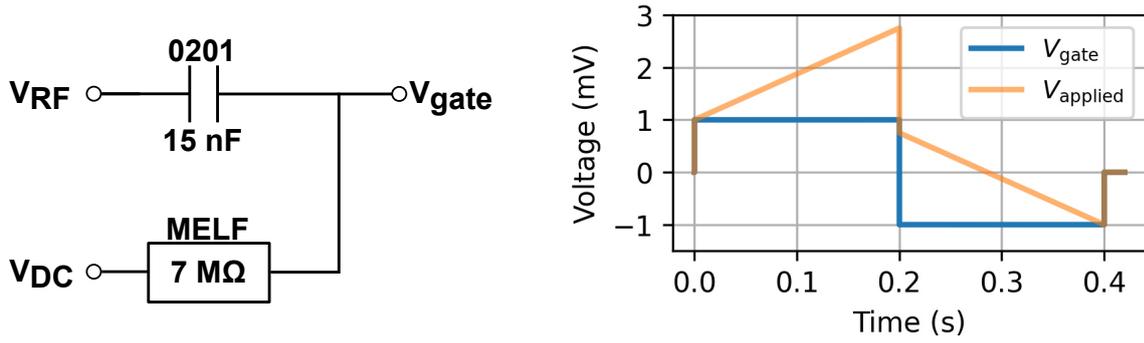


Figure 3.7: Bias-tee on the standard PCB and compensation principle. **Left** Resistive bias-tee used in our setup. The RF path is capacitively coupled to the gate through a 15 nF 0201 capacitor, while the DC bias is supplied via a 7 M Ω MELF resistor. **Right** A shaped voltage pulse applied at V_{RF} (orange) produces an approximately square pulse at the device gate (blue). The gate waveform is constructed to have zero time-integral. Data are from an LTspice simulation with typical parameters used in our experiments.

return to 0 V at the end of the sequence, allowing sequential operation. Without this inverse section, the bias-tee could introduce slow drifts of the quantum-dot potentials and reduce readout fidelity [124].

The choice of τ is a compromise. Larger τ (increased R and/or C) reduces the needed pre-emphasis α and preserves RF headroom which is a useful consideration because external drive voltages must be $\sim 44\times$ larger to overcome the -33 dB cryogenic attenuation. However, a larger τ slows DC operations: a DC step applied through the bias-tee follows $V_{\text{gate}}(t) = V_{\text{step}} [1 - e^{-t/\tau}]$, reaching 95% of its value only after 3τ (here $3\tau \approx 315$ ms). In practice, coarse DC ramps are performed faster (we seek qualitative features rather than exact voltages), whereas for precise setpoints we wait $\geq 5\tau$.

Experimental Determination of the Compensation

The compensation factor α is calibrated separately for each bias-tee, as component tolerances and mK-temperature behaviour differ. We track a sharp charge transition over time (up to 200 ms) using a square readout pulse with linear pre-emphasis, $V_{RF}(t) = V_0 (1 + \alpha t)$. During calibration we *omit* the inverse segment to avoid biasing the result; applying an incorrect α could otherwise produce charge accumulation. Instead, after each readout window we set the RF drive to 0 V and insert a delay of 5τ (here ≈ 525 ms) to allow $> 99\%$ discharge before the next iteration.

Figure 3.8 compares a poor calibration (left), where the charge transition drifts during the 100 ms trace, with a good calibration (right), where the working point is stable over 200 ms. Further details on the readout protocol used are given in Sec. 4.5.

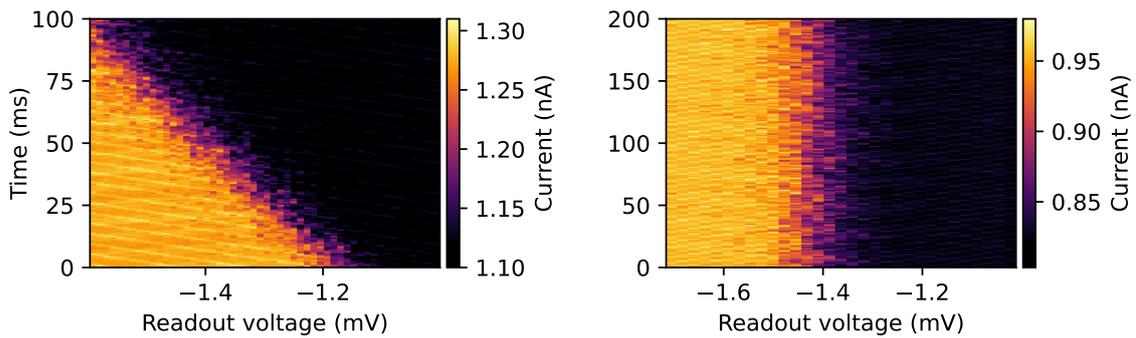


Figure 3.8: Experimental bias-tee compensation calibration. The device working point is stepped from 0 V to various readout voltages; the time-dependent current is colour-coded. A constant α is applied to the readout pulses. After each acquisition the voltage returns to 0 V, and the next iteration is delayed by 5τ to fully discharge the bias-tee. Ten averages are shown. **Left** Incorrect α leads to a systematic drift of the charge transition during the 100 ms window. **Right** Proper calibration stabilises the working point over 200 ms.

Pulse Construction Recipe

- Enforce a zero-DC condition at the gate: $\int V_{\text{gate}}(t) dt = 0$.
- Ensure any additional inverse pulse does not interfere with qubit operations.
- Instantaneous (fast) edges transmit through the bias-tee; slow (DC) components must be pre-emphasised.
- Use a linear pre-emphasis: $V_{\text{RF}}(t) = V_0 (1 + \alpha t)$; calibrate α per channel.
- During α calibration, insert long delays ($\geq 5\tau$) between averages to guarantee full discharge of the bias-tee.

3.7 FPGA-Based Control and Data Acquisition (ADwin Pro II)

Qubit DC control and readout were executed with an ADwin-Pro II system fitted with a Pro-CPU-T12 processor module [125]. The T12 is built around a Xilinx Zynq system-on-chip (SoC) that combines a dual-core ARM Cortex-A9 CPU (1 GHz) with a field-programmable gate array (FPGA) fabric (“programmable logic”). An FPGA is a reconfigurable digital device that implements logic directly in hardware (counters, state machines, timers, direct memory access engines, etc.) and therefore provides deterministic, low-jitter timing that is independent of the host PC operating system. In our setup, hard real-time sequencing runs on the T12 (FPGA/CPU) with 1 GB on-board memory for buffering.

Fast digitisation is provided by a Pro II-AIn-F-4/16 card (16-bit resolution, ± 10 V range), whose per-channel analogue-to-digital converters (ADCs) operate up to 4 MS/s with an analogue bandwidth of roughly 0-600 kHz and large on-module burst memory (~ 256 MB) for continuous high-rate capture. Analogue control voltages are generated with 16-bit digital-to-analogue converter (DAC) output boards (± 10 V range, μs -scale settling), which

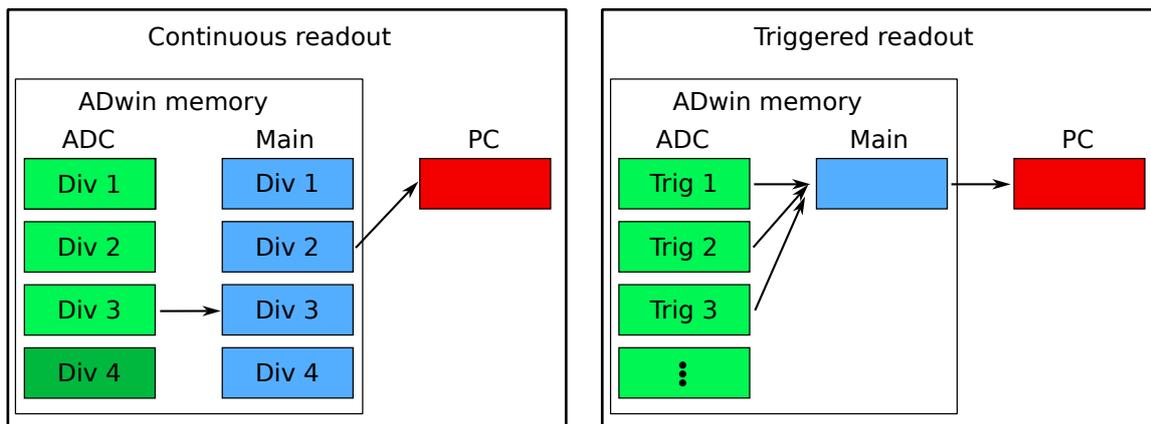


Figure 3.9: ADwin readout implementation. **Left** Continuous mode. Both the ADC buffer and main memory are divided into four segments (Div 1-4). The ADC currently writes to Div 4 (dark green) and wraps cyclically. Using status flags, the process copies the oldest safe segment (here Div 3) to the corresponding main-memory segment (blue), while another main-memory segment (Div 2) is concurrently transferred to the PC. The total rolling buffer spans 1 s. **Right** Triggered burst mode. Each external trigger from the AWG (within a pulse train) initiates a burst that is written contiguously in ADC memory. After the train completes, the AWG pauses while the ADwin drains the ADC to main memory and then to the PC. Averaging and post-processing of the 4 MS/s data are performed on the PC while the next train is being recorded.

connect to the Pro bus for synchronous, hardware-timed updates. Real-time experiment logic is implemented in ADbasic processes that coordinate triggering, buffered acquisition, and voltage setting. On the host side, Python drivers wrap these ADbasic processes and present an interface to the Spin Suite; in particular, the real-time engine `spin_excite` that orchestrates parallel execution across instruments, and streams buffered data during runtime [116]. This combination of ADbasic for deterministic timing and Python/Spin Suite for high-level experiment description and data handling provides robust, hardware-timed control together with flexible scripting.

Within this thesis, the DC output control originally developed by Sven Fuhrmann [126] was restructured and optimised. In addition, time-aligned input sampling was implemented and used for spin-qubit readout with DC digitisation at up to 4 MS/s. The complete ADwin Pro II driver stack, including the Python interface and the ADbasic programs running on the FPGA system, is available on GitHub [127]. The following sections briefly describe the time-domain processes used for the 4 MHz DC readout, and present two software/hardware approaches for increasing the effective DAC voltage resolution of the ADwin system.

Continuous Data Acquisition

Continuous acquisition is used primarily for noise spectroscopy, where a long time trace is Fourier transformed and the resulting spectrum analysed (see Sec. 4.2.3). The AIn-F-4/16 input card supports on-board averaging of the raw 4 MS/s stream, allowing effective output rates of 2 kS/s, 10 kS/s, 100 kS/s, 1 MS/s, or the full 4 MS/s.

Table 3.1: ADwin process overview. “Continuous” indicates a periodically repeating process with a fixed time window; “External” indicates triggering by an external signal (e.g. from the AWG). The rightmost column lists the associated Spin Suite modules [116] used within the qkit measurement framework [120].

No.	Job	Priority	Trigger	Spin Suite
1	Outputs: Ramping (parallel), Oversampling; Inputs: single samples (averaged)	Lowest, always running	Continuous	Spin_tune (sync)
2	Inputs: Continuous readout	High, only started when needed	Continuous	Spin_watch (async)
3	Inputs: Burst measurement	High, only started when needed	External	Spin_excite (sync'd async)
4	Readout of process 3 after one average	Low, only started when needed	Continuous	Spin_excite (sync'd async)

A practical challenge is avoiding read-write contention while data are being continuously acquired, moved to main memory, and transferred to the host. We therefore partition both the ADC buffer and the ADwin main memory into four equal segments that operate as a circular pipeline (Fig. 3.9, left). The writer pointer advances segment-by-segment; a segment can be read only if the writer is not currently on that segment and not on the subsequent one. Hence, at any time two segments are safe to read and ship to the next stage. With a total buffer depth corresponding to 1 s of data, ideal, contention-free transfers occur every 0.25 s.

Triggered Burst Data Acquisition

All time-domain measurements presented in Chs. 4 to 6 use a burst readout that acquires time-aligned, externally triggered data at up to 4 MS/s for pulse trains as long as ~ 25 s (Fig. 3.9, right). A digital trigger from the AWG to the ADwin starts each readout window; a typical pulse train contains < 1000 such windows. Once a train finishes, the AWG idles while the ADwin drains the ADC buffer to main memory and then transfers the complete dataset to the PC. On the host, traces are averaged (or filtered) to customary effective rates (e.g. 100 kS/s); the PC signals completion, upon which the ADwin re-arms and triggers the AWG to launch the next train.

The control handshake is thus twofold: the ADwin initiates new pulse trains by triggering the AWG, whereas the AWG issues intra-train triggers that mark the readout windows. To decorrelate the acquisition from periodic disturbances and reduce systematic timing biases, the start time of each train (i.e. each average in a pulse-parameter sweep) is randomly dithered by the ADwin within a window of 10 ns to 1 ms.

An overview of the processes implemented in this thesis and running in parallel on the ADwin system is given in Table 3.1.

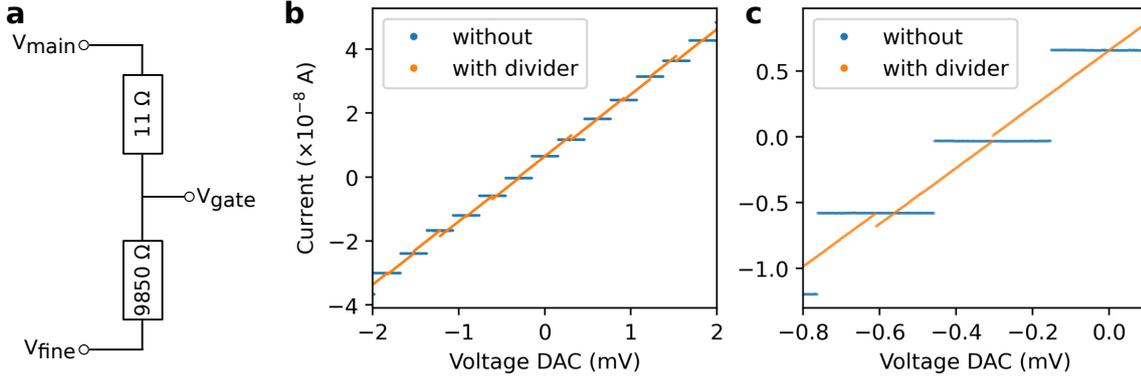


Figure 3.10: Two-DAC voltage divider. **a** Schematic: a resistive divider (illustrated with $11\ \Omega$ and $9.85\ \text{k}\Omega$) combines two ADwin DAC outputs. V_{main} sets the coarse level and V_{fine} provides sub-LSB trimming. The resulting gate voltage follows Eq. (3.7). **b,c** Measured DAC current transfer (through $50\ \text{k}\Omega$, IV-converter, integration time $0.1\ \text{s}$): raw 16-bit steps (blue) versus the divider-enhanced resolution (orange). Unequal DAC step spacing can cause slight overlaps when changing V_{main} , hence V_{main} should be kept constant during fine sweeps while changing V_{fine} over a broad range.

Voltage Divider Implementation

The ADwin 16-bit DACs provide a $\pm 10\ \text{V}$ range, i.e. a least significant bit (LSB) of $\Delta V_{\text{DAC}} = 20\ \text{V}/2^{16} \approx 305\ \mu\text{V}$. For experiments that rely purely on precise DC control, we increase the effective voltage resolution at the device without reducing the available amplitude by combining two DACs through a resistive divider placed between their outputs (rather than to ground). The gate node then sees a weighted average of the two DAC voltages,

$$V_{\text{gate}} = \frac{V_{\text{main}} R_{\text{fine}} + V_{\text{fine}} R_{\text{main}}}{R_{\text{main}} + R_{\text{fine}}}, \quad (3.7)$$

so that increments on V_{fine} are attenuated by the factor $R_{\text{main}}/(R_{\text{main}} + R_{\text{fine}})$ (sub-LSB trimming), while the full gate span remains near $\pm 10\ \text{V}$ as V_{main} sweeps.

In our implementation we use $R_{\text{main}} = 11\ \Omega$ and $R_{\text{fine}} = 9.85\ \text{k}\Omega$, giving an attenuation factor of $\approx 1/896$ for V_{fine} . Thus the fine step at the gate is $\Delta V_{\text{gate}} \approx 305\ \mu\text{V}/896 \approx 0.34\ \mu\text{V}$. Real DACs exhibit nonlinearity meaning their raw step sizes are not perfectly uniform; to avoid step overlap artefacts, V_{main} is held fixed while V_{fine} is swept over the desired range (see Fig. 3.10b,c).

The divider presently resides at room temperature; ideally it should be mounted inside the cryostat (4 K or mK stage, subject to heat load) to reduce Johnson-Nyquist noise, which scales as $V_{\text{RMS}} \propto \sqrt{RT}$ [128, 129]. Cooling-power constraints (watts at 4 K, $< 1\ \mu\text{W}$ at mK in a Sionludi [130]) limit the allowable dissipation of resistive networks at the cold stages.

Implementation Notes

- **SNR:** Lower resistance leads to lower pickup and Johnson-Nyquist noise ($V_{\text{RMS}} \propto \sqrt{R}$).

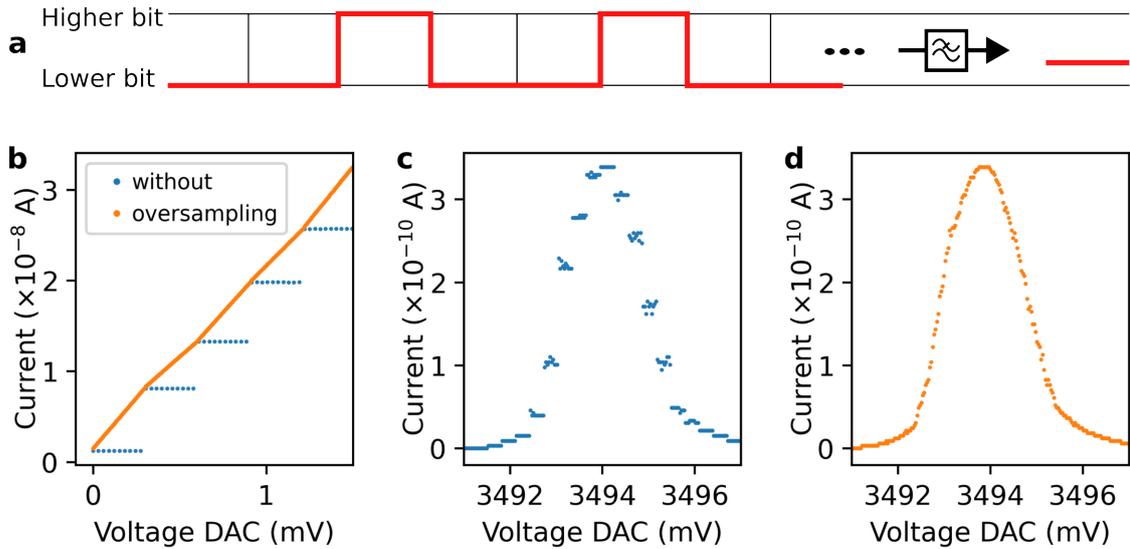


Figure 3.11: Oversampling implementation. **a** Principle: the DAC output voltage toggles between two adjacent bits in time segments (here: two “low” segments and one “high” per period). After the two-stage low-pass (cutoff ~ 250 Hz), the sample sees the time average, i.e. a stable voltage between the two bits. **b-d** Measurement with a $50\text{ k}\Omega$ load and IV conversion (integration time 0.1 s): raw 16-bit steps (blue) versus oversampled outputs (orange). Slope variations due to non-linearities in individual bit steps. **c** Coulomb oscillation of an SET without oversampling shows discrete plateaus; **d** the same sweep with oversampling factor 20 eliminates visible plateaus.

- **DAC current:** The total divider resistance is $R_{\Sigma} = R_{\text{main}} + R_{\text{fine}}$. The highest current $I_{\text{max}} = \Delta V / R_{\Sigma}$ has to be supplied by the DAC. In our case $I_{\text{max}} \approx 2.0\text{ mA}$, comfortably below the ADwin’s recommended 5 mA .
- **Power dissipation:** The divider dissipates $P = \frac{(V_{\text{main}} - V_{\text{fine}})^2}{R_{\Sigma}}$; verify thermal margins, especially if installed cold.
- **Calibration:** Measure $V_{\text{gate}}(V_{\text{main}}, V_{\text{fine}})$ once to extract effective gains and correct for resistor tolerances and DAC non-linearities; in operation, keep V_{main} fixed and step V_{fine} (use ≥ 10 fine-LSBs per target step to average out non-linearities).

Oversampling Method

Oversampling of the ADwin DAC outputs provides a purely software based implementation to increase effective voltage resolution at the device node. The DAC is toggled at high rate (up to 500 kHz) between two adjacent voltage bits, and the two-stage low-pass on the gate line ($\sim 250\text{ Hz}$ cutoff) averages this pulse-density-modulated waveform to a DC level between the bit steps (Fig. 3.11). Let V_L and V_H be the voltages of the two neighbouring DAC steps and let the duty cycle spent at V_H be $D \in [0, 1]$. Provided the toggle frequency is much higher than the line cutoff, the gate sees

$$V_{\text{gate}} \approx (1 - D) V_L + D V_H,$$

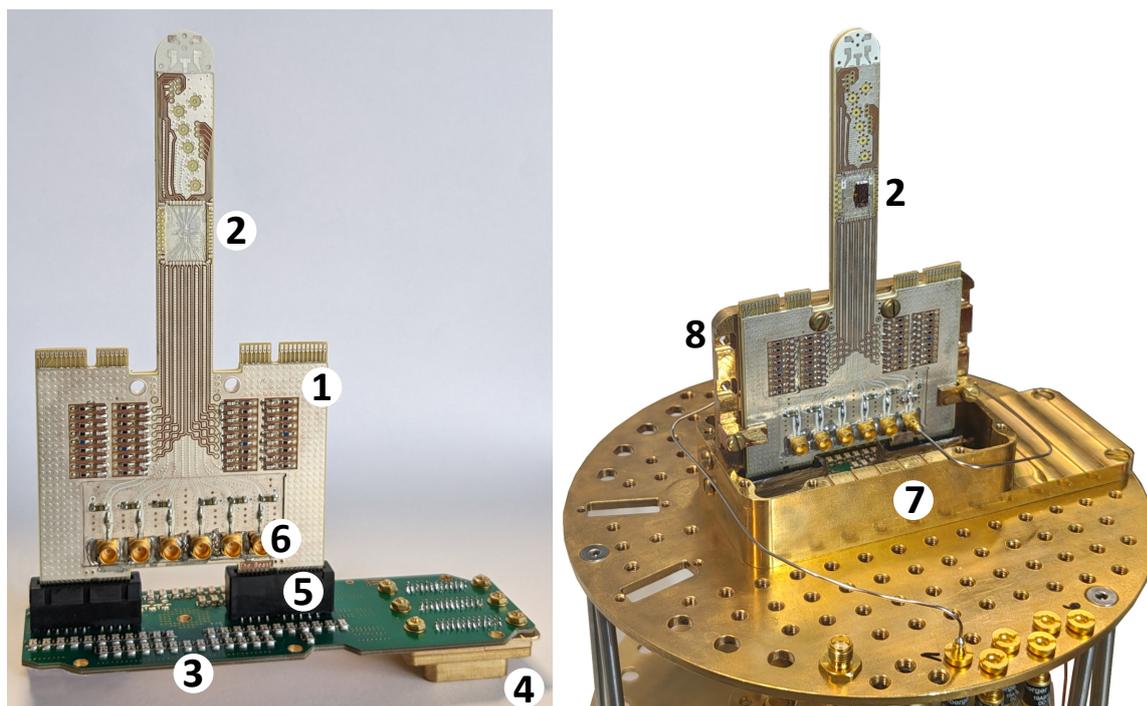


Figure 3.12: The Beast and Connector Board. **Left** (1) The Beast PCB; (2) chip area; (3) connector board with first low-pass filtering stage; (4) micro-sub-D connector to the DC lines in the cryostat; (5) PCIe connector for DC contact between The Beast and the connector board; (6) SMP connector for RF contact. **Right** A chip (2) glued to The Beast and mounted on the mK stage of the cryostat. Two RF cables are shown installed as an example. The lower mK shield (7) is in place, while the upper part of the shield is removed to reveal the assembly. A resting frame (8) provides both mechanical stability and thermal anchoring for The Beast.

so that with a duty-cycle granularity of $1/N$ (“oversampling factor” N) the minimum increment is approximately $(V_H - V_L)/N \approx \text{LSB}/N$. Compared to a hardware voltage divider, this method yields unique (non-overlapping) voltages and requires no additional components; in practice, small residual slope changes remain due to DAC non-linearities, as visible in Fig. 3.11b.

3.8 Novel Printed Circuit Board: The Beast

From the outset it was clear that the three RF lines available on the standard PCB (see Sec. 3.2) are insufficient for advanced silicon spin-qubit experiments. With one RF line typically reserved for the Rabi drive, only two gates remain for fast pulsing which is too restrictive for double-dot operation, where two plunger gates and the interdot barrier ideally require independent, high-bandwidth control [131]. To overcome this limitation, we developed a more capable interface PCB, referred to as *The Beast* which was made by the Swiss PCB manufacturer Optiprint [132].

It provides up to 69 DC contacts and 12 RF contacts to the chip, integrates dedicated bias-tee locations and tank-circuit nodes for RF reflectometry [133] directly on the board,

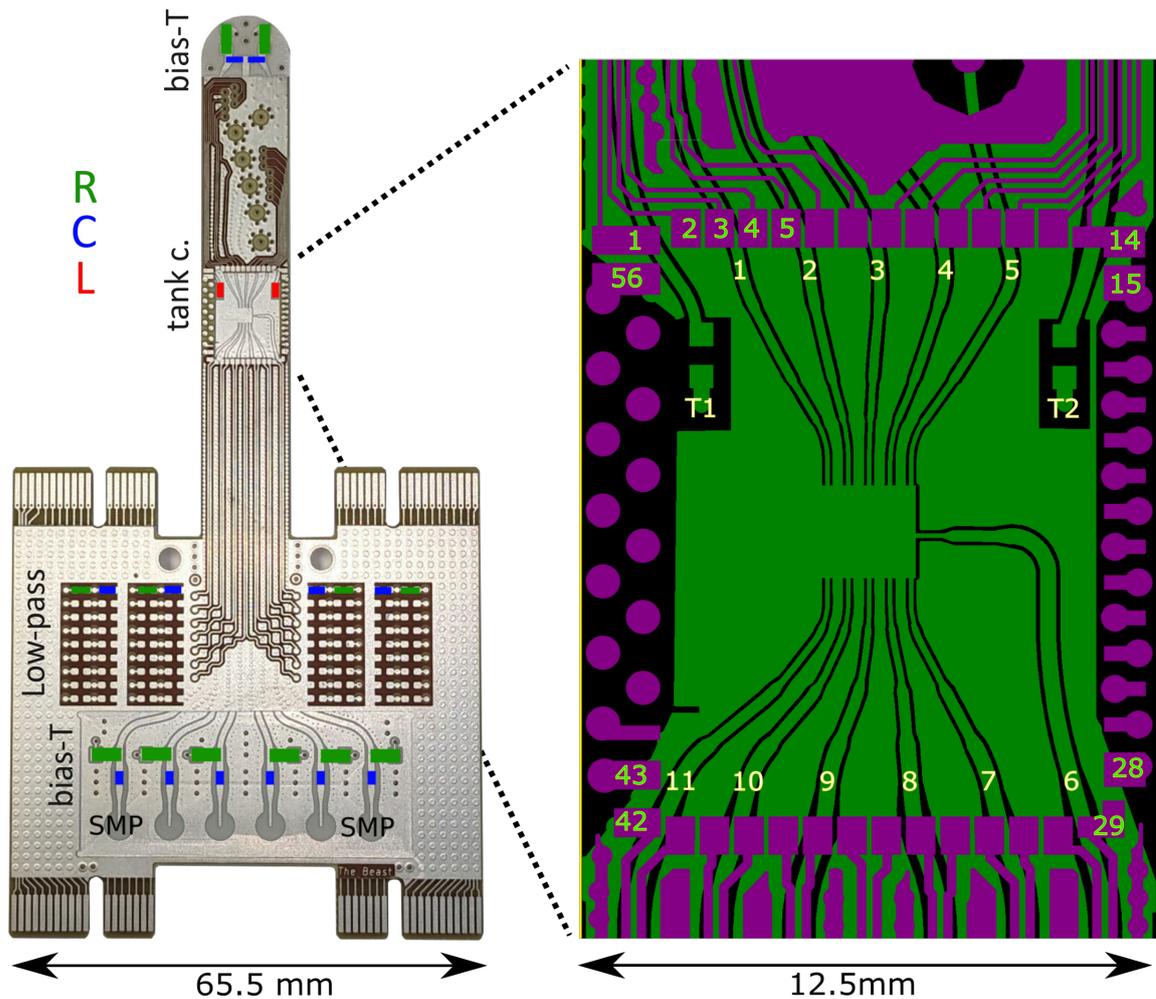


Figure 3.13: Surface-mount components and arrangement of DC/RF lines at the sample space. **Left** Solder locations for SMD components forming the second low-pass filter stage and the on-board bias-tees. Positions of the SMP connectors and the inductor for the tank circuit are indicated. The two bias-tees of the tank-circuit readout are placed at the top edge of the PCB. **Right** Zoom of the sample area. The top layer provides DC bond pads (purple) on all sides. The RF layer beneath (green) routes eleven CPWs that converge centrally and are initially shorted to ground (opened during sample integration). Labels T1 and T2 mark the pads for the SMD inductors of the tank-circuit readout.

and includes a second-stage DC low-pass filter network. Wire bonds connect the chip pads to the surrounding DC/RF pads; to facilitate reliable bonding, the top copper is finished with an immersion silver coating. The central chip region accommodates devices up to $9 \times 10 \text{ mm}^2$.

RF connectivity is provided via twelve SMP connectors distributed on the front and back sides of The Beast, feeding coplanar waveguides (CPWs) that route microwave signals to the sample region. DC connectivity is made through two 36-pin PCIe connectors that mate to a dedicated connector board (Fig. 3.12), which (i) hosts the first low-pass filtering stage, (ii) interfaces to the cryostat wiring via a Micro-D connector, and (iii) is rigidly fastened to

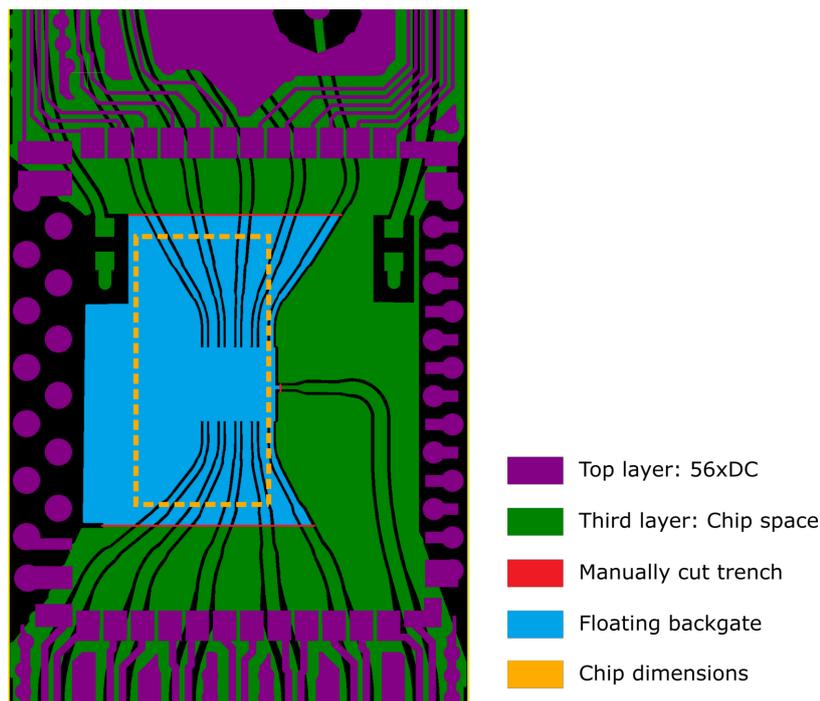


Figure 3.14: Chip placement on The Beast. CPW terminations are created at positions compatible with the chip size by manually cutting the CPW centre conductor with a scalpel. A global back-gate (blue) can be realised by cutting a trench along the red lines to break the local ground connection.

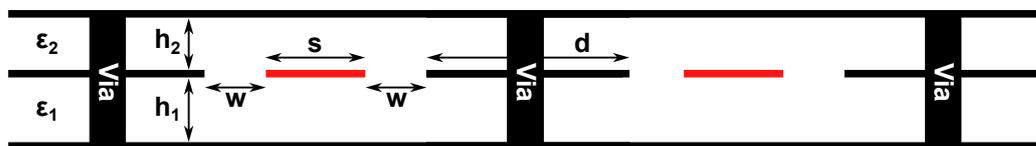


Figure 3.15: Schematic of the coplanar waveguides (CPWs) used on The Beast. Two laterally spaced CPWs (centre conductors in red) are sandwiched between two ground planes (black; one above, one below). The dielectric above the CPW has thickness h_2 and relative permittivity ϵ_2 ; the dielectric below has h_1 , ϵ_1 . Each CPW uses centre conductor width s and gap w . The separation between the two CPWs is d . All ground planes are stitched together with a via fence to maintain a low-impedance reference.

the mK plate for effective thermal anchoring of the DC lines. A complete schematic of the DC path is provided in Fig. B.5 (Appendix).

This expanded I/O capability makes The Beast particularly suited to experiments requiring multiple high-frequency control channels together with a large number of DC biases, while maintaining the mechanical robustness and thermal performance needed for stable millikelvin operation.

Fabrication details (layer stack, networks, storage/handling, breakout PCB, pin assignment, and routing) are compiled in Appendix B.

Microwave lines

In our setup all microwave transmission lines are designed for a characteristic impedance $Z_0 = 50 \Omega$. When a load Z_L is connected, the voltage reflection coefficient Γ is

$$\Gamma = \frac{V^-}{V^+} = \frac{Z_L - Z_0}{Z_L + Z_0},$$

and, for a lossless line, the time-averaged power delivered to the load is

$$P_L = \frac{|V^+|^2}{Z_0} (1 - |\Gamma|^2),$$

using the standard travelling-wave convention [121]. Thus, matching the on-board CPWs to $Z_L = 50 \Omega$ ($\Gamma = 0$) eliminates reflections and maximises power transfer from the cryogenic coax to the chip.

To achieve the matching condition on The Beast, we tune the CPW dimensions (s, w, h_1, h_2) for the chosen dielectrics (ϵ_1, ϵ_2) (see Fig. 3.15). For the inner-core routing of the board we used $\epsilon_1 = \epsilon_2 = 3.5$, $h_1 = 0.42$ mm, $h_2 = 0.30$ mm, $s = 0.35$ mm and $w = 0.19$ mm, yielding a calculated impedance of approximately 50Ω . In other regions of the PCB the CPW is tapered to accommodate launches and fan-outs while preserving the target impedance. The calculation method for a CPW sandwiched between ground planes is given in Sec. D.1.

To minimise crosstalk between neighbouring CPWs, we keep the ground region between them sufficiently wide. A practical rule-of-thumb is $d \geq 2s + w$ for the ground sheet between two CPWs (see Fig. 3.15) [134].

Via stitching

Vias which are mechanically drilled holes that are subsequently metallised (galvanically plated) on the inside to form a conductive barrel, provide vertical interconnects between copper layers (through, blind, or buried). In multilayer RF PCBs, these vias must stitch all ground planes together densely to avoid parasitic stubs and unintended propagation paths. We therefore implement (i) longitudinal via fences flanking each CPW, and (ii) transverse stitching to connect all ground planes. The two opposite via fences form a narrow “channel” that behaves as a short rectangular waveguide. The dominant TE_{10} mode—i.e., the lowest-order transverse-electric waveguide mode whose cutoff is set primarily by the fence-to-fence spacing a —must remain below cutoff. Accordingly, we choose

$$a < \frac{\lambda}{2} = \frac{c}{2f\sqrt{\epsilon_r}},$$

with λ the wavelength in the dielectric and f the highest frequency of interest [121]. For our target up to 40 GHz in $\epsilon_r = 3.5$, $\lambda \approx 4.0$ mm, so $a < 2.0$ mm. Along each fence, the via-to-via pitch is kept well below a wavelength (in dielectric) to maintain a solid RF ground; as a design guideline we use a spacing $\lesssim \lambda/8$ (here ≈ 0.5 mm) [135]. In the implemented layout this results in >3000 stitching vias distributed around the CPWs and ground planes.

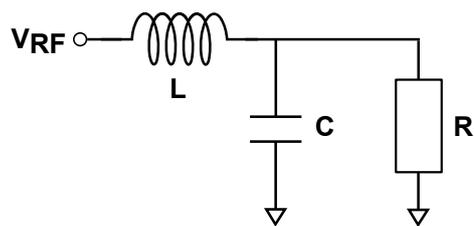


Figure 3.16: Exemplary lumped-element tank circuit used for RF-reflectometry readout. The device contributes a capacitance C (and a resistance R), while a discrete inductor L is placed close to the chip to minimise stray capacitance parallel to C .

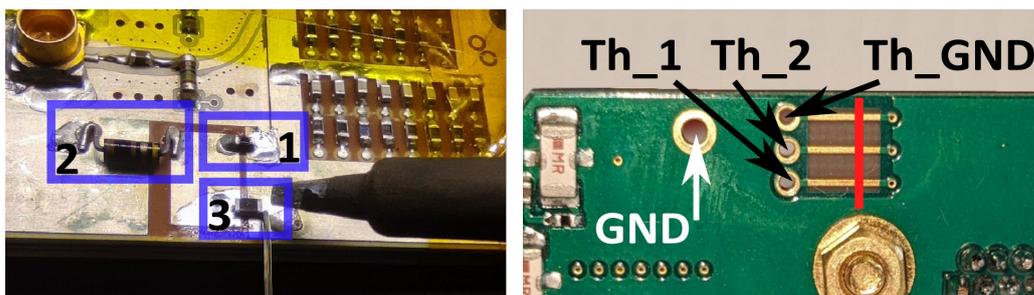


Figure 3.17: Thermometers mounted on The Beast. **Left** Three thermometers are connected in parallel from a central pad to ground: (1) Panasonic SC $22\ \Omega$, (2) Allen-Bradley C $100\ \Omega$, and (3) Vishay RuOx $1\ \text{k}\Omega$. The RuOx sensor is sensitive to excessive heating and is therefore soldered last, using the minimal possible thermal budget. **Right** Connector board showing the three DC lines dedicated to thermometry. By default, these lines are connected to the DC wiring of the cryostat. If the connection is interrupted at the location marked by the red line, the through-holes indicated by the black and white arrows can be used to attach separate thermometry wiring.

Tank Circuit

The Beast provides dedicated landing pads (Fig. 3.13, T1/T2) to mount SMD inductors and realise on-board LC resonators for RF reflectometry. A single $50\ \Omega$ feedline can service two independent tank circuits by frequency multiplexing: each device gate is connected to the common RF line through its own resonator tuned to a distinct resonance $f_0^{(i)}$, while DC biasing of each gate is supplied through its own resistive bias-tee.

Changes in the device's effective capacitance or resistance alter the resonator impedance (equivalently its resonance frequency and linewidth), producing a measurable change in the amplitude and phase of the RF signal reflected.

This architecture supports μs -scale integration times for spin readout [136], detailed design strategies are reviewed in [137].

Thermometry

The Beast includes a compact, low-thermal-mass on-board thermometer formed by wiring three thermistors in parallel. Using three dedicated DC lines plus ground enables a four-wire (Kelvin) measurement to eliminate lead resistance. Each element has a monotonic,

negative temperature coefficient $R_i(T)$ (NTC behaviour), but with a different useful range. By choosing room-temperature resistances that differ by orders of magnitude, the total parallel resistance

$$R_{\text{tot}}(T) = \left(\frac{1}{R_1(T)} + \frac{1}{R_2(T)} + \frac{1}{R_3(T)} \right)^{-1}$$

is dominated by the thermistor with the strongest T -dependence in the current temperature window, providing useful sensitivity from room temperature down to the millikelvin range. The three sensors and their typical ranges are:

- Panasonic SC 22 Ω : ~ 150 -400 K,
- Allen-Bradley C 100 Ω : ~ 1 -150 K,
- Vishay RuO_x 1 k Ω : ~ 10 mK-1 K.

A practical calibration procedure is to log R_{tot} throughout a controlled cool-down (and warm-up) while measuring temperature with a well-calibrated reference thermometer, and then fit the resulting data to obtain the mapping $T(R_{\text{tot}})$.

The three dedicated DC lines can alternatively be repurposed to drive a laser diode (e.g., for chip reset/illumination) or a micro-heater at the sample. Accordingly, the on-board thermometry traces are implemented with increased width to accommodate the higher currents required for these functions.

4 The SiMOS Spin Qubit: Definition and Characterisation

A central goal of this work is the integration of a gate-defined silicon spin qubit with a single-molecule magnet (SMM) into a hybrid quantum platform. As a prerequisite, the spin qubit must be reliably defined, tuned, and read out with high fidelity. This chapter therefore introduces the device layout and gate-stack design of the SiMOS spin-qubit structures used in this thesis and reviews the essential operating principles. The presentation begins with the lateral gate geometry and the elements of silicon band structure relevant to qubit definition. The charge-sensing toolbox based on a proximal single-electron transistor (SET) is then established, including Coulomb-diamond characterisation and charge-noise spectroscopy, followed by quantitative procedures for extracting gate lever arms and the electron temperature.

The second part concentrates on spin physics: the spin-selective readout mechanism and its analysis in single-shot “blip” traces are introduced. The longitudinal relaxation time T_1 is quantified and coherent control using electron spin resonance (ESR) and rapid adiabatic passage (RAP) for robust state transfer are described. The chapter concludes with Ramsey and Hahn-echo spectroscopy to determine the inhomogeneous dephasing time of the qubit which requires precise state preparation, phase-accurate manipulation, and faithful single-shot readout.

Two SiMOS devices from successive fabrication generations are studied, denoted devices A and B (Tab. 4.1). In both devices the lateral gate geometry is essentially unchanged while the key process improvement at imec is an increased thermal-oxide thickness from 8 nm in device A to 20 nm in device B (Tab. C.1) which was implemented to reduce strain at the ^{29}Si - SiO_x interface and thereby improve interface quality. The basic charge characterisation presented in this chapter uses data from device B, reflecting its improved tunability and immediate relevance. A complementary assessment of the charge-noise spectrum for device A is provided in Appendix C.2. Coherent single-qubit control experiments are shown of device A. Both devices were fabricated on nominally identical silicon substrates with a residual ^{29}Si concentration of approximately 400 ppm.

4.1 Device Layout

The term metal-oxide-semiconductor (MOS) derives from the field-effect transistor architecture widely used in integrated circuits [59]. In practice, the “metal” gate electrode

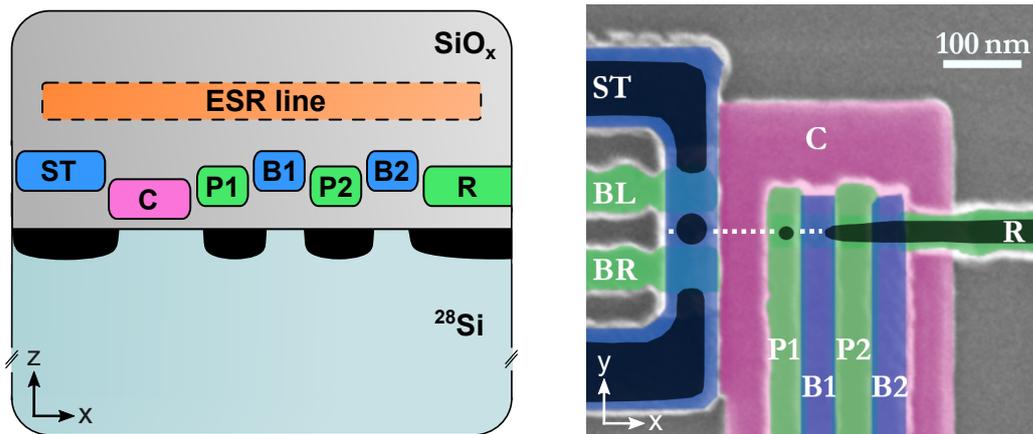


Figure 4.1: Device layout and heterostructure of the SiMOS spin-qubit device. **Left** Schematic cross-section (taken along the dotted white line in the right panel) showing three overlapping gate layers on SiO_x above isotopically enriched ^{28}Si . Under operating positive bias, electron accumulation regions (black) form at the Si/SiO_x interface, defining a quantum dot beneath the plungers P1 and P2 and an adjacent electron reservoir beneath R. B1 and B2 tune the tunnel couplings to the reservoir and across the interdot barrier, respectively. An on-chip ESR line above the oxide provides a transverse radio-frequency magnetic field for spin control. **Right** False-colour SEM image (top view) of the gate pattern. ST, BL and BR form the proximal SET charge sensor. C is a confinement gate that electrostatically separates the sensor from the qubit region and provides lateral confinement along y . The dotted white line marks the cross-section shown at left.

may be realised by highly doped polycrystalline silicon rather than a true metal, both in conventional MOSFETs and in the devices considered here.

The qubit heterostructure comprises a natural-silicon substrate capped by an epitaxial ^{28}Si layer of thickness 100 nm with a residual ^{29}Si concentration of ≈ 400 ppm. A thermally grown SiO_x layer provides vertical insulation and forms the MOS interface. Highly doped polysilicon gates patterned in three overlapping layers control the electrostatics. Details of electron accumulation at the interface are discussed in Sec. 2.1.

Figure 4.1 combines a cross-sectional schematic and a false-colored top-down scanning-electron-microscope (SEM) image of the gate layout. Precise layer heights of the individual polysilicon and SiO_x layers are given in Tab. C.1. The device separates into two functional regions: an SET used for charge sensing, and a double-quantum-dot (DQD) region used for qubit operation. These regions are electrostatically isolated by the centre confinement gate C (pink). Although the DQD geometry supports two dots, throughout this chapter the left dot is operated while the right dot is merged with the reservoir.

The SET consists of a top gate ST that controls its electrochemical potential and two barrier gates BL and BR that define the tunnel barriers to source and drain. The DQD region comprises the confinement gate C, two plunger gates P1 and P2, two barrier gates B1 and B2, and a reservoir gate R. The electron reservoir accumulated under R serves as an electron source and drain for the dots. The plunger/barrier array in the qubit region has a gate pitch of 80 nm (centre-to-centre spacing of the plungers). The three gate layers are arranged as follows: the bottom layer contains C (pink), which separates and laterally confines the SET and DQD regions. The middle layer contains BL and BR of the SET together with P1,

Table 4.1: Overview of devices.

Name	Platform	Enrichment	Specifier	ESR line	Experiments
Device A	²⁸ SiMOS	400 ppm	AL102977_D03_Die7_SD5	TiN	Single dot
Device B	²⁸ SiMOS	400 ppm	AL210563_D09_Die5_SD5	Al	Double dot

P2 and R (green). The top layer contains ST of the SET and the quantum-dot barrier gates B1 and B2 (blue). This layered architecture shapes the three-dimensional electrostatic potential, enabling precise control of tunnel couplings and dot energies.

In operation, the SET acts as a sensitive electrometer. ST tunes the SET island potential while BL and BR adjust the barrier transparencies. The SET typically operates in the many-electron regime [98], containing thousands of electrons between BL and BR. The confinement gate C is held at low or slightly negative potential to repel carriers and suppress unintended conduction between the sensor and qubit regions. Within the DQD, the plungers P1 and P2 set the single-particle energies and hence the electron occupation, whereas B1 and B2 control the tunnel rates to the reservoir and across the interdot barrier. Suitable combinations of plunger and barrier voltages permit reliable single-electron loading, manipulation, and readout. Operating voltages for devices A and B are summarised in Appendix C.3.

Microwave control of the qubit is delivered by a coplanar waveguide (CPW) that terminates on top of the qubit region in a shorted strip acting as an on-chip antenna. The resulting radio-frequency current generates a transverse magnetic field used to drive electron-spin-resonance (ESR) transitions. Device A employs a TiN ESR line, whereas device B uses Al. Device identifiers and materials are listed in Table 4.1.

4.2 The Single-Electron Transistor

Single-electron charge detection on a quantum dot coupled to a reservoir requires an electrometer capable of resolving individual electron tunnelling events. Each electron transfer must generate a signal with sufficient fidelity to register charge rearrangements on the chip. The single-electron transistor (SET) provides this functionality [138, 139]. It is itself a quantum dot tunnel-coupled to source and drain leads and, at its working point, conducts currents in the pA to nA range. Strong capacitive coupling between the SET and the target dot is achieved by placing them in close lateral proximity. In the gate design of devices A and B this spacing is set by the confinement gate C (Fig. 4.1) and is approximately 50 nm.

The SET potential landscape is highly tunable and is sketched in Fig. 4.2. The barrier gates BL and BR deplete the 2DEG accumulated by the top gate ST, producing lateral electrostatic confinement and forming a quantum dot with quantised levels μ_N . The energy required to add one electron to the dot is the addition energy

$$E_{\text{add}} = \mu_{N+1} - \mu_N = E_C + \delta E, \quad (4.1)$$

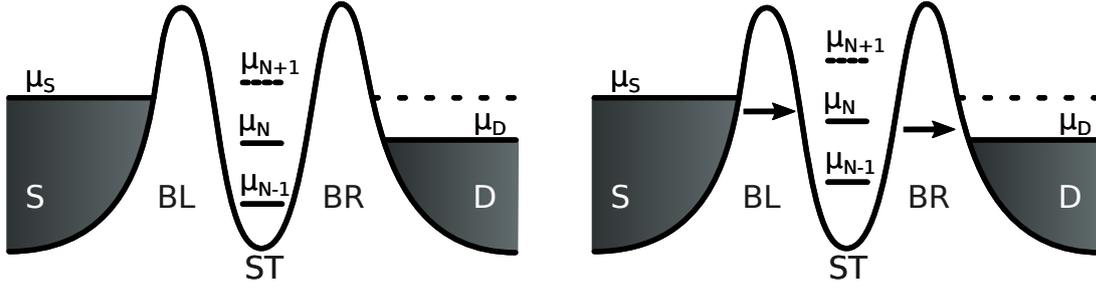


Figure 4.2: Electrochemical potential diagram of an SET region. The central quantum dot is defined by the topgate ST and both barrier gates BL and BR and coupled to source (S) and drain (D). The electron occupation of the dot is given by N . The potentials of source (μ_S) and drain (μ_D) are offset by an energy bias eV_{SD} . The dot electrochemical potential μ_N for the $N - 1 \rightarrow N$ transition is separated by the addition energy E_{add} . **Left** Coulomb-blockade regime: no current flows as the discrete dot levels are not between the energy window spanned by μ_S and μ_D . **Right** Transport regime: a tunnelling current flows from S to D (arrows) as μ_N lies between μ_S and μ_D .

where δE is the energy level spacing for two discrete quantum levels, which is zero when an electron is added to the same spin-degenerate level. The charging energy E_C accounts for the electrostatic repulsion between the added electron and those already on the island. It is determined by the total capacitance C_Σ to the surrounding environment [140]:

$$E_C = \frac{e^2}{2C_\Sigma} \quad (4.2)$$

with e being the elementary electron charge. Control of the dot electrochemical potential μ via a gate voltage V_G is described by the lever arm α . Using the capacitive coupling C_G between gate and dot gives

$$\alpha = \frac{C_G}{C_\Sigma} = \frac{\Delta\mu}{e\Delta V_G}. \quad (4.3)$$

4.2.1 Coulomb Blockade

Changing the SET top-gate voltage V_{ST} shifts the island levels by αV_{ST} with the lever arm α . Applying a source-drain bias opens a bias window. Tunnelling occurs only if at least one level lies within this window (Fig. 4.2). Only one electron can occupy the accessible level μ_N at a time. Once it tunnels to the drain, the level becomes available for the next event, producing a measurable current. If no level lies within the window, the island charge is fixed and the current is blocked, resulting in a phenomena called Coulomb blockade [139].

The conductance can be tuned between blockade and transmission by V_{ST} , yielding Coulomb oscillations as shown in Fig. 4.3. The voltage spacing between adjacent peaks, ΔV , follows

$$\Delta V = \frac{E_{add}}{e\alpha}.$$

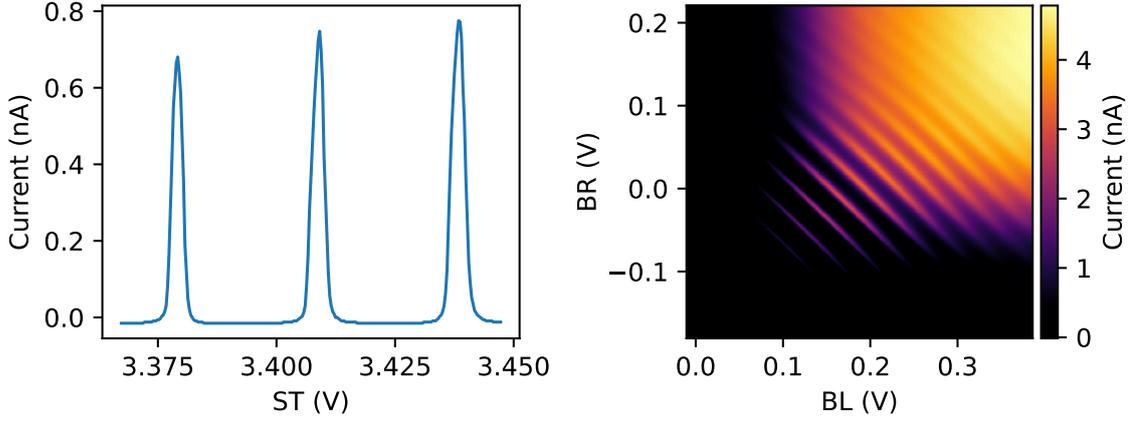


Figure 4.3: Coulomb oscillations of the SET used for charge sensing (device B). **Left** SET current versus top-gate (ST) voltage at fixed barrier voltages. The periodic peaks reflect successive single-electron charging of the SET island. The spacing converts to an addition energy via the gate lever arm. **Right** SET current versus the two barrier-gate voltages. Diagonal ridges arise from capacitive cross-coupling of both barriers to the island potential, while their individual voltages tune the tunnel transparencies. Reducing either barrier voltage eventually suppresses conduction and pinches off the SET. For charge sensing, a barrier working point lies in the regime of isolated Coulomb peaks.

Within a parallel-plate approximation the relevant capacitance scales as $C \propto A$, so $E_C \propto 1/A$ via Eq. (4.2), and hence $\Delta V \propto 1/A$. A smaller lateral area A therefore yields larger peak spacings. This scaling is used in two-dimensional SET maps (e.g. Fig. 4.3) to estimate island size and, via mutual-capacitance signatures from parasitic dots, their approximate position relative to the sensor.

The minimal full width at half maximum (FWHM) of a Coulomb peak, V_{FWHM} , is set by lifetime and thermal broadening. Lifetime broadening due to finite tunnel rate Γ gives $V_{\text{FWHM}} \propto \hbar\Gamma/(e\alpha)$ [140], while thermal broadening of the Fermi distribution yields $V_{\text{FWHM}} \propto k_B T/(e\alpha)$ [63], with electron temperature T and Boltzmann constant k_B . The Fermi function imposes a fundamental limit on level discrimination, justifying our cryogenic measurement setup (Sec. 3.1).

Experimentally, the lever arm α of the SET top gate ST on the island electrochemical potential (see Eq. (4.3)) is obtained from a bias-spectroscopy measurement of the SET current versus source-drain bias V_{SD} and ST voltage. An example is shown in Fig. 4.4. Three regions are evident: a diamond-shaped area of (near) zero current where the SET is in Coulomb blockade [139, 140], and regions of positive/negative current at finite bias. The blockade region is well described by the constant-interaction model (further details see [98, 141]).

The lever arm follows from the diamond dimensions as

$$\alpha = \frac{\Delta V_{\text{SD}}}{2 \Delta V_{\text{ST}}} \quad (4.4)$$

where ΔV_{SD} is the diamond height (tip-to-tip in bias) and ΔV_{ST} the width of a single diamond in gate voltage. Practically, the edges are extrapolated from their slopes near

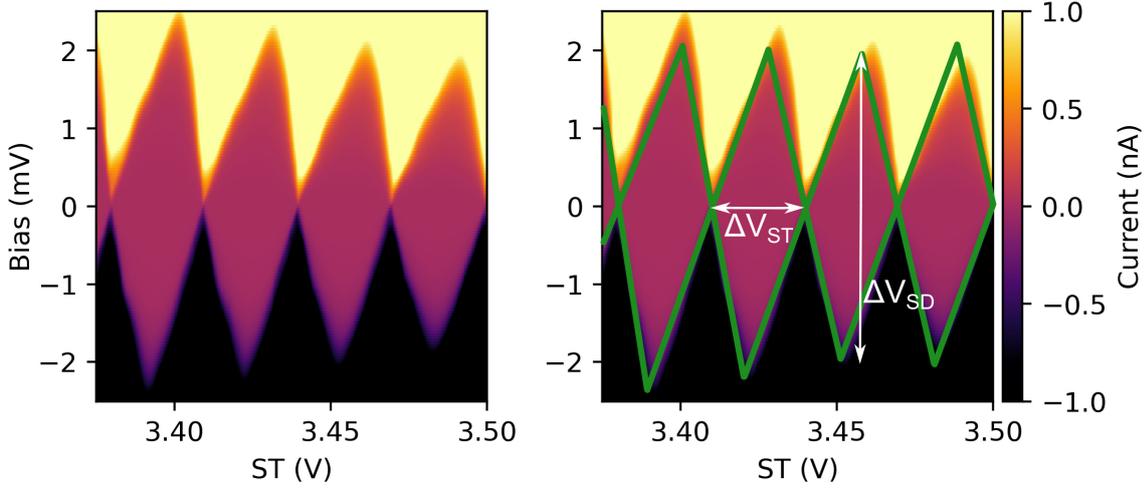


Figure 4.4: Coulomb diamonds of the SET (device B). **Left** SET current through the island formed under ST as a function of V_{SD} and V_{ST} . **Right** Same data with linear guides (green) along the diamond edges (zero current regime). The diamond height ΔV_{SD} and width ΔV_{ST} are indicated by white arrows.

Table 4.2: SET lever arm extraction. Systematic errors of 0.5 mV on ΔV_{SD} and 1 mV on ΔV_{ST} give a combined lever arm $\alpha = (7.0 \pm 0.3 \pm 0.9)\%$ (statistical \pm systematic).

	Diamond 1	Diamond 2	Diamond 3	Diamond 4	Average
ΔV_{SD} (mV)	4.4	4.2	3.9	4.1	4.2
ΔV_{ST} (mV)	29.9	29.4	29.4	30.1	29.7
α (%)	7.4	7.1	6.6	6.8	7.0

$V_{SD} = 0$ to determine the crossing points (green guidelines in Fig. 4.4, right). The extracted values are summarised in Tab. 4.2. Systematic uncertainties of 0.5 mV on ΔV_{SD} and 1 mV on ΔV_{ST} were assigned to reflect the limited precision of the graphical construction. Averaging over four diamonds yields $\alpha = (7.0 \pm 0.3 \pm 0.9)\%$, where the first uncertainty is statistical and the second systematic.

4.2.2 Charge Sensing

Coulomb oscillations of the SET can be used to sense charge on a nearby dot. A charge rearrangement on the target dot induces an offset charge Δq on the SET island, which appears as a break in the oscillation pattern. Equivalently, it looks like a shift in the SET top-gate voltage,

$$\Delta V_{ST} = \frac{\Delta q}{C_{ST}},$$

with C_{ST} the capacitance between the ST gate and the SET island [142]; see Fig. 4.5, left.

To maximise the step height and hence the signal-to-noise ratio (SNR), the SET is biased on a steep conductance flank. There the transconductance dI_{SET}/dV_{ST} is large, so a given ΔV_{ST} produces a larger current step (Fig. 4.5, right).

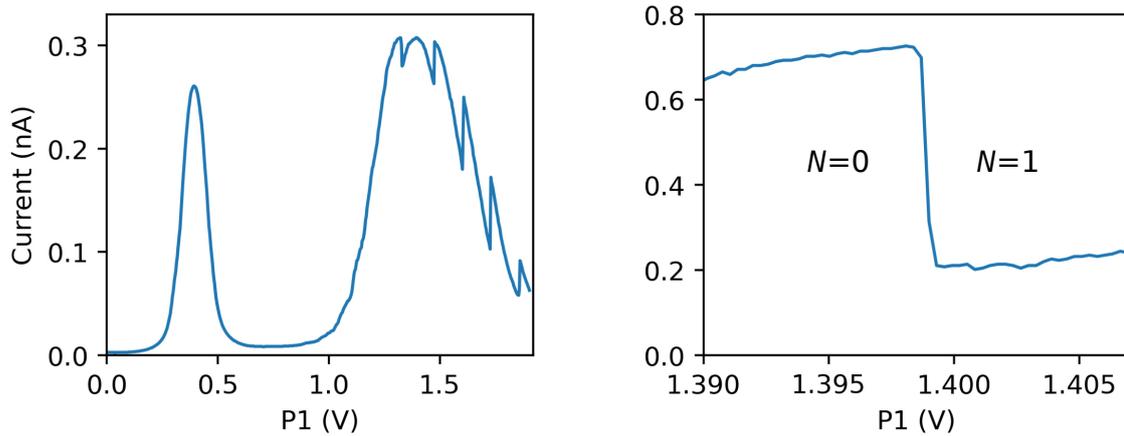


Figure 4.5: Charge sensing of a neighbouring quantum dot using the SET (device B). **Left** SET current versus the adjacent plunger P1 voltage. Periodic Coulomb oscillations arise from capacitive cross-coupling of P1 to the SET island. Superimposed discrete current steps for $P1 > 1.3$ V mark successive single-electron charging events of the target dot as its electrochemical potential crosses the reservoir Fermi level. **Right** A single transition with the SET biased on a conductance flank (optimised ST) voltage. The enhanced transconductance yields a large step, providing a suitable working point for high-sensitivity charge sensing. Charge occupancy N of the dot is indicated.

Increasing the source-drain bias V_{SD} raises the SET current and can further increase the step height, but at the cost of peak broadening, which reduces dI_{SET}/dV_{ST} . There is therefore a bias “sweet spot”. In addition, excessive current can heat the device and series resistors, elevating the electron temperature (Fig. 4.8, right), which has been shown to shorten the spin-relaxation time T_1 in SiMOS spin qubits [143].

For the single-qubit measurements presented in this chapter, SET sensing targeted the $0 \leftrightarrow 1$ charge transition of the dot. The SET was operated on the left-hand conductance flank. This choice makes the $N=1$ state correspond to a lower SET conductance which is beneficial as that state occupies most of the pulse cycle thereby reducing average sensor current and mitigating heating.

For charge sensing, two practical V_{SD} regimes are used: (i) High-bias tuning (~ 1 mV-2 mV): the Coulomb diamonds broaden, leaving little or no fully blockaded region, so the sensor is not “blind” over wide ranges of dot-gate voltage and discrete charge steps are readily detected during initial gate sweeps. (ii) Low-bias time-domain sensing (~ 0.1 mV): the peaks are sharper and the local transconductance dI_{SET}/dV_{ST} on the chosen flank is larger, improving step visibility and signal-to-noise ratio (SNR) while minimising dissipation.

4.2.3 Charge Noise Spectroscopy

Charge noise denotes low-frequency fluctuations of the electrostatic environment that shift dot energies and tunnel barriers. Prominent microscopic sources include defects at the Si/SiO_x interface or within the oxide that stochastically capture and emit electrons [144, 145]. Charge noise degrades qubit coherence and gate fidelity and has been identified

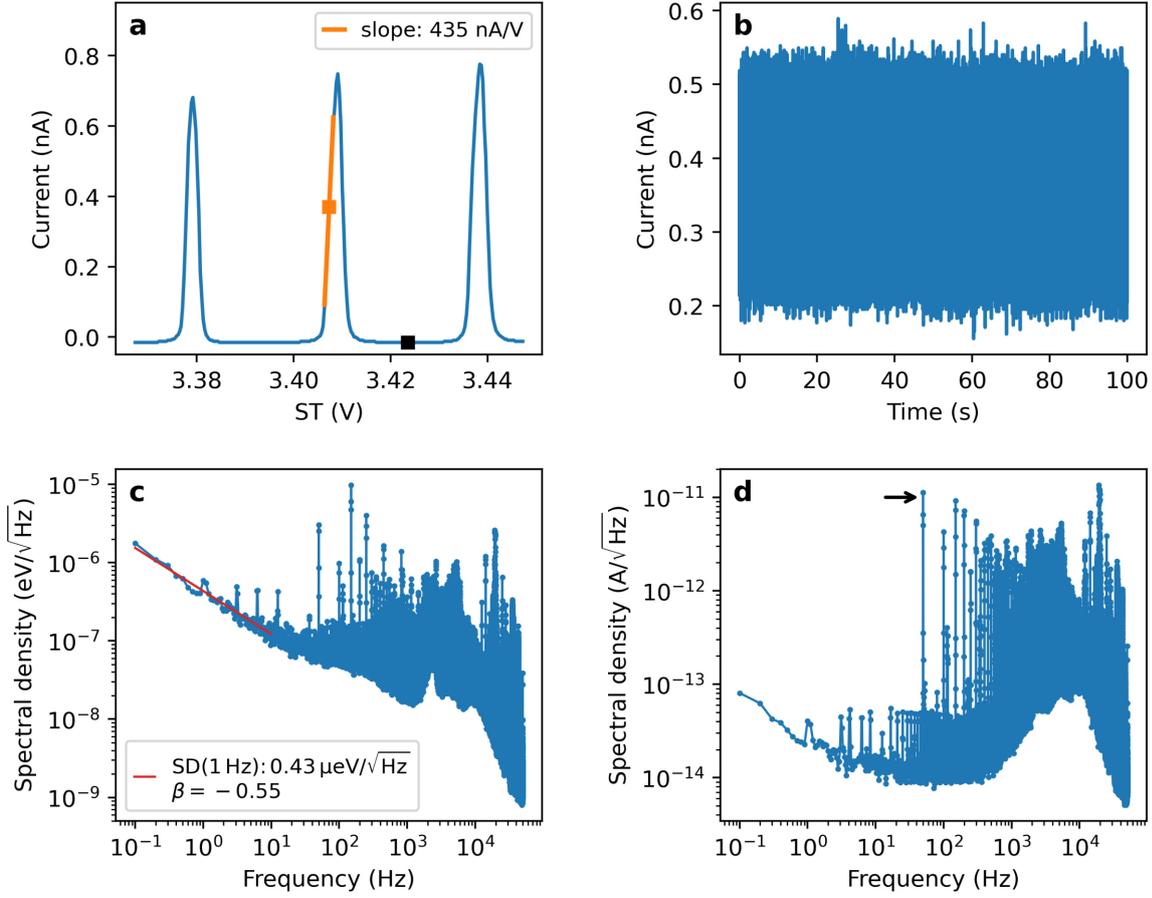


Figure 4.6: Spectral noise density (SD) of the SET (device B). **a** Coulomb oscillations of the SET with two measurement positions indicated. The linear fit on the left peak flank (orange) sets $dI_{\text{SET}}/dV_{\text{ST}}$ used in **c**. The black dot marks a charge-noise-insensitive point used to estimate the setup noise floor in **d**. **b** Time trace of the SET current sampled at 100 kHz. **c** Spectral noise density of the trace in **b**, averaged with Welch’s method (10 Hann windows). The fit is performed over 0.1 Hz to 10 Hz. **d** Noise floor recorded at the charge-noise-insensitive point (black dot in **a**). The mains peak at 50 Hz is indicated (arrow). The IV-converter bandwidth of ~ 24 kHz appears as a roll-off of the noise floor above that frequency.

as a principal limitation for SiMOS qubits [2, 146]. Its impact is typically more severe in devices with micromagnets that convert voltage noise into qubit-frequency noise via engineered field gradients. By contrast, the intrinsic spin-orbit coupling in SiMOS without micromagnets is comparatively weak [54], and therefore charge noise generally has a weaker impact in such SiMOS devices.

To compare devices, the noise magnitude is quantified via current spectroscopy of the SET. First, the local transconductance at the working point is obtained from the flank of a Coulomb peak (Fig. 4.6 a); here

$$\frac{dI_{\text{SET}}}{dV_{\text{ST}}} = 291 \text{ nA/V}.$$

A 100 s time trace of the SET current is then recorded at sampling rate $f_s = 100$ kHz using the newly implemented continuous readout mode of the ADwin ProII input card (Sec. 3.7).

Table 4.3: Charge noise spectroscopy of device B. A complete analysis for device A is provided in Appendix C.2.

Measurement	Spectral noise density at 1 Hz	β
1	$(0.39 \pm 0.04) \mu\text{eV}\sqrt{\text{Hz}}$	-0.43 ± 0.06
2	$(0.30 \pm 0.03) \mu\text{eV}\sqrt{\text{Hz}}$	-0.42 ± 0.06
3	$(0.43 \pm 0.02) \mu\text{eV}\sqrt{\text{Hz}}$	-0.55 ± 0.02
average	$(0.37 \pm 0.04) \mu\text{eV}\sqrt{\text{Hz}}$	-0.47 ± 0.04

The highest frequency of the spectrum is given by the Nyquist frequency $f_{\text{Ny}} = f_s/2$ [147]. The power spectral density (PSD) is estimated with Welch’s method [148, 149], using ten equal segments and a Hann window to reduce end discontinuities of the time trace segments. Welch’s routine returns a PSD in units of A^2/Hz converting to an energy amplitude spectral noise density (SD) referenced to the ST gate and yields

$$\text{SD}(f) = \frac{e \alpha}{dI_{\text{SET}}/dV_{\text{ST}}} \sqrt{\text{PSD}(f)}, \quad (4.5)$$

with lever arm α and elementary charge e . This provides SD in $\mu\text{eV}/\sqrt{\text{Hz}}$.

The low-frequency spectrum is empirically described by a power law,

$$\text{SD}(f) = A f^\beta,$$

consistent with an ensemble of two-level fluctuators (TLFs) with a broad distribution of time constants [150]. Fitting over 0.1 Hz to 10 Hz gives a charge-noise amplitude at 1 Hz of $(0.37 \pm 0.04) \mu\text{eV}/\sqrt{\text{Hz}}$ and an exponent $\beta = -0.47 \pm 0.04$ (Table 4.3), comparable to or slightly lower than previous measurements on this device generation [151]. An exponent $\beta > -1$ indicates a spectrum shallower than ideal $1/f$ (“pink”) noise, which can arise from a non-uniform TLF ensemble and/or partial admixture of a white-noise floor. Similar sub-unity exponents have been reported in silicon qubits [146, 152].

To verify that the measurement is not limited by the instrumentation, a “noise-floor” spectrum was taken at a charge-insensitive point between Coulomb peaks (Fig. 4.6d). As expected, narrow spurs at 50 Hz and its harmonics originate from the setup. The 50 Hz line sits about three orders of magnitude above the white floor in this configuration. When the sensor is parked on the Coulomb-peak flank (Fig. 4.6c), the same 50 Hz line is only two orders of magnitude above the local floor, indicating an overall elevation of the background due to genuine charge noise. Consequently, the spectral density near 1 Hz is not dominated by the electronic setup but by device charge noise. While the analysis above pertains to device B, a comprehensive treatment for device A is provided in Appendix C.2.

4.3 Plunger Gates: Lever Arm Extraction

While the SET lever arm can be extracted directly from Coulomb diamonds, the same procedure is not applicable to the plunger gates P1 and P2 because no current flows

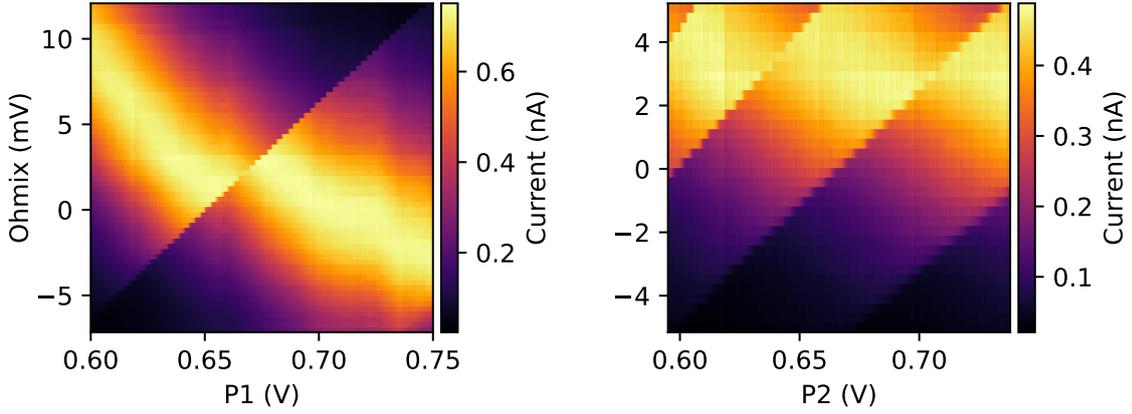


Figure 4.7: Lever-arm extraction of P1 and P2 of device B from current maps of the SET sensor (device B). SET tunnelling current as a function of the qubit plunger voltage and the reservoir potential V_{Ohmics} . The ohmic bias is generated by a room-temperature voltage divider (Sec. 3.7). Charge-addition lines appear as sharp, linear features while their slope gives the lever arm via Eq. (4.6). **Left** P1: $\alpha = (12.5 \pm 0.2)\%$. **Right** P2: $\alpha = (7.5 \pm 0.2)\%$.

through the quantum dots themselves. Instead, the lever arms are determined from a two-parameter map of the SET current versus the plunger voltage V_P and the electrochemical potential of the ohmic 2DEG channel under the reservoir gate R (labelled V_{Ohmics}). The electron reservoir is tunnel-coupled to the dot, so sweeping V_{Ohmics} shifts directly μ_{Ohmics} , while sweeping V_P shifts the dot's electrochemical potential μ_{dot} . Single-electron additions appear as sharp current steps that align along straight lines in the (V_P, V_{Ohmics}) plane. Along a charge-transition line the condition $\mu_{\text{dot}}(V_P) = \mu_{\text{res}}(V_{\text{Ohmics}})$ holds, which gives the plunger lever arm

$$\alpha = \frac{dV_{\text{Ohmics}}}{dV_P}. \quad (4.6)$$

Hence, the slope of a transition line directly yields the plunger lever arms of the device which is shown in Fig. 4.7.

This plunger-reservoir slope method offers several advantages over lever-arm extraction from Coulomb diamonds. First, it is a true zero-current measurement for the sensed dot: no transport flows through the dot, so the applied reservoir bias drops fully across the ohmic 2DEG (no ambiguity from series resistances or bias division in the DC lines). Second, it avoids the subjective step of defining Coulomb blockade by a current threshold, which can distort the apparent diamond shape and is strongly affected by electron temperature and lifetime broadening. In this method, the lever arm follows directly from the geometric slope of charge-addition lines, independent of thresholds. Third, the measurement can be performed at small source-drain biases and over narrow gate ranges, reducing dissipation and heating and mitigating bias-induced changes to tunnel barriers or dot size.

4.4 Electron Temperature Determination

The newly commissioned cryostat Sionludi L (“Idefix”) and the sample holder (see Secs. 3.1 to 3.3) were characterised via the electron temperature at the qubit chip. The effective electron temperature is governed primarily by the thermalisation of the wiring and filters to the cryostat’s cold stages.

Charge sensing of the transition width of a neighbouring quantum dot provides an estimate of the electron temperature of the coupled reservoir. The measurement principle is sketched in Fig. 4.8a: thermal broadening of the reservoir Fermi edge opens an energy window within which electrons tunnel into or out of the dot. Sweeping the dot level through this window changes the average dot occupancy according to the Fermi-Dirac distribution. An adjacent SET, biased on a conductance flank, reads out this average occupancy capacitively, without driving current through the dot itself. The sensor signal is essentially independent of the dot-reservoir tunnelling rate Γ , so the measurement remains operational even when Γ is made very small.

Experimentally, the SET is biased on a conductance flank and the tunnelling current is recorded while sweeping the dot plunger voltage over a narrow window across the $0 \leftrightarrow 1$ charge transition. The step in SET current follows the Fermi-Dirac function and is fitted as [139, 153]

$$I_{\text{SET}}(V) = \frac{I_A}{\exp\left(\frac{\alpha(V - V_0)}{k_B T_e}\right) + 1} + I_C, \quad (4.7)$$

where I_A and I_C depend on the SET working point, V_0 is the centre of the transition, α is the plunger lever arm to energy (extracted in Sec. 4.3), and T_e is the electron temperature.

To ensure that the extracted width reflects thermal broadening (not lifetime broadening), the dot-reservoir coupling is tuned such that

$$h\Gamma \ll k_B T_e.$$

At the same time, it is convenient to choose Γ faster than the data-acquisition bandwidth so that time-averaged occupancies are measured rather than resolving individual tunnelling events. For typical temperatures around $T_e \sim 100$ mK, tunnelling rates $\Gamma \sim 1$ MHz comfortably satisfy $h\Gamma \ll k_B T_e$ and provide smooth, well-averaged steps.

Applying this procedure with ~ 2 s averaging per point yields an electron temperature $T_e = (126 \pm 3)$ mK (Fig. 4.8b). This value is comparable to that obtained on a separate setup (“Obelix”), $T_e = (118 \pm 9)$ mK [154].

Although the sensed dot carries no transport current, the biased charge sensor can still emit phonons or photons and induce back-action heating [153, 155, 156]. This is evident in Fig. 4.8c, where increasing the SET current from 43 pA to 775 pA approximately doubles the extracted T_e . The most likely origin is local dissipation in the sensor and its reservoirs (plus shot-noise back-action), rather than series-line heating. In practice, the SET bias is therefore kept as low as possible consistent with adequate signal-to-noise. The achievable

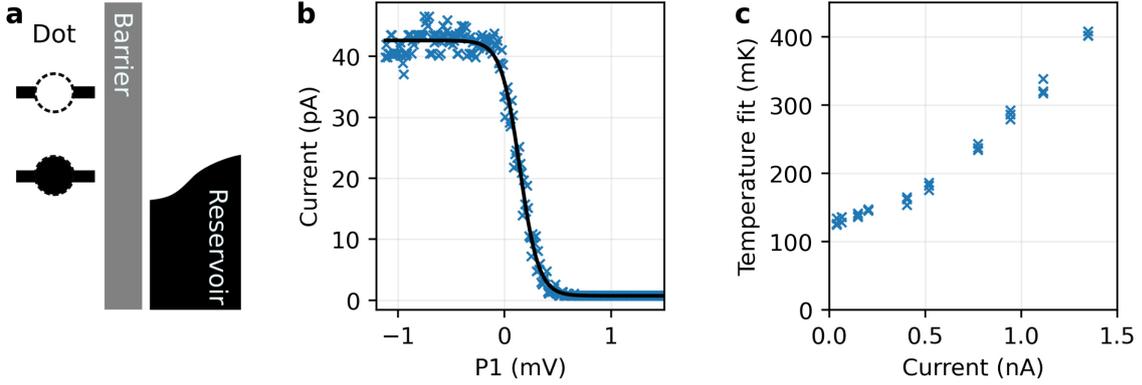


Figure 4.8: Electron temperature measurement (device B). **a** Schematic of a dot that is occupied by a single electron (black dot). The dot is tunnel-coupled to an electron reservoir with a Fermi-Dirac-broadened occupation. Sweeping the plunger changes the average dot occupancy, which is sensed by a nearby SET. **b** SET current across the $0 \leftrightarrow 1$ transition versus V_{P1} (offset by 0.652 V). Each point is an average over ~ 2 s. A room-temperature voltage divider is used for V_{P1} (Sec. 3.7). A fit to Eq. (4.7) gives $T_e = (126 \pm 3)$ mK with an SET current of 43 pA. **c** Dependence of extracted T_e on SET current, fitted using Eq. (4.7). Higher sensor current increases T_e due to sensor-induced heating.

sensor current is ultimately limited by the IV-converter trade-off: higher transimpedance gain lowers the measurement bandwidth and vice versa.

4.5 Spin-Selective Readout

After characterising the electron charge transition, the next step involves introducing the electron spin. Consequently, a readout protocol for selective spin tunnelling is implemented by spin-to-charge conversion through energy-selective tunnelling. In a static magnetic field B_{ext} , the lowest orbital splits into spin states separated by the Zeeman energy $E_Z = g\mu_B B_{\text{ext}}$. The plunger gate is tuned such that the reservoir electrochemical potential μ_{Res} lies between the spin-resolved dot addition potentials,

$$\mu_{\downarrow} < \mu_{\text{Res}} < \mu_{\uparrow},$$

while the charging energy E_C forbids double occupancy. Under these conditions a spin-up electron can tunnel out to the reservoir, whereas a spin-down electron remains confined. A nearby SET converts the resulting change in dot charge into a measurable current signal. This readout scheme is illustrated in Fig. 4.9 and is commonly referred to as Elzerman readout [157].

High-fidelity operation relies on a clear energetic hierarchy. Thermal and lifetime broadening must be small compared with the Zeeman splitting,

$$k_B T \ll E_Z \quad \text{and} \quad h\Gamma \ll E_Z,$$

to suppress thermally activated and lifetime-broadened errors [98, 157]. In addition, the spin-up electron should leave the dot before relaxing to the spin-down state on a timescale

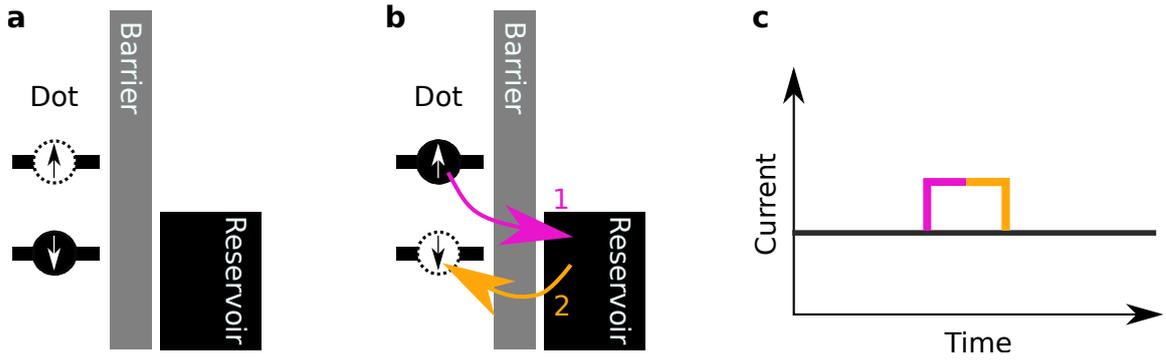


Figure 4.9: Spin-to-charge conversion (energy-selective tunnelling) for single-shot spin readout. A single-electron quantum dot is tunnel-coupled to an electron reservoir. A static magnetic field Zeeman-splits both spin states by E_Z . The plunger gate sets the dot electrochemical potentials μ_\downarrow and μ_\uparrow such that the potential of the reservoir sits in between ($\mu_\downarrow < \mu_{\text{Res}} < \mu_\uparrow$), while the charging energy E_C forbids double occupancy. In this configuration a spin-up electron can leave the dot whereas a spin-down electron remains confined. A nearby SET charge sensor converts this transient charge motion into a measurable current change. **a** Spin-down initial state: the ground state (filled black) lies below the reservoir Fermi level E_F . With $k_B T, h\Gamma \ll E_Z$, neither thermal activation nor lifetime broadening drives tunnelling on the measurement timescale, so no charge rearrangement occurs and the SET current remains steady. **b** Spin-up initial state: the electron tunnels out to the reservoir at rate $\Gamma_{\text{out}}^\uparrow$ (1). The dot then refills from the reservoir into the spin-down ground state at rate $\Gamma_{\text{in}}^\downarrow$ (2). This out-in sequence produces a transient change of the dot charge that is detected by the SET. The readout is intrinsically destructive of the initial spin state. **c** Representative time traces of the SET current with no blip (spin-down, black) and a single blip (coloured) indicating the out-in tunnelling sequence for spin-up.

T_1 , i.e. $\Gamma \gg 1/T_1$. Also, the subsequent refilling into the spin-down ground state should occur on a timescale that is resolvable by the bandwidth of the readout electronics. For the single-qubit experiments discussed in this work, the external magnetic field $B_{\text{ext}} = 0.66$ T. The readout conditions therefore read

$$\Gamma \ll 18.5 \text{ GHz} \quad \text{and} \quad T_e \ll 890 \text{ mK}$$

which is satisfied by the experimental parameters $\Gamma \leq 1$ kHz and $T_e < 300$ mK (for sensor currents < 1 nA; see Sec. 4.4). In practice, the plunger P1 sets the alignment ($\mu_\downarrow, \mu_\uparrow$) relative to E_F , while a single barrier gate B1 tunes Γ .

Blip evaluation

A projective measurement of the qubit state yields one of the basis states, but a qubit can be prepared in a superposition $a|\uparrow\rangle + b|\downarrow\rangle$. The spin-up probability $|a|^2$ is therefore obtained statistically from repeated single-shot measurements.

Figure 4.10a,b show time traces of the SET current recorded at the spin-to-charge conversion working point (Fig. 4.9). To generate a substantial spin-up population, a π -pulse is applied prior to readout (Sec. 4.8). With the chosen tunnel coupling between dot and reservoir, blips associated with genuine spin-up events occur at the beginning of the readout window (typically $t < 10$ ms), as the initially prepared spin-up electron promptly

tunnels out and the dot refills with spin-down. Blips appearing later in the window are consistent with thermally activated events in which a spin-down electron is thermally promoted to escape to the reservoir. Such events (false positives) are distributed uniformly in time, whereas genuine spin-up events cluster near the start of the readout.

Automated post-processing thresholds the SET current at a discrimination level I_{th} to classify outcomes (Fig. 4.10a). Each time series is binarised sample-by-sample into 0 ($\leq I_{\text{th}}$) and 1 ($> I_{\text{th}}$). A single-shot measurement is counted as “spin up” if any sample crosses threshold within the readout window. The binarised raster in Fig. 4.10c visualises this classification. Averaging the binary outcomes over all shots yields the measured spin-up fraction. In the dataset shown, 191 out of 200 traces contain at least one threshold crossing, corresponding to a spin-up fraction of 95.5%.

The characterisation in this chapter establishes the essential operating toolbox: a quantitative voltage-to-energy conversion via lever arms, a resolved electron temperature and charge-noise floor, and a high-fidelity spin-selective readout channel with tunable tunnel rates. With the device tuned to stable single-electron occupancy and the readout working point defined, the stage is set for time-domain control of the spin qubit.

4.6 Spin Relaxation Time

The longitudinal spin-relaxation time T_1 quantifies the return of an excited spin to the thermal ground state. In gate-defined silicon dots the dominant relaxation channels are phonon-assisted processes enabled by spin-orbit and spin-valley mixing [41, 70, 98]. At fixed temperature the rate typically increases with magnetic field [158]. Experimentally, T_1 is extracted by preparing a non-equilibrium spin population, letting it evolve for a variable wait time, and reading out the spin state in single shot.

The non-equilibrium spin population is achieved by the Elzerman pulse sequence [157], a three-stage sequence shown in Fig. 4.11. First, during the *empty* stage both spin levels are pulsed above the reservoir Fermi energy E_F so the dot is emptied. Next, during the *load* stage the plunger is pulsed well below E_F , allowing an electron to enter with either spin resulting in a spin-up fraction of $\approx 50\%$. The dot is then held for a controlled interval t_{wait} during which an initially loaded spin-up electron ($|\uparrow\rangle$) may relax to $|\downarrow\rangle$. Finally, during the *readout* stage the dot is placed in the energy-selective configuration $\mu_{\downarrow} < \mu_{\text{Res}} < \mu_{\uparrow}$. If the dot contains $|\uparrow\rangle$ at the start of the read window, the electron tunnels out and the dot refills in $|\downarrow\rangle$, producing a characteristic SET current jump which is called “blip”. If the dot is already in $|\downarrow\rangle$, no tunnelling occurs.

An arbitrary waveform generator (AWG) applies the Elzerman pulse sequence to the plunger gate. The corresponding waveform is shown in Fig. 4.12. Because the plunger is delivered through a bias-T, the empty and load pulses are DC-compensated to maintain a zero time-average (Sec. 3.6), and a short discharge pulse resets the bias-T between averages. In the right panel the read voltage level is scanned while keeping the rest of the sequence

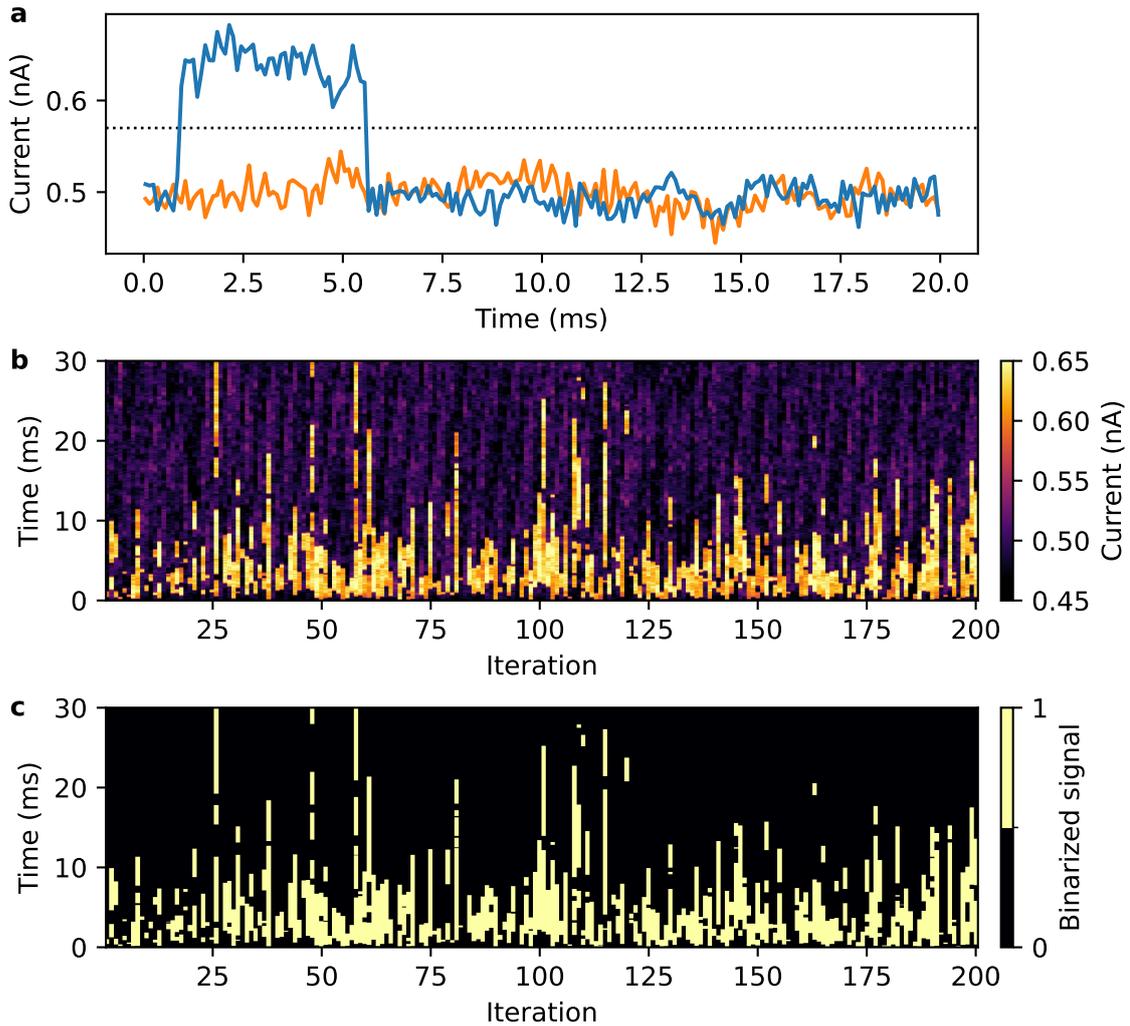


Figure 4.10: Single-shot Elzerman spin readout and threshold-based classification (device A). Traces are sampled at 10 kHz within a fixed readout window. **a** Representative SET-current traces with a discrimination threshold I_{th} (dotted). A spin-down outcome produces no rise above threshold over the window (orange), whereas a spin-up outcome yields a transient tunnelling “blip” above I_{th} (blue) as the electron tunnels out and the dot refills in the ground state. **b** Raster plot of 200 single-shot acquisitions, recorded after π -pulse initialisation (Sec. 4.8). Bright streaks correspond to blips associated with spin-up events, dark traces indicate spin-down. **c** Binarised outcome map obtained by applying the threshold from **a** to each time trace in **b**: a value of 1 (yellow) marks any sample crossing I_{th} within the readout window and is counted as a spin-up detection, while a completely 0 (dark) trace corresponds to spin-down. This representation is used to compute state probabilities which is in this case 95.5%.

fixed. The resulting horizontal “spin-tail” identifies the voltage range where the energy-selective condition $\mu_{\downarrow} < \mu_{\text{Res}} < \mu_{\uparrow}$ is satisfied. Biasing near the left edge of it maximises the $|\uparrow\rangle$ blip rate, whereas biasing near the right edge suppresses thermally activated false positives. In practice, the read level is swept and the setting that optimises visibility (maximising the $|\uparrow\rangle$ blip count while minimising thermal background) is selected.

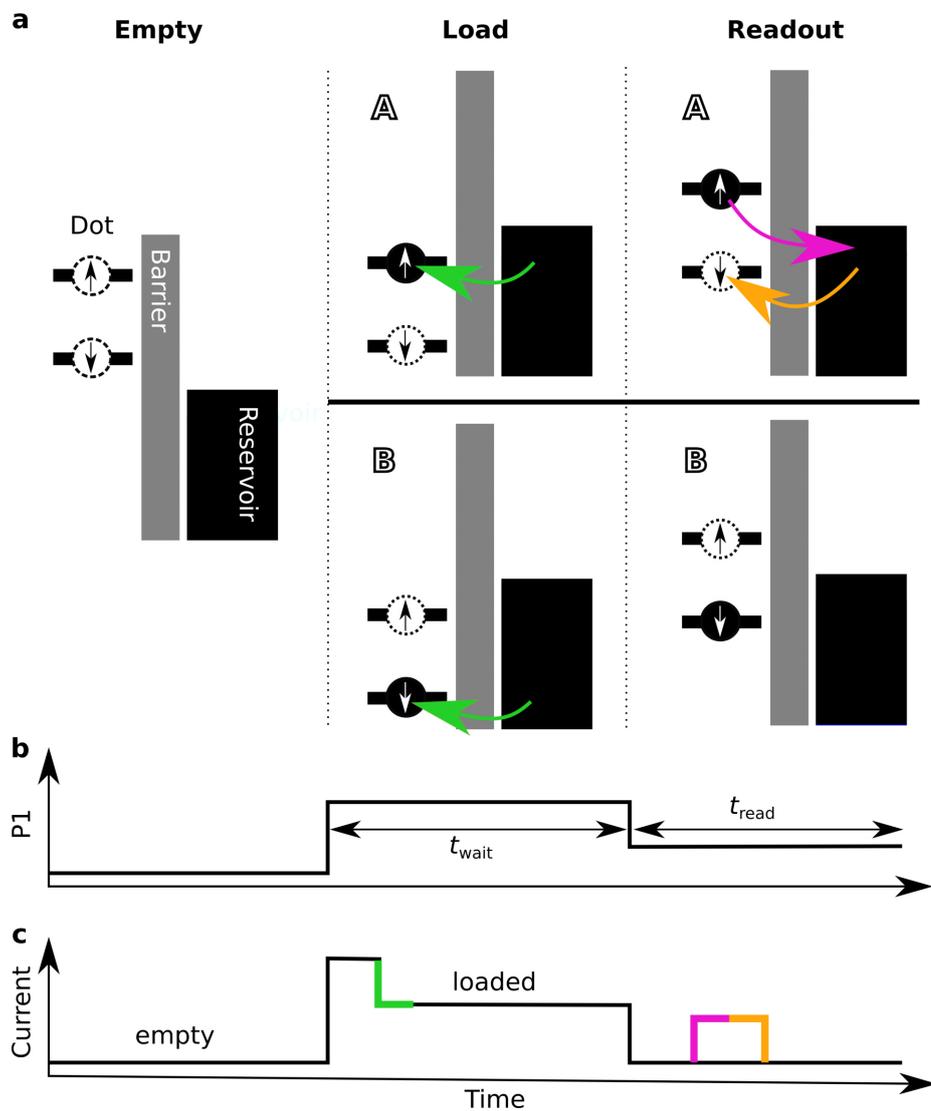


Figure 4.11: Schematic of the three-stage Elzerman pulse sequence [157] used for spin-relaxation measurements. **a** Alignment of the Zeeman-split dot potentials ($\mu_{\downarrow}, \mu_{\uparrow}$) in the single electron regime relative to the reservoir Fermi energy E_F . **b** Corresponding plunger P1 waveform and an exemplary SET-current response (**c**). *Empty stage:* Both dot levels are pulsed above E_F , fully depleting the dot. *Load stage:* A fast voltage pulse brings the dot levels well below E_F , allowing an electron to enter the dot (green arrow). Either spin may load at random: case A loads $|\uparrow\rangle$ (top row), case B loads $|\downarrow\rangle$ (bottom row). The levels are then held for a wait time t_{wait} to probe relaxation of $|\uparrow\rangle \rightarrow |\downarrow\rangle$. *Readout stage:* The dot is positioned in the energy-selective configuration $\mu_{\downarrow} < \mu_{Res} < \mu_{\uparrow}$. If the dot contains $|\uparrow\rangle$ (case A), the electron tunnels out (magenta arrow) and the dot quickly refills with $|\downarrow\rangle$ (orange arrow), producing a transient blip in the SET current within the readout window t_{read} . If the dot contains $|\downarrow\rangle$ (case B), no tunnelling occurs and no blip is observed. The lower traces illustrate the applied plunger P1 voltage sequence with t_{wait} and t_{read} marked. The corresponding SET-current response shows a loading step (green) followed, in case A only, by a readout blip (magenta/orange).

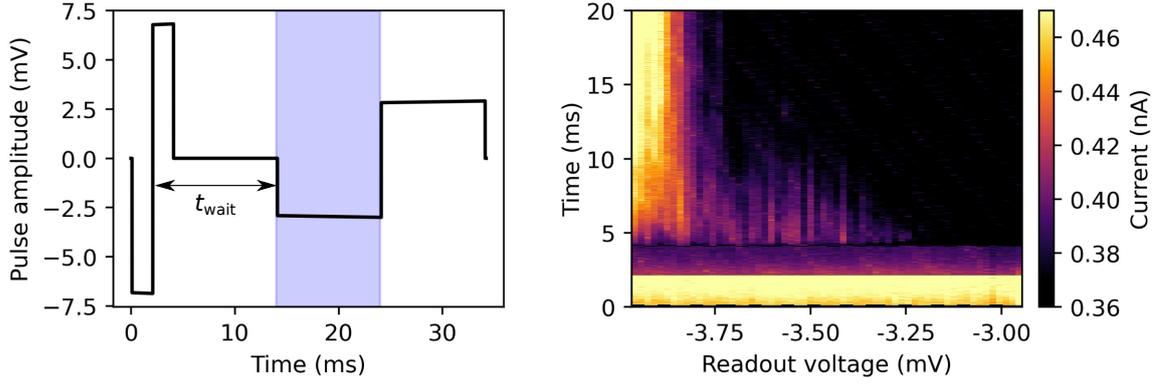


Figure 4.12: Implementation and calibration of the Elzerman pulse sequence (device A). **Left** Applied waveform to the plunger gate P1. The baseline (0 mV) corresponds to a loaded position with both spin levels well below the reservoir Fermi level E_F . Each cycle begins with an empty pulse (2 ms) followed by a load pulse (2 ms). Both are bias-T-compensated to keep the long-term charge average at zero (Sec. 3.6). After a wait period t_{wait} at the load level, the voltage is stepped to a read level where the SET current is digitised (blue shaded window). A discharge pulse at the end resets the bias-T prior to the next averaged cycle. **Right** Measured SET current for the full pulse sequence of **a** while sweeping the read level (horizontal axis). Shown for $t_{\text{wait}} = 2$ ms and $B_{\text{ext}} = 1$ T. Each trace is averaged 50 times. A 2 ms empty phase (high current) is followed by a 2 ms load/wait phase (low current). The subsequent horizontal “spin-tail” indicates the range of read voltages ($[-3.75 \text{ mV}, -3.25 \text{ mV}]$) satisfying the energy-selective condition $\mu_{\downarrow} < \mu_{\text{Res}} < \mu_{\uparrow}$. The “spin-tail” decays over ≈ 5 ms. The optimum read level is chosen near the right edge of the tail, where spin-up events yield clear blips while thermal false positives are minimised.

For each value of t_{wait} a large number of single-shot traces is acquired and classified using a fixed current threshold (Sec. 4.5). The measured spin-up fraction P_{\uparrow} is defined as the fraction of traces containing at least one above-threshold excursion within the read window. In the regime relevant here ($k_B T \ll E_Z$) the equilibrium spin-up population during readout is negligible. The decay of the non-equilibrium $|\uparrow\rangle$ population is then well described by a single exponential [98, 157],

$$P_{\uparrow}(t_{\text{wait}}) = P_{\uparrow}^0 \exp\left(-\frac{t_{\text{wait}}}{T_1}\right) + c, \quad (4.8)$$

where P_{\uparrow}^0 denotes the spin-up loading probability (set by the details of the load step). Thermal events which offset P_{\uparrow} are captured in c .

Figure 4.13 plots the measured spin-up fraction versus t_{wait} . A fit to Eq. (4.8) yields a relaxation time

$$T_1 = (117 \pm 9) \text{ ms} \quad \text{at} \quad B_{\text{ext}} = 0.66 \text{ T}.$$

The extracted T_1 confirms that relaxation is slow compared with the single-shot readout time, validating the readout hierarchy $\Gamma \gg 1/T_1$ assumed throughout this chapter.

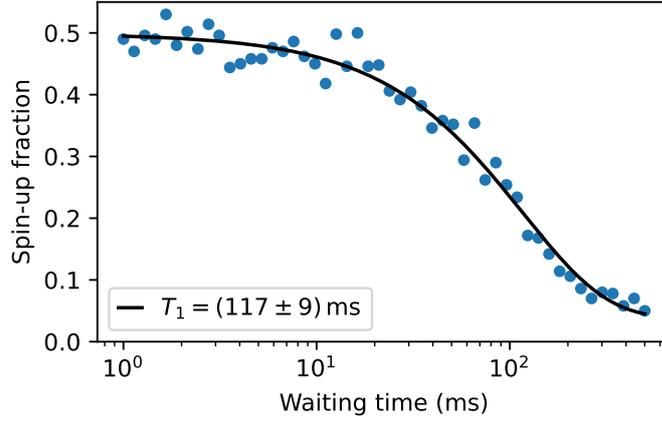


Figure 4.13: Spin relaxation time measured with the Elzerman sequence (device A). The spin-up fraction is plotted versus the waiting time in the load stage. Each data point represents the fraction of single-shot traces that contain a readout blip (500 shots averaged per point). The solid line is a fit to an exponential decay (Eq. (4.8)). The extracted relaxation time is $T_1 = (117 \pm 9)$ ms at $B_{\text{ext}} = 0.66$ T.

4.7 Rapid Adiabatic Passage

Locating the qubit Larmor frequency ω_0 is the first step toward coherent control, yet it is non-trivial in the present setup. The static field is generated by a home-built 3D vector magnet whose spatial inhomogeneity and small position shifts between cooldowns cause run-to-run variations of ω_0 . To search efficiently over a broad frequency window while maintaining sensitivity to a narrow resonance (typical linewidth < 0.5 MHz), rapid adiabatic passage (RAP) is employed to scan the spectrum in coarse segments of ~ 20 MHz.

RAP provides a robust spin inversion by sweeping the microwave frequency through resonance slowly enough that the state adiabatically follows the instantaneous eigenstate, yet fast enough to outrun relaxation ($T_1 \approx 117$ ms for the data in this chapter). In the rotating-wave approximation (RWA) introduced in Sec. 2.1.4, the driven spin is described by the time-dependent two-level Hamiltonian

$$H_{\text{RWA}}(t) = \frac{\hbar}{2} \begin{pmatrix} \Delta\omega(t) & \Omega_{\text{R}} \\ \Omega_{\text{R}} & -\Delta\omega(t) \end{pmatrix}, \quad \Delta\omega(t) \equiv \omega_0 - \omega_{\text{MW}}(t), \quad (4.9)$$

where Ω_{R} is the on-resonance Rabi frequency set by the transverse drive, and $\Delta\omega(t)$ is the (chirped) detuning. In the rotating frame, the state precesses about an effective field whose polar angle θ between the quantisation axis of the qubit and the microwave field evolves as [159]

$$\theta(t) = \arctan\left(\frac{\Omega_{\text{R}}}{\Delta\omega(t)}\right). \quad (4.10)$$

A linear chirp from large negative to large positive detuning rotates θ by π , and, provided the evolution is adiabatic, the spin follows this axis yielding deterministic inversion. A convenient adiabaticity criterion near resonance is [160]

$$\left|\frac{\partial}{\partial t}(\Delta\omega)\right| \ll \Omega_{\text{R}}^2 \quad (4.11)$$

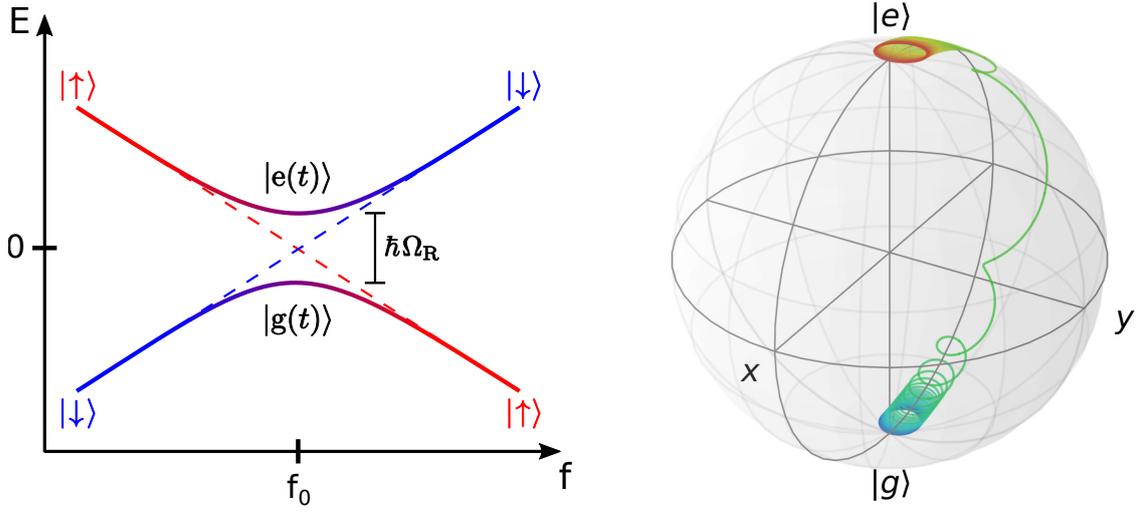


Figure 4.14: Working principle of rapid adiabatic passage (RAP). **Left** Energy diagram of both spin-down (blue) and spin-up (red) state versus drive frequency. As the microwave is linearly chirped through f_0 , the system traverses an avoided crossing of size $\hbar\Omega_R$. For $|\frac{\partial}{\partial t}(\Delta\omega)| \ll \Omega_R^2$ (Eq. (4.11)) the state adiabatically follows the upper (or lower) branch, producing inversion. **Right** Simulated spin trajectory on the Bloch sphere in the rotating frame during a RAP sweep, showing inversion from $|g\rangle$ to $|e\rangle$ and the small nutation about the time-varying effective field [169]. Taken from [160].

which sets an upper bound on the chirp rate. Under these conditions the dynamics are well captured by Landau-Zener theory [88]: starting in $|\downarrow\rangle$, the spin-up probability after a single passage is

$$P_{\uparrow} = 1 - \exp\left(-\frac{\pi \Omega_R^2}{2 |\frac{\partial}{\partial t}(\Delta\omega)|}\right) = 1 - \exp\left(-\frac{\pi^2 f_R^2}{|\frac{\partial}{\partial t}(\Delta f)|}\right) \quad (4.12)$$

with $f_R = \Omega_R/2\pi$ and $\Delta f = \Delta\omega/2\pi$. Hence, larger Ω_R or a slower sweep exponentially suppress non-adiabatic transitions. RAP was developed in nuclear magnetic resonance (NMR) experiments [113, 161, 162] and has since been adapted to spin qubit control [163–168].

To start RAP with the on-chip ESR line, the first step is to find safe microwave power levels. A continuous-wave (CW) tone at fixed frequency is applied while monitoring the SET. The onset of thermal broadening of the Coulomb peaks serves as a marker for excessive dissipation. At 18.5 GHz, noticeable broadening appears for source powers $\gtrsim 3$ dBm (Fig. 4.15).

In the triggered RAP sequence, however, the RF duty cycle is typically $\sim 10^{-2}$ (e.g., a ~ 200 μ s chirp within a 20 ms cycle), reducing the time-averaged power at the device by ~ 20 dB relative to CW. This permits the use of higher burst powers with source settings around 20 dBm, without appreciable steady-state heating, while providing a sufficiently large on-chip microwave field to satisfy the adiabatic condition in Eq. (4.11).

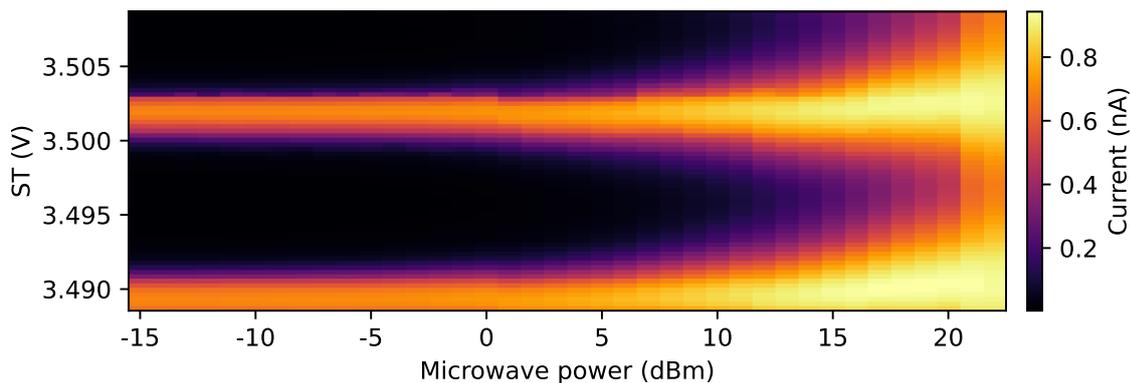


Figure 4.15: Microwave-power induced heating of the on-chip ESR line at 18.5 GHz (device A). Shown is the SET current over topgate (ST) voltage and microwave power. A continuous tone broadens the SET Coulomb peaks once the applied power exceeds ~ 3 dBm. Microwave tone supplied by a Rohde&Schwarz SMW200A signal generator.

The RAP waveform comprises (i) a frequency chirp generated by an AWG (baseband span ≤ 100 MHz) and up-converted with single-sideband IQ mixing to ~ 18.5 GHz (Sec. 3.5), and (ii) synchronised gate pulses on plunger P1 to set the loaded and readout level. To limit spurious events, the microwave burst is terminated ~ 1 ms before the charge-readout window to allow the device to re-thermalise. The microwave source supports hardware gating (RF blanking), whereby an external trigger enables or disables the RF output on demand. This suppresses the continuous background power that would otherwise leak continuously through the IQ mixer (Sec. 3.5), thereby mitigating unwanted device heating and thus background events during readout periods.

Figure 4.16a shows the timing of the microwave tone in respect to the plunger pulse sequence. For coarse frequency spectroscopy (Fig. 4.16b), the static field is swept while the microwave carrier is fixed. This avoids the frequency-dependent transmission variations of the microwave line over a very wide span. A typical coarse scan uses a $210 \mu\text{s}$ chirp across a 15 MHz window at 18.5 GHz with 20 dBm source power, averaging 100 shots per point. During coarse frequency spectroscopy, the readout level on P1 is set to yield a controlled $\sim 20\%$ thermal background, providing headroom for slow working-point drift. For final optimisation the readout level is retuned to minimise false positives and the sweep range and power are refined to maximise visibility ($\sim 80\%$, Fig. 4.16c). Once the resonance is narrowed down to ~ 10 MHz, a discrete frequency ESR search (next section) is used to determine ω_0 precisely.

RAP thus serves two roles in this work: (i) a robust, power- and detuning-tolerant spin-inversion technique for initialisation and control, and (ii) an efficient, high-dynamic-range spectroscopic tool to localise the Larmor frequency across broad search windows despite magnetic field changes by cooldown-to-cooldown variations.

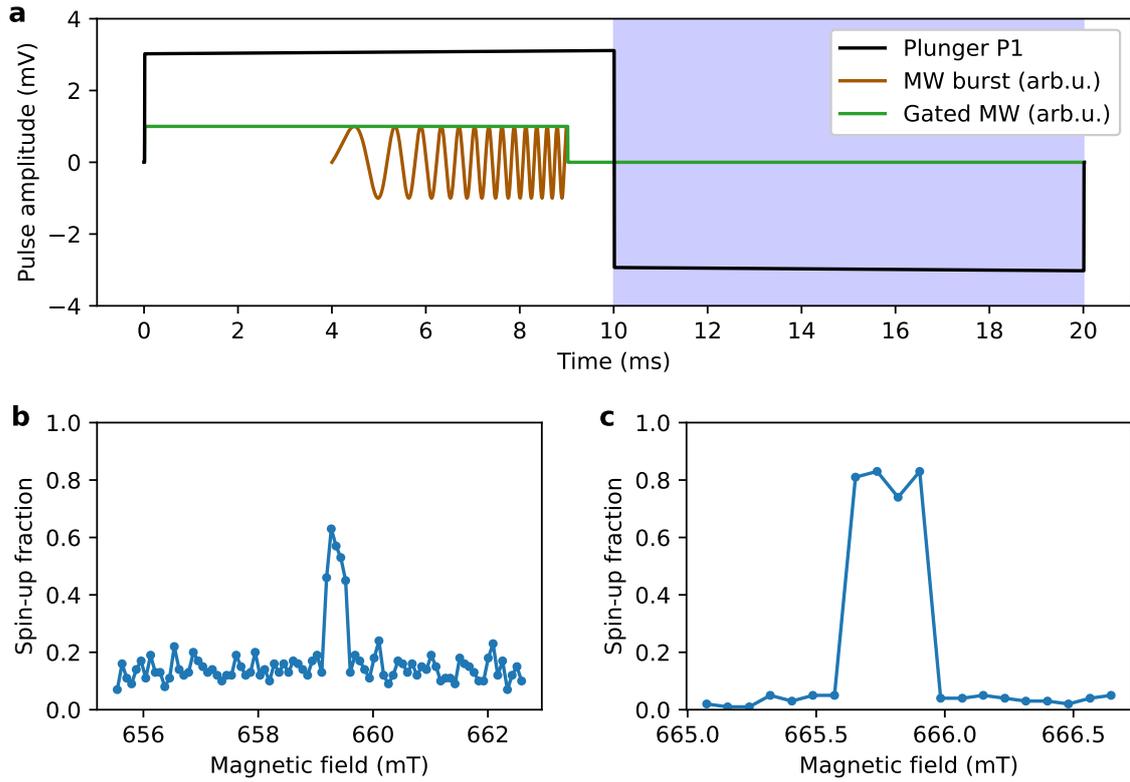


Figure 4.16: Experimental implementation of RAP (device A). **a** Pulse sequence used for spin manipulation. The baseline (0 mV) corresponds to a loaded state with $\mu_{\downarrow}, \mu_{\uparrow} < \mu_{\text{Res}}$. The microwave chirp (brown) is exaggerated in length and frequency for visibility and is always ended ~ 1 ms before the charge-readout window (blue). A hardware RF-gate trigger (green) suppresses mixer leakage during readout, greatly reducing background events. **b** Coarse resonance search via magnetic-field sweep at fixed carrier frequency: chirp length $210 \mu\text{s}$, span 15 MHz centred at 18.5 GHz, source power 20 dBm, 100 averages per point. **c** RAP with the same chirp length but reduced span (10 MHz) and optimised read voltage of P1 at 18.5 GHz and 22 dBm.

4.8 Rabi Oscillations

Tuning the microwave drive to the qubit's Larmor frequency ($\omega_{\text{MW}} = \omega_0$) drives coherent Rabi oscillations between the spin eigenstates (Sec. 2.1.4). In practice, the oscillations decay due to fluctuations of the effective magnetic field experienced by the qubit. Two prominent noise sources are: (i) current noise in the superconducting magnet and electromagnetic pickup on the wiring, which translate into fluctuations of the static field B_0 and (ii) stochastic nuclear-spin flips of residual ^{29}Si nuclei, which generate a fluctuating Overhauser field that adds to B_0 [33]. Both mechanisms produce a time-dependent detuning $\Delta\omega(t) = \omega_0(t) - \omega_{\text{MW}}$. The resulting detuning jitter dephases the qubit state, reducing its phase coherence. Complete dephasing leads to an ensemble-averaged spin-up probability of $P_{\uparrow} = 1/2$.

The characteristic timescale for the decay of phase coherence is denoted T_2 . Since the pure basis states $|0\rangle$ and $|1\rangle$ do not dephase, coherence loss occurs only when the qubit occupies a superposition, e.g. $|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\phi}|1\rangle)$ on the equator of the Bloch sphere. Under

resonant driving the state cycles between eigenstates and superpositions, so the decay of Rabi oscillations is described by a drive-specific coherence time T_2^R , which generally differs from the free-induction time T_2^* obtained from a Ramsey experiment.

Under a coherent drive with on-resonance Rabi frequency Ω_R and detuning $\Delta\omega$, the excited-state probability in the rotating-wave approximation after spin-down initialisation is

$$P_{\uparrow}(t) = \frac{\Omega_R^2}{\Omega_{\text{eff}}^2} \sin^2\left(\frac{\Omega_{\text{eff}}}{2} t\right) \exp\left(-\frac{t}{T_{2,\text{Rabi}}}\right), \quad \Omega_{\text{eff}} = \sqrt{\Delta\omega^2 + \Omega_R^2}, \quad (4.13)$$

where Ω_{eff} is the magnitude of the effective rotation vector in the rotating frame [170, 171]. The exponential envelope captures dephasing from noise frequencies near the drive frequency [170, 172]. Introducing a detuning ($\Delta\omega \neq 0$) tilts the rotation axis and increases the precession rate to Ω_{eff} , while reducing the maximum swing of the population. The corresponding increase of the energy splitting $\hbar\Omega_{\text{eff}}$ between both spin states under detuned driving is also visible in the left panel of Fig. 4.14.

Single-Tone Spectroscopy and Resonance Identification

At fixed magnetic field ($B_0 = 0.66$ T), the drive frequency is swept while applying a long microwave burst (longer than $T_{2,\text{Rabi}}$), so that the detected spin-up fraction approaches a steady $\sim 50\%$ on resonance. To avoid heating during readout, the output of the microwave is switched off 1 ms before readout according to the waveform shown in Fig. 4.16a. Figure 4.17 shows a representative sweep near ~ 18.485 GHz. If the RF source is not hardware-gated, a second, spurious peak is observable. This peak is caused by continuous leakage of the local oscillator (LO) through the IQ mixer and excites the qubit during the readout window. The two peaks are offset by the intermediate frequency (IF, here 100 MHz). The difference in origin of the two peaks is evident in the associated single-shot readout traces: LO leakage produces recurring blips throughout readout, whereas a pulsed drive yields blips only at the beginning of the readout window.

Continuous LO leakage was used to track the resonance over long times, as in Fig. 4.18. Frequency wander within a ~ 0.5 MHz window over 12.5 h is observed, including occasional abrupt jumps and slower drifts. In this setup, stable operation therefore requires periodic re-tuning of the drive frequency (approximately hourly). The drift can be reduced by improved magnetic shielding by μ -metal wrapped around the cryostat to suppress environmental field fluctuations [173, 174].

A rough estimate of T_2^* can be obtained from the ESR linewidth in the low-power limit, where power broadening is negligible. Reducing the drive power narrows the peak until the width is set by inhomogeneous dephasing. In the data of Fig. 4.19, the 20 dBm trace shows clear power broadening, while the 2 dBm trace does not narrow further but loses contrast. The peak that does not exhibit power broadening can be used for precise frequency determination and for estimating T_2^* . The fit of the peak to a Gaussian line shape yields a

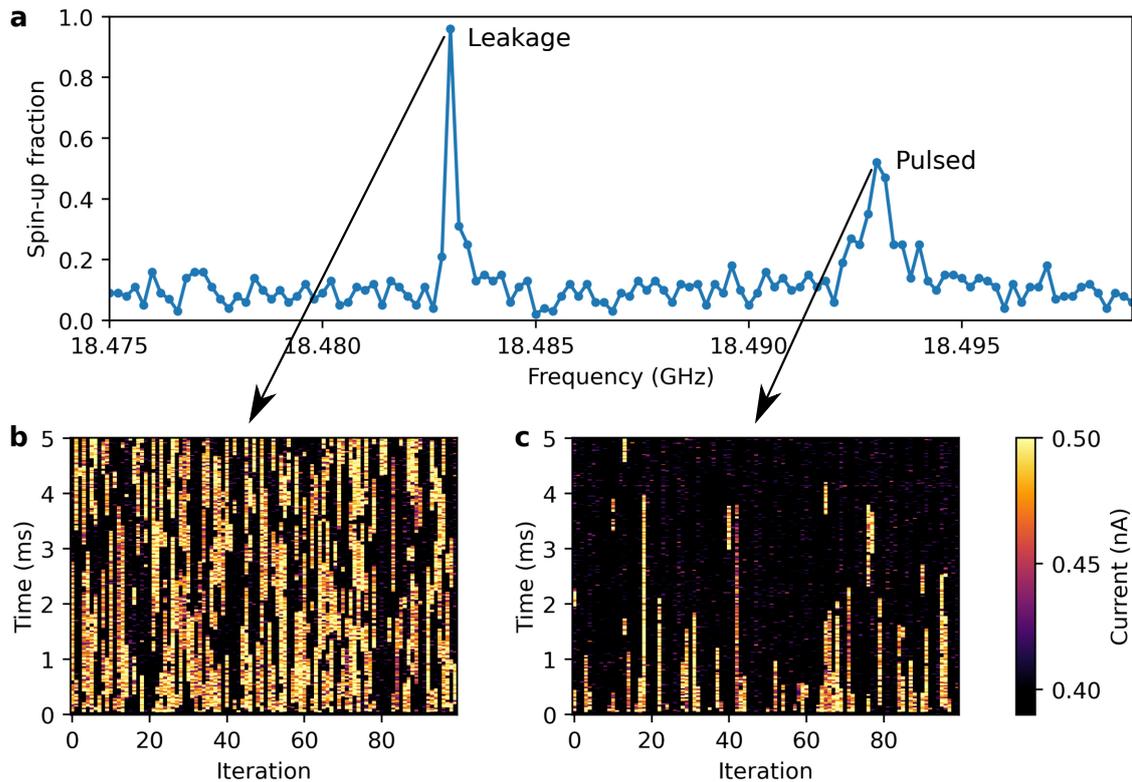


Figure 4.17: Single-tone spectroscopy using a 100 MHz AWG tone mixed with an 18.483 GHz LO tone (device A). The pulse duration is 419 μs . The source's gated mode was disabled, so continuous LO leakage reached the device. **a** Spin-up fraction versus mixed frequency. Two peaks appear: the left peak ("Leakage") arises from continuous LO leakage that excites the qubit during readout; the right peak ("Pulsed") corresponds to the intended, pulsed drive. **b** Raw data of the readout: Continuous excitation by LO leakage drives the spin throughout the readout window, yielding a near-100% spin-up fraction. **c** With pulsed excitation, blips appear only at the start of spin-selective readout.

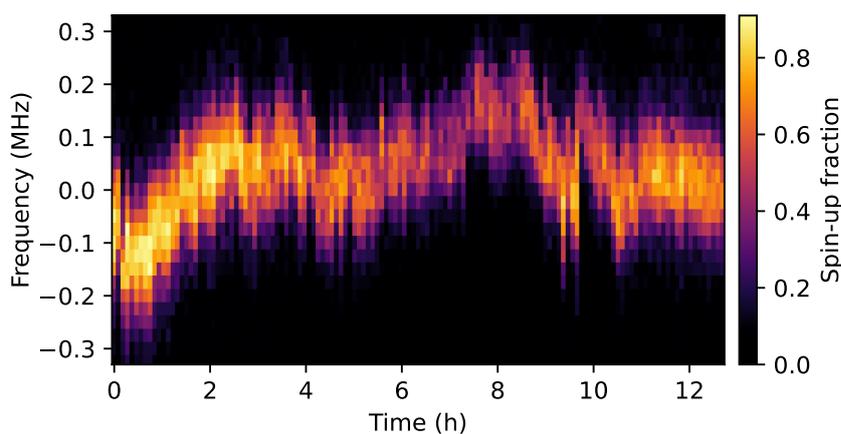


Figure 4.18: Resonance tracking using the mixer's CW leakage (device A). Shown is the spin-up fraction over frequency (offset from 18.5 GHz) over the time of 12.5 h. The resonance shifts within this time ~ 0.5 MHz. Each pixel averages 200 shots.

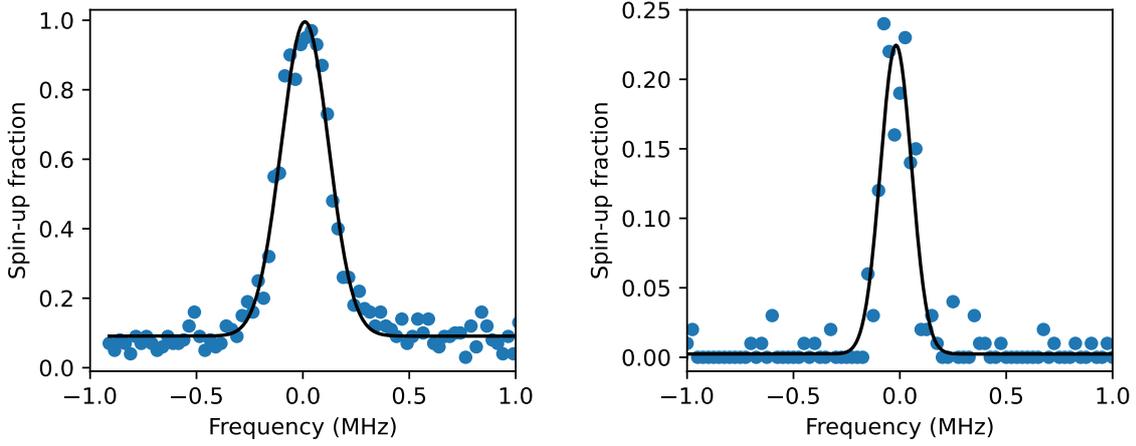


Figure 4.19: Gauss fit (black line) of the qubit resonance (device A). The microwave tone is applied for a pulse duration of $15 \mu\text{s}$ at frequencies offset by 18.485 GHz . Each trace was recorded over 10 min with 100 averages per point. **Left** FWHM = $(263 \pm 7) \text{ kHz}$ at 20 dBm (power broadened). **Right** FWHM = $(172 \pm 5) \text{ kHz}$ at 2 dBm (near low-power limit).

full width half maximum (FWHM) of $(172 \pm 7) \text{ kHz}$. The free-induction time relates to the linewidth as [98, 175]

$$T_2^* = \frac{2\sqrt{\ln 2}}{\pi \text{FWHM}}$$

and yields $T_2^* = (3.1 \pm 0.1) \mu\text{s}$, consistent with Ramsey measurements (Sec. 4.9) for this device.

Rabi Oscillations

Stepping the pulse duration of a resonant microwave drive produces coherent Rabi oscillations [98, 176] which are shown in Fig. 4.20. The oscillation frequency is set by the transverse microwave magnetic field amplitude generated by the on-chip antenna. Writing the field as $B_1(t) = B_1 \cos(\omega_{\text{MW}} t + \phi)$, and noting that a linearly polarised field contributes only its co-rotating component to the Rabi oscillations, the Rabi frequency reads

$$\Omega_{\text{R}} = \frac{g \mu_{\text{B}} B_1}{2\hbar}. \quad (4.14)$$

Using the measured Rabi frequency of $f_{\text{R}} = \Omega_{\text{R}}/2\pi \approx 421 \text{ kHz}$ (with $g \approx 2$) implies that the ESR antenna produces a field amplitude of $B_1 \approx 30 \mu\text{T}$.

Under resonant driving, the state vector precesses about an axis in the equatorial plane of the Bloch sphere, enabling full population transfer $|0\rangle \leftrightarrow |1\rangle$. Introducing a detuning $\Delta\omega \neq 0$ tilts the rotation axis toward z , increases the precession rate to $\Omega_{\text{eff}} = \sqrt{\Delta\omega^2 + \Omega_{\text{R}}^2}$, and reduces the oscillation amplitude to $\Omega_{\text{R}}^2/\Omega_{\text{eff}}^2$ (Eq. (4.13)). The effect is sketched in Fig. 4.21a and observed experimentally as the characteristic chevron pattern in Fig. 4.21b,c.

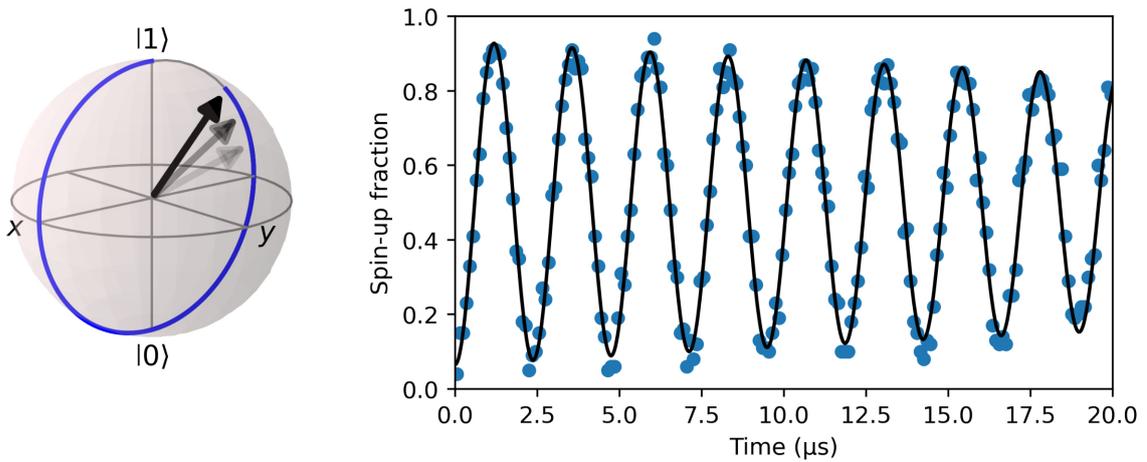


Figure 4.20: Coherent qubit control showing Rabi oscillations (device A). **Left** Bloch-sphere illustration of resonant Rabi driving. By stepping the microwave pulse duration, the state vector is rotated around the drive axis (y -axis) to reach specific target states (black arrows). **Right** Measured Rabi oscillations at a microwave frequency of 18.486 GHz and a source power of 25 dBm, with 100 single shots averaged per point and the microwave source operated in gated mode. A fit (black) of the data to Eq. (4.13) yields a Rabi frequency $f_R = (421.6 \pm 0.3)$ kHz and a Rabi dephasing time $T_{2,\text{Rabi}} = (84 \pm 13)$ μs .

Slow drifts and occasional jumps of the resonance frequency are visible because the maps were recorded over many hours (16.0 h and 22.5 h, respectively).

Fitting the Rabi oscillations provides precise π - and $\pi/2$ -pulse durations, enabling preparation of arbitrary spin superposition states via time-calibrated rotations on the Bloch sphere which are the basic building blocks for subsequent quantum-control experiments.

4.9 Ramsey Experiment

The Ramsey pulse sequence [177] is the standard tool to quantify dephasing from low-frequency noise. A $\pi/2$ pulse, calibrated from Rabi measurements, prepares a superposition $|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\phi}|1\rangle)$ on the equator of the Bloch sphere. After a free-evolution interval τ , the qubit state is projected onto the measurement basis (the z axis) by the second $\pi/2$ pulse (Fig. 4.22b). During the free-evolution interval on the equator, slow fluctuations of the qubit's Larmor frequency, caused by magnetic noise, lead to shot-to-shot variations of the accumulated phase (Fig. 4.22a). Ensemble averaging then causes the transverse component to decay, driving the measured spin-up probability toward $1/2$. The corresponding inhomogeneous dephasing time T_2^* of this decay sets an upper bound on how long phase coherence can be maintained without error mitigation techniques.

During the free evolution interval τ of the Ramsey pulse sequence, noise frequencies much faster than $1/\tau$ are averaged out and so don't contribute to the decay time T_2^* [172]. The pulse sequence is therefore dominated by quasi-static noise frequencies smaller than $1/\tau$.

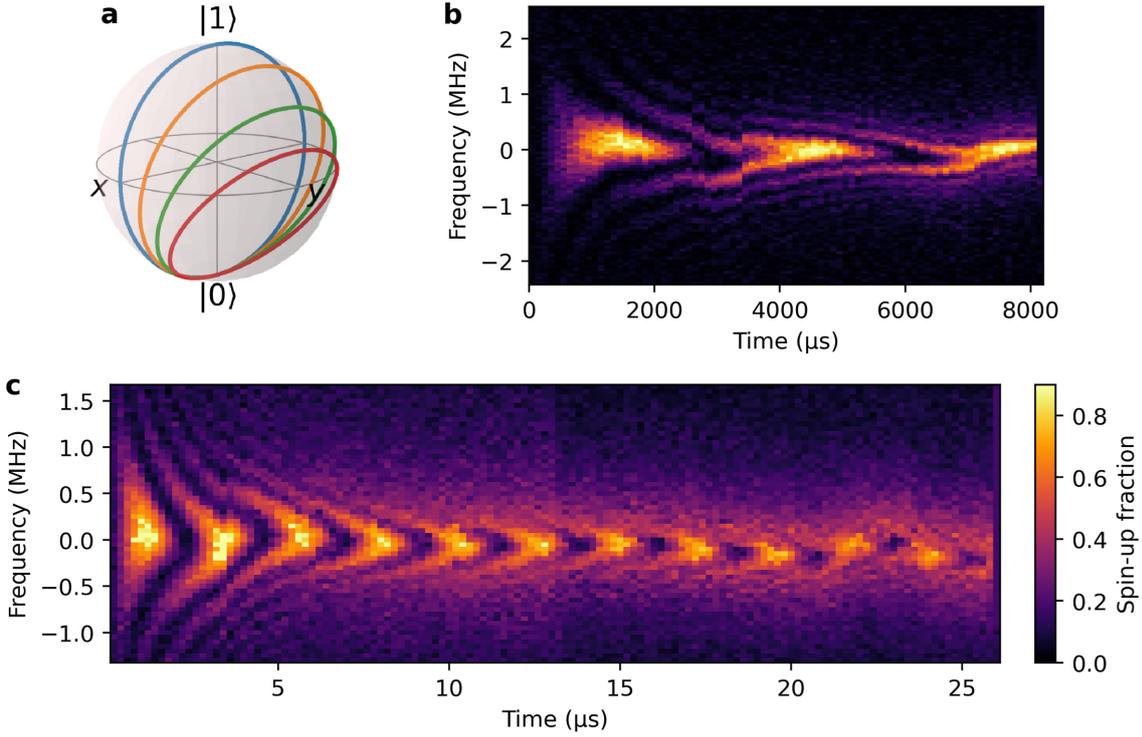


Figure 4.21: Rabi oscillations as a function of drive detuning (device A). **a** Bloch-sphere representation. For a resonant drive (blue), the rotation axis is in y direction and the state undergoes full-amplitude Rabi oscillations at Ω_R . With finite detuning $\Delta\omega \neq 0$ (increasing from orange to green to red), the rotation axis tilts toward z , the oscillation frequency increases to Ω_{eff} given in Eq. (4.13), and the oscillation amplitude is reduced by $\Omega_R^2/\Omega_{\text{eff}}^2$. **b,c** Chevron maps of spin-up probability (shared colorbar) versus pulse duration and microwave frequency for two centre frequencies (offsets 18.485 GHz in **b** and 18.486 GHz in **c**) at ~ 25 dBm. The characteristic chevron pattern shows increasing oscillation frequency and decreasing contrast with increasing detuning. In **c**, the oscillations decay visibly toward a steady $\sim 50\%$ spin-up level as coherence is lost during driving. Recorded over a time span of 16.0 h (**b**) and 22.5 h (**c**).

For robust fitting, we intentionally apply a small detuning $\Delta\omega$ to the drive frequency ω_{MW} so that the Ramsey signal oscillates at $\Delta\omega/2\pi$ while its envelope decays. The data are well described by a Gaussian envelope [39, 178] which is indicative of quasi-static noise. The spin-up probability for spin-down initialisation is given by [98]

$$P_{\uparrow}(t) = A \cos(\Delta\omega t + \phi) \exp\left[-\left(\frac{t}{T_2^*}\right)^2\right] + c, \quad (4.15)$$

where A is the amplitude, ϕ the phase offset set by pulse timing, and c a constant background of thermally excited blip events.

Figure 4.22c shows Ramsey oscillations at $\Delta f = 739$ kHz, yielding $T_2^* = (3.1 \pm 0.1) \mu\text{s}$. A frequency-time map (Fig. 4.22d) displays the characteristic Ramsey fringes. The dephasing times and qubit parameters are discussed at the end of this chapter in Sec. 4.11.

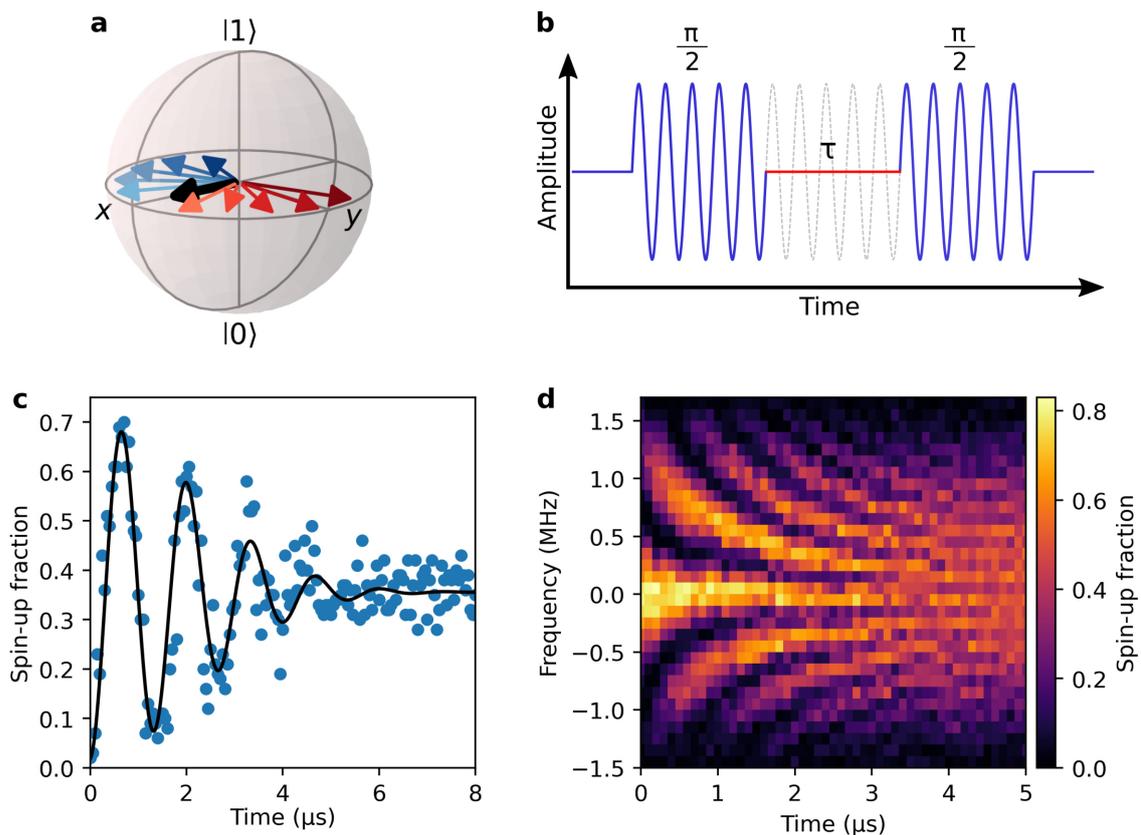


Figure 4.22: Ramsey pulse-sequence measurement (device A). **a** Bloch-sphere schematic: Shown is the qubit state of the ensemble (black arrow) after an initial $\pi/2$ -pulse that brings the qubit state to the equator. Each single shot (indicated by coloured arrows) accumulates phase during the free-evolution time τ . Positive phase accumulation is shown in deeper red while more negative phase accumulation is shown in deeper blue. This fan out of individual states leads to a decay of the average transverse Bloch vector amplitude of the ensemble. **b** Ramsey pulse sequence: two in-phase $\pi/2$ -pulses separated by τ . The second pulse maps the dephased superposition back onto the measurement axis. **c** Measured spin-up probability versus τ at a fixed detuning. A fit to Eq. (4.15) yields $T_2^* = (3.1 \pm 0.1) \mu\text{s}$ at $\Delta f = 739 \text{ kHz}$ (qubit frequency $\sim 18.48 \text{ GHz}$). Each point averages 100 single shots, total acquisition time $\sim 21 \text{ min}$. **d** Ramsey fringes: spin-up probability versus detuning of the Ramsey sequence and τ (frequency axis offset by 18.48 GHz). The oscillatory pattern with a decaying envelope is characteristic of Ramsey interferometry. Measured with a microwave power of 25 dBm .

4.10 Hahn-Echo

Reliable quantum control requires mitigation of dephasing to extend the number of coherent operations. A standard technique to suppress slow environmental field fluctuations is the Hahn-echo pulse sequence [17, 98, 179]. Conceptually, a refocusing π pulse reverses the phase accumulated during the first free-evolution period: if the superposition $|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\phi}|1\rangle)$ acquires a phase shift $+\Delta\phi$ over $\tau/2$, the π pulse maps subsequent evolution to an additional phase shift $-\Delta\phi$, so that quasi-static (slow) detuning noise cancels at time τ . A Hahn-echo therefore filters noise with frequencies $< 1/\tau$, while very

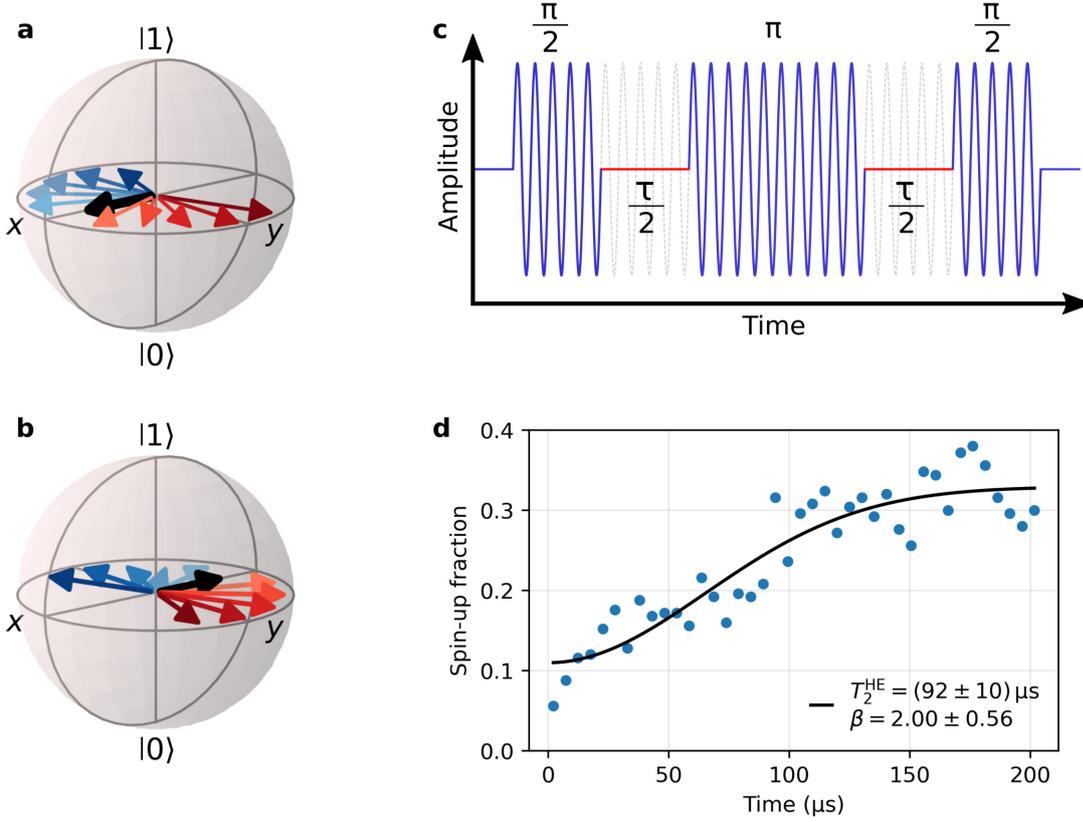


Figure 4.23: Hahn-echo coherence measurement (device A). **a** Bloch-sphere sketch after an initial $\pi/2$ pulse around the y axis. The state is brought on the equator and, during the first free-evolution interval $\tau/2$, individual Bloch vectors (colored arrows) fan out due to quasi-static detuning noise, reducing the average transverse Bloch vector amplitude (black arrow). **b** After a π pulse around the y axis at $t = \tau/2$ the phase accumulation of all trajectories is reversed. For quasi-static detuning, the accumulated phase in the second $\tau/2$ cancels that of the first, so the ensemble refocuses and coherence is recovered at total time τ . **c** Hahn-echo pulse sequence consisting of a $\pi/2$ pulse, an idle evolution time $\tau/2$, a refocusing π pulse, followed by another $\tau/2$ idle period and a final in-phase $\pi/2$ pulse that maps the recovered transverse qubit vector onto the measurement axis. **d** Measured spin-up fraction versus total evolution time τ at a drive frequency of 18.50 GHz and source power ~ 19 dBm. Each point averages 250 single shots. The solid line is a fit to the stretched-exponential model in Eq. (4.16). Fitted parameters are indicated in the panel.

fast noise averages out within each interval. The remaining decoherence is dominated by noise around $1/\tau$.

The complete sequence starting with an initial $\pi/2$ to bring the state to the equator, refocusing π at $\tau/2$, and a final $\pi/2$ to project onto the measurement axis is shown in Fig. 4.23c. In the experiment, the spin-up fraction is recorded as a function of the total free-evolution time τ (Fig. 4.23d), and the decay is fit with a phenomenological stretched exponential [61, 178]:

$$P_{\uparrow}(t) = A \exp \left[- \left(\frac{t}{T_2^{\text{HE}}} \right)^{\beta} \right] + c, \quad (4.16)$$

where A is the contrast, β is the stretching exponent that reflects the underlying noise spectrum, and c accounts for a small offset from background events. The fit yields $T_2^{\text{HE}} = (92 \pm 10) \mu\text{s}$ and $\beta = 2.0 \pm 0.6$. The echo therefore extends coherence by a factor of ~ 30 relative to the Ramsey value $T_2^* = (3.1 \pm 0.1) \mu\text{s}$.

4.11 Discussion of Qubit Characterisation Results

This section summarises the single-qubit characterisation and compares the measured parameters with state-of-the-art SiMOS devices controlled by ESR. Plausible noise sources limiting the observed coherence are identified, and targeted experiments are outlined to disentangle their respective contributions.

The spin-lattice relaxation time is $T_1 = (117 \pm 9) \text{ms}$. Although second-scale T_1 values were observed on other samples during this work, the value for device A is much larger than the timescales relevant here: with a readout duration of 5 ms, decay during readout is small ($< 4\%$). Moreover, T_1 does not limit the homogeneous dephasing time ($T_2 \leq 2T_1$ [172]) as the reported T_2 is three orders of magnitude smaller than this upper bound.

A Rabi frequency $f_{\text{R}} = (421.6 \pm 0.3) \text{kHz}$ is determined which is limited by the available microwave generator power. Together with a driven-decay time $T_2^{\text{R}} = (84 \pm 13) \mu\text{s}$, this yields a Rabi quality of $f_{\text{R}} T_2^{\text{R}} \approx 35$ cycles before the envelope drops to $1/e$. Compared with literature for similar ESR-SiMOS architectures (see Tab. 4.4), both f_{R} and T_2^{R} lie toward the lower end of reported ranges. However, increases in available drive power or adjustments to cryostat attenuation should make higher f_{R} accessible.

Ramsey and Hahn-echo measurements give $T_2^* = (3.1 \pm 0.1) \mu\text{s}$ and $T_2^{\text{HE}} = (92 \pm 10) \mu\text{s}$, respectively. In the comparison of Tab. 4.4, two published T_2^* values are similar to the ones reported here, while several exceed it by roughly an order of magnitude. A similar trend holds for echo coherence. This demonstrates a stronger noise influence in the shown measurements than in many recent reports.

The large separation between the observed T_2^* and T_2 (by a factor of ~ 30) indicates that low-frequency detuning noise is the dominant influence on the qubit. This is consistent with the filter-function picture: the Ramsey sequence is highly sensitive to low-frequency noise set by the free-evolution time τ (roughly $f \lesssim 1/\tau$), whereas the Hahn-echo sequence rejects the DC component and is primarily sensitive to noise at frequencies of order $1/\tau$ [172, 180]. Moreover, an echo exponent of $\beta = 2.0 \pm 0.6$ (Gaussian envelope) is consistent with dephasing limited by a steep low-frequency spectrum, often modelled as $1/f$ [178, 180].

Noise Sources

- (i) Residual ^{29}Si nuclei (spin 1/2) produce an Overhauser field that adds vectorially to the applied field and fluctuates on a wide range of timescales due to nuclear dipolar

Table 4.4: Comparison of qubit parameters for device A with other publications of SiMOS devices controlled by ESR, ordered chronologically by publication year. Uncertainties of the parameters are mostly not reported and thus are omitted for clarity. In addition to performance metrics, the table lists the residual ^{29}Si concentration and the static magnetic field B used for qubit manipulation. Blank entries indicate values not reported.

Parameter	This work	Ref.[101]	Ref.[146]	Ref.[181]	Ref.[62]	Ref.[100]	Ref.[53]	Ref.[58]
T_1 (s)	0.117		1		1.6			9.5
f_R (MHz)	0.421			0.33	0.900	1.5		0.658
T_2^R (μs)	84	380		18.6				
T_2^* (μs)	3.1	120	33	3.3	24	2.3	8.1	40.6
T_2^{HE} (ms)	0.092	1.2	0.401			0.012		1.9
^{29}Si (ppm)	400	800	800	800	800	50		400
B (T)	0.66	1.4	1.4	0.45	0.61	0.23	0.82	0.05-0.50
Year	2025	2014	2018	2019	2022	2022	2024	2025

flip-flops [182, 183]. This appears as quasi-static detuning noise changing in each single-shot measurement.

- (ii) External magnetic fields from the measurement environment (e.g. from mains transformers) introduce narrowband components at 50 Hz and its harmonics.
- (iii) Current noise in the vector magnet power supply introduces also magnetic field noise at 50 Hz and its harmonics. This can come from ground-loop pickup or residual supply ripple.
- (iv) Gate voltage noise modulates the vertical electric field E_z at the dot position. Through spin-orbit coupling at the Si/SiO_x interface this produces a Stark shift of the electron g -tensor, so that fluctuations of E_z translate into Larmor-frequency noise and dephasing [39, 101, 146].
- (v) Charge noise (trap occupancy changes) at the Si/SiO_x interface acts as an effective gate-voltage noise [146]. This also leads to dephasing by the already mentioned Stark shift.

Given that device A is fabricated using isotopically enriched silicon with residual ^{29}Si concentration of approximately 400 ppm (far below natural silicon's 4.7% ^{29}Si concentration [184]), hyperfine-induced inhomogeneous broadening should be strongly suppressed. Together with the literature values in Tab. 4.4 (including higher T_2^* at higher ^{29}Si concentration), this suggests that other noise sources are the dominant limitations here. At present the available data do not single out a unique mechanism. In the following, targeted experiments are outlined to discriminate between the candidate noise sources.

Experiments to Identify the Dominant Noise Source

The influence of current noise in the vector magnet can be tested using a different coil setup with a changed current-to-field conversion factor. Because the required DC current for a given field then differs, the same absolute supply noise maps to a different field noise, providing a clear diagnostic. Ideally, experiments are performed with a persistent-mode magnet that, once energised, is independent of the current source and therefore cannot contribute additional field noise. This effectively removes the vector magnet as a noise source. Magnetic noise from the laboratory environment can be reduced and tested by installing μ -metal shielding around the cryostat.

The Overhauser-field contribution from residual ^{29}Si can be tested by characterising newer device generations with lower residual concentrations. If a systematic change in coherence time is observed with concentration, the nuclear bath is implicated as a source of the low-frequency noise.

Gate noise can be varied by adjusting the filtering. The present two-stage low-pass with a 250 Hz cutoff appears reasonable but can be made more stringent to assess its impact. Alternatively, battery powered voltage sources can be employed to eliminate ground loop effects. By reducing voltage noise on the gates the Stark contribution $\partial\omega/\partial V_{\text{gate}}$ is effectively modified. If charge noise intrinsic to the chip is responsible for electrostatic shifts, a small plunger-voltage offset, which is compensated by nearby confinement gates, can reduce the vertical field E_z at fixed electron occupancy. This weakens spin-orbit coupling and thereby reduces the Stark shift [146] which should lower the associated dephasing contribution.

In conclusion, this chapter provided the introduction to single qubit measurements and a characterisation of single-qubit parameters of device A, which have also been compared to published literature values. The same qubit is then used in the next chapter (Ch. 5) for sensing applications.

5 Quantum Magnetometry with Single-Molecule Magnets Coupled to Spin Qubits

Building on the framework of the hybrid system developed in Ch. 2, this chapter presents the experimental magnetometry of single-molecule magnets (SMMs) using an industrial silicon metal-oxide-semiconductor (SiMOS) spin qubit as a nanoscale sensor. The content is adapted from and extends the results reported in “Hybrid Quantum Systems: Coupling Single-Molecule Magnet Qudits with Industrial Silicon Spin Qubits” [185].

The SMM in this implementation is $[\text{TbPc}_2]^-$, as described in Sec. 2.2. Canonical behaviours of the electronic spin state including magnetic hysteresis, angular anisotropy and temperature-dependent relaxation of $[\text{TbPc}_2]^-$ are demonstrated via shifts in the qubit’s resonance frequency.

The present SiMOS process at imec is not yet optimised for hybrid operation (see chip optimisations in Ch. 7), resulting in a vertical spacing of $\gtrsim 250$ nm between qubit and magnetic material. At such distances the dipolar field from a single molecule is insufficient for reliable detection as has been discussed before in Sec. 2.3. Therefore a dilute crystal comprising 5% $[\text{TbPc}_2]^-$ and 95% diamagnetic YPc_2 is employed, which amplifies the stray field at the silicon qubit while preserving the intrinsic SMM character by minimising intermolecular interactions.

The measurements rely on a protocol that enables continuous and robust tracking of the qubit resonance: Rapid adiabatic Passage during Spin-Selective tunnelling (RPSS). In RPSS, a single frequency-chirped microwave pulse is applied concurrently during spin-selective readout, providing high-throughput resonance detection over a 100 MHz window without repeated pulse recalibration. First, the readout statistics and the visibility metric used throughout the analysis are established, and the calibration of the three-axis vector magnet that permits controlled field-orientation sweeps at fixed magnitude is summarised. The magnetometry results obtained with RPSS are then presented, including hysteresis along the easy axis, angular anisotropy and thermal relaxation dynamics of the $[\text{TbPc}_2]^-$ ensemble, followed by a brief discussion of sensitivity and constraints of the current implementation.

Measurement Principle

The underlying detection principle for sensing the SMM stray field with the silicon qubit relies on field changes caused by switching events of the $[\text{TbPc}_2]^-$ molecules. These events modify the total magnetic field \mathbf{B}_{tot} seen by the qubit. The external field \mathbf{B}_{ext} and the SMM stray field at the qubit site \mathbf{B}_{SMM} add vectorially,

$$\mathbf{B}_{\text{tot}} = \mathbf{B}_{\text{ext}} + \mathbf{B}_{\text{SMM}},$$

and set the Larmor frequency of the silicon qubit,

$$f_0 = \frac{g \mu_B}{h} |\mathbf{B}_{\text{tot}}|,$$

with Landé factor $g \approx 2$ and Bohr magneton μ_B . In the implemented configuration, the SMM produces a frequency shift of about 34 MHz on a baseline resonance near 18.4 GHz.

The external magnetic field supplied by a 3D vector magnet is set to $B_{\text{ext}} = 0.66$ T providing a sufficiently large Zeeman splitting to enable the energy-selective (Elzerman [157]) readout of both qubit spin states (Sec. 4.5). \mathbf{B}_{ext} is aligned along a hard axis of the molecule (perpendicular to the easy axis) to minimise perturbation of the intrinsic switching behaviour.

Figure 5.1 shows the $[\text{TbPc}_2]^-$ crystal positioned on the SiMOS device using a wire-bonder as a micromanipulator under ambient conditions. The crystal has an edge length of approximately $50 \mu\text{m}$ and adheres to the chip via a thin layer of Apiezon N grease, which also provides thermal anchoring to the cryostat cold stage. The magnetic easy axis is readily aligned because it coincides with a crystal face. As shown in Fig. 2.6, the molecular symmetry axis is parallel to a unit-cell vector, offering a clear crystallographic reference for alignment.

A slight lateral displacement of the crystal along the chip surface (x -axis) is supposed to enhance the projection of the SMM stray field onto \mathbf{B}_{ext} . Denoting by θ the angle between \mathbf{B}_{ext} and \mathbf{B}_{SMM} at the qubit location, the field magnitude is

$$B_{\text{tot}} = \sqrt{B_{\text{ext}}^2 + B_{\text{SMM}}^2 + 2B_{\text{ext}}B_{\text{SMM}} \cos(\theta)}.$$

Placing the crystal edge above the qubit so that the bulk is laterally offset targets $\theta \approx 0^\circ$ (or 180°), thereby maximising the influence of B_{SMM} on B_{tot} . In this case B_{SMM} adds linearly to B_{ext} , maximising the resonance shift. This is illustrated schematically in Fig. 5.2.

5.1 Rapid Adiabatic Passage during Spin-Selective Tunnelling (RPSS)

Detecting small shifts of the silicon-qubit resonance f_0 necessitates a control-and-readout sequence that is both sensitive and robust. Rapid Adiabatic Passage (RAP; see Sec. 4.7) is

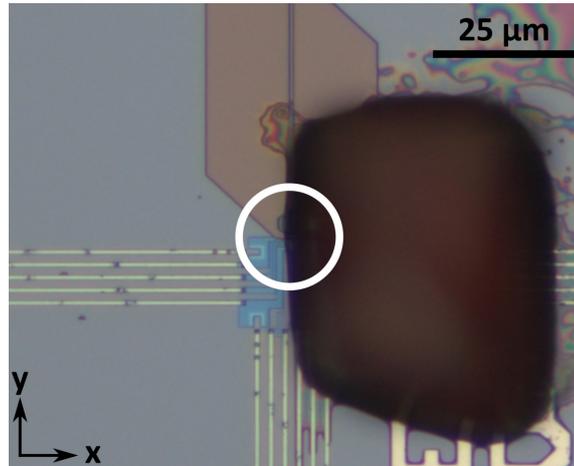


Figure 5.1: SMM crystal placement. Optical image of a dilute crystal (dark brown) comprising 5% $[\text{TbPc}_2]^-$ and 95% diamagnetic YPc_2 positioned on the SiMOS device. Apiezon N grease provides thermal contact and adhesion. The molecular easy axis is oriented out of the plane (z -direction). The spin qubit is located at the centre of the white circle, with its quantisation axis set by an external magnetic field B_{ext} applied along the x -axis. The on-chip microwave antenna (light brown) enters from the top, providing an oscillating magnetic field along the y -axis. DC gate lines (light green) connect from the left, right and bottom edges.

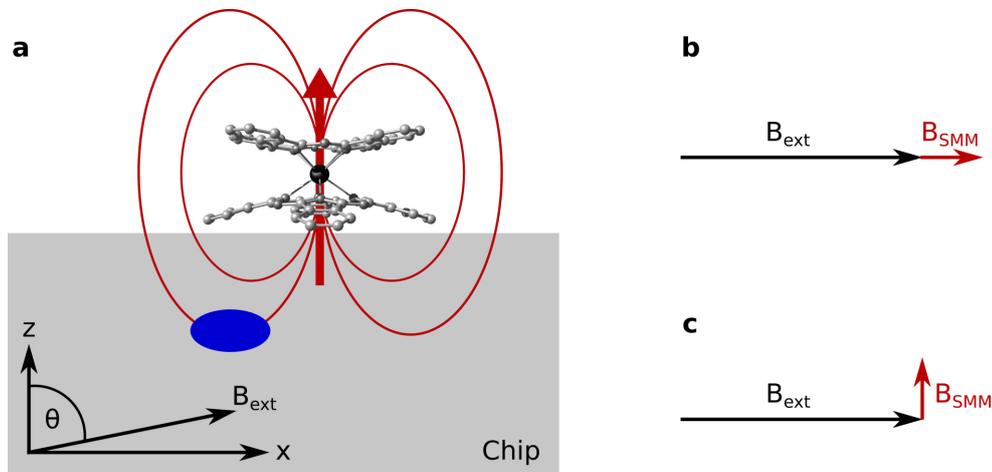


Figure 5.2: Magnetic-field alignment and lateral displacement of the SMM. **a** Schematic of the alignment. A qubit (blue) beneath the chip surface (grey) is coupled to a $[\text{TbPc}_2]^-$ single-molecule magnet positioned above. The molecular easy axis (red arrow) is normal to the chip, and the SMM dipolar field lines are shown in red. Only one molecule is shown for simplicity, while in the experiment a diluted ensemble is used. An external magnetic field B_{ext} of fixed magnitude is rotated by an angle θ in the x - z plane. A lateral offset of the SMM along x is supposed to enhance the x component of the stray field at the dot position. **b** With lateral displacement as in **a**, at $\theta = 90^\circ$ the SMM field B_{SMM} is collinear with B_{ext} , giving $B = B_{\text{ext}} + B_{\text{SMM}}$ and a maximised frequency shift. **c** Without lateral displacement, at $\theta = 90^\circ$ the SMM field is orthogonal to B_{ext} , so $B = \sqrt{B_{\text{ext}}^2 + B_{\text{SMM}}^2}$ and the effect on the qubit resonance is much smaller.

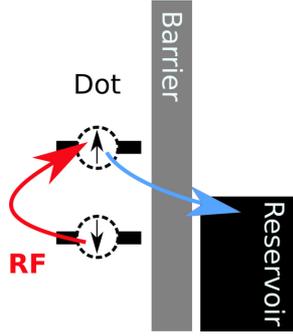


Figure 5.3: Simultaneous spin excitation and readout in the Elzerman configuration. A quantum dot with Zeeman-split spin states is tunnel-coupled to a reservoir through a barrier. The reservoir Fermi level is biased between the spin-dependent energy levels of the quantum dot, so with fixed gate voltages the dot is naturally initialised in $|\downarrow\rangle$. A radio-frequency (RF) pulse drives $|\downarrow\rangle \rightarrow |\uparrow\rangle$ (red arrow) by RAP. The $|\uparrow\rangle$ state can then tunnel to the reservoir with characteristic time $\sim 1/\Gamma$ (blue arrow), where Γ is the tunnel rate. The resulting lateral charge motion is detected by a nearby electron transistor. Subsequently, a $|\downarrow\rangle$ electron reloads from the reservoir, reinitialising the qubit.

used to invert the spin and thereby locate f_0 : the microwave frequency is swept across the resonance at fixed amplitude. When the sweep rate of the detuning Δf satisfies the adiabatic condition

$$\left| \frac{\partial \Delta f}{\partial t} \right| \ll 2\pi f_R^2, \quad (5.1)$$

with f_R the Rabi frequency, the state follows the instantaneous eigenstate and full spin inversion is obtained [113, 162]. The corresponding inversion probability is captured by Landau-Zener theory [88] (see Eq. (4.12)).

The appeal of RAP is that full spin inversion does not rely on precise π -pulse timing, instead satisfying the adiabatic criterion is sufficient. In contrast, conventional electron-spin resonance (ESR; Sec. 4.8) demands accurate pulse control. Because f_0 must be identified within a ~ 100 MHz window, and the microwave transfer function varies with frequency due to the signal chain, an ESR-based approach would require re-calibrating the π -pulse at each frequency point. This creates substantial overhead and forces fine frequency steps smaller than the typical linewidth ($\gtrsim 172$ kHz; see Fig. 4.19) to track the resonance reliably. RAP avoids both issues and permits also much coarser frequency stepping without re-calibration.

In standard qubit operation (Secs. 4.7 and 4.8), the dot levels are pulsed well below the reservoir Fermi level during microwave excitation to suppress thermally assisted tunnelling to the reservoir that may be enhanced by ohmic heating. Manipulation is then followed by a dedicated readout window, typically 5 ms to 10 ms in our setup. Resolving f_0 across 100 MHz with 1 MHz spacing would therefore take $\gtrsim 0.5$ s per single average.

To reduce the acquisition time substantially, RAP is performed *during* spin-selective readout. Specifically, the dot is kept at the Elzerman readout point [157] while the microwave frequency is chirped. In this RPSS mode, a 100 MHz scan lasts only about 0.15 s, cutting

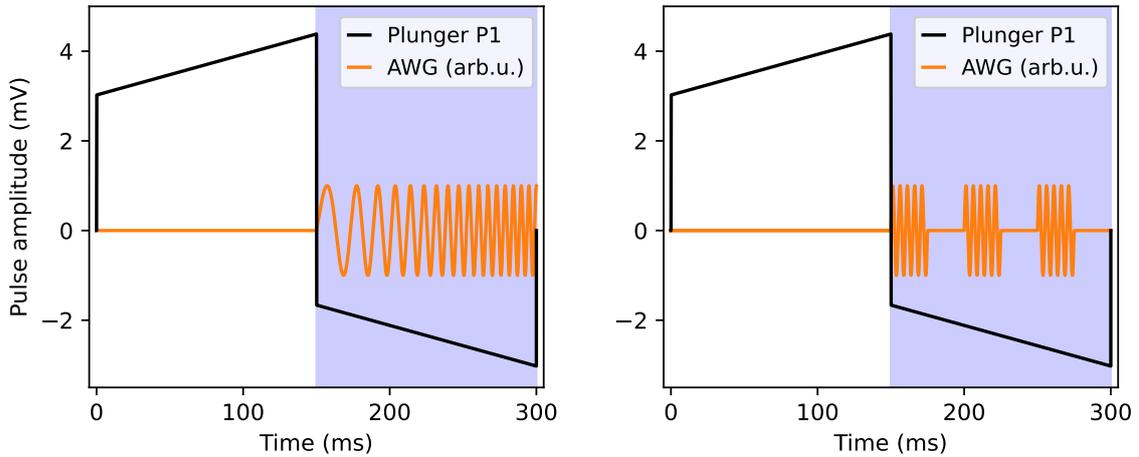


Figure 5.4: Pulseforms for RPSS and ESR tracking. **Left** RPSS: a microwave chirp (orange) is applied while the plunger (black line) is held at the readout point (blue shaded region). A small linear ramp compensates bias-tee discharge, and an initial inverted segment removes residual charge to enable seamless repetition. **Right** ESR tracking during Elzerman readout: identical plunger timing while discrete fixed-frequency tones from the AWG are used to drive ESR. Mapping the blip time to the discrete ESR frequency yields f_0 .

the resonance-tracking time by more than a factor of three. The pulse sequence is detailed in Fig. 5.4.

The RPSS scheme is illustrated in Fig. 5.3. The reservoir Fermi level is positioned between the Zeeman-split qubit states $|\downarrow\rangle$ and $|\uparrow\rangle$. The $|\downarrow\rangle$ state is stable, whereas $|\uparrow\rangle$ tunnels to the reservoir and produces a characteristic SET response (“blip”). During RPSS, a frequency chirp is applied in the readout configuration of the qubit. Whenever the momentary chirp frequency crosses f_0 , RAP transfers population to $|\uparrow\rangle$, which promptly tunnels. Correlating the blip time with the chirp frequency yields f_0 . In this way, f_0 is extracted from a continuous frequency sweep rather than by iteratively stepping the microwave frequency and performing separate readouts at each point.

During RPSS operation, the device is held effectively at a fixed energy alignment between dot and reservoir. In practice, the plunger potential was applied via a high-frequency line (total attenuation ≈ 33 dB, i.e. $\sim 1/44$ in voltage) to provide fine resolution. Because the RF drive is capacitively coupled through a resistive bias-tee, the programmed pulse (Fig. 5.4) is combined with a small compensating ramp to counteract bias-tee discharge (Sec. 3.6). An inverted pre-segment (during $t \leq 150$ ms) is used to remove residual charge on the bias-tee and allow continuous repetition. While the plunger was pulsed in this implementation, RPSS can in principle be executed without any plunger pulsing. The main requirement being sufficient DC voltage resolution which can be attained by additional voltage dividers (Sec. 3.7).

The plunger voltage P1 that places the quantum dot at the energy-selective readout point is determined using an Elzerman pulse sequence [157] (Fig. 4.11). In calibration measurements of the RPSS technique, the AWG applies a frequency chirp of 30 MHz during a 150 ms readout, centred on the expected resonance f_0 . The chirp is upconverted with a

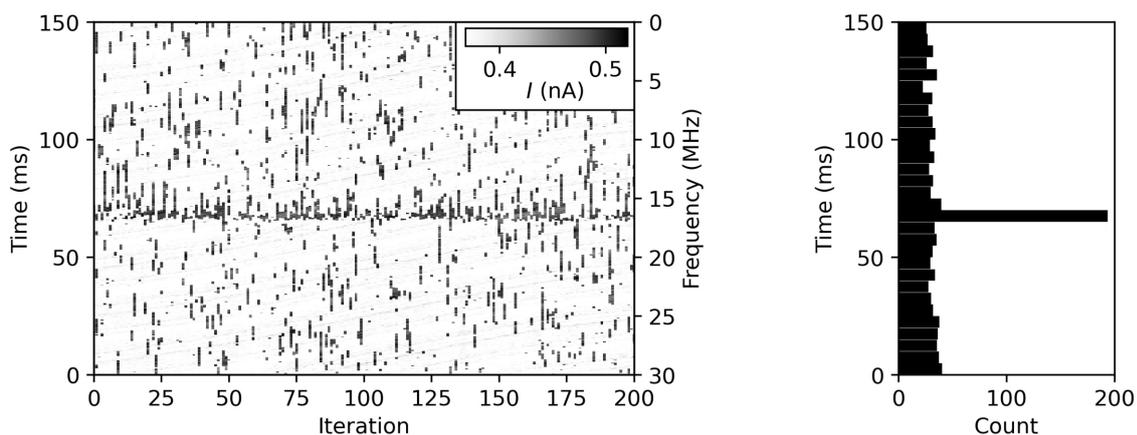


Figure 5.5: Spin-resonance detection using RPSS. **Left** Raw data from 200 single-shot measurements, each acquired during a 150 ms chirp of 30 MHz (right axis) offset by 18.57 GHz, corresponding to a sweep rate of 0.2 GHz s^{-1} . The sweep proceeds from high to low frequency (due to frequency mixing) at a source power of 8 dBm. The tunnelling rate is $\approx 300 \text{ Hz}$. **Right** Histogram of blip start times from the left panel, converted to frequency via the sweep rate, showing a pronounced peak at the Larmor resonance. A bin width of 1 MHz is used. The bin at resonance contains 97% spin-up events, whereas the average background in off-resonant bins is 16%. The visibility is defined as the difference between these on- and off-resonant spin-up fractions.

single-sideband mixer (Sec. 3.5) using a microwave source at 18.6 GHz and 8 dBm output power (setup attenuation in Sec. 3.4). The resulting sweep rate is 0.2 GHz s^{-1} .

The left panel of Fig. 5.5 shows 200 consecutive single-shot traces of the SET current. High (low) current corresponds to an unoccupied (occupied) dot. Each trace is thresholded and binarised (see Sec. 4.5) to identify individual tunnelling events (blips) that signify spin-up electrons escaping to the reservoir. A line of blips appears at a consistent time of $\approx 68 \text{ ms}$ across traces, indicating a well-defined resonance at 18.569 GHz. On average, each trace contains 5.5 blips, of which 4.7 arise from thermally excited ground-state electrons and constitute a background signal.

To extract the resonance frequency, the start times of all blips are histogrammed (right panel of Fig. 5.5) and mapped to frequency using the known sweep rate. A fixed bin width of 1 MHz is chosen so the resonance is captured within a single bin. The resonance visibility is defined as the difference between the on-resonance spin-up fraction (fraction of traces with a blip start in the maximum bin) and the average spin-up fraction in the neighbouring off-resonant bins. This metric is used throughout this chapter to quantify signal contrast. The bin with the highest count is taken as the qubit resonance, which converts to magnetic field via the electron gyromagnetic ratio $\gamma_e \approx 28 \text{ GHz T}^{-1}$ [186], hence 1 MHz corresponds to $\sim 36 \mu\text{T}$.

Robust Blip Identification

The starting times of the blips were identified with an advanced blip-detection tool, the Python library Denoise [187], developed by Noah Gläser [188]. It offers two key advantages

for the binarisation of SET-current traces. First, it automatically and robustly determines the HIGH/LOW levels and the discrimination threshold, remaining stable even under slow working-point drift and avoiding manual tuning. Second, rather than simple pointwise thresholding, it applies total-variation binarisation with a step penalty that suppresses noise-induced toggling. This prevents a single physical blip from being fragmented into multiple short events, preserves sharp leading edges, and yields accurate blip start times thereby improving the precision of the extracted resonance frequency by RPSS.

Alternative Cumulative Blip Count

An alternative to binning is to accumulate the extracted blip start times into a cumulative count as a function of time (and, via the known sweep rate, frequency). Fitting this cumulative distribution (Fig. 5.6) with a cumulative distribution function (CDF) [189] of a normal distribution, augmented by a linear background, locates the point of maximum slope, which corresponds to the qubit resonance responsible for the blip events. The fit model for the number of accumulated blips is

$$N_{\text{blips}} = a \text{CDF}\left(\frac{f - f_0}{\sigma}\right) + b f + c, \quad (5.2)$$

where f is the momentary chirp frequency, f_0 the Larmor frequency, σ the standard deviation, a the on-resonance amplitude, and b, c parameterise a linear background due to thermally activated blips.

For the cumulative blip count shown in Fig. 5.6, the fit yields $|a| = 162$, i.e. the number of on-resonance blips. With 200 averages this corresponds to a visibility of 81%. Overall, 1107 blips occurred (given by c), corresponding to about 4.7 thermal-background blips per single-shot measurement. While the cumulative fit returned a nominally finer frequency uncertainty of $\sigma = 0.2$ MHz, it proved less robust in automated fitting, so binned histograms were adopted for the main analysis.

RPSS Visibility Analysis

In calibration measurements the dependence of the RPSS resonance visibility on microwave power and sweep rate was investigated. The left panel of Fig. 5.7 shows that, for a sweep rate of 0.2 GHz/s, the visibility increases with drive power and saturates at approximately 89% around 3 dBm. For powers $\gtrsim 5$ dBm the background rate rises and suppresses the net signal, which can be attributed to thermally activated tunnelling that empties the dot before the spin is driven into resonance. In that situation the microwave drive still crosses f_0 , but the quantum dot is already unoccupied so neither excitation nor a subsequent blip occurs, leading to a dead time of the sensor.

To extract the Rabi frequency, the power dependence was fitted with a Landau-Zener model Eq. (4.12), using $f_{\text{R}} \propto 10^{P_{\text{MW}}/20}$, where P_{MW} denotes the applied microwave output

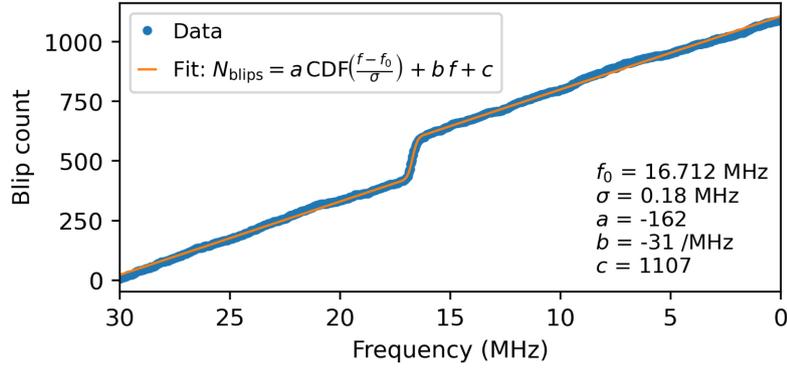


Figure 5.6: Cumulative blip count analysis. The same dataset as in the left panel of Fig. 5.5 is analysed without 1 MHz binning: all blip start times are cumulatively summed over time and mapped to frequency using the sweep rate. The data are fitted with Eq. (5.2) (normal CDF plus linear background) to account for resonance-driven and thermal blips. The fit returns a resonance at $f = 16.7$ MHz offset from 18.57 GHz, with $\sigma = 0.2$ MHz, providing finer nominal frequency resolution than the 1 MHz-binned approach.

power in dBm. This yields $f_R = (17.8 \pm 0.3)$ kHz at 8 dBm. The fit (black line) was restricted to $P_{MW} \leq 5$ dBm because the increased background at higher powers reduces the apparent visibility in a way not captured by Eq. (4.12).

Adiabaticity for the data in the left panel of Fig. 5.7 is assessed via the condition Eq. (5.1). At the applied power level of 8 dBm it holds that $2\pi f_R^2 \approx 1.98$ GHz/s, so the applied sweep rate of 0.2 GHz/s satisfies the adiabatic criterion. Consistently, evaluating Eq. (4.12) predicts spin inversion close to unity. The histogram in the right panel of Fig. 5.5, acquired at identical parameters, returns a visibility of 81 %, indicating full inversion in combination with a non-optimal readout fidelity.

The right panel of Fig. 5.7 plots the visibility as a function of sweep rate at fixed drive power of 8 dBm. The Landau-Zener model Eq. (4.12) predicts about 96 % inversion even at 1 GHz/s, yet the measured visibility decreases sharply. The reason is that the binning is fixed in frequency, not in time. The time spent in one frequency bin is

$$\Delta t = \frac{\Delta f_{\text{bin}}}{\left| \frac{\partial}{\partial t} \Delta f \right|},$$

so larger sweep rate yields a shorter Δt for each 1 MHz bin. With a tunnel rate $\Gamma \approx 300$ Hz the mean blip onset time is $1/\Gamma \approx 3.3$ ms. If $\Delta t < 1/\Gamma$, many tunnelling events occur outside the sampling window of a bin, fewer blips are registered inside a single bin, and the visibility appears reduced. Slower sweeps increase Δt and the blip visibility, but require longer acquisition times and larger waveform memory.

The dot-reservoir coupling also introduces a trade-off. Stronger coupling reduces the timing jitter of the blip onset and therefore lowers the standard deviation of the f_0 estimate under RPSS. At the same time it shortens the blip duration. Very short blips are attenuated by the low-pass filtered readout and can fall below the detection threshold, which reduces the overall blip count. In practice, a suitable operating point balances minimal scatter in f_0

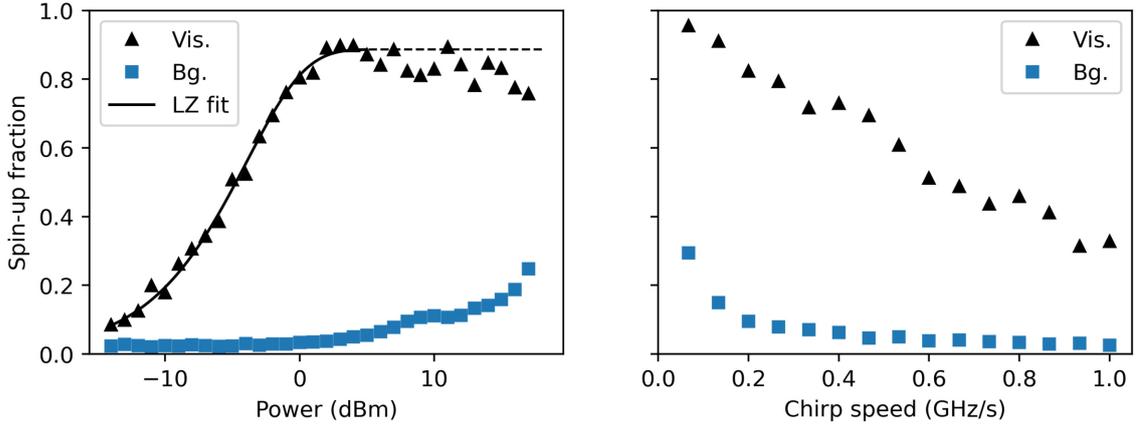


Figure 5.7: Analysis of the RPSS visibility. **Left** Resonance visibility (black triangles, “Vis.”) and background (blue squares, “Bg.”) as a function of applied microwave power at a fixed sweep rate of 0.2 GHz/s. The power dependence (with P_{MW} in dBm) is fitted for $P_{MW} \leq 5$ dBm (black line) using a Landau-Zener based model, $P_{\uparrow} = A \left(1 - \exp\left(-\pi^2 (f_R^{8 \text{ dBm}} \sqrt{10^{(P_{MW}-8)/10}}) / 0.2 \text{ GHz/s}\right) \right)$, yielding $f_R^{8 \text{ dBm}} = (17.8 \pm 0.3) \text{ kHz}$ and $A = (89 \pm 1) \%$. **Right** At fixed power (8 dBm) the finite tunnel rate reduces the visibility with increasing sweep rate: a 1 MHz frequency bin spans a progressively shorter time interval, decreasing the probability of a blip falling into the bin. In both panels the microwave frequency is centred at 18.57 GHz.

with a sufficiently high detection probability. A coupling rate of about 300 Hz provides such a compromise.

For the magnetometry measurements on TbPc_2 presented in Sec. 5.3, the chirp span was extended to 100 MHz to capture the full hysteresis loop. With the pulse length fixed at 150 ms the corresponding sweep rate is 0.67 GHz/s, which reduces the RPSS visibility to below 50 %. The contrast remained sufficient to localise the qubit resonance.

The field sensitivity of RPSS was evaluated from full resonance spectra acquired with 50 averages taking 15 s. With a histogram bin width of $\delta f = 1 \text{ MHz}$, the corresponding field increment is $\delta B = \delta f / \gamma_e = 1 \text{ MHz} / 28 \text{ GHz T}^{-1} \approx 36 \mu\text{T}$. This yields a field sensitivity of $\delta B \sqrt{T} \approx 36 \mu\text{T} \sqrt{15 \text{ s}} \approx 138 \mu\text{T} / \sqrt{\text{Hz}}$. The frequency window used for tracking is 100 MHz, which maps to a sensitive magnetic-field span of 3.6 mT. This resolution is sufficient to follow the dynamic magnetic response of the $[\text{TbPc}_2]^-$ ensemble.

Alternative to RPSS: Rabi Excitation during Spin-Selective Tunnelling

Using the protocol of simultaneous spin readout and manipulation (Fig. 5.3), both continuous-frequency sweeps (RAP) and discrete-frequency excitations (Rabi experiments) can be implemented and are shown here as a second way to track the qubit resonance. In the discrete case a sequence of fixed-frequency microwave pulses is applied, each targeting a specified detuning. Spin excitation occurs only when the applied frequency matches the qubit resonance, and for pulse durations $t \gg T_2$ the time-averaged spin-up fraction is limited to a maximum of 50%. The pulse form is shown in the right panel of Fig. 5.4.

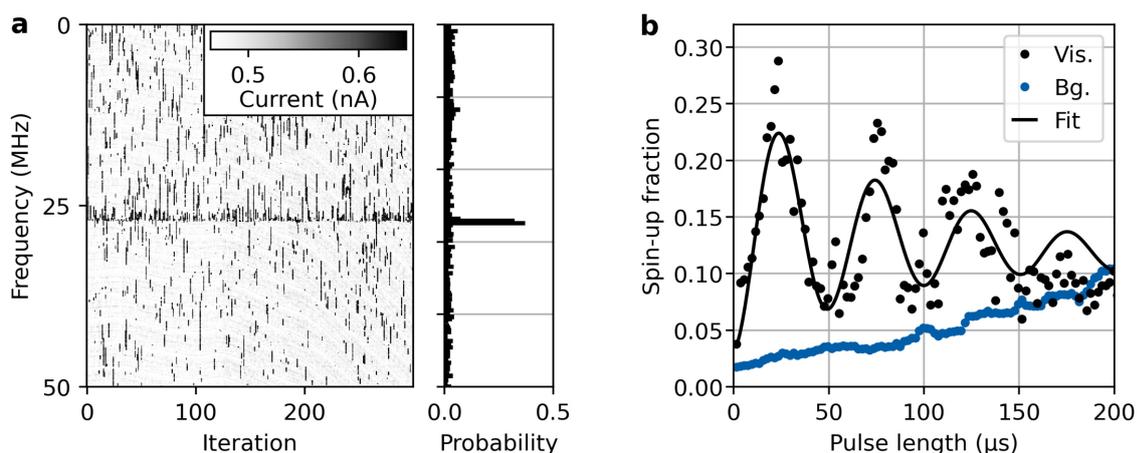


Figure 5.8: Discrete-frequency excitation during spin-selective tunnelling. **a**, Left: Raw data from 300 single-shot measurements acquired with a discrete-frequency protocol. Each 200 ms readout trace comprises of 1,000 intervals with a length of 200 μs , each consisting of a pulsed microwave excitation followed by an idle period, repeated across a set of frequency steps with 50 kHz spacing. Shown are representative traces using 24 μs microwave on-time during the 200 μs interval, converted to frequency (offset by 18.485 GHz) using the known time-to-frequency mapping. Right: Histogram of extracted blips from the left panel with a bin width of 300 kHz. **b** Maximum visibility of the histogram in panel **a** as a function of RF pulse duration (black points). Coherent Rabi oscillations are observed with a fitted frequency of (19.7 ± 0.3) kHz (black line). The background count (blue points) increases with pulse length due to the larger microwave duty cycle within the fixed measurement window, resulting in reduced visibility.

As a demonstration, a measurement using 1,000 frequency steps over a 50 MHz range with 50 kHz spacing was performed. Each step comprised a 200 μs interval containing a microwave pulse followed by an idle period to allow tunnelling events to the reservoir, producing a total single-shot measurement time of 200 ms. The resulting data (Fig. 5.8a) show a pronounced resonance, consistent with resonant microwave driving of the qubit.

A potential advantage of this method is the compatibility with calibrated π -pulses. With appropriate optimisation across the full frequency range, such pulses could in principle deliver ideal visibility while reducing measurement time relative to RAP. A broadband calibration of this type was not pursued in the present study.

Coherent control within this framework was verified by varying the excitation pulse duration to resolve Rabi oscillations. As shown in Fig. 5.8b, oscillations at approximately 19.7 kHz are observed. The approach is constrained by increased thermal background at extended pulse durations. A higher microwave duty cycle raises the probability of thermal blip events, which introduces uncertainty in the extracted spin-up fraction and limits the achievable fidelity.

In summary, the calibration of π -pulses across a wide frequency range could enable full inversion at lower background levels than RPSS, the required calibration overhead is substantial. The adiabatic RPSS approach remains more robust and provides full spin inversion over a broad range of sweep rates without the calibration overhead. Accordingly, RPSS is adopted for the magnetometry measurements that follow.

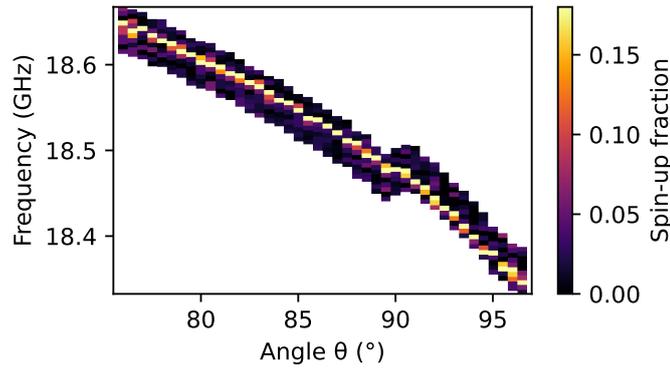


Figure 5.9: Resonance frequency versus magnetic-field angle θ for the initial, uncalibrated vector-magnet. The colour scale encodes the spin-up fraction. A simple feedback updates the baseband microwave frequency after each column, which enables rapid tracking of the resonance across a wide span. The residual frequency drift with angle originates from coil misalignment. After applying the calibration angles, the resonance trace is substantially flattened (see Fig. 5.10).

5.2 Calibration of the External Vector Magnet

A custom-built 3D vector magnet (Fig. 3.1), wound from NbTi superconducting wire, controls the magnetic field at the chip location. The objective in the present experiment is to vary the field direction while keeping the total field amplitude constant, which stabilises the qubit resonance frequency during angular sweeps. With this approach, the direction of the external field can be used to manipulate the magnetisation of an SMM, while the resulting magnetisation changes are detected as shifts of the qubit resonance frequency.

The magnetic field can be tuned within the yz plane up to a magnitude of 0.28 T. A simultaneous compensation field is applied along x so that the total field modulus remains approximately constant at $B_{\text{const}} = 0.66$ T. For polar scans in the yz plane the field is parameterised by the in-plane angle θ (from $+z$ towards $+y$) and the projected in-plane amplitude B_{pol} .

Small coil misalignments relative to the Cartesian axes are accounted for by fixed correction angles: $\Delta_{xy} = 0.75^\circ$ for the x - y coil plane and $\Delta_{xz} = -2.6^\circ$ for the x - z coil plane. The correction angles were determined from broad parameter sweeps. The y - z configuration is taken as orthogonal for simplicity, which is adequate because the x component dominates in the data shown and exceeds the yz components by more than a factor of two.

Current-to-field conversion factors and alignment corrections are obtained with a feedback procedure that adjusts the baseband microwave frequency after each acquisition, thereby keeping the qubit resonance centred inside a 100 MHz detection window (Fig. 5.9).

Assuming a constant total field $B_{\text{const}} = 0.66$ T and an in-plane component B_{pol} at angle θ in the yz plane, the required coil-set fields are

$$\begin{aligned} B_x(B_{\text{pol}}, \theta) &= \sqrt{B_{\text{const}}^2 - B_{\text{pol}}^2} + B_{\text{pol}} [\cos \theta \tan(\Delta_{xz}) + \sin \theta \tan(\Delta_{xy})] \\ B_y(B_{\text{pol}}, \theta) &= B_{\text{pol}} \sin \theta \cos(\Delta_{xy})^{-1} \\ B_z(B_{\text{pol}}, \theta) &= B_{\text{pol}} \cos \theta \cos(\Delta_{xz})^{-1} \end{aligned}$$

where (B_x, B_y, B_z) denote the applied field components in the chip-fixed Cartesian frame and the factors Δ_{xy} and Δ_{xz} implement the measured coil tilts. After calibration, angular sweeps proceed with minimal resonance drift, which enables the magnetometry measurements presented in the next section.

5.3 Quantum Magnetometry Experiments on TbPc_2

This part uses a SiMOS spin qubit as a cryogenic sensor to investigate the magnetisation dynamics of a $[\text{TbPc}_2]^-$ ensemble. The qubit, introduced in Ch. 4, senses local magnetic-field variations through shifts of its Zeeman splitting, which are tracked using the RPSS protocol introduced earlier in this chapter.

The measurements capture key magnetic signatures of the ensemble, including magnetic hysteresis [48], angular anisotropy [190], and thermally assisted spin relaxation [44]. A static field of 660 mT is used to set a stable Zeeman splitting for qubit control and readout. The easy axis of the $[\text{TbPc}_2]^-$ crystal is aligned with the device z axis, which is straightforward since the molecular symmetry axis follows the unit-cell direction (Fig. 2.6). During the measurements, the field magnitude is kept constant while its direction is varied by up to $\pm 25^\circ$ about the x axis, thereby tuning the field projection along the easy axis and enabling controlled switching of the TbPc_2 electronic spin states.

Detection of Magnetic Hysteresis

The molecular response along the easy axis is probed by rotating the external field in the xz plane and tracking the qubit resonance frequency. Instead of plotting the resonance versus angle, Fig. 5.10a displays the resonance frequency as a function of the corresponding z -component of the field, B_z , which is aligned with the easy axis of the $[\text{TbPc}_2]^-$ ensemble. A pronounced hysteresis loop is observed at the base temperature of 48 mK. The centre of the loop is shifted by -28.3 mT (labelled z_0), which matches a small residual z field from the x coil. From the loop a remanent magnetisation of 0.4 mT and a coercive field of 12.8 mT are obtained.

At an elevated temperature of 210 mK the hysteresis loop is broader and the edges are smoother, consistent with thermal activation in the Ising ground doublet (Sec. 2.2.1). This temperature dependence indicates that the observed hysteresis is dominated by the molecular ensemble rather than by instrumental effects.

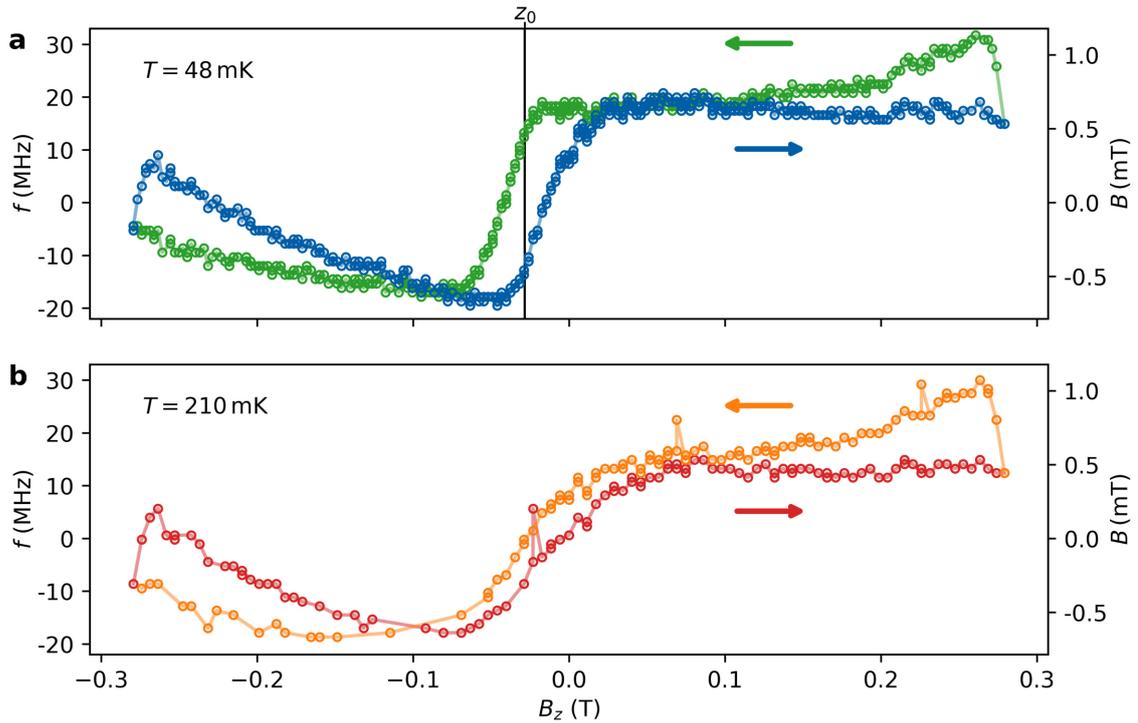


Figure 5.10: Silicon spin-qubit magnetometry of a $[\text{TbPc}_2]^-$ ensemble. The qubit resonance frequency is shown with an offset of 18.435 GHz versus the field component along the $[\text{TbPc}_2]^-$ easy axis (z). Data are taken during angular sweeps in the xz plane ($\theta = 0$) at 48 mK (a) and 210 mK (b). Each point averages 50 single shots with a bin width of 1 MHz. A g -factor of 2 [186] converts frequency to field up to a constant offset. Forward and backward sweeps show clear magnetic hysteresis linked to reversal of the ensemble spin state ($J = \pm 6$). Artefacts at the sweep reversal points (± 0.280 T) come from the vector magnet. A small residual slope for $B_z < -0.1$ T is due to imperfect coil calibration.

At both temperatures, reproducible features appear immediately after the sweep reversals at ± 0.280 T. Following each reversal, the inferred field exhibits a positive step of order 0.5 mT, and the qubit resonance shifts accordingly. These steps can be attributed to field-history effects of the NbTi type-II superconducting coils, most plausibly involving the redistribution of pinned flux vortices when the sweep direction is changed [191]. As the field is swept away from the reversal point, the vortex configuration relaxes and the additional field contribution decays back towards its baseline value. Because the vector magnet is mounted on the 4 K stage, these artefacts are expected to be largely independent of the sample temperature which is consistent with the similar turnaround features observed at 48 mK and 210 mK in Fig. 5.10. A control measurement performed without the $[\text{TbPc}_2]^-$ crystal (Fig. 5.11) retains the reversal artefacts while eliminating the central hysteresis loop, further supporting this interpretation.

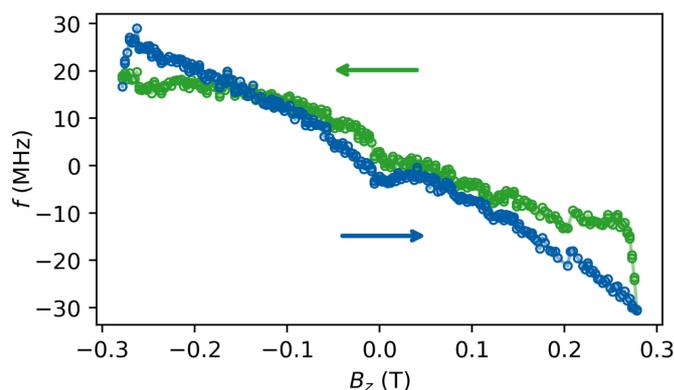


Figure 5.11: Background measurement without the $[\text{TbPc}_2]^-$ crystal using RPSS at 48 mK. The qubit resonance (offset by 18.435 GHz) is plotted versus B_z during angular sweeps in the xz plane. Arrows mark the sweep direction. A residual tilt comes from imperfect coil compensation. Coil-induced artefacts remain at the turnaround points, while the central hysteresis is absent.

Hyperfine Transition Broadening

Electronic spin reversal of $[\text{TbPc}_2]^-$ is expected to occur at four distinct values of B_z , corresponding to the four hyperfine branches ($m_I = -3/2, -1/2, 1/2, 3/2$) shown in Fig. 2.7. Micrometre-scale SQUID (μSQUID) magnetometry [192] resolves this four-step pattern clearly in the absence of a transverse field (left panel of Fig. 5.12). The crystal is not perfectly aligned at the level of individual molecules, and small molecule-to-molecule tilts of the easy axis are expected. When a transverse field is applied, these tilts translate into a distribution of effective switching fields, which broadens the hyperfine steps. This broadening is apparent for $H_\perp = 400$ mT in the right panel of Fig. 5.12. In the hysteresis loop measured with the SiMOS qubit (Fig. 5.10), the transverse field component is larger, $B_\perp \approx 0.66$ T. The associated broadening is therefore stronger, and the four hyperfine steps are not resolved.

Angular-Dependent Magnetisation Measurements

The $[\text{TbPc}_2]^-$ crystal is positioned on the chip such that its easy axis is aligned with the device z axis. This alignment can be verified experimentally by sweeping the magnetic-field direction in the yz plane and observing the characteristic magnetic anisotropy of the ensemble. The spherical-trajectory settings are calibrated and corrected as described in Sec. 5.2, ensuring that the total field magnitude remains constant while only the direction is varied. The sweeps are performed along radial trajectories, so that the absolute value of B_y or B_z increases as the angle moves away from the origin. The resulting data are shown in Fig. 5.13.

A sharp resonance frequency shift is observed along the z -axis, indicating that the easy axis of the $[\text{TbPc}_2]^-$ ensemble is aligned closely within the xz -plane of the 3D vector magnet. This angular dependence confirms both the pronounced magnetic anisotropy of

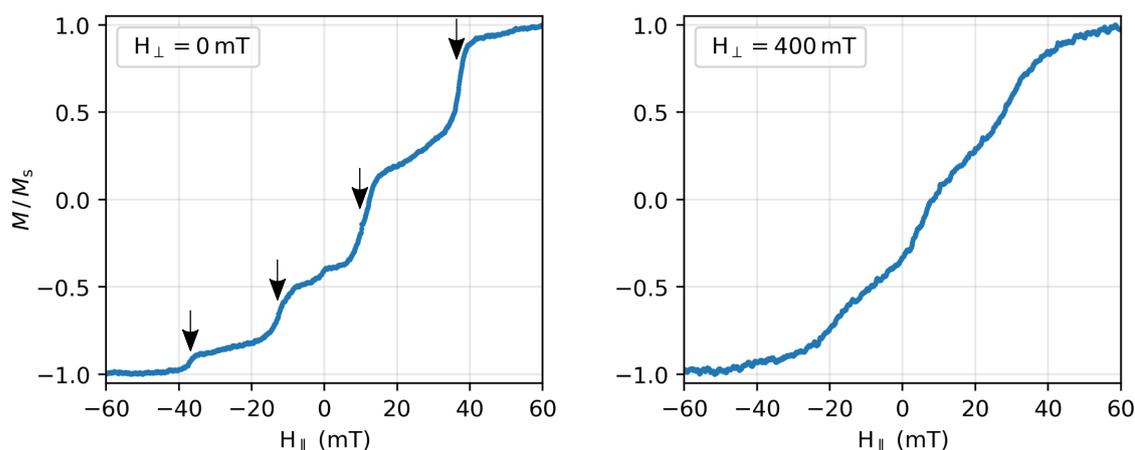


Figure 5.12: μ SQUID measurements of the same $[\text{TbPc}_2]^-$ crystal performed in a Sionludi L at 25 mK [193]. Normalised magnetic moment M as a function of the field H_{\parallel} along the easy axis. Sweep direction from left to right. Two cases of perpendicular field are shown: **left** $H_{\perp} = 0$ mT and **right** $H_{\perp} = 400$ mT. **Left** Four hyperfine steps, indicated by arrows, are visible without a transverse field. **Right** Increasing H_{\perp} broadens these features because small misalignments of individual molecules lead to a wider distribution of switching fields.

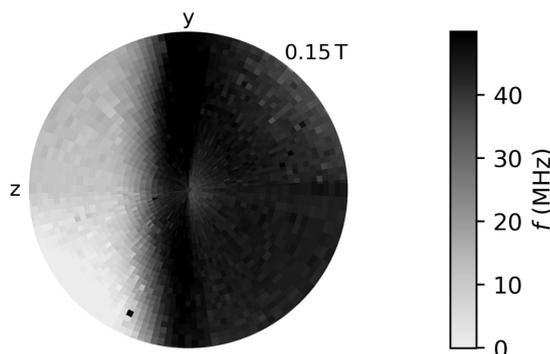


Figure 5.13: Polar plot of the resonance frequency at 48 mK during angular sweeps in the yz plane. The radial axis spans 0.00 T to 0.15 T. The sweep direction runs from the centre outward. A sharp transition along the z direction confirms alignment with the $[\text{TbPc}_2]^-$ easy axis. The qubit resonance frequency is shown with an offset of 18.420 GHz.

the molecule and demonstrates that their magnetisation produces a measurable dipolar contribution to the local field sensed by the qubit.

Thermal Relaxation Dynamics

A key feature of SMMs is the strong temperature dependence of their magnetisation relaxation [44]. To access this behaviour, the time evolution of the $[\text{TbPc}_2]^-$ ensemble is monitored via the qubit resonance frequency. The ensemble is first saturated at -0.280 T along z . The field is then returned to the setpoint z_0 , which corresponds to $B_z = -28.3$ mT

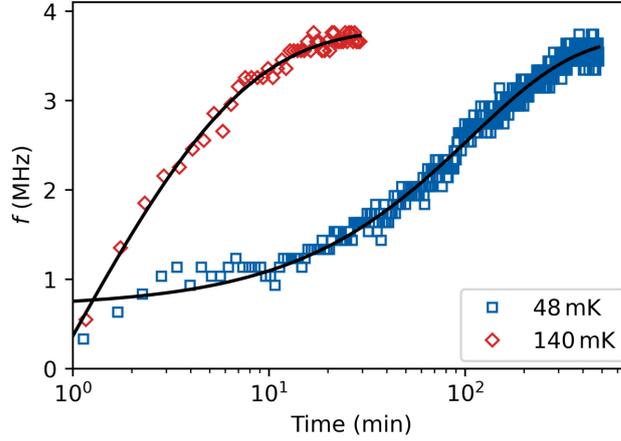


Figure 5.14: Time-resolved relaxation of the $[\text{TbPc}_2]^-$ ensemble after saturating at -0.280 T. The field is swept quickly to the setpoint z_0 from Fig. 5.10, and the qubit resonance is then monitored over time. A small drift of the vector-magnet setting during the run slightly shifted the effective z_0 , which lowers the visibility and changes the absolute frequency slightly. Each point averages 100 acquisitions (30 s per point) taken with a sweep rate of 0.267 GHz/s and a frequency bin width of 0.1 MHz. Stretched-exponential fits (Eq. (5.3)) give $\tau_{48\text{ mK}} = (107 \pm 2)$ min, $\beta_{48\text{ mK}} = 0.81 \pm 0.02$, and $\tau_{140\text{ mK}} = (0.8 \pm 0.3)$ min, $\beta_{140\text{ mK}} = 0.46 \pm 0.06$. The resonance frequency is plotted with an offset of 18.46 GHz.

because a small residual z -component is produced by the nominal x coil (Fig. 5.10). At this setpoint, the resonance is tracked continuously using repeated RPSS acquisitions. Each data point averages 100 single-shot measurements, giving a time step of 30 s per point. For the relaxation traces, the RPSS histogram bin width is reduced to $\delta f = 0.1$ MHz. This finer binning is used because the resonance shifts span only a narrow range and the spectra are acquired over a reduced 40 MHz window. The corresponding field increment is

$$\delta B = \frac{\delta f}{\gamma_e} = \frac{0.1 \text{ MHz}}{28 \text{ GHz T}^{-1}} \approx 3.6 \mu\text{T},$$

which yields an effective field sensitivity of

$$\delta B \sqrt{T} \approx 3.6 \mu\text{T} \sqrt{30 \text{ s}} \approx 20 \mu\text{T}/\sqrt{\text{Hz}}.$$

Data are taken at 48 mK and 140 mK and are shown in Fig. 5.14. In both cases the resonance drifts by about 3 MHz, which reflects a change of the ensemble magnetisation over time. The relaxation is much faster at the higher temperature, as expected for thermally activated dynamics. The data are fitted with a stretched exponential,

$$f(t) = A \exp[-(t/\tau)^\beta] + C, \quad (5.3)$$

where A is the amplitude, τ the characteristic time, β the stretch exponent, and C a constant offset. This form captures the distributed relaxation expected in SMM ensembles [194], where variations in local environments result in a broad spectrum of relaxation rates. At 48 mK the fit yields $\tau = (107 \pm 2)$ min with $\beta = 0.81 \pm 0.02$, which indicates moderately distributed relaxation dynamics only, as β is close to 1. At 140 mK the relaxation speeds

up to $\tau = (0.8 \pm 0.3)$ min and the exponent drops to $\beta = 0.46 \pm 0.06$. The smaller β points to a broader spread of relaxation rates at higher temperature, which is consistent with thermally assisted access to more relaxation pathways (Sec. 2.2.1) [194, 195]. These measurements demonstrate that the SiMOS qubit can track the magnetisation dynamics of a $[\text{TbPc}_2]^-$ ensemble over timescales spanning minutes to hours, establishing a robust foundation for more advanced hybrid quantum sensing schemes.

Conclusion

This chapter demonstrates a practical hybrid sensing architecture that combines an industrial SiMOS spin qubit with a $[\text{TbPc}_2]^-$ molecular spin ensemble. The qubit acts as a sensitive, CMOS-compatible interface that resolves resonance shifts on the order of 0.1 MHz induced by changes in the ensemble magnetisation. Stable resonance tracking at a fixed working point is enabled by the RPSS protocol, which combines rapid adiabatic passage with spin-selective tunnelling and operates without continuous π -pulse recalibration. The protocol is not tuned for ultimate sensitivity. Instead, it is designed for robust operation in the present regime where an ensemble signal produces MHz-scale frequency shifts, consistent with the measured remanent field of 0.4 mT.

The measurements resolve characteristic signatures of the electronic spin states of the ensemble magnetisation, including pronounced hysteresis loops, a clear angular anisotropy consistent with the molecular easy axis, and thermally activated relaxation dynamics. The relaxation times extracted for the $[\text{TbPc}_2]^-$ ensemble are $\tau = (107 \pm 2)$ min at 48 mK and $\tau = (0.8 \pm 0.3)$ min at 140 mK. Hyperfine-resolved steps associated with the nuclear spin are not observed, which is attributed to small molecule-to-molecule tilts of the easy axis in the crystal combined with the transverse field required for qubit operation, leading to a broadened distribution of nuclear switching fields.

The dominant limitation of the present device is the comparatively large vertical separation between the qubit and the molecules, on the order of 250 nm or more. This distance weakens the dipolar coupling and reduces the induced qubit frequency shifts. Reducing the separation to below 10 nm through an improved gate stack and process flow (see Ch. 7) would substantially enhance the coupling and could enable detection of individual $[\text{TbPc}_2]^-$ molecules.

5.4 Outlook: Single-Molecule Sensing

The next step is to move from ensemble measurements to single-molecule sensing. With an improved gate stack of the chip (Ch. 7), the qubit-molecule spacing can be reduced, which increases the magnetic coupling. For a spacing of about 18 nm the expected shift of the qubit resonance from a single switching event of $[\text{TbPc}_2]^-$ is on the order of 160 kHz (Sec. 2.3).

Resolving a constant frequency shift of that size can be accomplished with either a Rabi pulse sequence or a Ramsey pulse sequence as introduced in Sec. 4.8 and Sec. 4.9. To compare both approaches, the measured spin-up fraction after each sequence is evaluated for parameters close to the experiment. For a simple estimate, a frequency shift $\Delta f = 100$ kHz from an SMM switching event and an evolution time $\tau = 1 \mu\text{s}$ are assumed for both sequences. Dephasing is neglected and the visibility is set to one for this proof-of-principle comparison.

For a Rabi experiment starting from the initial state $|\downarrow\rangle$, the probability to measure the spin-up state after a drive pulse of duration τ is

$$P_{\uparrow}(\tau) = \frac{\Omega^2}{\Omega^2 + \Delta\omega^2} \sin^2\left(\frac{\tau}{2}\sqrt{\Omega^2 + \Delta\omega^2}\right), \quad \Omega = 2\pi f_{\text{R}}, \quad \Delta\omega = 2\pi\Delta f. \quad (5.4)$$

For $f_{\text{R}} = 500$ kHz, the probability difference between on-resonance driving and a detuned drive at $f_0 + \Delta f$ evaluates to $\Delta P_{\uparrow} = P_{\uparrow}(f_0) - P_{\uparrow}(f_0 + \Delta f) \approx 3.9\%$ at $\tau = 1 \mu\text{s}$ (a π -pulse duration). Increasing the pulse time to $3\pi/2$ ($\tau = 1.5 \mu\text{s}$) raises the contrast to $\Delta P_{\uparrow} \approx 6.4\%$. Depending on the amount of averages, such changes in the spin-up fraction can still be comparable to the measurement noise. A larger ΔP_{\uparrow} is therefore desirable, as it reduces the number of averages required and improves the robustness of the measurement.

A Ramsey sequence, on the other hand, depends on the accumulated phase $\Delta\phi = 2\pi\Delta f\tau$ while the qubit is on the equator of the Bloch sphere. With $\tau = 1 \mu\text{s}$ and $\Delta f = 100$ kHz the readout

$$P_{\uparrow}(\tau) = \frac{1}{2} [1 + \cos(\phi_0 + \Delta\phi)] \quad (5.5)$$

yields $\Delta P_{\uparrow} \approx 9.8\%$ at the fringe top ($\phi_0 = 0$), which is much larger than for the Rabi case. The Ramsey response can be increased further by operating at a slightly detuned working point ($\phi_0 = \pi/2$), which gives $\Delta P_{\uparrow} \approx 29.4\%$ for the same parameters. This higher detuning sensitivity shows that a phase-accumulating pulse like Ramsey is better suited to reveal small shifts of the qubit resonance.

The maximal phase-accumulation time is limited by T_2^* , so $\tau < T_2^*$. For the present single qubit $T_2^* \approx 3.1 \mu\text{s}$ (Sec. 4.9). Extending T_2^* allows longer phase accumulation and therefore a larger ΔP_{\uparrow} at fixed Δf . Figure 5.15 shows simulations for both Rabi and Ramsey with line cuts at $\tau = 1 \mu\text{s}$ and $\tau = 19 \mu\text{s}$ in panels b and d, clearly illustrating that increasing τ steepens the slope of spin-up fraction versus detuning and thus improves sensitivity. This highlights the trade-off between spectral precision and measurement time in qubit-based magnetometry.

In practice, improving T_2^* requires suppressing low-frequency magnetic noise to which Ramsey measurements are particularly susceptible. Dynamical-decoupling sequences such as Hahn echo or Carr-Purcell-Meiboom-Gill (CPMG) suppress quasi-static detuning noise and thereby extend the coherence time [180]. However, precisely because these protocols refocus static or slowly varying frequency shifts, they are not directly sensitive to a step-like change in f_0 caused by an SMM switching event and are therefore not suitable for the sensing task considered here. A genuine improvement of T_2^* can be achieved by further ^{28}Si enrichment, which reduces the Overhauser field from residual ^{29}Si nuclei and

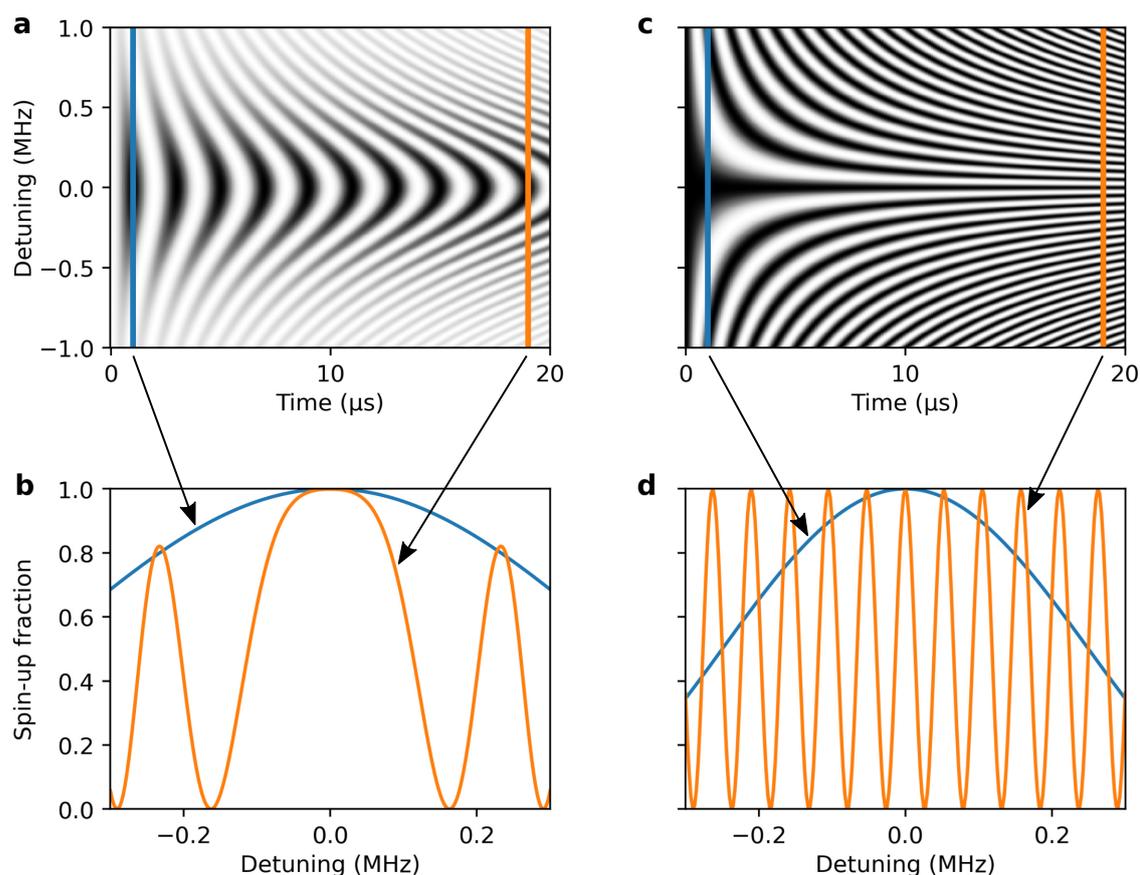


Figure 5.15: Sensitivity of pulse protocols for Larmor frequency determination. **a, c** Time-resolved spin dynamics under a Rabi pulse (**a**) and a Ramsey sequence (**c**), simulated with a Rabi frequency of 500 kHz in an ideal, decoherence-free system. The spin-up fraction is shown in grayscale (white = 0 %, black = 100 %). Vertical lines mark representative readout times at 1 μ s (blue) and 19 μ s (orange). **b, d** Spin-up probability as a function of detuning taken at those readout times. The Ramsey sequence (**d**) exhibits sharper and more closely spaced fringes than the Rabi protocol (**b**), reflecting higher sensitivity to resonance shifts.

narrows the inhomogeneous linewidth. Additional gains are expected from reducing slow field drift and current ripple of the vector magnet by using low-noise, well-filtered current supplies. Such improvements are pivotal for resolving the subtle field shifts produced by individual molecular spin transitions, enabling a transition from ensemble averages to single-molecule readout in future experiments.

6 Quantum Sensing Using Double Quantum Dot States

After the single-dot sensing experiments presented in Ch. 5, this chapter investigates how a double quantum dot can be employed for magnetic-field sensing. The focus is not on realising universal two-qubit logic, but on exploiting double-dot spin physics to obtain readout schemes and signal dependencies that are advantageous for sensing applications.

In a double dot, an additional spin-readout mechanism becomes available based on Pauli spin blockade (PSB) [196]. For sensing applications, a key advantage of PSB is that the readout mechanism operates over a wide range of external magnetic fields, from very small fields up to the tesla regime [197, 198]. This contrasts with the single-dot readout used in Ch. 5, which used a transverse magnetic field of 0.66 T to enable spin-selective tunnelling. With PSB-based readout this constraint is removed, enabling readout at very low magnetic fields (even on the order of the geomagnetic field) [197], which is beneficial for sensing applications.

Moreover, single-dot sensing schemes are primarily sensitive to the *absolute* Zeeman energy of a single spin. In a double-dot configuration the relevant quantity can instead be a magnetic-field *gradient* between the two dots. This gradient is encoded in coherent singlet-triplet₀ (ST₀) oscillations [17], whose frequency depends on the Zeeman energy difference between the two qubit sites. A SMM placed asymmetrically with respect to the double dot would generate a local dipolar field that differs between the two dots, thereby inducing a gradient that can potentially be characterised through changes in the ST₀ oscillation frequency. This opens a route to sensing schemes that are largely insensitive to global changes of the absolute magnetic field and associated environmental noise, yet remain highly sensitive to local field gradients induced by nearby SMMs.

6.1 Introduction to Double Dot Systems

The same device layout used for the single-dot measurements presented earlier can be tuned to form a double quantum dot that is tunnel coupled to a common electron reservoir. For all measurements in this chapter, device B (Tab. 4.1) is used. The gate layout, including the regions in which electrons are accumulated (shaded in black), is shown in the left panel of Fig. 6.1. Each dot is controlled by its own plunger gate (P1 and P2), which tunes the electrochemical potential of the dot and thereby enables loading of electrons from

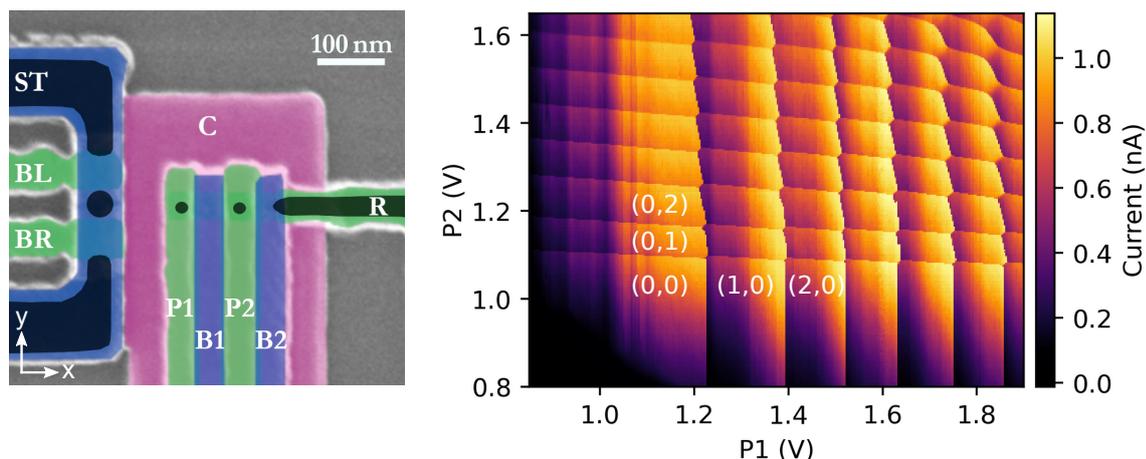


Figure 6.1: Double-dot device and charge stability map. **Left** False-colour SEM image (top view) of the multi-layer gate stack. The gate layer closest to the Si/SiO₂ interface is shown in pink, the intermediate layer in green, and the upper layer in blue. In operation, the regions shaded in black host accumulated electrons forming the two quantum dots and the nearby sensor. The in-plane directions x and y are indicated. **Right** Charge stability diagram showing the sensor current as a function of the two plunger gate voltages $P1$ and $P2$. Electron occupations (n_1, n_2) of the left and right dot are indicated in parentheses. During the measurement, the sensor top gate (ST) is linearly compensated to keep the sensor close to its optimal working point while the plunger voltages are varied.

the reservoir. While dot 2 (on the right) is directly tunnel coupled to the reservoir, the left dot can only exchange electrons via the right dot. Sweeping both plunger voltages produces the charge stability diagram shown in the right panel of Fig. 6.1. Charge states are labelled (n_1, n_2) , with n_1 (n_2) being the number of electrons in the left (right) dot. The experiments described in this chapter use two electrons confined in the double dot for qubit manipulation. Once the desired electron number is loaded, the tunnel coupling to the reservoir is no longer required for the coherent manipulation itself. In initial tests, the tunnel coupling to the reservoir was therefore turned off (by ramping down gates B2 and subsequently R), resulting in a fixed charge configuration. However, under these conditions (confinement gate C at -0.2 V) the charge occupation proved insufficiently stable: single-electron jumps in and out of the dot via residual charge puddles typically occurred on the timescale of about one event per day. To mitigate this, the reservoir is kept weakly connected in the final measurements, ensuring that the double-dot electron occupation is stabilised by exchange with the reservoir on long timescales.

The region around the $(2, 0) \leftrightarrow (1, 1)$ charge transition is shown in a zoomed-in view in Fig. 6.2. A high-symmetry line, referred to as the detuning axis, is indicated in white and crosses both the $(2, 0)$ and $(1, 1)$ charge regions. Moving along this axis keeps the total electron number of the double dot constant while changing the relative electrochemical potentials of the two sites. This is highly convenient, and all double-dot experiments presented in this chapter rely on voltage changes of $P1$ and $P2$ along this detuning line. The detuning ε denotes the energy difference between the ground state energies of the left and right dot while keeping the sum of their electrochemical potentials constant. Qubit initialisation is performed at point A in the $(2, 0)$ configuration, and this point is chosen

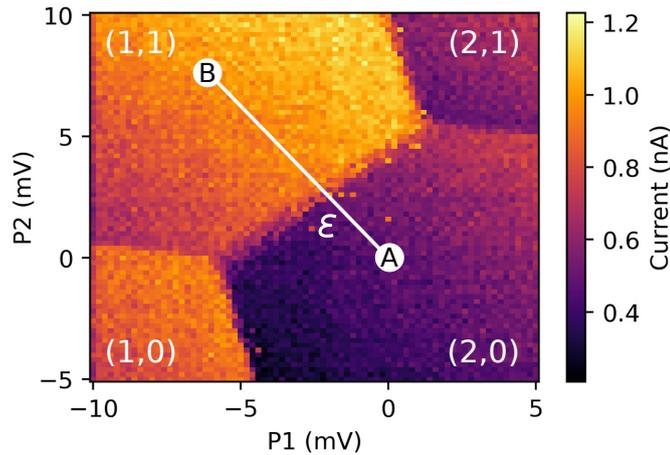


Figure 6.2: Charge stability diagram around the $(2, 0) \leftrightarrow (1, 1)$ charge configuration. The dot occupations (n_1, n_2) are indicated in parentheses. The white line marks the symmetry line of the detuning axis ϵ connecting points A and B. Measurements typically start at point A for initialisation, and negative detuning is applied along this axis to move into the $(1, 1)$ charge state towards point B. The map was acquired in “video mode”: the AWG applies voltage pulses to the capacitive branch of the bias-tee (see Sec. 3.6). For each pixel, the gate voltage is pulsed from the reference point A to the corresponding coordinate in $(P1, P2)$ space, and the sensor current is integrated for $20 \mu\text{s}$. Static offsets are supplied to the resistive side of the bias-tees, with DC voltages $P1 = 1.6 \text{ V}$ and $P2 = 1.0 \text{ V}$.

to correspond to $\epsilon = 0 \text{ eV}$. A negative detuning pulse towards point B moves the system into the $(1, 1)$ charge region.

Experimentally, the conversion from gate voltage to detuning energy is obtained using a combined lever arm of P1 and P2. Sweeps along the detuning line are characterised by a ratio $P2/P1 = -1.25$. Using the single-plunger lever arms $\alpha_1 = 12.5\%$ and $\alpha_2 = 7.5\%$ of the used device, measured in Sec. 4.3, the overall lever arm along the detuning axis is

$$\alpha_\epsilon = \alpha_1 + 1.25 \alpha_2 = 12.5\% + 1.25 \times 7.5\% \approx 21.9\% .$$

6.2 The Singlet-Triplet Qubit

The main appeal of the ST_0 qubit for the quantum-computing community is that two-axis qubit control can be achieved purely by electrical gate pulses [199, 200], without the need for high-frequency magnetic or electric fields. This renders the ST_0 encoding particularly attractive for scalable architectures [201].

By changing the detuning parameter ϵ , the charge configuration $(1, 1)$ can be transferred to $(2, 0)$, effectively moving both electrons into the same dot. In this situation the spatial wavefunctions overlap strongly and the electrons experience an exchange interaction J , often referred to as Heisenberg coupling [202]. This interaction is directly controlled by the detuning and can therefore be tuned electrically. Including the Zeeman energy of

each spin, the resulting two-spin Hamiltonian (written in terms of dimensionless spin operators) reads [202]

$$H = g_1 \mu_B \mathbf{B}_1 \cdot \mathbf{S}_1 + g_2 \mu_B \mathbf{B}_2 \cdot \mathbf{S}_2 + J \mathbf{S}_1 \cdot \mathbf{S}_2. \quad (6.1)$$

The first two terms describe the Zeeman energies $E_{Z,1}$ and $E_{Z,2}$ of the individual spins, while the last term captures the exchange coupling. It is useful to distinguish two limiting regimes set by the competition between the Zeeman-energy difference $\Delta E_Z = E_{Z,1} - E_{Z,2}$ and the exchange energy J . The eigenstates of Eq. (6.1) can then be discussed in the following limits [17]:

- **Weakly coupled ($J \ll \Delta E_Z$):** In this regime the natural basis is the product-state basis $\{ |\uparrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\downarrow\downarrow\rangle \}$. The Zeeman gradient dominates the level structure, and the exchange term provides only a small perturbative shift to the energies of the antiparallel states $|\uparrow\downarrow\rangle$ and $|\downarrow\uparrow\rangle$.
- **Strongly coupled ($J \gg \Delta E_Z$):** In this limit the eigenstates are well described by the singlet-triplet basis $\{ |T_+\rangle, |T_0\rangle, |T_-\rangle, |S\rangle \}$, with

$$\begin{aligned} |T_+\rangle &= |\uparrow\uparrow\rangle, \\ |T_0\rangle &= \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle), \\ |T_-\rangle &= |\downarrow\downarrow\rangle, \\ |S\rangle &= \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle). \end{aligned}$$

where $|S\rangle$ denotes the singlet and $|T_\pm\rangle, |T_0\rangle$ the triplet states. In this regime the exchange coupling J sets the singlet-triplet₀ energy splitting.

With the polarised triplet states $|T_+\rangle = |\uparrow\uparrow\rangle$ and $|T_-\rangle = |\downarrow\downarrow\rangle$ energetically split off by the Zeeman energy, the remaining two states $|S\rangle$ and $|T_0\rangle$ form a convenient two-level subspace. These states are used to encode the logical qubit states, for example $|0\rangle \equiv |S\rangle$ and $|1\rangle \equiv |T_0\rangle$. This choice is also natural from the perspective of readout: in the large- J regime relevant for Pauli spin blockade, the eigenstates close to the $(1, 1) \leftrightarrow (0, 2)$ transition are predominantly singlet- or triplet-like.

A Bloch-sphere representation of the ST_0 qubit is shown in the left panel of Fig. 6.3. In the $\{|S\rangle, |T_0\rangle\}$ basis the effective two-level Hamiltonian can be written as [17, 199]

$$H_{ST_0} = \frac{1}{2} J(\varepsilon) \sigma_z + \frac{1}{2} \Delta E_Z \sigma_x, \quad (6.2)$$

where $J(\varepsilon)$ generates rotations about the Bloch z -axis and ΔE_Z generates rotations about the x -axis. The product states $|\uparrow\downarrow\rangle$ and $|\downarrow\uparrow\rangle$ correspond to points on the equator of this Bloch sphere, representing equal-weight superpositions of $|S\rangle$ and $|T_0\rangle$ with different relative phases. In contrast to $J(\varepsilon)$, which can be tuned over orders of magnitude by adjusting the detuning ε , ΔE_Z is typically set by the device and magnetic-field configuration and is therefore always present during operation. Consequently, the qubit undergoes continuous precession about the Bloch x -axis, which must be accounted for when designing and calibrating gate sequences in traditional ST_0 qubit experiments [17, 200].

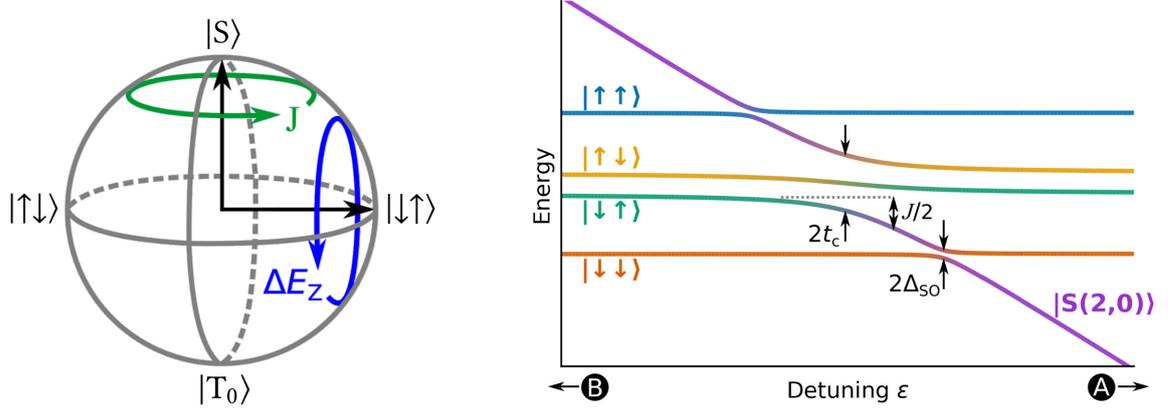


Figure 6.3: ST_0 qubit Bloch sphere and energy diagram. **Left** Bloch-sphere representation of the ST_0 qubit in the $\{|S\rangle, |T_0\rangle\}$ basis. The exchange interaction J generates rotations about the z -axis (green arrow), while the Zeeman energy difference ΔE_z generates rotations about the x -axis (blue arrow). The product states $|\uparrow\downarrow\rangle$ and $|\downarrow\uparrow\rangle$ lie on the equator. **Right** Energy diagram of the Hamiltonian Eq. (6.3) as a function of detuning ε , obtained from a QuTiP [203] simulation. The coloured lines denote the eigenenergies. The four product spin states reside in the $(1, 1)$ charge configuration, while the $|S(2, 0)\rangle$ branch (purple) disperses strongly with detuning. Avoided crossings are set by the tunnel coupling t_c and the spin-orbit-induced coupling Δ_{SO} . Couplings and level spacings are adjusted for visibility and are not to scale. Measurement points A and B from Fig. 6.2 are indicated along the detuning axis as a guide to the operating points.

The first landmark experiments on the ST_0 qubit were carried out by Petta et al. in 2005 [17], who demonstrated coherent control between $|S\rangle$ and $|T_0\rangle$ in a GaAs double-dot device, including Ramsey and Hahn-echo pulse sequences. More recently, capacitive coupling between two ST_0 qubits has been realised, enabling two-qubit operations in this encoding [201, 204].

Including the doubly occupied singlet state $|S(2, 0)\rangle$, the two-electron Hamiltonian in the basis $\{|\uparrow\uparrow\rangle, |\uparrow\downarrow\rangle, |\downarrow\uparrow\rangle, |\downarrow\downarrow\rangle, |S(2, 0)\rangle\}$ can be written as [205, 206]

$$H = \begin{pmatrix} E_z & 0 & 0 & 0 & \Delta_{SO} \\ 0 & \Delta E_z/2 & 0 & 0 & t_c \\ 0 & 0 & -\Delta E_z/2 & 0 & -t_c \\ 0 & 0 & 0 & -E_z & \Delta_{SO} \\ \Delta_{SO} & t_c & -t_c & \Delta_{SO} & U - \varepsilon \end{pmatrix} \quad (6.3)$$

where E_z is the average Zeeman energy, ΔE_z the Zeeman-energy difference between the two dots, ε the interdot detuning, and U the on-site Coulomb energy of the $(2, 0)$ configuration. The parameter t_c denotes the spin-conserving interdot tunnel coupling between $|S(1, 1)\rangle$ and $|S(2, 0)\rangle$, while Δ_{SO} describes spin-flip tunnelling between $|\uparrow\uparrow\rangle$ or $|\downarrow\downarrow\rangle$ and $|S(2, 0)\rangle$ [207]. These couplings are treated as effective tunnel matrix elements capturing the underlying microscopic processes. In a SiMOS device, Δ_{SO} originates from spin-orbit interaction induced by the strong vertical electric field and the atomistic Si/SiO_x interface, which together generate Rashba- and Dresselhaus-like spin-orbit contributions [208, 209]. The same interface-induced spin-orbit interaction leads to a pronounced g -factor anisotropy [209], as characterised in Sec. 6.4. It has been shown that ΔE_z can be changed

by orders of magnitude by varying the angle of the applied external magnetic field [209, 210].

The energy diagram of the full Hamiltonian in Eq. (6.3) as a function of detuning is shown in the right panel of Fig. 6.3. The four $(1, 1)$ spin states appear as approximately horizontal lines in this representation, while the $|S(2, 0)\rangle$ state (purple) disperses strongly with detuning energy ε and crosses the $(1, 1)$ manifold. Three avoided crossings become visible, set by the tunnel coupling t_c and the spin-orbit-induced coupling Δ_{SO} . The voltage points A and B of the charge stability diagram (introduced in Fig. 6.2) are indicated on the energy diagram as a guide to the operating points used in the experiments.

Phenomenologically, the detuning dependence of the exchange interaction can be modelled as

$$J(\varepsilon) \approx J_0 \exp\left(\frac{\varepsilon}{\varepsilon_0}\right), \quad (6.4)$$

where J_0 sets the overall energy scale and ε_0 is the characteristic detuning scale [211].

Landau-Zener Processes

For qubit initialisation and readout, transitions across the avoided crossings in the energy diagram are exploited. The transition probabilities are described by Landau-Zener theory [88, 89], as discussed previously in the context of rapid adiabatic passage (Sec. 4.7). Near a given avoided crossing, the dynamics can be reduced to an effective two-level system with instantaneous eigenenergies $E_1(\varepsilon)$ and $E_2(\varepsilon)$, which are swept in time by changing the detuning $\varepsilon(t)$. The relevant sweep rate is

$$v = \frac{d}{dt}(E_1(\varepsilon) - E_2(\varepsilon)) \approx \dot{\varepsilon}, \quad (6.5)$$

The approximation $v \approx \dot{\varepsilon}$ is valid when the diabatic level separation $E_1 - E_2$ depends approximately linearly on detuning in the vicinity of the anticrossing (so that $d(E_1 - E_2)/d\varepsilon \approx 1$). In situations where the level spacing exhibits strong curvature as a function of ε , the full derivative $d(E_1 - E_2)/d\varepsilon$ should be retained when evaluating v [212]. The probability to undergo a non-adiabatic transition at the avoided crossing is then given by the Landau-Zener formula [212]:

$$P_{\text{LZ}} = \exp\left(-\frac{2\pi\Delta^2}{\hbar v}\right). \quad (6.6)$$

where Δ is the coupling matrix element that sets the minimum energy gap at the anticrossing. By controlling v during the passage through the anticrossing, the adiabaticity of the sweep can be tuned to realise a desired transition probability into the excited state. For $v \ll 2\pi\Delta^2/\hbar$ the passage is adiabatic with respect to the perturbation Δ and $P_{\text{LZ}} \rightarrow 0$, whereas for $v \gg 2\pi\Delta^2/\hbar$ the passage is diabatic and $P_{\text{LZ}} \rightarrow 1$.

Initialisation of the Singlet-Triplet Qubit

Initialisation of the ST_0 qubit relies on controlled Landau-Zener passages starting from the ground state $|S(2, 0)\rangle$ at position A, indicated in Figs. 6.2 and 6.3. A detuning pulse from point A towards point B along the detuning line moves the system into the (1, 1) charge region. Two different initialisations of the (1, 1) configuration are used: (i) initialisation into the product state $|\downarrow\uparrow\rangle$ and (ii) initialisation into the singlet $|S(1, 1)\rangle$.

For both protocols it is essential to cross $|S(2, 0)\rangle \leftrightarrow |\downarrow\downarrow\rangle$ diabatically, such that the system remains on the singlet branch. In the Landau-Zener description this corresponds to a sweep rate v that satisfies $v \gg 2\pi \Delta_{S0}^2/\hbar$, so that the non-adiabatic transition probability at this anticrossing approaches unity.

Initialisation into $|\downarrow\uparrow\rangle$: The detuning pulse is chosen adiabatic with respect to the charge anticrossing between $|S(2, 0)\rangle$ and $|S(1, 1)\rangle$, whose gap is set by $2t_c$. This requires $v \ll 2\pi t_c^2/\hbar$, ensuring that $|S(2, 0)\rangle \rightarrow |S(1, 1)\rangle$. As the detuning is further increased into the regime $J(\varepsilon) \ll \Delta E_Z$, the instantaneous ground state in the $S_z = 0$ subspace continuously evolves from $|S(1, 1)\rangle$ towards the product state $|\downarrow\uparrow\rangle$. If the sweep remains slow compared to the spin mixing driven by ΔE_Z , i.e. $v \ll 2\pi \Delta E_Z^2/\hbar$, the system follows this lowest eigenstate adiabatically and ends in $|\downarrow\uparrow\rangle$ [186, 213]. At this detuning point the exchange energy $J(\varepsilon)$ is negligible and $|\downarrow\uparrow\rangle$ is an eigenstate of the Hamiltonian, so the state is stationary.

Initialisation into $|S(1, 1)\rangle$: Here, a semi-adiabatic sweep is employed. The detuning is again changed adiabatically with respect to the charge anticrossing (gap $2t_c$), so that $|S(2, 0)\rangle$ maps to $|S(1, 1)\rangle$ as before, implying $v \ll 2\pi t_c^2/\hbar$. However, the sweep is made diabatic with respect to the slower $S \leftrightarrow T_0$ basis rotation driven by the Zeeman-energy difference ΔE_Z [17]. This requires $v \gg 2\pi \Delta E_Z^2/\hbar$, so that once in the (1, 1) sector the state does not have time to follow the instantaneous eigenbasis towards the product state $|\downarrow\uparrow\rangle$ but instead is effectively “frozen” in $|S(1, 1)\rangle$.

In the regime $J \ll \Delta E_Z$ the state $|S(1, 1)\rangle$ is not an eigenstate of the Hamiltonian Eq. (6.1), but can be expressed as a superposition of the eigenstates $|\downarrow\uparrow\rangle$ and $|\uparrow\downarrow\rangle$, which are split by ΔE_Z . The resulting energy difference between the two eigenstates leads to a relative phase accumulation in time and thus to coherent oscillations between $|S(1, 1)\rangle$ and $|T_0(1, 1)\rangle$. In general, the ST_0 oscillation frequency is given by [214]

$$\Omega(\varepsilon) = \frac{1}{\hbar} \sqrt{\Delta E_Z^2 + J(\varepsilon)^2}, \quad (6.7)$$

which interpolates between gradient-dominated oscillations ($J \ll \Delta E_Z$) and exchange-dominated oscillations ($J \gg \Delta E_Z$).

6.3 Pauli Spin Blockade Readout

Readout of the ST_0 qubit state is performed via Pauli spin blockade (PSB) in the strongly coupled regime $J \gg \Delta E_Z$ in the (2, 0) charge region close to point A, indicated in Figs. 6.2

and 6.3. In this regime the spin state is converted into a charge configuration of the double dot. Owing to the Pauli exclusion principle, the total two-electron wavefunction must be antisymmetric under particle exchange. As a consequence, the lowest-energy two-electron state in $(2, 0)$ is a singlet, whereas a $(2, 0)$ triplet requires an antisymmetric spatial component and therefore forces one electron into a higher-lying valley or orbital state [98]. The resulting singlet-triplet energy splitting enables the blockade mechanism: while the $|S(1, 1)\rangle$ state can tunnel into the $|S(2, 0)\rangle$ configuration, the triplet state $|T_0(1, 1)\rangle$ remains in $(1, 1)$, because accessing the corresponding $(2, 0)$ triplet would require an additional energy. By detecting whether the system resides in $(2, 0)$ or $(1, 1)$ using a nearby SET charge sensor, the spin state of the qubit can be inferred. During readout, the detuning sweep rate v is chosen adiabatic with respect to the interdot tunnel coupling t_c and diabatic with respect to ΔE_Z . This choice maps any $|S(1, 1)\rangle$ population reliably onto $|S(2, 0)\rangle$ while keeping $|T_0(1, 1)\rangle$ in the $(1, 1)$ configuration, thereby preserving the singlet-triplet populations prepared during the manipulation stage [17].

Due to the orbital and valley structure of the quantum dots (Fig. 2.3), two additional electrons are loaded into the left dot to fill the lowest valley state. The relevant charge transition for PSB readout is therefore not $(2, 0) \leftrightarrow (1, 1)$ but $(4, 0) \leftrightarrow (3, 1)$. In this configuration the singlet-triplet splitting between $|S(4, 0)\rangle$ and $|T_0(4, 0)\rangle$ is set by the comparatively robust orbital level spacing, rather than by the strongly varying valley splitting [42, 215, 216]. As a result, the energy window over which PSB-based readout can be performed is more stable and possibly enlarged [42, 217]. For notational simplicity, these two additional “core” electrons are omitted in the following discussion.

Experimentally, PSB is first characterised by preparing a mixture of singlet and triplet states in the $(1, 1)$ configuration and probing how this mixture converts to $(2, 0)$ under a detuning sweep. The pulse sequence, shown in the left panel of Fig. 6.4, starts by preparing $|S(1, 1)\rangle$ via a detuning sweep from point A to point B (Figs. 6.2 and 6.3) in 10 ns, chosen to satisfy the mentioned inequalities for v . The system is then held in $(1, 1)$ for a dwell time $t_{\text{dwell}} = 1 \text{ ms} \gg T_2^{\text{ST}_0}$ to allow $|S(1, 1)\rangle$ to dephase fully into an incoherent mixture of S and T_0 (the relevant coherence time $T_2^{\text{ST}_0}$ will be extracted in Sec. 6.4). Finally, the detuning is swept back towards point A while the SET current is recorded.

The resulting data, shown in the right panel of Fig. 6.4, reveal two current plateaus corresponding to the $(1, 1)$ (low current) and $(2, 0)$ (high current) charge configurations. As the detuning is varied along the time axis, singlet states can adiabatically tunnel to $(2, 0)$, whereas triplet states remain blockaded in $(1, 1)$ until the detuning is large enough to reach the avoided crossing $|T_0(1, 1)\rangle \leftrightarrow |T_0(2, 0)\rangle$, where charge transfer eventually occurs. The optimal PSB readout point lies between the $|S(1, 1)\rangle \leftrightarrow |S(2, 0)\rangle$ and $|T_0(1, 1)\rangle \leftrightarrow |T_0(2, 0)\rangle$ anticrossings, labelled S and T in Fig. 6.4, defining the readout window used in the following experiments.

After this initial test, a similar pulse scheme is used (Fig. 6.5a), where the detuning is pulsed directly to the readout point R instead of sweeping through it as before. For each cycle the sensor current is recorded for generally $100 \mu\text{s}$, and a $10 \mu\text{s}$ segment within this window is integrated to obtain a single current value. This value is then used to discriminate

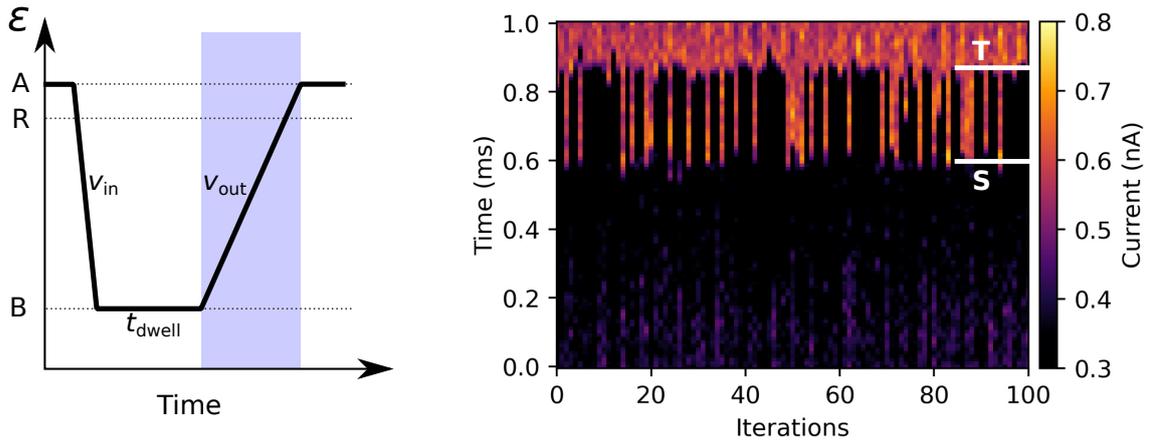


Figure 6.4: Detuning sweep over the $(1, 1) \leftrightarrow (2, 0)$ transition. **Left** Schematic detuning pulse starting at the initialisation point A in the $(2, 0)$ region. The detuning is ramped to point B in the $(1, 1)$ region with sweep rate v_{in} preparing a singlet state, held there for a dwell time t_{dwell} , and then swept back towards A with sweep rate v_{out} . During the final part of the return ramp (blue shaded region) the SET current is recorded. The readout point R lies between A and B along this segment. For clarity the pulse shape is shown without the additional compensation of bias-tee effects discussed in Sec. 3.6. **Right** Colour plot of the SET current as a function of time and iteration index for 100 repetitions of the pulse. Low current corresponds to the $(1, 1)$ configuration, while high current indicates $(2, 0)$. Two distinct transition bands (labelled S and T) appear when the decohered mixture of $|S(1, 1)\rangle$ or $|T_0(1, 1)\rangle$ tunnels to $(2, 0)$, providing a clear signature of PSB. The data were taken with $v_{in} = 220$ keV/s, $t_{dwell} = 1$ ms $\gg T_2^{ST_0}$ and $v_{out} = 2.2$ eV/s at nominally zero applied magnetic field (ambient geomagnetic field).

between an unblocked singlet, which tunnels to $(2, 0)$, and a blockaded triplet $|T_0(1, 1)\rangle$. By varying the readout detuning and compiling a histogram of the resulting current values, the map in Fig. 6.5b is obtained. It reveals a range -0.45 meV $< \epsilon < -0.12$ meV where both charge states occur, identifying the PSB window. The width of this window directly reflects the energy splitting of 0.33 meV between $|T_0\rangle$ and $|S\rangle$ in the configuration with both electrons on the left dot (experimentally $(4, 0)$). A specific readout point is chosen within this window at $\epsilon = -0.29$ meV, and the corresponding current histogram is shown in Fig. 6.5c. The two peaks associated with singlet (S) and triplet (T_0) outcomes are clearly separated and are distinguished using a threshold current $I_{th} = 0.34$ nA. A current below I_{th} is assigned to a singlet outcome, whereas a current above I_{th} is interpreted as a triplet.

One of the main advantages of PSB readout is that it does not require a local electron reservoir. Whereas Elzerman readout (Sec. 4.5) relies on tunnelling between the qubit dot and a reservoir, PSB readout only uses interdot charge transfer. This is particularly beneficial for scaling, as the fabrication overhead associated with reservoirs, additional gate electrodes and ground connections can be reduced significantly [186, 218].

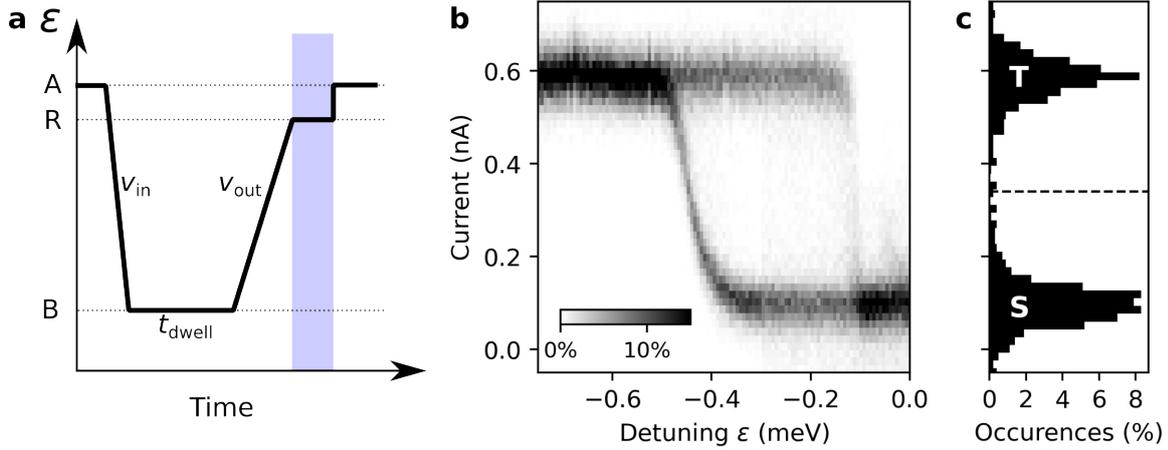


Figure 6.5: Determination of the PSB readout point. **a** Detuning pulse starting at the initialisation point A in the $(2, 0)$ charge configuration. The detuning is ramped with sweep rate v_{in} to point B in $(1, 1)$, followed by a dwell time t_{dwell} . The detuning is then ramped back with sweep rate v_{out} to the readout point R, where the SET current is integrated over a fixed time window (blue shaded region). After readout the detuning is returned to point A. For clarity the pulse shape is shown without the additional compensation of bias-tee effects discussed in Sec. 3.6. **b** Histogram of the recorded SET current as a function of readout detuning. The greyscale indicates the relative occurrence (in %) of a given current value during a $10 \mu\text{s}$ integration window. The SET is calibrated such that high current corresponds to the $(1, 1)$ charge configuration and low current to $(2, 0)$. For detunings $\epsilon < -0.45 \text{ meV}$ the system remains in $(1, 1)$, while for $\epsilon > -0.12 \text{ meV}$ it relaxes to $(2, 0)$. Between these limits a blockade region is observed where the readout detuning lies within the PSB window and both $|S(2, 0)\rangle$ and $|T_0(1, 1)\rangle$ contribute. **c** Line cut of panel **b** at $\epsilon = -0.29 \text{ meV}$. Two well-separated peaks are visible, corresponding to T_0 (upper peak) and S (lower peak) outcomes. The dashed line marks the threshold current $I_{\text{th}} = 0.34 \text{ nA}$ used to discriminate between S and T_0 . Parameters: $v_{\text{in}} = v_{\text{out}} = 220 \text{ keV/s}$; $t_{\text{dwell}} = 200 \mu\text{s}$; external magnetic field $B = 15 \text{ mT}$ along \hat{x} ; 1000 averages per detuning point.

6.4 Singlet-Triplet Oscillations for Sensing

The driving mechanism for singlet-triplet (ST_0) oscillations is the Zeeman energy difference between the two electrons in the left and right quantum dot. A SMM placed asymmetrically with respect to the double dot generates a local field gradient ΔB . This modifies the Zeeman energies of the two spins and produces a difference $\Delta E_Z = \hbar f_{\text{ST}}$ (in the limit $J \rightarrow 0$), which in turn drives ST_0 oscillations at frequency f_{ST} [17, 199]. Under the assumption $\mathbf{B}_1 \parallel \mathbf{B}_2$, the Zeeman energy difference can be written as

$$\Delta E_Z = E_{\uparrow\downarrow} - E_{\downarrow\uparrow} = \mu_B (g_1 B_1 - g_2 B_2) = \mu_B (g \Delta B + B \Delta g), \quad (6.8)$$

where B and g are the average magnetic field and average g -factor over the two dots, while $\Delta B = B_1 - B_2$ and $\Delta g = g_1 - g_2$ denote the corresponding differences. This expression makes clear that both a magnetic-field gradient and a g -factor difference can contribute to ΔE_Z . In the literature, ΔB generated by on-chip micromagnets [200, 219] or by an Overhauser-field gradient in GaAs [17, 220] has been identified as the dominant drive term for ST_0 dynamics. In contrast, experiments in SiMOS quantum dots have established Δg as the main contribution, originating from the Stark shift induced by vertical electric fields

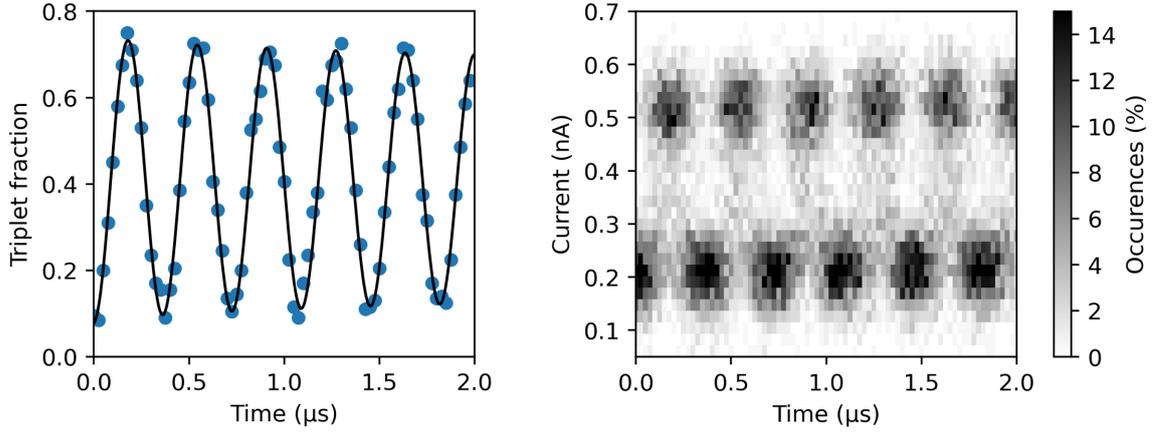


Figure 6.6: ST_0 oscillations. **Left** Measured triplet fraction as a function of the evolution time in the (1, 1) configuration. The data points (blue) are fitted with a cosine function (black line), yielding a frequency $f_{ST} = (2.749 \pm 0.005)$ MHz. **Right** Histogram of the single-shot SET readout traces underlying the left panel, obtained by integrating the SET current for 10 μ s. A current threshold of $I_{th} = 0.36$ nA is used to discriminate singlet and triplet outcomes. The data were taken at an external magnetic field of $B = 0.3$ T applied along the x direction.

on the plunger gates together with interface-induced spin-orbit coupling and valley-orbit mixing [67, 209, 221].

For sensing applications that target local magnetic-field gradients, a g -factor gradient is undesirable, since it produces a linear dependence $\Delta E_Z \propto \Delta g B$ and therefore makes the ST_0 oscillation frequency sensitive to changes in the global magnetic field rather than purely to ΔB . As discussed below, Δg exhibits a pronounced dependence on the angle of the applied field, which offers a means to minimise it. In this optimised configuration the ST_0 oscillation frequency becomes, to good approximation, a direct measure of the magnetic-field gradient across the double dot. The corresponding dependence of f_{ST} on ΔB is then

$$\frac{df_{ST}}{d(\Delta B)} = \frac{\mu_B g}{h} \approx 28 \text{ GHz/T}, \quad (6.9)$$

which sets the frequency sensitivity of the sensing technique.

The measurement of ST_0 oscillations employs the pulse sequence already introduced in Fig. 6.5, with a fixed readout detuning. By varying the dwell time in the (1, 1) configuration, coherent oscillations of the triplet fraction are observed, as shown in Fig. 6.6. A cosine fit to the data yields an oscillation frequency of $f_{ST} = (2.749 \pm 0.005)$ MHz at an external magnetic field of $B = 0.3$ T applied along the x direction.

In the industrial device studied no micromagnetic structure is present that could deliberately generate a field gradient, and the residual ^{29}Si concentration is reduced to ≤ 400 ppm by isotopic enrichment. The Overhauser-field gradient is therefore expected to be strongly suppressed, so that $\Delta B \approx 0$ and the ST_0 oscillations are attributed to a g -factor difference

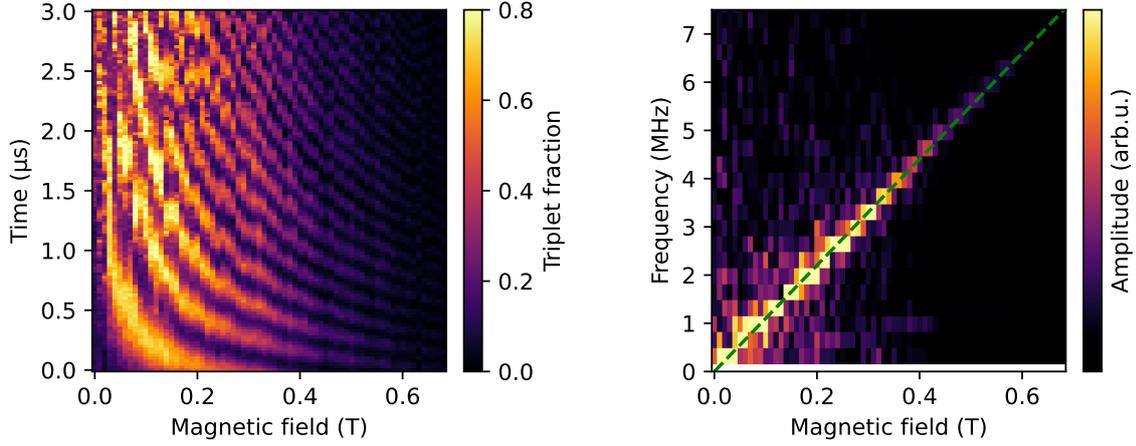


Figure 6.7: ST_0 oscillations as a function of magnetic field. **Left** Colour map of the triplet fraction as a function of evolution time in the (1, 1) configuration and external magnetic field applied along the x direction. The field dependence of the visibility is attributed to a detuning shift of the optimal readout point caused by the magnetic-field dependence of the Δ_{SO} anticrossing. **Right** Fourier transform of the time-domain data in the left panel, plotted as a function of frequency and magnetic field. The dominant oscillation peak disperses linearly with field and is fitted with Eq. (6.10), yielding $df_{ST}/dB = (10.90 \pm 0.07)$ MHz/T and $f_0 = (0.05 \pm 0.03)$ MHz.

Δg between the two dots. This interpretation is supported by the dependence of the ST_0 frequency on the external magnetic field shown in Fig. 6.7. A linear fit of the form

$$f_{ST}(B) = (df_{ST}/dB) B + f_0 \quad (6.10)$$

yields $df_{ST}/dB = (10.90 \pm 0.07)$ MHz/T and $f_0 = (0.05 \pm 0.03)$ MHz. Interpreting the intercept as a residual Zeeman splitting due to a static field difference gives

$$\Delta B = \frac{\Delta E_Z}{\mu_B g} = \frac{f_0 h}{\mu_B g} = (2 \pm 1) \mu\text{T}, \quad (6.11)$$

which is negligible within the experimental uncertainty and consistent with the absence of a designed magnetic-field gradient in the device.

The slope df_{ST}/dB can be used to extract the g -factor difference,

$$\Delta g = \frac{df_{ST}}{dB} \frac{h}{\mu_B} \Rightarrow \frac{\Delta g}{g} = (3.89 \pm 0.03) \times 10^{-4}. \quad (6.12)$$

This result indicates that the observed ST_0 oscillations are driven predominantly by a small gradient in the effective g -factor across the double dot, corresponding to a relative mismatch of 0.039%. The extracted mismatch is consistent with literature reports on SiMOS double dots, where the effective g -factor difference is typically of order 10^{-4} to 10^{-3} [209, 210].

One of the main drawbacks of using ST_0 oscillations for sensing is the strong angular dependence of Δg . This is illustrated in Fig. 6.8, which shows the ST_0 oscillations as a

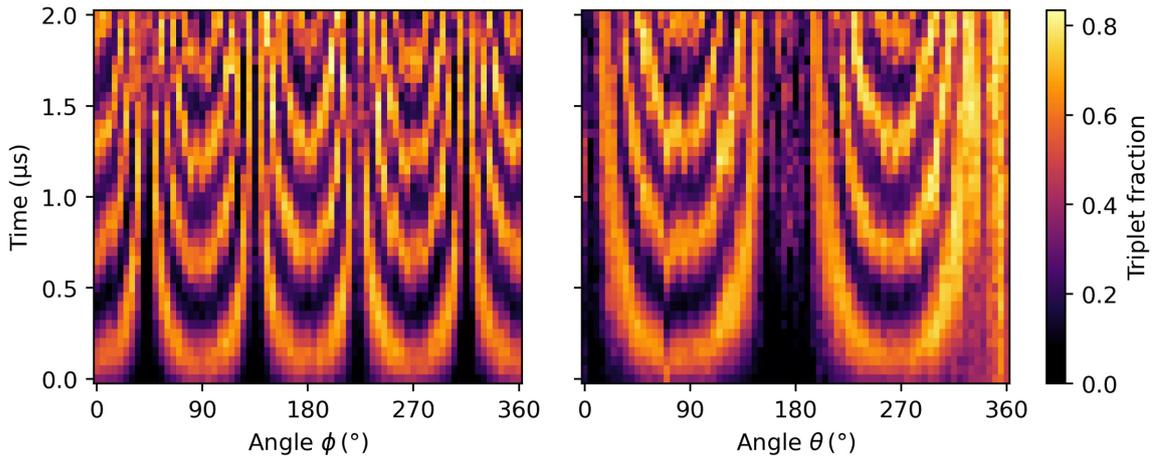


Figure 6.8: ST_0 oscillations as a function of the direction of the external magnetic field ($B = 150$ mT), parameterised in spherical coordinates. The colour scale (triplet fraction) is shared between both panels. **Left** Dependence on the azimuthal angle ϕ for fixed polar angle $\theta = 90^\circ$. Symmetric modulations of the ST_0 oscillation frequency f_{ST} are observed with a periodicity of 90° . The directions $\phi = 0^\circ$ (along x) and $\phi = 90^\circ$ (along y) correspond to the in-plane $\langle 110 \rangle$ crystal axes of the Si wafer. **Right** Dependence on the polar angle θ for fixed azimuth $\phi = 0^\circ$. Here f_{ST} exhibits a symmetric modulation with a periodicity of 180° . The direction $\theta = 0^\circ$ (along z) corresponds to the wafer normal (Si $\langle 001 \rangle$), while $\theta = 90^\circ$ (along x) again lies along a $\langle 110 \rangle$ in-plane axis.

function of the direction of the external magnetic field. The direction-dependent g -tensor of each dot causes the oscillation frequency f_{ST} to vary with angle in a manner that is symmetric with respect to the Si crystal axes ($\langle 001 \rangle$ and $\langle 110 \rangle$). A detailed discussion of the microscopic origin and modelling of this behaviour can be found in Refs. [208, 209, 216, 221]. At the same time, this angular dependence offers an opportunity: by choosing an operating angle (ϕ, θ) at or near a minimum of f_{ST_0} , one can realise a configuration where $\Delta g(\phi, \theta) \approx 0$ and the ST_0 oscillation frequency becomes, to a good approximation, purely sensitive to the magnetic-field gradient ΔB , which is the ideal regime for sensing experiments.

Overhauser-Field Background

Recording ST_0 oscillations over extended evolution times reveals a distinctly different behaviour from that observed in ESR Rabi experiments (Sec. 4.8). As shown in the left panel of Fig. 6.9, the ST_0 oscillations are initially coherent, but the phase rapidly becomes scrambled while the overall visibility remains high. In contrast, for single-spin Rabi oscillations (Sec. 4.8) fast noise during the evolution time usually leads to an exponential decay of the oscillation amplitude, while the phase of the averaged oscillations remains well defined.

In the ST_0 case, the observed dephasing is consistent with predominantly low-frequency noise. A natural contribution is slow magnetic-field noise from the residual ^{29}Si nuclei (400 ppm), which gives rise to a quasi-static Overhauser field during each experimental

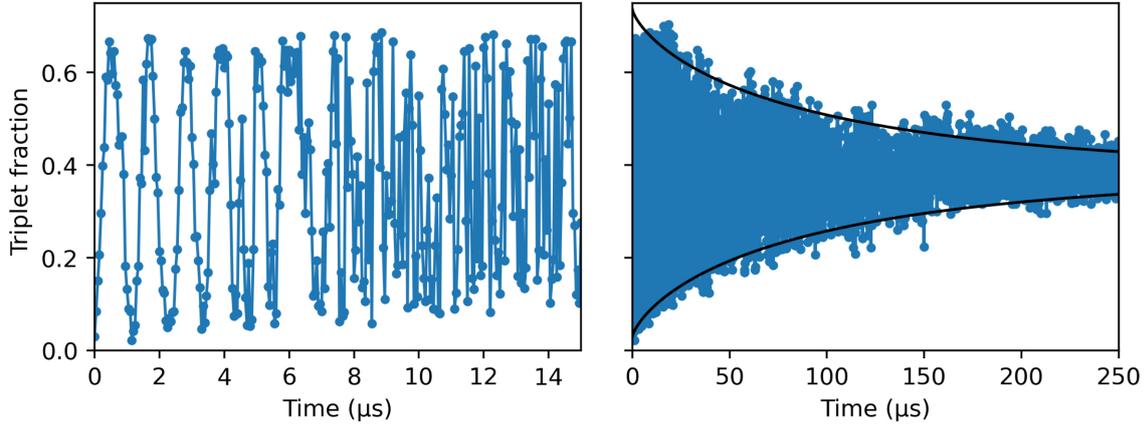


Figure 6.9: ST_0 oscillations measured over extended times. **Left** Triplet fraction as a function of dwell time in the (1, 1) configuration. The oscillations start out coherent at short times, but the phase becomes randomised at longer times while the visibility remains high, indicative of slow frequency fluctuations from a quasi-static Overhauser field. **Right** Same data plotted over a longer time window. The envelope is obtained by extracting local maxima and minima from windows of 30 neighbouring data points and fitting the upper and lower envelopes with $C \pm A \exp[-(t/T_2^{ST})^\beta]$, yielding $T_2^{ST} = (95 \pm 5) \mu\text{s}$ and $\beta = 0.74 \pm 0.04$. Each data point corresponds to an average of 500 single-shot measurements acquired over ≈ 10 s, with the full data set recorded over ≈ 14 h.

cycle and drifts on time-scales much longer than the ST_0 evolution and readout window [186, 222, 223]. In addition, low-frequency electrical noise can contribute via Stark shifts of the effective g -factor: fluctuations of the local electric field at the dot (due to gate voltage noise or charge rearrangements at the Si/SiO_x interface) translate into fluctuations of the ST_0 precession frequency [146, 209, 221]. Each data point in the time trace is obtained by averaging about 500 single-shot realisations over ~ 10 s, so that each point effectively samples a slightly different ST_0 frequency f_{ST} . This slow frequency wandering causes an apparent loss of phase coherence in the dataset, even though the underlying single-shot dynamics remain coherent. When the same data are viewed over a longer time window (right panel of Fig. 6.9), the decay of the oscillation envelope is well described by a stretched exponential of the form $C \pm A \exp[-(t/T_2^{ST})^\beta]$, yielding $T_2^{ST} = (95 \pm 5) \mu\text{s}$ and $\beta = 0.74 \pm 0.04$. A stretching exponent $\beta < 1$ indicates that the dephasing is governed by a broad distribution of slow noise processes rather than a single characteristic rate, which is consistent with slowly fluctuating Overhauser fields as well as slow Stark shift induced frequency drifts present during the measurement [224].

For sensing applications, the key figure of merit is not T_2^{ST} in isolation, but rather the choice of an evolution time that yields a stable and well-defined triplet fraction over the total measurement duration required by the sensing protocol. In practice, this means operating at dwell times where the ST_0 oscillations still appear coherent and the impact of Overhauser drift on the inferred signal is minimal.

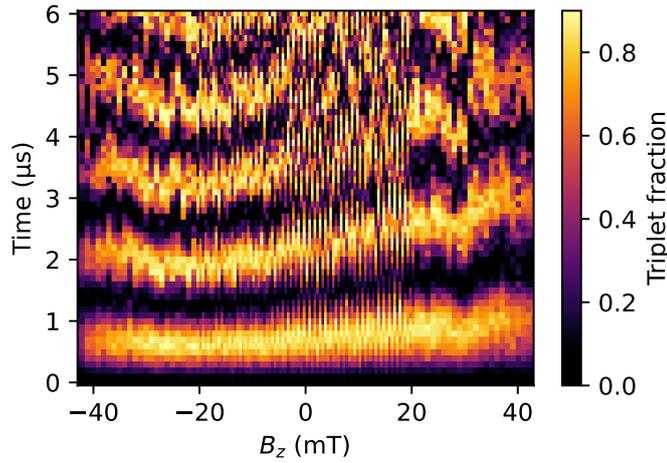


Figure 6.10: Sensing test with ST_0 oscillations. A diluted $TbPc_2$ crystal is placed on top of the device, intended to induce a magnetic-field gradient ΔB . Its easy axis is aligned along z . The colour map shows the triplet fraction as a function of dwell time and longitudinal magnetic field B_z , in the presence of a constant transverse field $B_x = 0.1$ T. Small changes in the ST_0 oscillation frequency with B_z are visible, but no clear signature that can be attributed to the $TbPc_2$ crystal is observed.

Proof-of-Principle Sensing with a $TbPc_2$ Ensemble

Sensing based on the ST_0 oscillation protocol was tested using the industrial SiMOS device. As discussed in Sec. 2.3, the present gate layout is not yet optimised for single-molecule sensing: the centre-to-centre separation between the double dot and the nominal SMM position exceeds 250 nm, such that the dipolar coupling to an individual $TbPc_2$ molecule is expected to be far below the detectable threshold. To enhance the magnetic signal, a 5% $TbPc_2$ crystalline ensemble (Sec. 2.2) was therefore deposited on top of the chip (Fig. 5.1).

A macroscopic $TbPc_2$ crystal is, however, not expected to generate a strong gradient ΔB across the ~ 80 nm double-dot separation, even though it produces a sizeable average stray field (Sec. 5.3). This is consistent with the measurements shown in Fig. 6.10. In this experiment a $TbPc_2$ crystal, with its easy axis aligned approximately along z , was positioned above the device. The triplet fraction was recorded as a function of dwell time in the (1, 1) configuration and longitudinal magnetic field B_z . Since $f_{ST} \approx 0$ for a purely z -oriented field at the chosen working point (where $\Delta g \approx 0$; see Fig. 6.8), additional transverse fields B_x were applied to render the ST_0 oscillations visible. The data shown in Fig. 6.10 correspond to $B_x = 0.1$ T. If the ensemble introduced a substantial gradient $\Delta B(B_z)$, additional features or shifts in the ST_0 oscillation frequency beyond the smooth background of g -factor anisotropy would be expected. Instead, only small, gradual changes of f_{ST} with B_z are observed, and no distinct signature can be unambiguously attributed to the $TbPc_2$ crystal.

These measurements therefore indicate that, in the current geometry, the field gradient produced by the $TbPc_2$ ensemble at the double dot is below the sensitivity of the ST_0 protocol. Realising ST_0 -based sensing of individual SMMs will require a reduced dot-

molecule separation that generates a larger and more controllable ΔB across the two dots. Gate-stack optimisations aimed at achieving this are discussed in Ch. 7.

Up to this point, all experiments have relied on gradients ΔE_Z and the exchange interaction $J(\varepsilon)$ to generate coherent dynamics in the ST_0 subspace. While the Zeeman-energy gradient has been characterised via ST_0 oscillations, the next step is to characterise the exchange interaction itself through exchange oscillations.

6.5 Exchange Oscillations

Beyond characterising $J(\varepsilon)$, exchange oscillations provide a direct verification of initialisation into the product state $|\downarrow\uparrow\rangle$. This constitutes an important step towards the proposed ESR experiments discussed in the next section. Demonstrating exchange-driven dynamics confirms that coherent control around both axes of the ST_0 Bloch sphere is available.

In the effective ST_0 Hamiltonian H_{ST_0} of Eq. (6.2), the exchange interaction $J(\varepsilon)$ generates rotations about the z -axis of the ST_0 Bloch sphere. In the regime $J(\varepsilon) \gg \Delta E_Z$, the eigenstates in the $S_z = 0$ subspace are $|S\rangle$ and $|T_0\rangle$. Conversely, in the opposite limit $J(\varepsilon) \ll \Delta E_Z$ the eigenstates are well approximated by the product states $|\uparrow\downarrow\rangle$ and $|\downarrow\uparrow\rangle$. The relation between both bases is

$$|\uparrow\downarrow\rangle = \frac{1}{\sqrt{2}}(|T_0\rangle + |S\rangle), \quad (6.13)$$

$$|\downarrow\uparrow\rangle = \frac{1}{\sqrt{2}}(|T_0\rangle - |S\rangle). \quad (6.14)$$

If the qubit is initialised in a product state in the regime $J(\varepsilon) \ll \Delta E_Z$ and the detuning is then changed diabatically such that $J(\varepsilon) \gg \Delta E_Z$, the initial product state becomes a coherent superposition of the eigenstates $|S\rangle$ and $|T_0\rangle$. The exchange splitting $J(\varepsilon)$ between these eigenstates leads to a relative phase accumulation in time and thus to coherent oscillations between $|\downarrow\uparrow\rangle$ and $|\uparrow\downarrow\rangle$ at frequency $\Omega(\varepsilon) = J(\varepsilon)/\hbar$.

The pulse sequence used to observe these exchange oscillations is illustrated in Fig. 6.11a. Initialisation into $|\downarrow\uparrow\rangle$ is achieved by an adiabatic ramp with sweep rate v_{in} from point A in the $(2, 0)$ configuration to point B in $(1, 1)$, where $J(\varepsilon) \approx 0$ and $|\downarrow\uparrow\rangle$ is, to good approximation, an eigenstate of H_{ST_0} . A fast pulse in detuning then moves the system close to the $(1, 1) \leftrightarrow (2, 0)$ charge transition, where $J(\varepsilon)$ is large, and the state is allowed to evolve for a variable dwell time t_{dwell} . This sudden turn-on of J (with respect to the spin dynamics) projects the product state into a coherent superposition of $|S\rangle$ and $|T_0\rangle$ and generates exchange oscillations. After the evolution period, the detuning is pulsed back to point B, effectively switching off $J(\varepsilon)$ again.

Readout is performed along the product-state axis by ramping adiabatically back to the readout detuning point R with $v_{\text{out}} = v_{\text{in}}$. During this ramp, the component of the state corresponding to $|\downarrow\uparrow\rangle$ maps onto the $|S(2, 0)\rangle$ branch, while the orthogonal component maps onto $|T_0(1, 1)\rangle$, which remains blockaded by PSB. The resulting singlet versus triplet

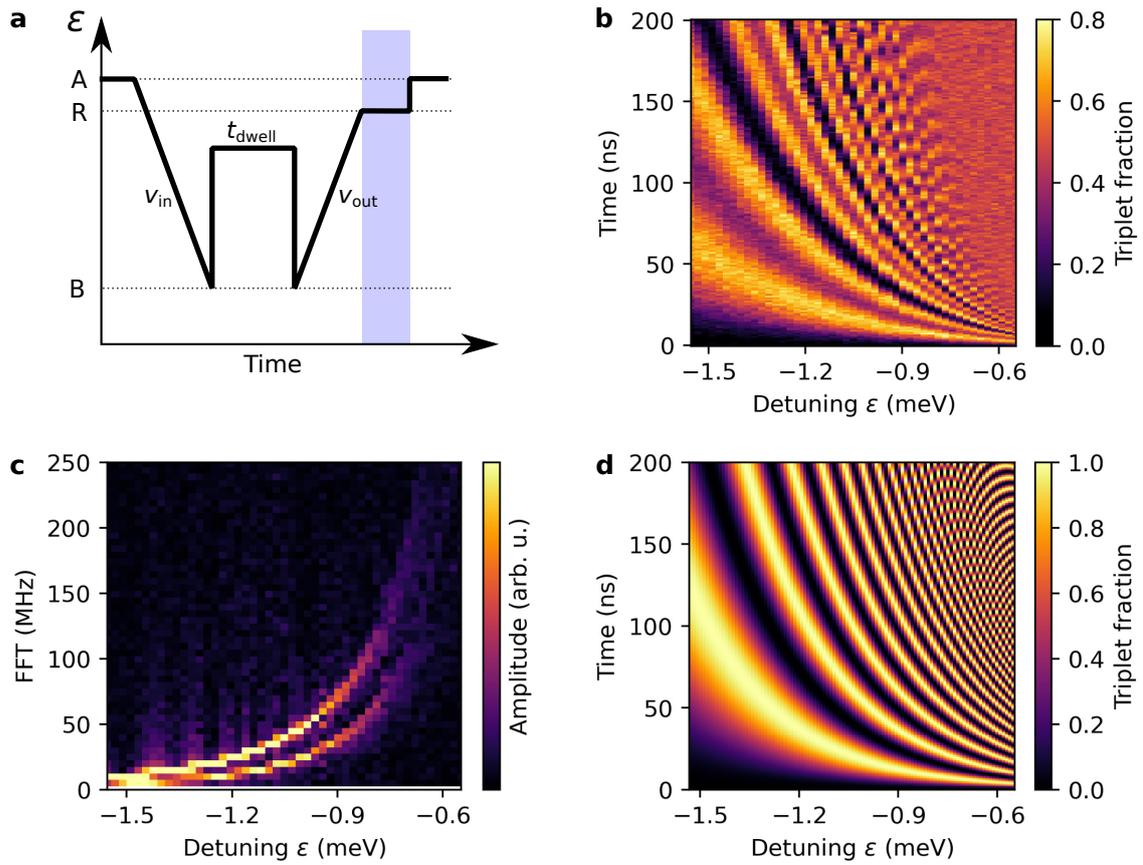


Figure 6.11: Exchange oscillations. **a** Schematic detuning pulse sequence. Starting from the initialisation point A in the $(2, 0)$ charge configuration, the detuning is ramped with sweep rate $v_{\text{in}} = 876 \text{ eV/s}$ to point B in $(1, 1)$. From there it is pulsed to a detuning value close to the $(1, 1) \leftrightarrow (2, 0)$ charge transition, where the exchange $J(\varepsilon)$ is large, and held for a variable dwell time t_{dwell} . The detuning is then pulsed back to B and subsequently ramped with sweep rate $v_{\text{out}} = v_{\text{in}}$ to the readout point R, where the SET current is integrated over a fixed time window (blue shaded region). Finally, the detuning is returned to point A. For clarity, the pulse shape is shown without the additional bias-tee compensation discussed in Sec. 3.6. **b** Measured triplet fraction (probability of the blockaded $|\uparrow\downarrow\rangle$ state) as a function of dwell time and detuning during t_{dwell} . The oscillation frequency increases approximately exponentially with detuning, and a pronounced beating pattern is visible, giving rise to dips in the oscillation maxima. The data were taken at an external magnetic field of $B_z = 30 \text{ mT}$. **c** Fourier transform of the data in **b**. Both branches are fitted with an exponential exchange model $J(\varepsilon) = J_0 \exp(\varepsilon/\varepsilon_0)$, yielding for the upper branch $J_0 = (2.3 \pm 0.2) \text{ GHz}$ and $\varepsilon_0 = (0.256 \pm 0.009) \text{ meV}$, and for the lower branch $J_0 = (2.4 \pm 0.3) \text{ GHz}$ and $\varepsilon_0 = (0.220 \pm 0.008) \text{ meV}$. **d** Numerical simulation of **b** using the exponential exchange model, performed with QuTiP and without including decoherence. In this simple model each detuning value gives rise to a single dominant oscillation frequency, and no additional frequency component is observed.

outcome probabilities therefore directly reflect the exchange-driven oscillations between the product states during t_{dwell} .

The experimentally acquired data are shown in Fig. 6.11b, which displays the triplet fraction as a function of dwell time for different detuning values during t_{dwell} . As the detuning is made more positive, the exchange interaction $J(\varepsilon)$ increases approximately exponentially and the oscillation frequency becomes faster. The dephasing of the oscillations is partic-

ularly rapid at the largest $J(\varepsilon)$, where $\partial J(\varepsilon)/\partial\varepsilon$ is largest and gate-voltage fluctuations (arising from charge noise or electrical noise) therefore have an enhanced impact on $J(\varepsilon)$ and directly modulate the oscillation frequency [178].

A further striking feature of the data is revealed by the Fourier transform shown in Fig. 6.11c. Two distinct frequency branches contribute to the exchange oscillations, with the upper branch appearing at approximately twice the frequency of the lower one. Both branches increase exponentially with detuning, with the upper branch following $J(\varepsilon)/h$ and reaching $J(-1.50 \text{ meV})/h \approx 6.5 \text{ MHz}$ and $J(-0.75 \text{ meV})/h \approx 121.4 \text{ MHz}$. A similar two-frequency structure has been reported in Ref. [211], where it was suggested that the double dot can occupy an incoherent mixture of two distinct singlet-triplet manifolds that are read out simultaneously, each with its own exchange energy. In that picture, the additional branch may arise from an excited valley or orbital configuration of the double dot. At present, the microscopic mechanism by which such an excited manifold becomes populated in this device remains unclear.

A QuTiP [203] simulation of the ST_0 Hamiltonian Eq. (6.2) is shown in Fig. 6.11d. Using a phenomenological exchange law $J(\varepsilon) = J_0 \exp(\varepsilon/\varepsilon_0)$ [178] and neglecting decoherence, the simulation reproduces the overall trend of increasing oscillation frequency with detuning. However, for each detuning value only a single dominant frequency component is obtained, and no additional frequency component is observed, underscoring that the simple single-manifold model does not capture the additional branch observed experimentally.

From the dominant (upper) frequency branch an exponential dependence

$$J(\varepsilon) = 2.3 \text{ GHz} \exp\left(\frac{\varepsilon}{0.256 \text{ meV}}\right) \quad (6.15)$$

is extracted. This parametrisation of $J(\varepsilon)$ is useful for choosing operating regimes such as $J \ll \Delta E_Z$ (gradient-dominated ST_0 sensing) or $J \gg \Delta E_Z$ (exchange-dominated dynamics). The observation of robust exchange oscillations and their exponential dispersion demonstrates that the protocol for preparing the product state $|\downarrow\uparrow\rangle$ is largely successful, even though the presence of a second frequency branch indicates that the initialisation is not perfectly confined to a single singlet-triplet manifold.

6.6 Towards ESR-Based Sensing with Pauli Spin Blockade

Spectroscopic sensing protocols based on tracking shifts of the qubit resonance frequency were discussed in Sec. 2.3 and Ch. 5. In the single-dot ESR experiments presented there, energy-selective tunnelling (Elzerman readout) requires a sufficiently large Zeeman splitting, such that the energy separation between $|\uparrow\rangle$ and $|\downarrow\rangle$ exceeds the relevant thermal broadening. This places a constraint on the external magnetic field and can be undesirable for sensing applications in which additional bias fields perturb the target system.

A double quantum dot offers an alternative route: PSB provides spin-to-charge conversion based on the large singlet-triplet energy separation in the $(2, 0)$ charge configuration [17,

98]. Importantly, this splitting is primarily set by orbital and valley physics and therefore persists even at very small magnetic fields. As a consequence, state preparation and readout remain feasible down to the low-field regime, in principle even at geomagnetic field levels. This makes PSB particularly attractive for ESR-based sensing in situations where external sensor fields should be minimised. The combination of coherent ESR control with PSB-readout has been demonstrated previously and provides clear experimental reference points [100, 218, 225].

The measurement procedure starts by preparing the antiparallel product state $|\downarrow\uparrow\rangle$ in the (1, 1) configuration, following the exchange-based preparation described in Sec. 6.5. During coherent control the system is pulsed deep into (1, 1) such that the exchange interaction is negligible, $J(\varepsilon) \approx 0$, and the product states are (to good approximation) eigenstates of the full two-spin Hamiltonian in Eq. (6.3). From $|\downarrow\uparrow\rangle$ two single-spin ESR transitions can be driven coherently,

$$|\downarrow\uparrow\rangle \leftrightarrow |\uparrow\uparrow\rangle, \quad |\downarrow\uparrow\rangle \leftrightarrow |\downarrow\downarrow\rangle. \quad (6.16)$$

Their frequencies are set by the local Zeeman energies of the two dots. Writing the static field components along the quantisation axis as $B_1 = B + \Delta B/2$ and $B_2 = B - \Delta B/2$ the two ESR frequencies read

$$f_{\text{ESR1}} = \frac{g_1\mu_B}{h} \left(B + \frac{\Delta B}{2} \right), \quad (6.17)$$

$$f_{\text{ESR2}} = \frac{g_2\mu_B}{h} \left(B - \frac{\Delta B}{2} \right). \quad (6.18)$$

Both transitions therefore respond to changes in the average field B , while their separation is governed by the Zeeman-energy difference ΔE_Z between the dots.

In the sensing concept pursued here, an asymmetrically placed SMM provides an additional local field contribution on one dot, thereby generating a measurable ΔB . Applying resonant radiofrequency tones to the on-chip antenna produces an oscillating drive field at the qubit, enabling coherent single-spin rotations and the implementation of Rabi and Ramsey pulse sequences (Sec. 5.4). By tracking the resulting oscillations as a function of drive frequency, f_{ESR1} or f_{ESR2} can be measured and monitored as the SMM magnetisation state changes and as the externally applied field direction is varied relative to the molecular easy axis.

After coherent control, the system is adiabatically pulsed to the PSB readout point near the (1, 1) \leftrightarrow (2, 0) transition, where the charge outcome depends on the spin configuration. For the chosen readout trajectory, $|\downarrow\uparrow\rangle$ is mapped onto the unblocked $|S(2, 0)\rangle$ branch, whereas the remaining product states stay in (1, 1) and remain blocked by PSB, yielding a charge signal that is detected with the SET. In the present context it is sufficient to read out the spin parity: antiparallel states (odd parity) versus parallel states (even parity). Parity readout is often advantageous compared with distinguishing S and T_0 , because the blockade decay rate of the even-parity outcomes Γ_{\pm} can be significantly slower than the decay rate of the T_0 state Γ_{T_0} [218]. Microscopically, for S - T_0 readout the unblocking of T_0

is dominated by efficient $S \leftrightarrow T_0$ mixing driven by a Zeeman-energy difference between the dots [98, 218, 226]

$$\Gamma_{T_0} \approx \Gamma_{\text{mix}}(S \leftrightarrow T_0) + \Gamma_{\text{spin-flip}}, \quad (6.19)$$

whereas the fully polarised triplets T_{\pm} (corresponding to $|\uparrow\uparrow\rangle$ and $|\downarrow\downarrow\rangle$) are protected by spin projection and typically require a genuine spin-flip process to unblock,

$$\Gamma_{\pm} \approx \Gamma_{\text{spin-flip}}. \quad (6.20)$$

Since in most devices $\Gamma_{\text{mix}}(S \leftrightarrow T_0) \gg \Gamma_{\text{spin-flip}}$, the parity-blockade decay can be orders of magnitude slower, directly improving discrimination in the charge sensor and the achievable readout fidelity [218].

Looking ahead, frequency tracking of f_{ESR1} and f_{ESR2} enables two TbPc₂ sensing experiments. First, ensemble magnetometry at reduced bias field: as observed in Sec. 5.3 and Sec. 6.4, TbPc₂ crystals predominantly provide an absolute stray field on the order of mT, while the field gradient ΔB at the scale of the double dot is negligible. In this regime, tracking f_{ESR1} or f_{ESR2} primarily measures changes in B . Relative to the single-dot readout used in Ch. 5, PSB-based readout removes the requirement of a large transverse bias field, which should reduce hyperfine-transition broadening and may enable the observation of sharp hyperfine-related features in the ensemble response. Second, single-molecule readout via a local gradient: for a single TbPc₂ placed asymmetrically on the double dot, the dominant signature is expected to be a local field contribution, i.e. a measurable ΔB with only a modest change in B . In this case, the frequency difference $\Delta f = f_{\text{ESR1}} - f_{\text{ESR2}}$ or tracking of a single ESR transition frequency provides a sensitive probe of magnetisation changes of the SMM.

In conclusion, ESR combined with PSB parity readout offers a promising route towards low-field sensing with double quantum dots, because spin-to-charge conversion no longer relies on a large Zeeman splitting to enable energy-selective tunnelling. Realising this experimentally requires (i) reliable preparation of $|\downarrow\uparrow\rangle$, (ii) sufficient spectral separation $\Delta E_Z/h$ to selectively address ESR1 and ESR2, and (iii) efficient delivery of the required drive fields in the low-frequency regime relevant for small bias fields. Once these technical challenges are overcome, ESR with PSB readout provides a powerful spectroscopic tool for detecting and tracking the magnetic state of single-molecule magnets.

7 Simulation-Driven Optimisation of Gate Stack Architectures for Spin Qubit Sensors

A central limitation for hybrid SMM-semiconductor sensors is the rapid $1/r^3$ decay of the SMM's dipolar field with sensor-sample separation r (see Sec. 2.3). To approach single-molecule sensitivity, the most effective lever is therefore geometric: reducing the physical distance between the SMM on the surface and the electron confined in the silicon quantum dot, while keeping the device in the few-electron regime with reliable tuneability.

In the industrial SiMOS device A (see Sec. 4.1, Tab. C.1), the vertical separation from a surface-deposited TbPc₂ molecule to the qubit is set by three stacked contributions above the channel: the gate stack (≤ 108 nm), the on-chip ESR antenna (≈ 100 nm), and a passivation layer (≈ 250 nm). The gate stack and antenna sits embedded within the passivation. This architecture is excellent for robustness, protection and yield in conventional qubit operation, but it is unfavourable for hybrid sensing since it places the molecule hundreds of nanometres away from the active quantum region.

The passivation layer is applied across an entire 300 mm wafer by the foundry and serves as a protection of the quantum region. For sensing, however, we intentionally seek controlled coupling to a spin on the surface. Requesting the omission of passivation at the foundry level is impractical. Instead, a post-processing route in our academic cleanroom is adopted and the passivation locally removed by a selective etch, thereby opening a window above the target dot (see Sec. 7.4). Once this window is opened, the passivation no longer contributes to the molecule-qubit separation within that region.

The ESR antenna placement is not an intrinsic constraint of the architecture. It need not be centred above the qubit: laterally offset designs have already been fabricated and characterised, showing only a modest reduction in drive speed. This follows from the near-field scaling of a straight conductor, whose microwave magnetic field decays as $B \propto 1/r$ [96]. Increasing the antenna-qubit distance r therefore lowers B and hence the Rabi rate. A shifted ESR line is thus entirely feasible, with the trade-off of slightly slower qubit manipulation. After passivation removal at the sensing site, the antenna no longer sets the minimum achievable SMM-qubit separation.

With passivation removed locally and the antenna positioned off-axis, the remaining spacer is the gate stack itself. This motivates a gate layout tailored to hybrid operation that creates a recess directly above the quantum dot. A practical route is to reduce the

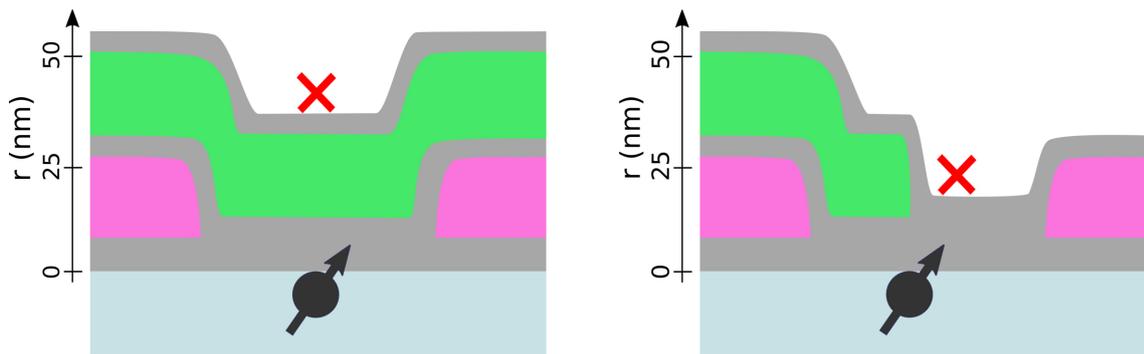


Figure 7.1: Gate-stack optimisation principle, the layer heights depicted in the figures are identical to device A in Sec. 4.1). **Left** In a standard stack, the minimum SMM-qubit distance (red cross to black spin) is set by the SiO_x thickness (grey) plus the plunger (green). The confinement gate C is pink; gates are n^{++} poly-Si; the qubit resides in isotopically enriched ^{28}Si (light blue). **Right** A shortened plunger gate opens a recess so an SMM can sit closer to the dot.

lateral extent of the plunger gate so that a pocket is opened above the dot (see Fig. 7.1, right). Since the trade-off is a reduced plunger lever arm, it is therefore useful to ask the following questions: (i) can shortened-plunger layouts still form few-electron dots at realistic biases; and (ii) can the confinement gates steer the dot toward the plunger edge so that the SMM can be positioned directly above it? Achieving this controlled lateral displacement towards the plunger is essential for minimising the SMM-dot separation. The nextnano [227] simulations below demonstrate that this is feasible.

In the purpose-designed layout of device A, once the local passivation has been removed and the antenna displaced, the remaining vertical separation is determined by the internal oxides of the gate stack and is approximately 18 nm. Using Eq. (2.14), a TbPc_2 at $r = 18$ nm produces a dipolar field change of about $5.6 \mu\text{T}$, corresponding via $\gamma_e \approx 28 \text{ GHz/T}$ to a qubit resonance shift of $\approx 160 \text{ kHz}$. Such shifts are resolvable by Ramsey spectroscopy and, more robustly, by dynamical-decoupling sequences such as Hahn-echo and CPMG (see Sec. 2.3). Further thinning is feasible: if the oxide is locally etched to achieve $r \approx 5$ nm, the expected frequency shift rises to $\approx 7.5 \text{ MHz}$, which is straightforward to detect.

The remainder of this chapter details the gate stack optimisation. It first outlines the simulation framework and validates the model against measurements on baseline devices. Two different shortened plunger architectures are compared, and their operating windows and dot-placement capabilities are quantified. The surface post-processing required to realise the recess is then described. Finally, initial device tests are presented for a sample fabricated in an imec 300 mm cleanroom for hybrid operation.

7.1 Computational Methodology and Simulation Framework

In this chapter, simulations are carried out with the Schrödinger-Poisson solver nextnano [227], which models semiconductor heterostructures from the nanometre to the micron scale.

It supports one-, two- and three-dimensional device geometries under equilibrium and non-equilibrium conditions, including external electric and magnetic fields, and accounts for strain prior to the electrostatic and quantum-mechanical solves. Material libraries cover group-IV and III-V (zincblende and wurtzite) semiconductors with arbitrary crystal orientation, doping profiles and multilayer stacks.

Our targets are SiMOS devices in which a two-dimensional electron gas (2DEG) is accumulated at the Si/SiO_x interface beneath patterned polysilicon gates. Within nextnano, the gate materials are treated as electrostatic Dirichlet boundaries (ideal conductors). Dielectrics enter via their relative permittivity, and silicon forms the quantum region in which bound states and subbands are computed. Unless stated otherwise, material parameters are taken from the native nextnano database to ensure consistency across designs.

Governing Equations and Working Principle of nextnano

The software treats quantum-confined semiconductor devices by solving the Schrödinger-Poisson system self-consistently [228]. For the devices considered here, lateral and vertical feature sizes are typically one to two orders of magnitude larger than the lattice constant (~ 0.5 nm [229]). At that scale, fully atomistic approaches (e.g. empirical tight binding) would require millions of atoms and basis orbitals and are therefore computationally prohibitive, whereas a continuum description retains the relevant physics at tractable calculation cost.

The continuum band-structure model employed is the $k \cdot p$ envelope-function framework [230]. Within this approach the crystal periodicity is absorbed into slowly varying envelope functions that obey effective Schrödinger equations. In this work, electrons in the conduction band are described by the single-band $k \cdot p$ (effective-mass) approximation, which captures the low-energy dispersion near the band minimum with a position-dependent effective mass of the electron and a conduction-band edge profile. Multiband $k \cdot p$ Hamiltonians (e.g. 6×6 for valence bands [231]) are available when band mixing or anisotropy must be resolved (notably for holes), but are not required for the electron states of interest in this study. The envelope $\Psi_{\alpha,E}(\mathbf{r})$ with the eigenenergy E and the conduction or valence band number α obeys

$$H_{\text{eff}} \Psi_{\alpha,E}(\mathbf{r}) = E \Psi_{\alpha,E}(\mathbf{r}) . \quad (7.1)$$

The effective mass Hamiltonian is given as

$$H_{\text{eff}} = -\frac{\hbar^2}{2} \nabla \cdot \left[m^{*-1}(\mathbf{r}) \nabla \right] + E_c(\mathbf{r}) + V_{xc}(\mathbf{r}) ,$$

where $m^*(\mathbf{r})$ is the effective mass, $E_c(\mathbf{r})$ the conduction-band edge, and $V_{xc}(\mathbf{r})$ an optional local-density exchange-correlation term. From the resulting wave functions, the local density of states is formed as

$$\rho(\mathbf{r}, E) = \sum_{\alpha} |\Psi_{\alpha,E}(\mathbf{r})|^2 .$$

Carrier densities of electrons follow from Fermi-Dirac occupation, with

$$f_{\text{FD}}(E, E_F, T) = \frac{1}{1 + \exp\left(\frac{E - E_F}{k_B T}\right)},$$

so that

$$n(\mathbf{r}) = \int \rho(\mathbf{r}, E) f_{\text{FD}}(E, E_F, T) dE$$

and the electrostatic potential satisfies Poisson's equation

$$\nabla \cdot (\epsilon_0 \epsilon_r(\mathbf{r}) \nabla \phi(\mathbf{r})) = e n(\mathbf{r}). \quad (7.2)$$

In its self-consistency cycle, nextnano begins with a pure band-edge quantum calculation, i.e. the first Schrödinger solve uses $E_c(\mathbf{r})$ to obtain an initial carrier distribution. Thereafter, the program alternates between updating the electrostatic potential $\phi(\mathbf{r})$ by solving Eq. (7.2) with the current carrier densities and boundary conditions, and recomputing the quantum states Eq. (7.1) using a total electron potential constructed from $E_c(\mathbf{r})$ and the self-consistent electrostatics

$$U_e^{(k)}(\mathbf{r}) = E_c(\mathbf{r}) - e \phi^{(k)}(\mathbf{r}).$$

The loop continues until the Poisson residual falls below the prescribed tolerance.

In the remaining (semi)classical regions - i.e. parts of the device where quantum confinement and wave-function penetration are negligible, nextnano does not solve the Schrödinger equation. Instead, carrier densities are obtained from a generalised Thomas-Fermi model [232] with Fermi-Dirac statistics that uses the local band edges and quasi-Fermi levels. These densities enter Poisson's equation Eq. (7.2) to update $\phi(\mathbf{r})$. This treatment captures drift of bands and electrostatics in bulk-like regions at low computational cost, while the explicitly quantum regions (thin channels/wells) provide the quantised subbands and wave-function shapes needed where confinement is essential.

Numerically, the nonlinear Poisson problem is handled by Newton-Raphson iteration [233] with an inexact line search for robustness and rapid convergence, while Schrödinger eigenproblems are solved with standard BLAS [234] and LAPACK [235] back-ends. Three-dimensional simulations are memory-bound; nextnano therefore supports anisotropic meshing, with typical in-plane spacings of 0.5 nm and adaptive refinement in z down to 0.01 nm near the interface, noting that the computational cost grows with the total number of grid points.

In the present, deliberately simplified simulations non-idealities at the Si/SiO_x interface were not included. Specifically, no fixed interface charge or trap density was incorporated into the Schrödinger-Poisson model. Dielectrics were treated as homogeneous and the interface as ideal. Consequently, absolute threshold voltages and lever arms may be globally offset relative to the experiment, an effect we absorb during tuning by adjusting gate voltages. This approximation is adequate for the aims of this chapter which were to compare gate layouts and demonstrate the feasibility of edge-localised dot formation.

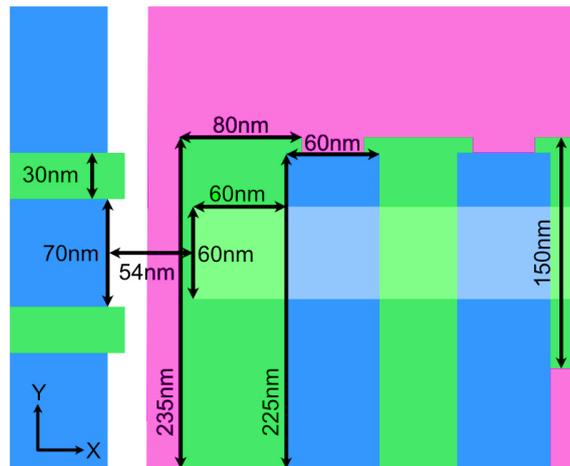


Figure 7.2: Gate layout of the imec SiMOS double-quantum-dot device with dimensions. The x-y plane is shown; different gate layers are colour-coded: bottom layer 1 (pink), layer 2 (green), top layer 3 (blue).

7.2 Baseline Device: Simulation of the Original Device

The simulations presented were carried out in collaboration with Niklas Ries as part of a Bachelor's thesis [236]; additional implementation details are documented there.

To keep runtimes tractable while retaining the relevant physics, the self-consistent Schrödinger-Poisson solver was restricted to the region that actually forms the quantum dot and its reservoir (under and around the gates P1, B1, and P2). Everywhere else on the chip, carrier densities were computed in a (semi)classical treatment using the generalised Thomas-Fermi model with Fermi-Dirac statistics, as outlined in Sec. 7.1. This hybrid partitioning concentrates quantum effort where quantisation is essential and avoids unnecessary cost in the SET regions.

The SET area (top gate and both barriers) was meshed with a coarser lateral grid to reduce memory and computation time. This produces visibly sharper edges in the simulated SET maps, but control runs with a uniformly fine grid yielded indistinguishable dot characteristics (position, electron density) in the qubit region, thereby justifying the mesh simplification for the SET.

Because the original imec design (Fig. 7.2) is available for laboratory testing, the simulated gate biases were derived from experimentally tuned devices and lightly adjusted for the numerical model. The voltages used in the simulation are listed in Appendix Tab. E.1. The resulting 2DEG density is shown in Fig. 7.3. By integrating the electron density in the quantum region beneath P1, we find an occupation of three electrons for the operating point considered.

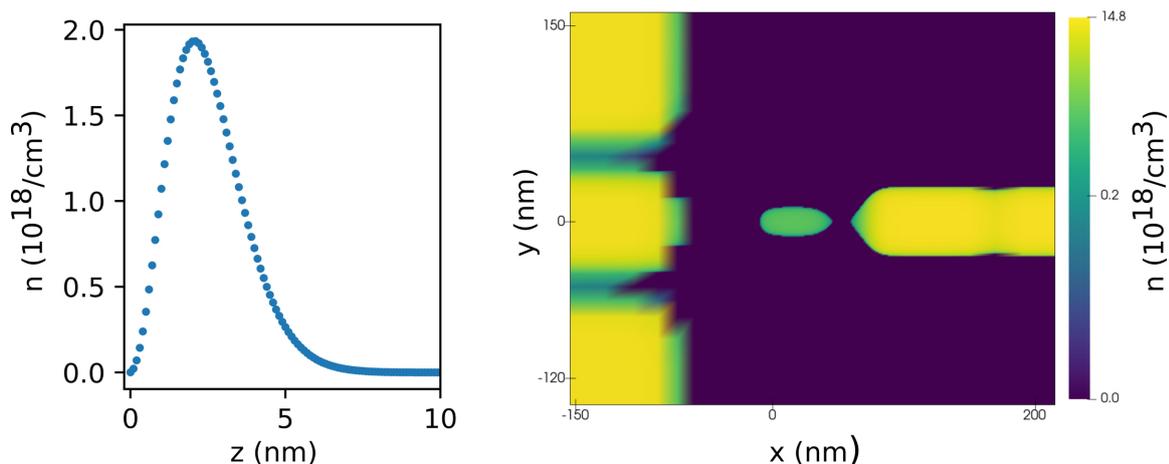


Figure 7.3: Electron density n simulated for a working point derived from experimentally tuned devices. **Left** Line cut along z (normal to the surface into the chip) through the dot centre; $z = 0$ nm is defined at the $\text{SiO}_x\text{-Si}$ interface. **Right** 2DEG density in the x - y plane over a 370 nm \times 300 nm window and integrated along z and on a logarithmic colour scale.

7.3 Gate-Layout Optimisation: Single- vs. Split-Plunger Designs

This section compares two gate layouts that are tailored for SMM sensing. While the eventual size and shape of the new layouts are of interest, the primary objective is to demonstrate feasibility: a quantum dot can be accumulated slightly off the plunger-edge so that its vertical separation from an SMM positioned above the recess is minimised.

As a baseline we use the original device employed experimentally in this thesis, with the simulated electron density discussed in Sec. 7.2. The optimised designs preserve the full gate stack and biases of the baseline, only the geometry of the left-dot plunger P1 is modified. To keep the parameter sweep both minimal and meaningful, a single geometric length of the plunger layout is chosen as the optimisation variable, while all other gate dimensions and voltages are held fixed.

A practical requirement is imposed on the plunger coupling: the modified P1 must retain a lever arm large enough to accumulate a dot with ≤ 10 electrons using plunger voltages ≤ 3.5 V. The 3.5 V ceiling reflects laboratory practice to avoid oxide breakdown between overlapping gates. Any geometry that cannot form a few-electron dot within this bias window is therefore ruled out.

Finally, the coupling to the reservoir is not explicitly simulated. With a functional barrier-gate layout in place, the tunnel rate can be tuned in experiment by adjusting the barrier gate, and is not expected to limit the conclusions of the present comparison.

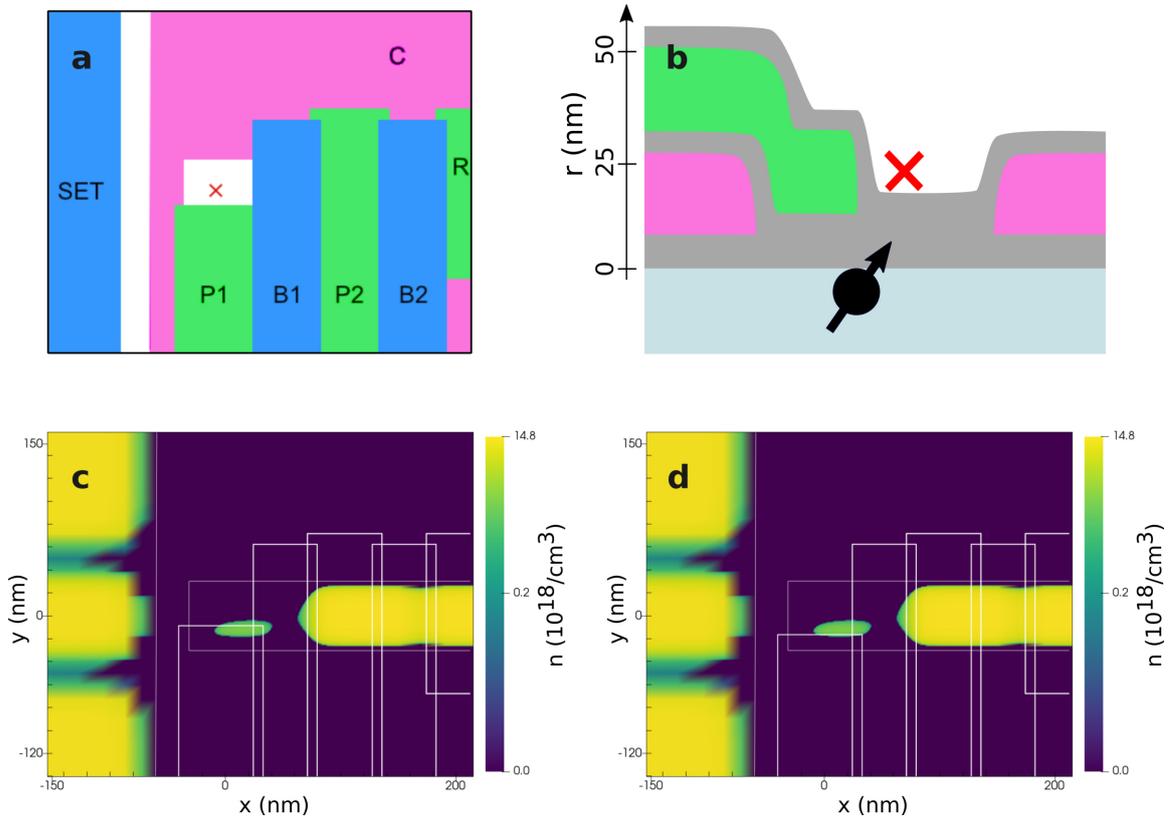


Figure 7.4: Reduced single-plunger P1 layout and simulated electron densities. **a** Top view of the gate layout showing the open recess where the SMM is intended to land (red cross). **b** Side view: the spin qubit (black) resides in isotopically enriched ^{28}Si (light blue); the plunger is green, the confinement gate C is pink, and oxides are grey. The recess allows the SMM (red cross) to sit closer to the dot. **c,d** 2DEG density in the x-y plane over a $370\text{ nm} \times 300\text{ nm}$ window and integrated along z and on a logarithmic colour scale. Gate footprint overlaid in white. Both cases accumulate ~ 3 electrons in the dot. In **c**, a 40 nm recess is obtained with $V_{P1} = 1.10\text{ V}$; the dot remains largely under P1. In **d**, a 50 nm recess with $V_{P1} = 1.80\text{ V}$ shifts the dot predominantly next to the plunger, i.e. at the recess edge.

Reduced Single-Plunger Layout

The single-plunger variant shortens P1 along the y-direction to open a recess above the dot (Fig. 7.4). In this geometry the simulated electron-density maps (panels c, d) show that a well-defined dot can be accumulated at the plunger edge once the recess is sufficiently wide. A recess of 50 nm which corresponds to a lateral overlap of the P1 edge with the C-gate aperture of 10 nm , reliably yields the desired off-edge localisation and is adopted as our baseline parameter choice.

The lateral positioning is enabled by the nearby confinement gate C, which is biased at -0.2 V and thereby electrostatically displaces the dot towards the centre of the C-gate gap. Both operating points illustrated yield a few-electron dot (~ 3 electrons) while keeping all other device dimensions and biases unchanged.

Split Double-plunger layout

The second layout employs two plunger fingers, P1 and P1M, which define a narrow slit (recess) between them (Fig. 7.5). Biasing the two plungers symmetrically accumulates a quantum dot in the centre of this recess. Although this requires one additional gate compared with the single-plunger variant, it provides deterministic lateral placement of the dot directly beneath the opening intended for the SMM.

The recess width controls the onset of unintended double-well behaviour: once it exceeds a critical value, the lateral confinement splits into two minima. In our simulations, recesses $\gtrsim 40$ nm produce two lateral density maxima (one near each plunger edge) separated by a shallow saddle; the bimodality is most evident on a linear colour scale (Fig. 7.6, left). Because such two-lobe profiles can evolve into a fully developed double dot if the saddle deepens or widens, we regard recesses < 30 nm as the safer operating regime, with smaller values preferable within fabrication limits.

A key advantage of the split layout is in-situ lateral tuning: applying slightly different voltages to P1 and P1M translates the dot along y . As illustrated in Fig. 7.6 (right), for a 50 nm recess the bias pair $V_{P1} = 1.10$ V, $V_{P1M} = 1.60$ V displaces the dot centre by ~ 15 nm along y . This tunability is attractive for optimising the SMM-dot coupling without altering the device geometry.

In conclusion, both proposed layouts can form a few-electron dot displaced towards the plunger edge at realistic biases. The split double-plunger offers clear advantages: deterministic accumulation of a dot in the recess, lateral tunability via asymmetric biases, and the possibility to maintain a larger lever arm with a relatively small recess (e.g. 20 nm), which also provides better electrostatic shielding of the 2DEG. These benefits come at the fabrication cost of one additional plunger gate.

7.4 Post-Processing: Buffered Oxide Etch Passivation Removal

With a gate layout tailored for strong SMM-qubit coupling in place, the remaining spacer between the SMM and the spin qubit is SiO_x . In addition, the foundry process at imec deposits a further 150 nm of SiO_x as a passivation layer to protect the device. For our experiments this additional passivation is detrimental because it increases the SMM-qubit separation. In this section we therefore summarise how the SiO_x passivation can be locally removed as an academic cleanroom post-process. The post-processing strategy was characterised in detail in the Master's thesis of Noah Gläser [188]; for exact procedures and recipes, please refer to that work.

Because aluminium bond pads and leads reside on the chip surface, a global passivation removal is infeasible as it would disrupt electrical connections. The objective is instead to open a $10\ \mu\text{m} \times 10\ \mu\text{m}$ window above each device so that an SMM can approach the

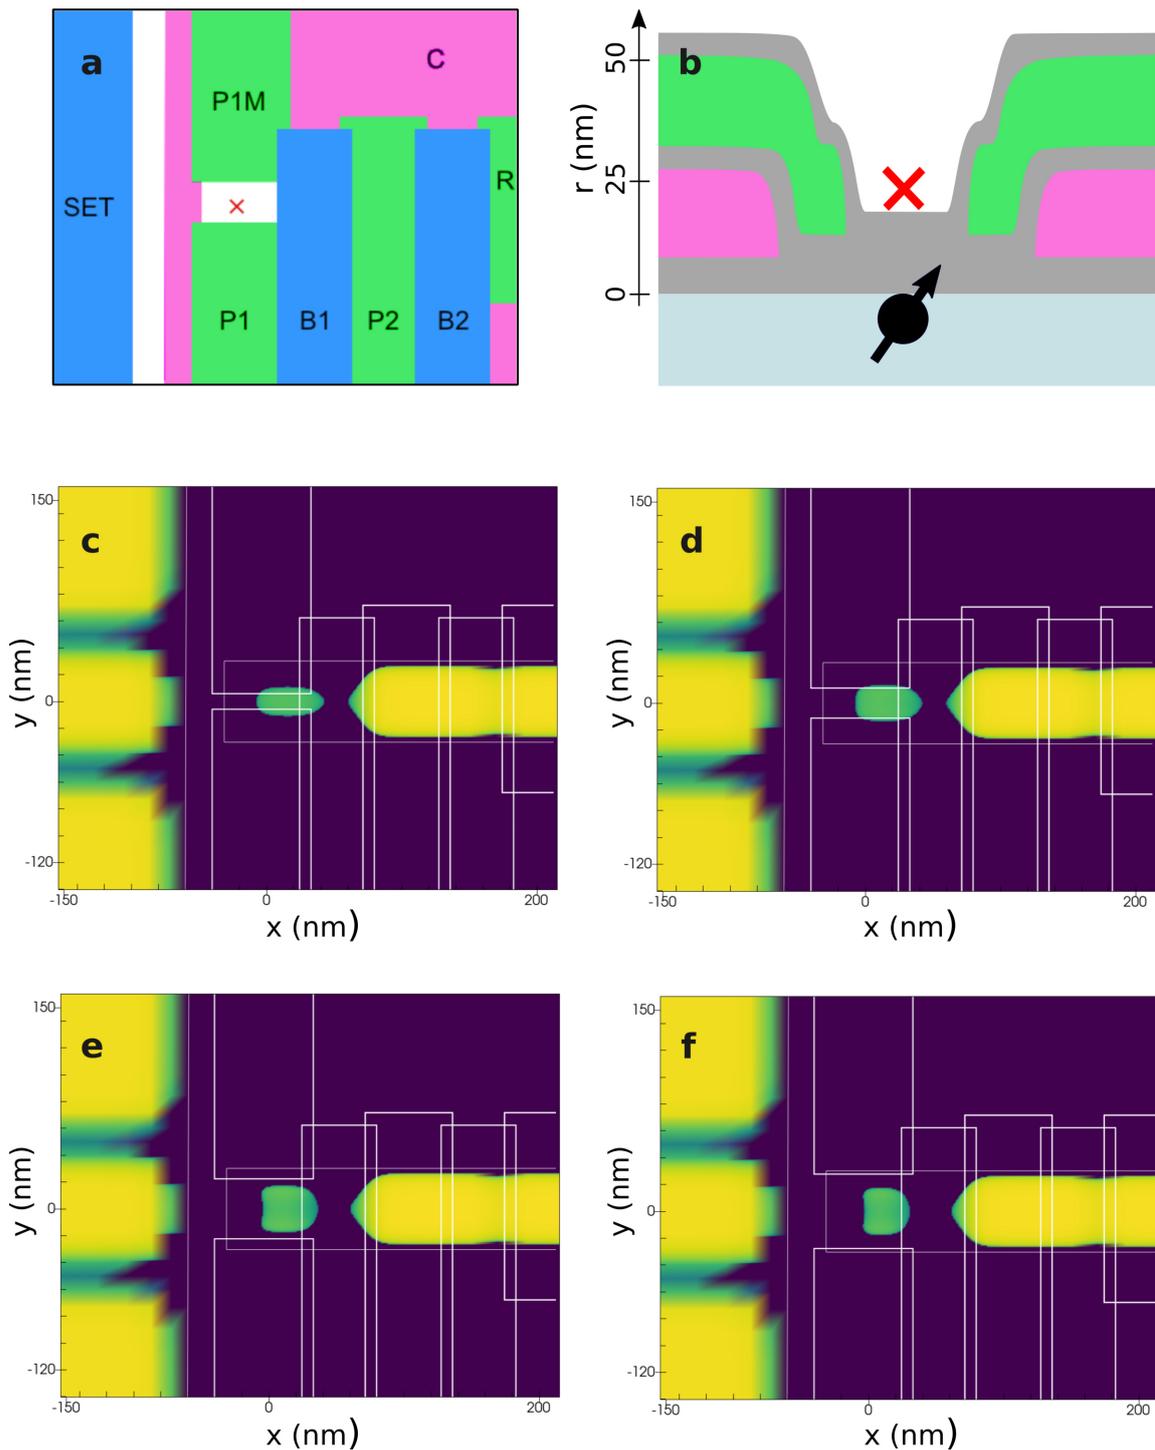


Figure 7.5: Split double-plunger layout and simulated electron densities. **a** Top view showing the open recess where the SMM is intended to land (red cross). **b** Side view: the spin qubit (black) resides in isotopically enriched ^{28}Si (light blue); plungers are green, the confinement gate C is pink, and oxides are grey. The recess allows the SMM (red cross) to sit closer to the dot. **c-f** 2DEG density in the x-y plane over a $370\text{ nm} \times 300\text{ nm}$ window, integrated along z and shown on a logarithmic colour scale (gate footprint overlaid in white). In all cases, the dot contains ~ 3 electrons. Symmetric plunger biases yield recess widths of **c** 10 nm at $V_{P1} = V_{P1M} = 0.70\text{ V}$, **d** 20 nm at 0.80 V, **e** 40 nm at 1.20 V, and **f** 50 nm at 1.60 V. The colour scale matches Fig. 7.4: blue indicates zero electron density and yellow corresponds to $14.8 \times 10^{18}\text{ cm}^{-3}$.

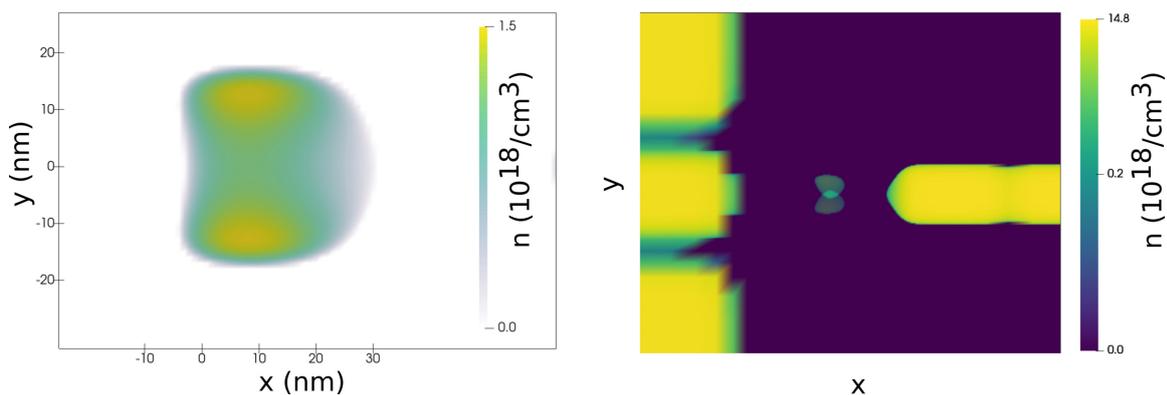


Figure 7.6: Asymmetric tuning with a split double-plunger. **Left** Electron density in the dot region plotted on a linear colour scale for the 50 nm recess case of Fig. 7.5 f; the density exhibits a two-lobe profile along y . **Right** 2DEG density in the x - y plane (370 nm \times 300 nm), integrated along z and shown on a logarithmic colour scale. For $V_{P1} = 1.10$ V, $V_{P1M} = 1.60$ V (overlaid with the inverted pair), the dot centre shifts by ~ 15 nm along y .

quantum region, while preserving the buried gate stack and, critically, the thermally grown gate oxide that provides inter-gate isolation.

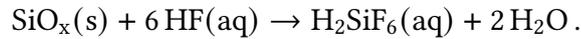
Semiconductor spin qubits are highly susceptible to inhomogeneities from nearby impurities [237, 238], so the ablation method must minimise damage and contamination. Wet chemical etching meets this criterion: particle energies are in the sub-eV regime with per-degree-of-freedom thermal energies $\sim k_B T \approx 0.03$ eV at room temperature, and the solid-liquid interface remains well thermalised to the bath. By contrast, prevalent dry-etch methods (plasma or reactive ion etching) employ ions with typical energies > 200 eV [239], and focused ion beams reach the keV regime [240]; this energy is deposited directly into the substrate, raising concerns of damage and contamination for sensitive qubit structures [241, 242]. Accordingly, wet chemical etching is adopted exclusively to preserve device integrity and materials purity.

Buffered oxide etchant (BOE) is a standard, SiO_x -selective wet etchant [243]. Although BOE is effective, it is a contact poison and requires strict handling protocols [244]. Other common etchants such as potassium hydroxide or piranha solution (sulphuric acid and hydrogen peroxide) do not, or only barely, attack SiO_x and are therefore unsuitable for this task [245].

To realise the local opening, the SiO_x passivation is removed only in a $10 \mu\text{m} \times 10 \mu\text{m}$ square centred above each device; all remaining areas are protected by photoresist. The resist-coated chip is submerged in BOE for a prescribed time and immediately rinsed in water to stop the reaction. Where SiO_x contacted the etchant, recesses remain whose depth depends primarily on the etch duration.

Etch chemistry and buffering.

The dissolution of SiO_x in hydrofluoric acid (HF) can be written as [246]



Fluoride ions (F^-) are continuously depleted at the interface and soluble fluorosilicate products form, which can transiently hinder access of fresh reactants. As the etch proceeds, the availability of active species diminishes and the rate slows unless buffered.

A buffered mixture (BOE) combines 49 % aqueous HF with 40 % aqueous ammonium fluoride (NH_4F) [245]. This has three main advantages: (i) NH_4F buffers and stabilises the HF_2^- concentration via the dissociation $\text{NH}_4\text{F}(\text{aq}) \rightleftharpoons \text{NH}_4^+ + \text{F}^-$ and subsequent complex formation $\text{HF} + \text{F}^- \rightleftharpoons \text{HF}_2^-$ -counteracting surface depletion and supporting a steady etch rate [246]; (ii) convenient etch rates at reduced free-HF concentration [247]; and (iii) HF_2^- , the active species in BOE, diffuses more slowly through resist than the smaller F^- ions that dominate in pure HF, mitigating under-mask attack [248]. In our case, “BOE” denotes a 7:1 mixing ratio of NH_4F solution to aqueous HF.

Process outcome and device integrity

In the final BOE test run a target etch depth of 130 nm was chosen. Using an etch duration of 77 s an average depth of 139 nm was obtained with a standard deviation of 2 nm across 27 etched devices [188]. An example of the etched window is shown in Fig. 7.7, where the opening is clearly visible in optical microscopy and atomic-force microscopy (AFM). One etched sample was subsequently tested at millikelvin temperatures: the qubit resonance was located by rapid adiabatic passage, demonstrating that the gate stack and ESR line remained fully functional after passivation removal.

In conclusion, buffered HF etching enables local removal of the SiO_x passivation in our academic cleanroom while preserving device operability. By reducing the oxide thickness above the qubit, the spacing between an SMM on the chip surface and the semiconductor spin qubit is minimised, directly supporting the hybrid-sensor architecture explored in this work.

7.5 Fabricated Device: Initial Tests

Devices implementing the split double-plunger layout were fabricated in the 300 mm wafer line at imec [34]. The left-dot plungers define a slit of 25 nm between P1 and P1M. The corresponding stability for a representative device in Fig. 7.8 shows that both plungers can independently-and jointly-accumulate electrons in the quantum dot, evidencing intact electrical isolation and effective control from each finger. Using the procedure described in Sec. 4.3, lever arms of 5.7% for P1 and 6.3% for P1M are extracted of the same device, with

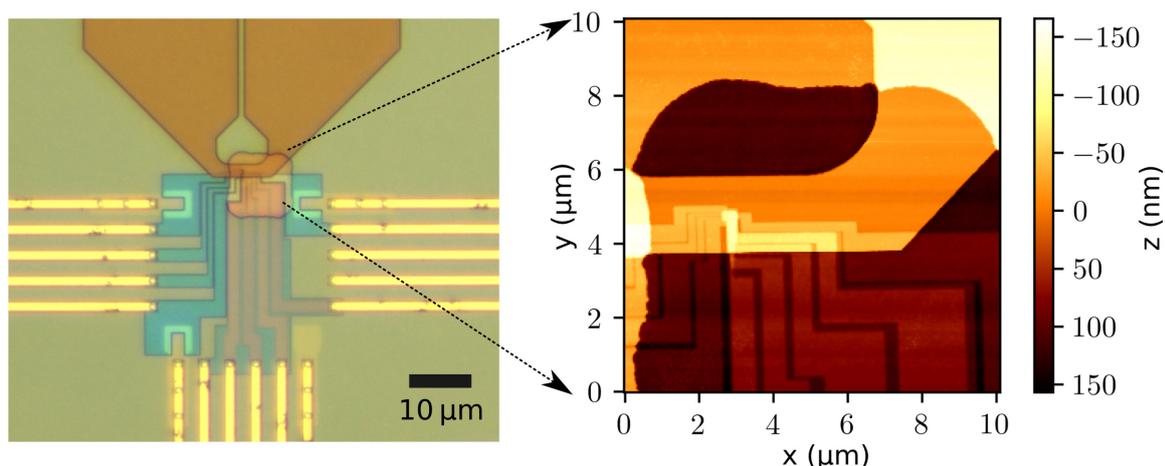


Figure 7.7: Enriched SiMOS device with a locally etched passivation window, imaged by optical microscopy and AFM. **Left** Optical micrograph of a qubit chip showing gate interconnects and ohmic contacts to the Si/SiO_x interface (left, right, bottom), where the 2DEG forms. The brown feature in the upper third is the ESR line; the qubit region lies beneath its taper. The etched opening appears as a rounded square centred below the taper. **Right** Atomic force micrograph of the etched window. The rounded-square edge is visible along the image perimeter. The orange ridge is the ESR line; near $x = 3 \mu\text{m}$, $y = 5 \mu\text{m}$ the metal gates converge to form the qubit layout, which is too small to resolve at this field of view.

statistical uncertainties below 0.1%. Employing both plungers for dot 1 yields a combined lever arm of 10.7% (P1&P1M).

In contrast, the right-dot plunger P2 employs a traditional (unslit) layout and exhibits a substantially larger lever arm of 17.5%. This increase reflects the absence of a slit above the channel, which maximises the gate capacitance.

The split double-plunger devices belong to a new generation of imec devices with a thinner SiO_x stack and several additional process optimisations, which together yield larger lever arms than reported earlier in this work. The top passivation remains present and must be removed to minimise the SMM-qubit spacing, as discussed in Sec. 7.4. Implementing this passivation-removal process is the next planned step for these devices.

This chapter established a pathway to reduce the SMM-qubit spacing while preserving few-electron operability. Using a self-consistent Schrödinger-Poisson framework (with semiclassical regions for efficiency), we validated a baseline SiMOS device and then evaluated gate layouts that create a recess above the dot. Both the shortened single-plunger and the split double-plunger architectures can position a stable few-electron dot at the plunger edge. The split double-plunger additionally provides in-situ lateral tuning via asymmetric biases and maintains a strong lever arm at modest recess widths, making it the preferred option for hybrid coupling. Complementarily, a buffered-oxide-etch post-process was demonstrated to open local windows in the passivation without degrading device functionality (Sec. 7.4). First imec-fabricated devices with the split layout exhibit independent plunger control. Taken together, these results show that the remaining spacer can be engineered down to the gate-oxide thickness, enabling SMM-induced frequency shifts in the MHz range for oxide thicknesses $< 10 \text{ nm}$, and motivate the next step of

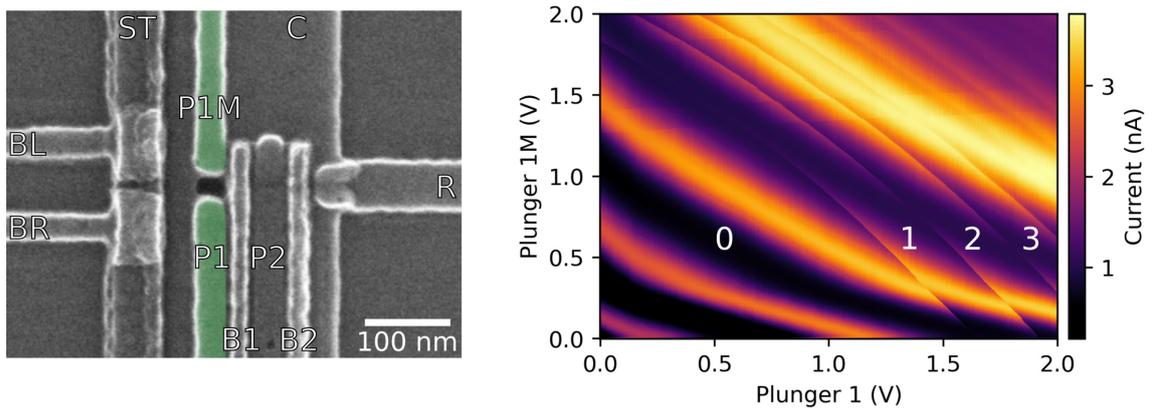


Figure 7.8: Custom-made split double-plunger device. **Left** Scanning electron micrograph of the gate layout with the two split plunger gates highlighted in green. Gate labels are indicated. **Right** Stability diagram of a split double-plunger device with dot occupation indicated by white numerals. Both plungers induce charge-addition lines (abrupt steps in the SET current), demonstrating that each gate can load electrons into the dot separately and in combination. The background modulation arises from Coulomb oscillations of the SET sensor.

applying local passivation removal to the new split-plunger devices and integrating surface SMMs.

8 Conclusion & Outlook

Single-molecule magnets (SMMs) provide chemically engineered, long-lived spin states that are attractive as qudits and quantum memories, but fast, local readout at the single-molecule level remains experimentally challenging. The central question addressed in this thesis is whether industrial silicon metal-oxide-semiconductor (SiMOS) spin qubits can act as viable quantum sensors for reading out SMMs. The hybrid architecture is motivated by the complementarity of the two subsystems: the SiMOS spin qubit provides an on-chip readout platform, whereas the SMM serves as a long-lived quantum memory and multilevel qudit. Overall, this work shows that SiMOS spin qubits already enable quantitative magnetometry of SMM ensembles, demonstrated here using terbium bis(phthalocyanine) in its anionic form $[\text{TbPc}_2]^-$. Progress towards single-molecule readout is currently limited primarily by device geometry, in particular the qubit-molecule separation imposed by the industrial gate stack.

Within this framework, the thesis first establishes the performance of the industrial SiMOS spin qubit as a reliable quantum sensor and uses quantitative estimates of the dipolar coupling to identify the geometric requirements for single-molecule sensitivity. On this basis, a first hybrid realisation is demonstrated in which an ensemble of $[\text{TbPc}_2]^-$ molecules is detected via qubit-based magnetometry. The thesis further develops double-dot sensing concepts that target both magnetic-field gradients and low-field absolute Zeeman splittings. Finally, self-consistent gate-stack simulations are used to design SiMOS devices with substantially reduced qubit-surface separation, providing a concrete route towards the single-SMM regime.

A dedicated experimental infrastructure based on a Sionludi L dilution refrigerator, equipped with a three-dimensional vector magnet and high-frequency wiring, was developed and optimised to support coherent control of industrial SiMOS spin qubits and magnetometry of SMMs. A central element of the conceptual framework is the choice of the SiMOS spin qubit as the sensing platform. Industrially fabricated SiMOS devices provide ultra-shallow quantum dots with excellent electrostatic control and reproducible device characteristics. In this context, the electron spin is not used primarily as a logical qubit for quantum computing. Instead, it serves as a quantum sensor whose resonance frequency shifts with the Zeeman energy set by the local magnetic field of nearby SMMs.

The suitability of the SiMOS device as a quantum sensor rests on reliable spin readout and coherent control of a single electron confined in the quantum dot. Spin-selective readout is achieved via a nearby single-electron transistor (SET), which converts the spin state into a detectable charge signal through energy-selective tunnelling. In this operating regime, the spin-lattice relaxation time is $T_1 = (117 \pm 9)$ ms, ensuring that relaxation

is negligible on the timescales of the readout and control sequences used throughout this work. Coherent control is established using Rabi, Ramsey, and Hahn-echo pulse sequences. Electrically driven electron spin resonance (ESR) yields a Rabi frequency of $f_R = (421.6 \pm 0.3)$ kHz, enabling rapid spectroscopy and resonance tracking. The measured inhomogeneous dephasing time, $T_2^* = (3.1 \pm 0.1)$ μ s, and Hahn-echo coherence time, $T_2^{\text{HE}} = (92 \pm 10)$ μ s, define the available evolution time for phase-sensitive sensing. The 30-fold increase in coherence under Hahn echo indicates that dephasing is dominated by slow, low-frequency detuning fluctuations, consistent with the filter-function picture in which Ramsey is highly sensitive to quasi-static noise whereas Hahn echo suppresses its DC component. Plausible contributions include Overhauser-field fluctuations from residual ^{29}Si nuclei, narrowband magnetic-field noise at 50 Hz and its harmonics from the environment and the vector-magnet current supply, and electric-field fluctuations that induce Stark shifts of the electron g -tensor at the Si/SiO_x interface. For the sensing protocols employed in this thesis, the relevant requirements are repeatable single-electron operation with spin-selective readout, coherent control with a Rabi rate that supports rapid resonance tracking, and coherence times that are sufficient for the chosen sensing sequences. The measured T_1 , f_R , T_2^* , and T_2^{HE} therefore provide a quantitative benchmark for the present sensor performance: they are adequate for robust ensemble magnetometry and resonance-tracking protocols, while also highlighting that longer T_2^* is required for high-resolution Ramsey-based sensing in the single-molecule limit.

The magnetic coupling between the SiMOS qubit and a TbPc₂ molecule is dominated by the dipolar interaction and thus decays as $1/r^3$ with the separation r between the molecular magnetic moment and the electron wavefunction. In the present industrial architecture, the gate stack and on-chip ESR antenna are encapsulated in a SiO_x passivation layer that fixes a minimum qubit-molecule spacing of 250 nm. At this distance, a magnetisation reversal of a single TbPc₂ molecule shifts the qubit resonance by only 60 Hz, far below the qubit's Fourier-limited linewidth of 3 kHz set by the measured echo coherence time. Quantitative coupling estimates show that approaching separations of a few tens of nanometres or less is necessary to reach frequency shifts in the hundreds of kHz, where they can be detected reliably by Ramsey-type pulse sequences. Importantly, both simulations and device considerations indicate that oxide thicknesses well below 10 nm should still yield functioning SiMOS qubits in the relevant few-electron regime, so such reduced separations are, in principle, technologically achievable. In the current generation of devices, this geometric constraint is overcome by using an ensemble of molecules: a dilute TbPc₂⁻ crystal in a diamagnetic YPc₂⁻ matrix is deposited on the chip surface, so that the net stray field from many molecules adds to produce a substantially larger, measurable resonance shift of 34 MHz with the existing industrial gate-stack geometry.

Magnetometry of a first hybrid architecture that couples an industrial SiMOS spin qubit to an ensemble of the single-molecule magnet [TbPc₂]⁻ is demonstrated using the Rapid adiabatic Passage during Spin-Selective tunnelling (RPSS) protocol developed in this work. A single frequency-chirped microwave pulse, applied concurrently with spin-selective readout, enables robust tracking of the qubit resonance. Depending on the required dynamic range, the protocol is operated either over a 100 MHz window with 1 MHz steps, or over a 30 MHz window with 0.1 MHz steps. This compact sensing scheme resolves

key magnetic signatures of the $[\text{TbPc}_2]^-$ ensemble: pronounced hysteresis loops when sweeping along the molecular easy axis, a clear angular anisotropy consistent with the easy-axis direction, and thermally activated magnetisation relaxation. The relaxation follows stretched-exponential dynamics with characteristic timescales of $\tau \approx 107$ min at 48 mK and $\tau \approx 0.8$ min at 140 mK. In the present configuration, a large transverse field of 0.66 T is required to maintain spin-selective readout of the SiMOS qubit. This field broadens the hyperfine-resolved switching steps of $[\text{TbPc}_2]^-$, due to a small distribution of molecular easy-axis orientations in the crystal, and prevents direct resolution of the Tb nuclear manifold. Nevertheless, the combination of RPSS-based resonance tracking and signal amplification by the molecular ensemble establishes a first quantitative benchmark for SiMOS-based quantum magnetometry of SMMs.

Double-dot sensing concepts based on a two-electron singlet-triplet (ST_0) qubit provide a complementary route to hybrid SMM-semiconductor magnetometry, motivated by the desire to operate at low absolute magnetic fields. This is attractive for molecular-spin readout because large transverse fields, needed for single-dot Elzerman readout, broaden field-sensitive features of the $[\text{TbPc}_2]^-$ ensemble and hinder direct access to its hyperfine structure. In one approach, the qubit is encoded in the $|S\rangle$ and $|T_0\rangle$ states, whose energy splitting depends predominantly on the difference in Zeeman fields across the two dots and, if the two dots have slightly different effective Landé g -factors, also on the absolute magnetic field. This makes the ST_0 qubit intrinsically sensitive to local field gradients while being comparatively insensitive to homogeneous field fluctuations. A TbPc_2 molecule placed asymmetrically with respect to the double dot therefore generates a differential Zeeman shift between the two dots, imprinting the molecular magnetisation onto the ST_0 precession frequency. In a second approach, ESR can be driven on one of the dots while readout proceeds via Pauli spin blockade between the (1, 1) and (0, 2) charge configurations, providing a means to measure the absolute Zeeman splitting at external magnetic fields close to zero. A key advantage of these double-dot schemes is that they do not rely on the large transverse field required for single-dot ESR with Elzerman readout. This relaxes an important constraint of the single-qubit architecture and provides a natural route to gradient-based and absolute-field sensing of SMMs in the low-field regime, where molecular hyperfine features are expected to be accessible. In the present device, coherent ST_0 oscillations are observed at (2.749 ± 0.005) MHz, set by the effective Zeeman-energy difference ΔE_Z between the dots. Pulsing into the strongly coupled regime yields exchange-driven oscillations reaching (123 ± 17) MHz. Together, these results confirm experimental access to both control axes of the ST_0 Hamiltonian via ΔE_Z and the tunable exchange interaction $J(\epsilon)$. A ΔE_Z gradient originating from an SMM has not yet been resolved in the present devices, because the qubit-molecule separation in the current industrial gate stack is too large for the resulting differential Zeeman shift to enter the detectable range.

Reaching single-molecule sensitivity requires increasing the dipolar coupling between the SiMOS qubit and the target SMM, which in practice is achieved most effectively by reducing the qubit-SMM separation set by the industrial gate stack. To address this geometric limitation, gate-stack optimisation for sensing applications is pursued using a self-consistent Schrödinger-Poisson solver implemented in nextnano. Different gate geometries are explored with the explicit goal of reducing the vertical separation between

the SiMOS qubit and an SMM while maintaining robust few-electron operation. A split double-plunger design is introduced, in which the plunger is implemented as two opposing finger gates that define a narrow gap above the dot. Simulations show that, for appropriate voltage biasing at the gates, the electron wavefunction can be steered into this gap and localised directly beneath the opening. For the present gate stack, the gap must remain below approximately 30 nm to avoid an inhomogeneous confinement potential in the quantum dot. These layouts can be combined with local removal of the SiO_x passivation layer in the opening. This passivation removal was implemented and tested via selective hydrofluoric wet etching. In the simulated structures, the combined approach reduces the effective qubit-surface separation from hundreds of nanometres to well below 10 nm, bringing the expected dipolar coupling into the MHz range for a single TbPc₂ molecule. A first generation of such sensing-optimised devices has been fabricated in the industrial cleanrooms of imec, implementing a split-plunger gap of 25 nm. This demonstrates that the proposed gate-stack modifications are compatible with the underlying SiMOS process flow and provides a concrete route towards experimentally accessing the single-SMM regime.

Taken together, the results of this thesis show that industrial SiMOS spin qubits can already function as practical quantum sensors for SMM ensembles, while also defining a realistic and technologically grounded pathway towards single-molecule sensitivity. In the current devices, ensemble magnetometry of TbPc₂⁻ demonstrates that the hybrid architecture is viable at the many-molecule level, and the main obstacle to reading out a single SMM is identified as geometric: the qubit-molecule separation set by the gate-stack and oxide thickness. Simulation-driven gate-stack optimisation, combined with the prospect of oxide thicknesses well below 10 nm, indicates that this limitation can be mitigated by realistic process modifications, bringing the expected dipolar coupling into a regime that is resolvable with established sensing protocols.

In the near term, hybrid measurements using double-dot Pauli spin blockade aim to demonstrate qubit operation in the new sensing-optimised SiMOS devices and extend ensemble experiments into a low-field regime where the Tb hyperfine structure can be resolved clearly. Progress towards the single-molecule level will rely on refined molecular placement strategies and on advanced sensing protocols that combine Ramsey-type sequences with PSB readout or ST₀-based gradient detection, exploiting both absolute-field and differential-field sensitivity. In the longer term, arrays of such sensors or mobile sensor qubits that can be shuttled on-chip could act as buses mediating interactions between multiple molecular spins. In this hybrid picture, TbPc₂ provides a chemically engineered, multilevel quantum memory. Combining such molecular states with an industrial silicon platform offers a route towards hybrid devices with on-chip readout of long-lived molecular degrees of freedom. The same approach also enables new opportunities in quantum sensing and in fundamental studies of nanoscale magnetism.

A Custom DC Line Fabrication

A.1 Fabrication of DC Lines

The following procedure describes the fabrication of a 24-line DC wiring assembly housed in a FeCuNi capillary terminated with a Fischer 00 series connector and a Micro-D connector (Comtronic CMKB-N1-B-25S-S). This wiring is used to route DC signals from room temperature to the mixing chamber stage in the Sionludi cryostat. The fabrication procedure described here, as well as the cabling of the newly commissioned Sionludi L cryostat, was carried out in collaboration with Viktor Adam [119]. Photographs of the completed DC lines, showing the connection from the mixing chamber stage to the sample PCB, are presented in Fig. A.1.

Materials

- Constantan wire, insulated, 80 μm diameter, $\sim 130 \Omega/\text{m}$, heat-resistant purple insulation.
- FeCuNi capillary, 70 cm length (Sionludi L) or 100-110 cm length (Sionludi XL), 1.0 mm inner diameter, 1.5 mm outer diameter.
- 0.25 mm diameter steel or copper pull-through wire, 2.2 m length.
- Apiezon N grease.
- Heat-shrink tubing, 7 mm length, two pieces per capillary end.
- Hermetic Fischer 00 series connector.
- Micro-D connector (Comtronic CMKB-N1-B-25S-S).
- 1.5 mm copper wire (bare, preferably silver-coated) for strain relief.
- Concentrated sulphuric acid (for insulation removal).
- Silver powder epoxy.

Procedure

1. **Wire preparation:** Cut insulated Constantan wire to the desired length: 70 cm for Sionludi L or 100-110 cm for Sionludi XL. Target a total line resistance of 70-100 Ω .
2. **Capillary preparation:** Cut FeCuNi capillary to the required length and carefully smooth the ends to prevent damage to the wire insulation during insertion.
3. **Wire bundling:** Fix two screwdrivers 90 cm apart. Anchor one end of the wire to one screwdriver, then wrap it 12 full turns around both screwdrivers, adding 3 extra turns. This yields 30 wires in total, 6 of which serve as spares.
4. **Tinning capillary ends:** Apply a small amount of tin using soldering grease to the ends of the capillary to aid later soldering. Leave 5 mm from each end untinned; tin the following \sim 1.5 cm to allow attachment of the strain relief.
5. **Heat-shrink placement:** Place two 7 mm heat-shrink tubes onto each capillary end for later protection of the wires.
6. **Pull-through wire setup:** Pass the 0.25 mm pull-through wire through the loop of bundled DC wires, bringing both ends together. Insert both ends into the capillary and wrap around a third screwdriver. Tension the pull-through wire and secure it.
7. **Grease application:** Apply Apiezon N grease evenly to the DC wires to enhance thermal contact within the capillary.
8. **Inserting wires into the capillary:** Remove the middle screwdriver and carefully slide the capillary over the DC wires, ensuring a small blob of grease emerges at the far end. If not, add more grease.
9. **Trimming ends:** Once the capillary is centred, cut off the wire ends that were in contact with the screwdrivers to remove any mechanically damaged sections.
10. **Insulation removal - first side:** Prepare a plastic guide plate (e.g. KF40 plastic blank flange) with \geq 30 small holes. Thread each wire through a hole, then immerse \sim 4 mm of the ends into concentrated sulphuric acid until the insulation disintegrates (\sim 1 h in cold acid, less if warm).
11. **Strain relief fabrication:** Form 1.5 mm copper wire into an L-shape (5 mm short leg, 20 mm long leg). Solder to the ground of both connectors in a position that avoids interference with tension reliefs or shielding boxes. Pre-heat connectors to \sim 100 $^{\circ}\text{C}$ for easier soldering.
12. **Micro-D connector termination:** Mechanically fix the capillary in the desired position relative to the connector. For each wire, pull until the stripped section can be inserted into the correct solder pin. Bend the uninsulated length into a V-shape to hide it inside the pin. This V-shape of exposed wire will clamp itself in the solder pin. Solder with a fine tip, ensuring minimal exposed conductor.

13. **Looping and strain relief:** Create a small loop between connector and capillary to better organise the wires inside the connector housing later. Pull back slightly so the heat-shrink at the capillary end can fit inside the shielding box.
14. **Heat-shrink and strain relief soldering:** Shrink the tubing at both ends to protect the wires. Quickly solder the strain relief to the capillary using soldering grease and $\sim 400^\circ\text{C}$ to minimise insulation damage.
15. **Excess wire handling:** Gather spare wires together and bind with fine string or copper wire (Wolfgang would say: “angel’s hair”).
16. **Preparation of opposite end:** Cut the extending DC lines to 3 cm and remove insulation as before (~ 4 mm via acid bath).
17. **Testing - initial:** Using a breakout box, verify there are no shorts between lines or to ground. Repair or replace faulty lines.
18. **Fischer connector termination:** Identify DC line 1 and solder it to pin 1 of the Fischer connector. Connect remaining lines accordingly. Solder the strain relief to the capillary and connect spare wires to ground.
19. **Testing - final:** Test all lines with breakout boxes for shorts or cross-connections. Correct as necessary.
20. **Final fitting:** File shielding boxes as required for connector fit. Adjust tension reliefs to ensure boxes close flush.
21. **Sealing:** Apply silver powder epoxy to seal the Fischer connector against radiation leaks at the cable entry point. Do *not yet* apply to the Micro-D connector, which will be frequently handled during installation.

Tips

- Test both breakout boxes for internal faults before use.
- Verify space inside shielding boxes before soldering strain reliefs.
- Use tweezers connected to a multimeter for easier continuity checks.
- As an alternative to acid stripping, melt insulation by rotating the wire end in a blob of hot solder, but avoid excessive heat that may dissolve conductor material.
- Remove residual insulation debris after the acid bath using plastic tweezers.
- For correct connector alignment, pre-shape a stiff copper wire inside the cryostat to determine the required angle between connectors.

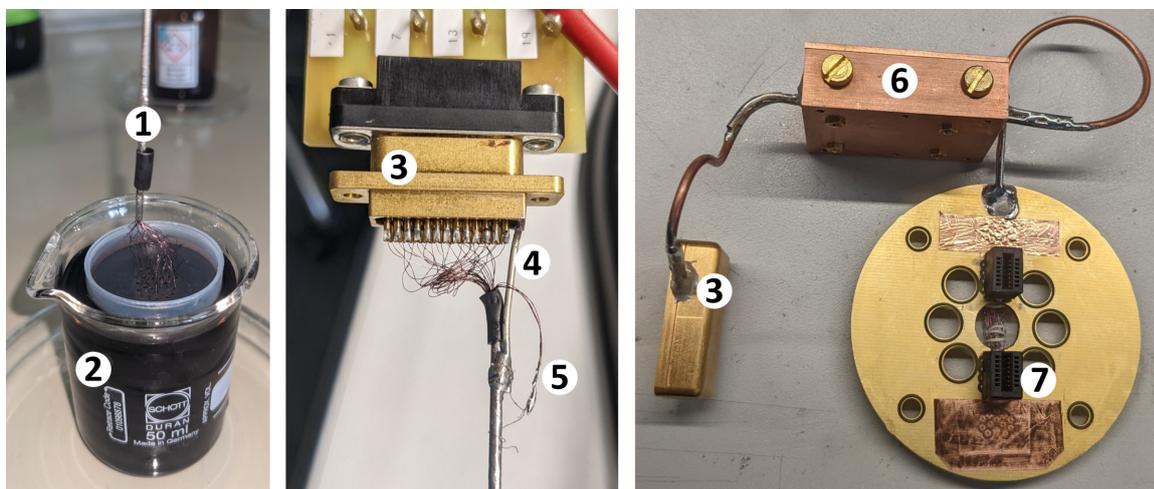


Figure A.1: DC wiring assembly above the mK stage. **Left:** FeCuNi capillary (1) with exposed Constantan wires prepared for insulation removal in concentrated sulphuric acid (2). A perforated plastic cap holds each wire in place at the acid surface. During this process, the heat-shrink tubing was inadvertently displaced from its intended position. **Middle:** Micro-D connector (3) with strain relief (4) and spare wires (5) for replacement in case of unwanted electrical contact. **Right:** Assembled mK-stage DC lines, showing silver epoxy sealing at the Micro-D connector (3), first-stage low-pass filter enclosed in a copper housing (6), and PCIe connector (7) for electrical connection to the sample PCB.

B Design and Fabrication: The Beast

Sample PCB and Connector Board

This appendix supplements the main text with fabrication details of the new sample PCB (The Beast) and the connector board mounted on the mK stage of the cryostat.

B.1 Layer Stack of The Beast

Figure B.1 shows the twelve-layer buildup. A key complexity is the via technology and associated lamination sequence. Of the 3703 vias, only 125 are through-vias; the remaining 3578 are blind or buried, which necessitated a multi-stage drill-plate-laminate process: nine drilling/plating steps and five sequential laminations, while maintaining layer-to-layer registration. The central sub-stack (layers In4–In7) is specified at ~ 1.4 mm overall thickness to mate reliably with the PCIe edge connectors used in the setup.

We use a combination of through-, blind-, and buried-vias to meet RF isolation and routing density goals while respecting manufacturing constraints [249]. Fabrication houses typically specify an aspect ratio (AR) limit for mechanically drilled vias, defined as

$$AR = \frac{\text{via diameter}}{\text{via depth}},$$

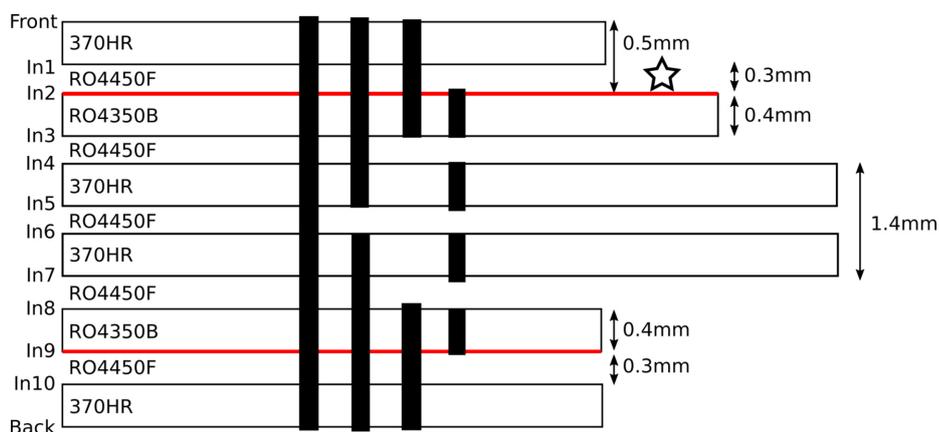


Figure B.1: Layer stack of The Beast. The twelve copper layers and dielectric materials are annotated at left; the chip location is marked by a star. Via types (through, blind, buried) are indicated as solid black interconnects. The CPW routing layers are highlighted in red. In total, 3703 vias are used in the design.

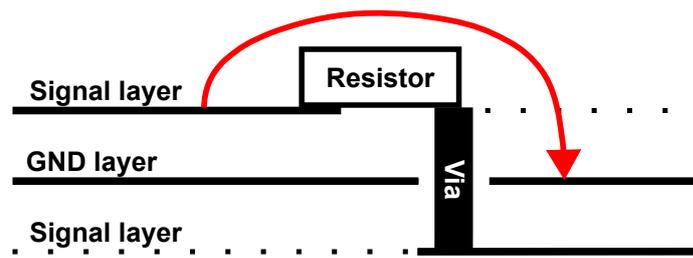


Figure B.2: Low-pass crosstalk reduction by buried routing. A top-layer DC trace (black) with an in-line resistor hands off to an inner signal layer through a via. A continuous ground plane between the top and inner layers screens the electric RF field (red), thereby reducing high-frequency leakage around the resistor.

(noting that some manufacturers define the inverse quantity). Small AR values are harder to drill and plate uniformly. While through-vias can meet AR limits of 0.1, shallow blind-vias (depth $\lesssim 0.3$ mm) are constrained by minimum drill and plating requirements, which effectively forces a larger finished diameter (AR = 2) and correspondingly larger pad/antipad geometries [249]. Laser-drilled μ -vias (typical depth ~ 150 μ m) were not employed here due to the layer thickness.

Vias are also used deliberately to suppress unintended RF coupling across the lumped low-pass stages. The short copper spans to and from an in-line resistor act as a small shunt capacitance, potentially bypassing high-frequency components that the filter is meant to attenuate. To mitigate this, the signal is dropped using a plated via to an inner routing layer immediately beneath a solid ground plane, forming an effective shield against fringing fields and reducing high-frequency crosstalk (see Fig. B.2). In The Beast this practice is used wherever space allows; the transition via is placed slightly away from the SMD pad to avoid interference during soldering.

B.2 Accessories for The Beast

Two auxiliary tools are supplied with The Beast: (i) shorting connectors for ESD safety during handling and storage, and (ii) a breakout PCB for rapid continuity checks of the DC wiring and the first low-pass filter stage.

The PCIe shorting connectors (Fig. B.3, left) tie all pins to a single node so that the entire interface remains at one potential. A dedicated pin routes this node through a series resistor (order megaohms) to a chassis or lab ground, providing a controlled discharge path. For sample transfer, best practice is to install the shorting connector on the top of The Beast while it is mated to the connector board, and to keep it fitted whenever it is transported or stored. For storage, a custom 3D-printed case (Fig. B.3, right) protects the bonded sample; the shorting connector should remain connected.

The breakout card (Fig. B.4) plugs into the connector board at the mK stage and exposes test points that map the PCIe pins to accessible pads. This tool is used to confirm the integrity of the DC lines from room temperature to the sample interface.

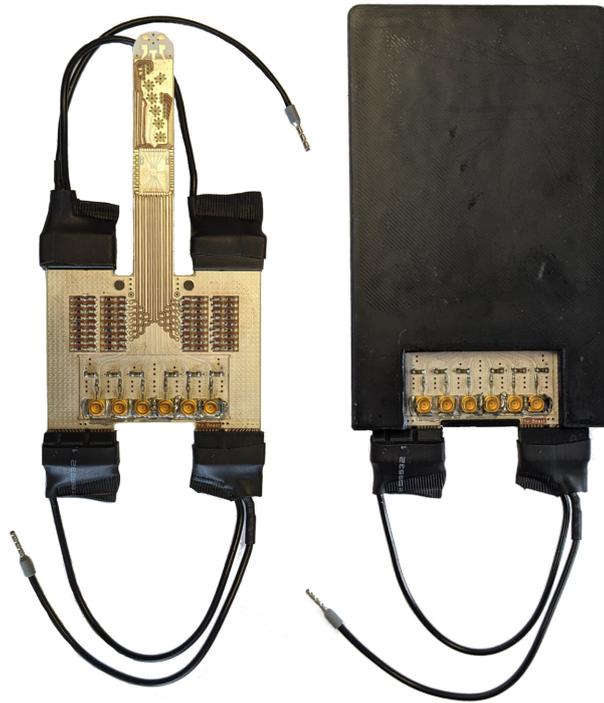


Figure B.3: Accessories for storage and ESD protection. **Left:** External PCIe shorting connectors: all pins are shorted together to a common potential; one additional pin provides a bleed path through a series megaohm-range resistor to safely discharge the assembly. **Right:** 3D -printed storage case for a wire-bonded sample; The Beast is fixed to the base with two screws.

B.3 Assembly & Soldering Reference

This section provides the practical information needed to assemble and solder The Beast PCB and its mK-stage connector board. The complete electrical network and filter topology are shown in Fig. B.5. Please observe the manufacturer’s recommended port labelling for the Mini-Circuits LFCN-80+ filters during placement (input/output orientation as per datasheet).

We implement two DC line classes:

- **“Fast” lines** (for ohmic/readout): two-stage RC low-pass with a -3 dB point near 59 kHz, plus an LFCN-80+ section to suppress RF in the 200–4500 MHz range.
- **“Slow” lines** (for gate biasing): two-stage RC low-pass with a -3 dB point near 222 Hz, followed by a resistive bias-tee for RF pulsing if needed. The bias-tee uses a 10 M Ω MELF resistor (DC feed) and a 15 nF 0603 capacitor (RF coupling), yielding a DC-path cutoff of ~ 1 Hz.

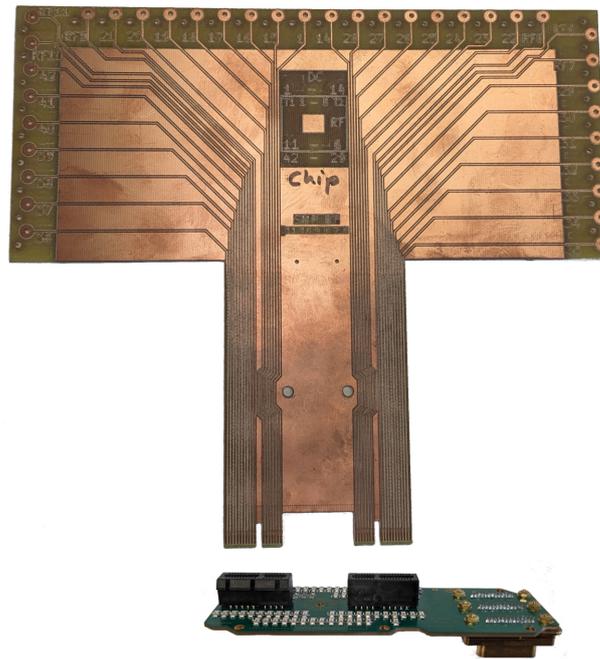


Figure B.4: Breakout PCB for the connector board. The copper-coloured breakout card (top) plugs into the green connector board (bottom) to verify end-to-end DC connectivity from the room-temperature harness to the PCIe connectors at the mK stage.

Pin maps and routing views

For soldering and inspection, the SMD pin numbering on The Beast is listed in Fig. B.6; the connector-board SMD pin map is given in Fig. B.7. Layer-by-layer routing overviews are provided in Fig. B.8-Fig. B.12.

Three-Dimensional Vector Magnet

To enable magnetic-field control in arbitrary directions at the sample, The Beast integrates a compact three-axis (3D) vector magnet. The field at the chip is the vector sum of three orthogonal coil contributions, with B_x aligned to the symmetry axis of the cold finger.

The axial (x) coil is an elliptically wound solenoid made from superconducting NbTi wire, chosen for high critical current and low loss at cryogenic temperatures. The transverse (y, z) axes are implemented as compact split pairs in a Helmholtz-like configuration to provide a uniform field in the central sample volume while keeping inductance and footprint small. The full assembly was modelled and optimised in COMSOL Multiphysics for field strength and homogeneity.

At the sample position, the design is intended to deliver up to 1.0 T along x and 0.3 T along each transverse axis (y, z). The coil set is thermally anchored to the elevated 4 K

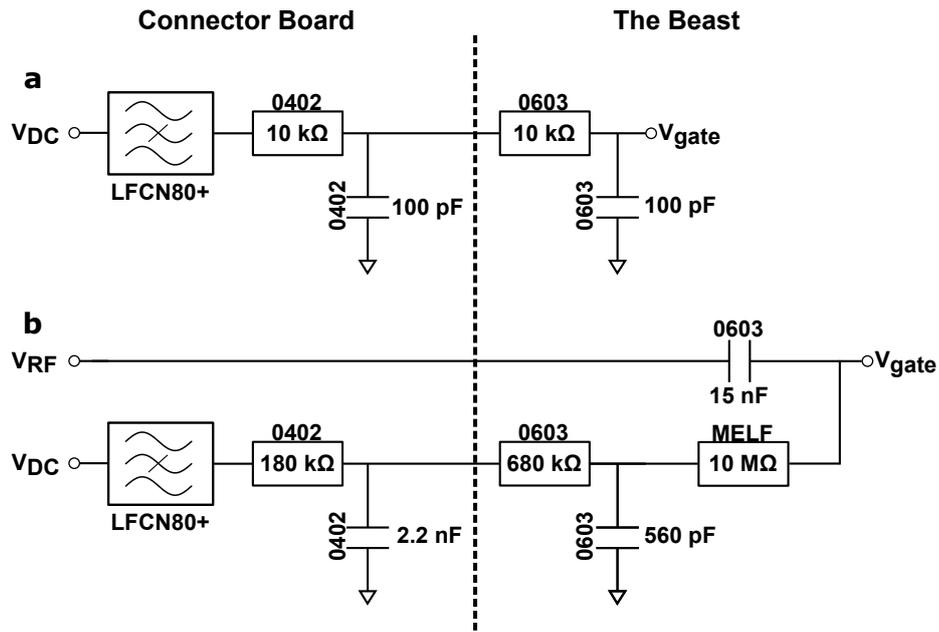


Figure B.5: Electrical circuit diagram of The Beast and the connector board. **a** Two-stage low-pass arrangement for ohmic/readout lines ($-3 \text{ dB} \approx 59 \text{ kHz}$) with an additional Mini-Circuits LFCN-80+ section providing strong suppression across 200-4500 MHz. The first RC stage is on the connector board; the second RC stage is on The Beast. **b** Two-stage low-pass for a slow gate line ($-3 \text{ dB} \approx 222 \text{ Hz}$) feeding a resistive bias-tee (DC feed 10 MΩ MELF, RF coupling 15 nF 0603), giving a DC-path cutoff of $\sim 1 \text{ Hz}$.

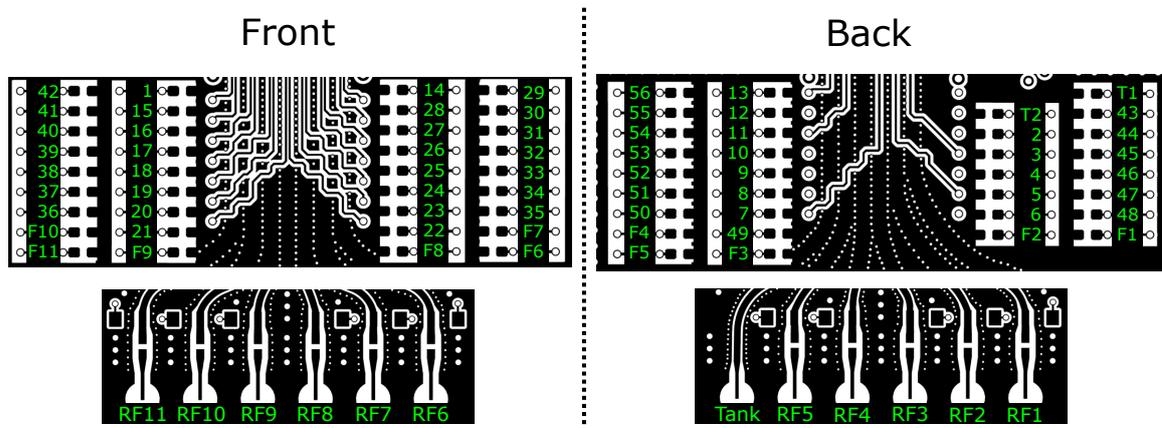


Figure B.6: SMD pin assignment on The Beast (front and back).

stage. At the time of writing, fabrication of the vector magnet was still in progress, and the quoted performance figures refer to the design specifications.

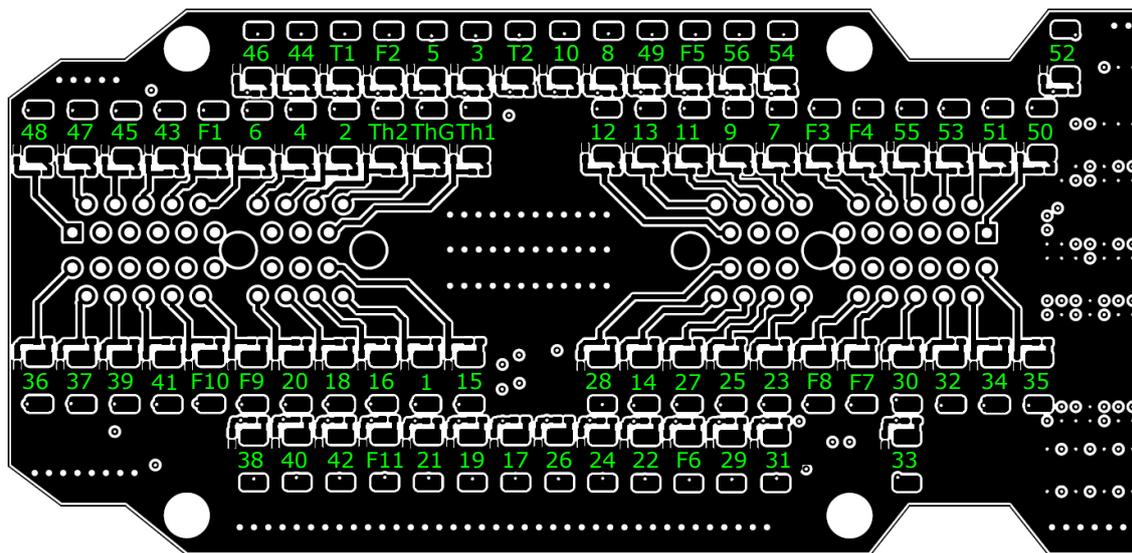
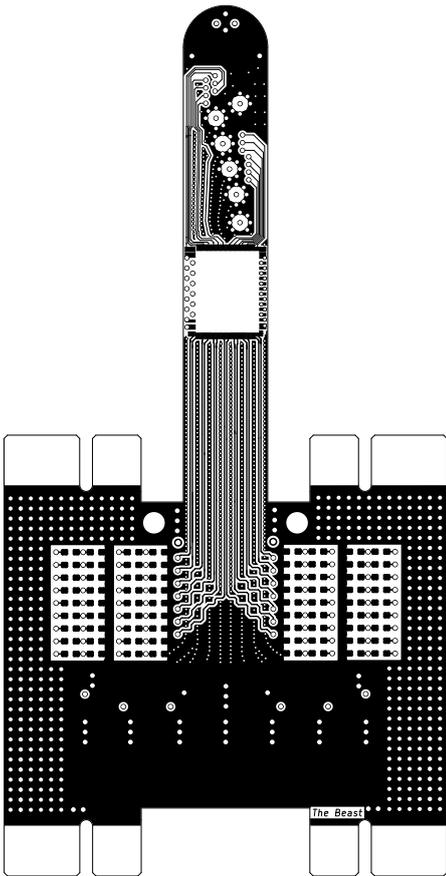
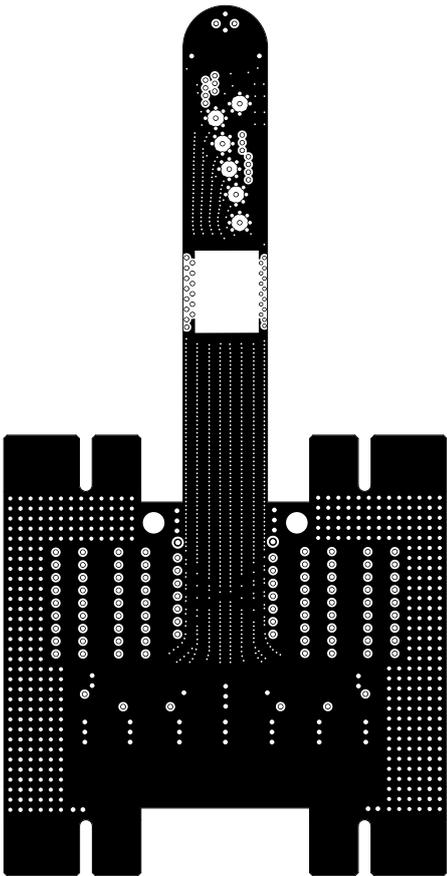


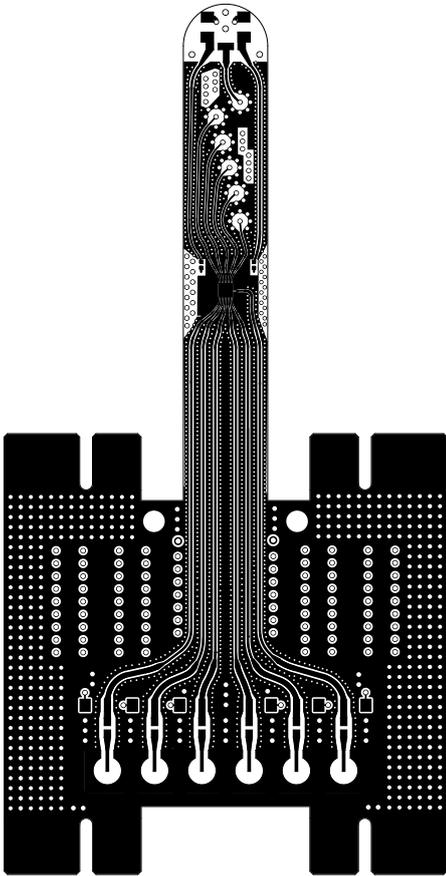
Figure B.7: SMD pin assignment on the connector board. The LFCN-80+ is placed above the numeric labels; the signal path runs toward the board centre where the PCIe connectors are located. F1-F11 and T1, T2 are attributed to the bias-tees of the RF lines and tank circuits. ThG, Th1, Th2 are attributed to the thermometry.



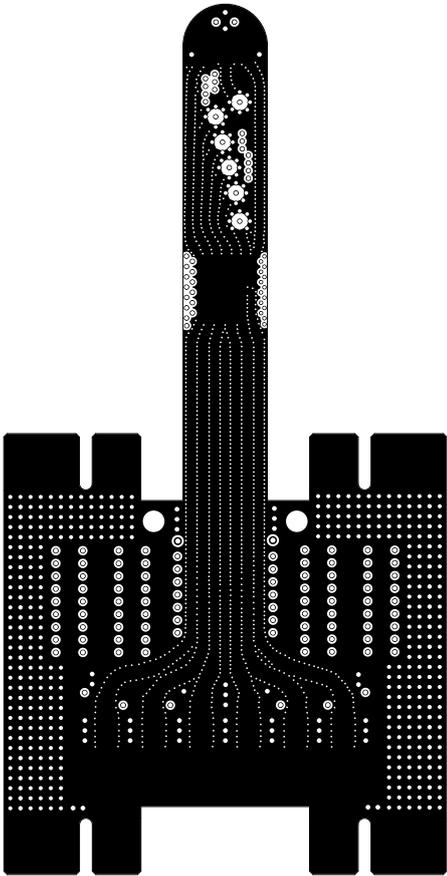
(a) Front layer



(b) Inner layer 1

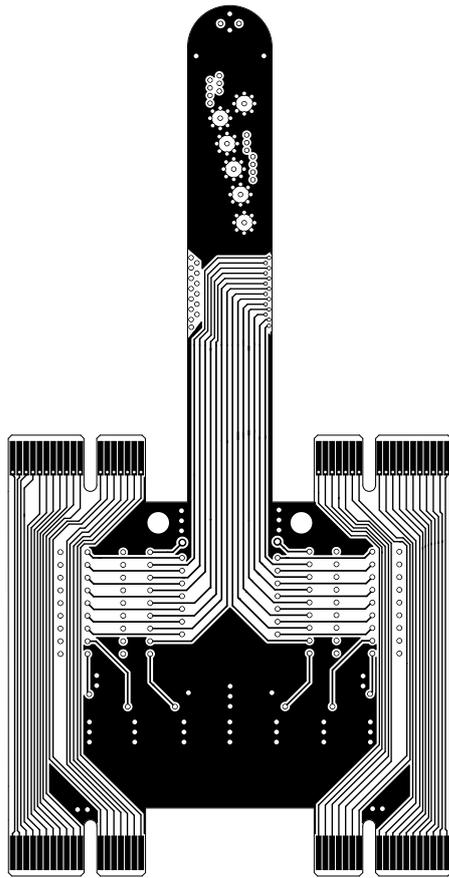


(c) Inner layer 2

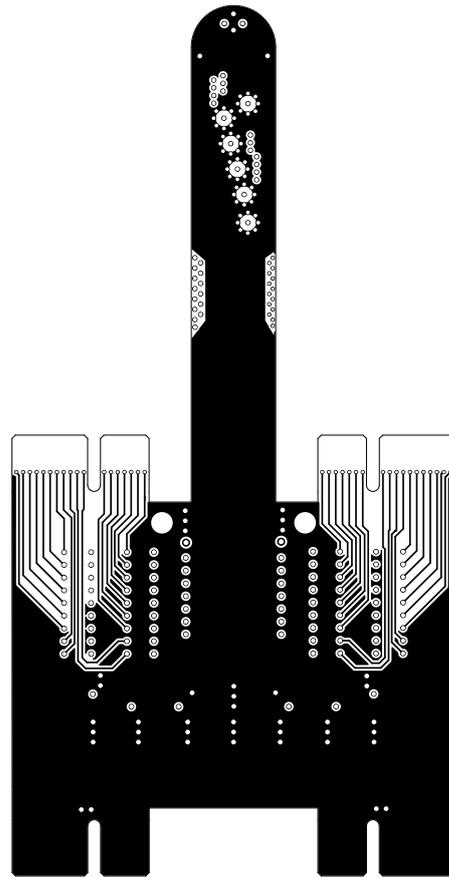


(d) Inner layer 3

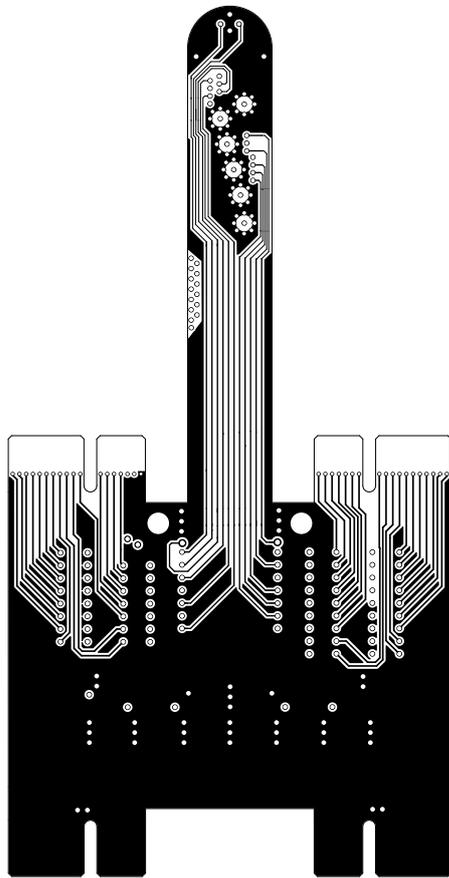
Figure B.8: Copper layers of The Beast.



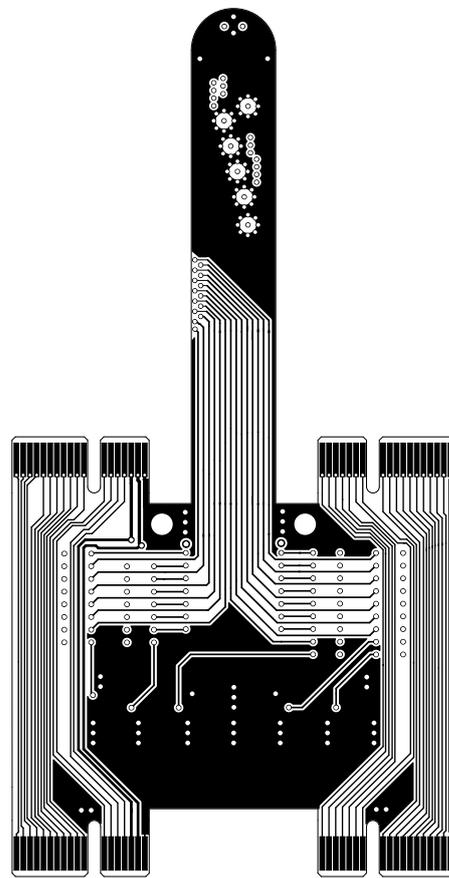
(e) Inner layer 4



(f) Inner layer 5

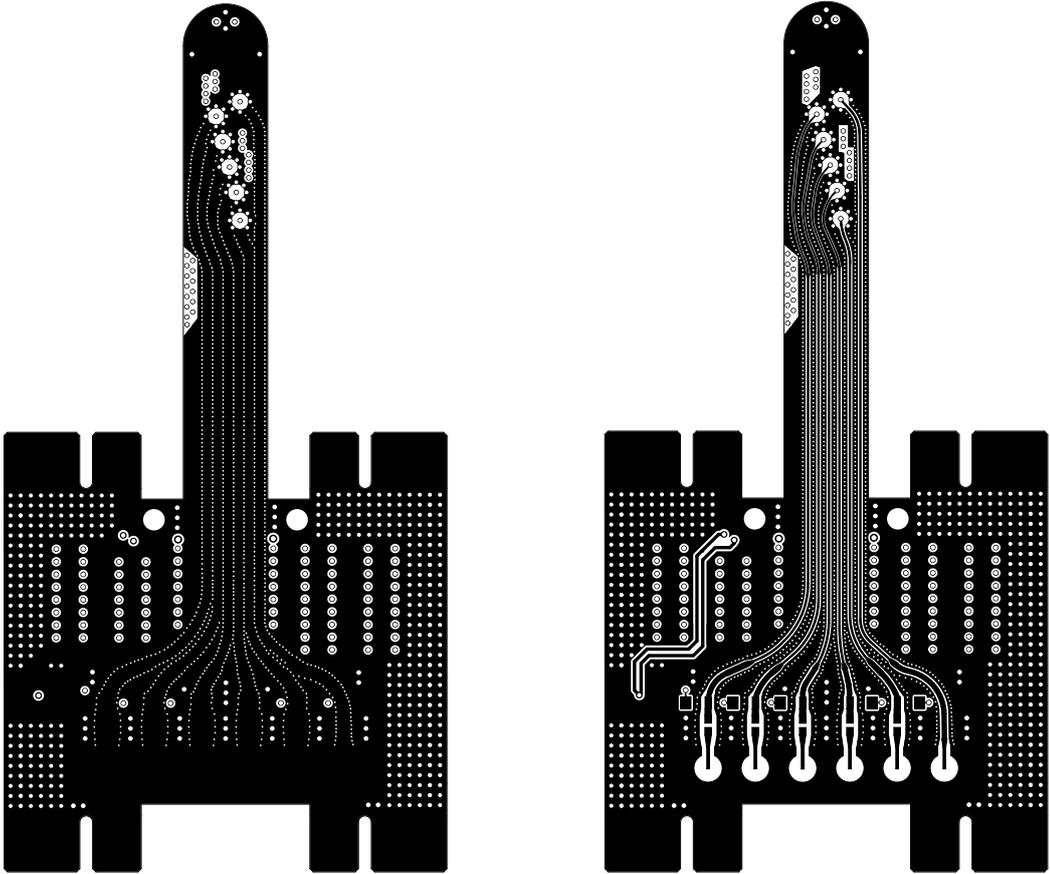


(g) Inner layer 6



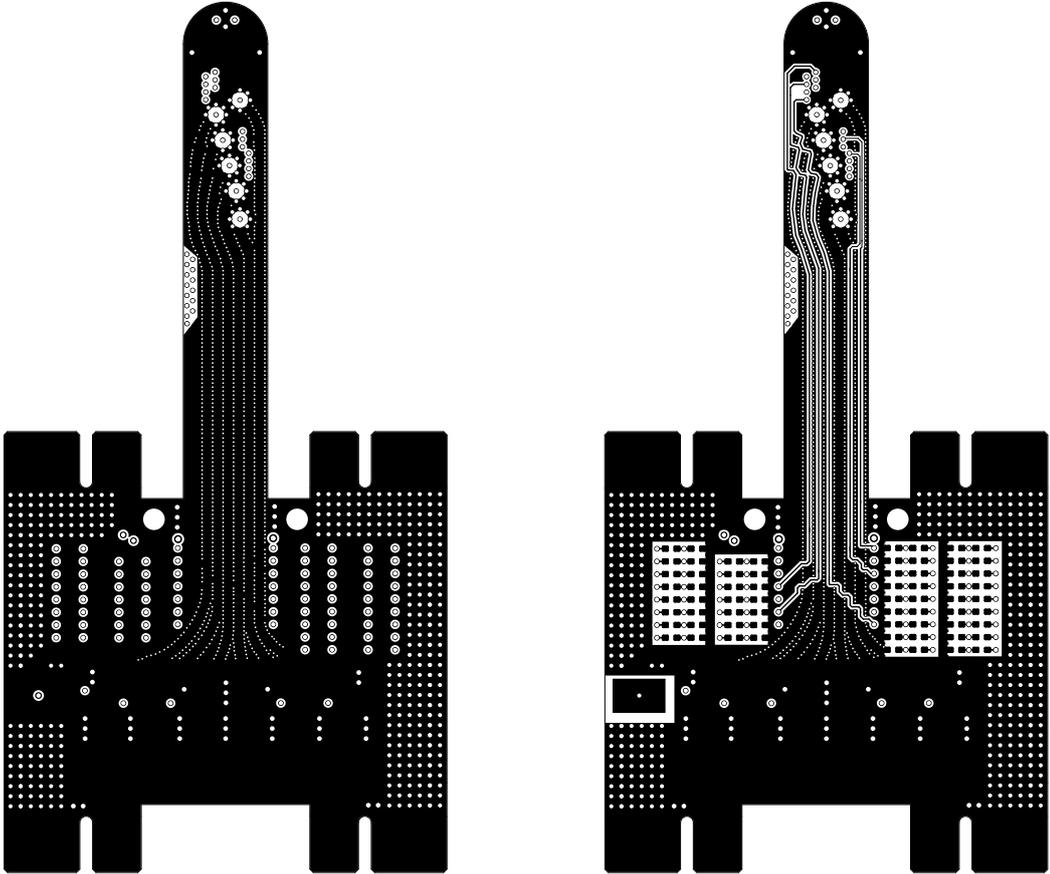
(h) Inner layer 7

Figure B.9: Copper layers of The Beast.



(i) Inner layer 8

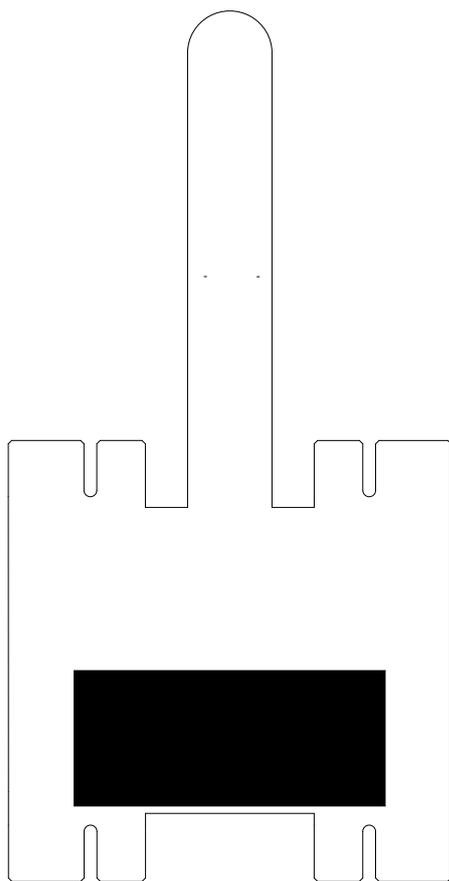
(j) Inner layer 9



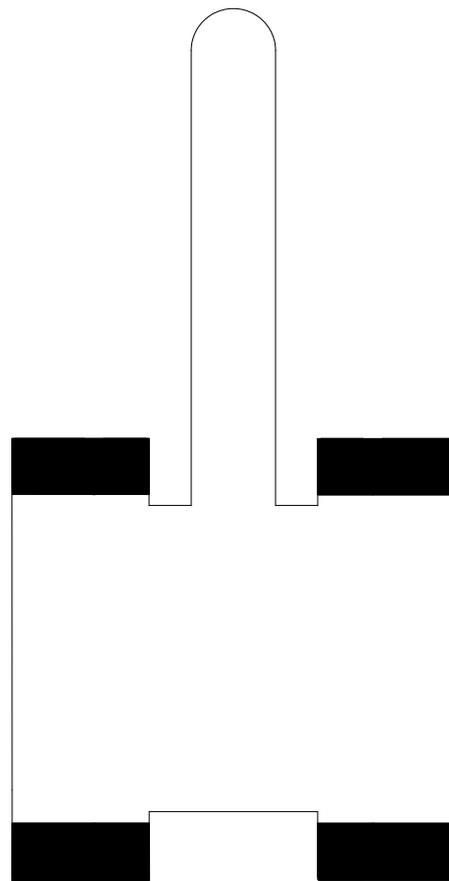
(k) Inner layer 10

(l) Back layer

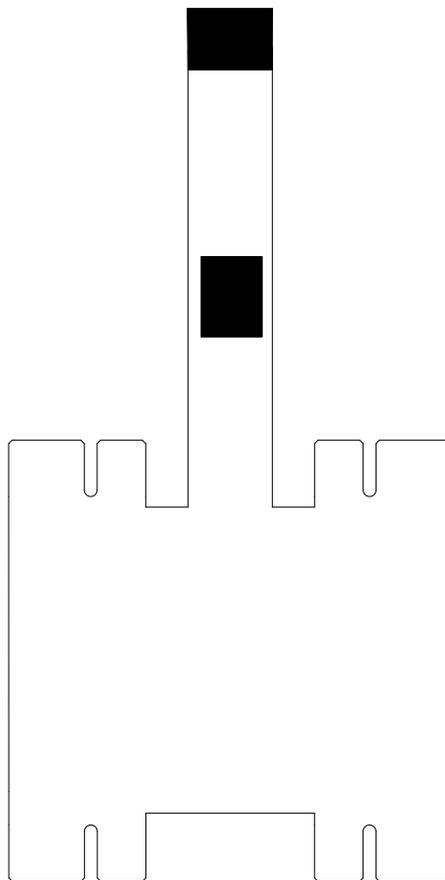
Figure B.10: Copper layers of The Beast.



(m) Cut-outs: front-In1; In10-back



(n) Cut-outs: front-In3; In8-back



(o) Cut-outs: front-In1

Figure B.11: Copper layers of The Beast.

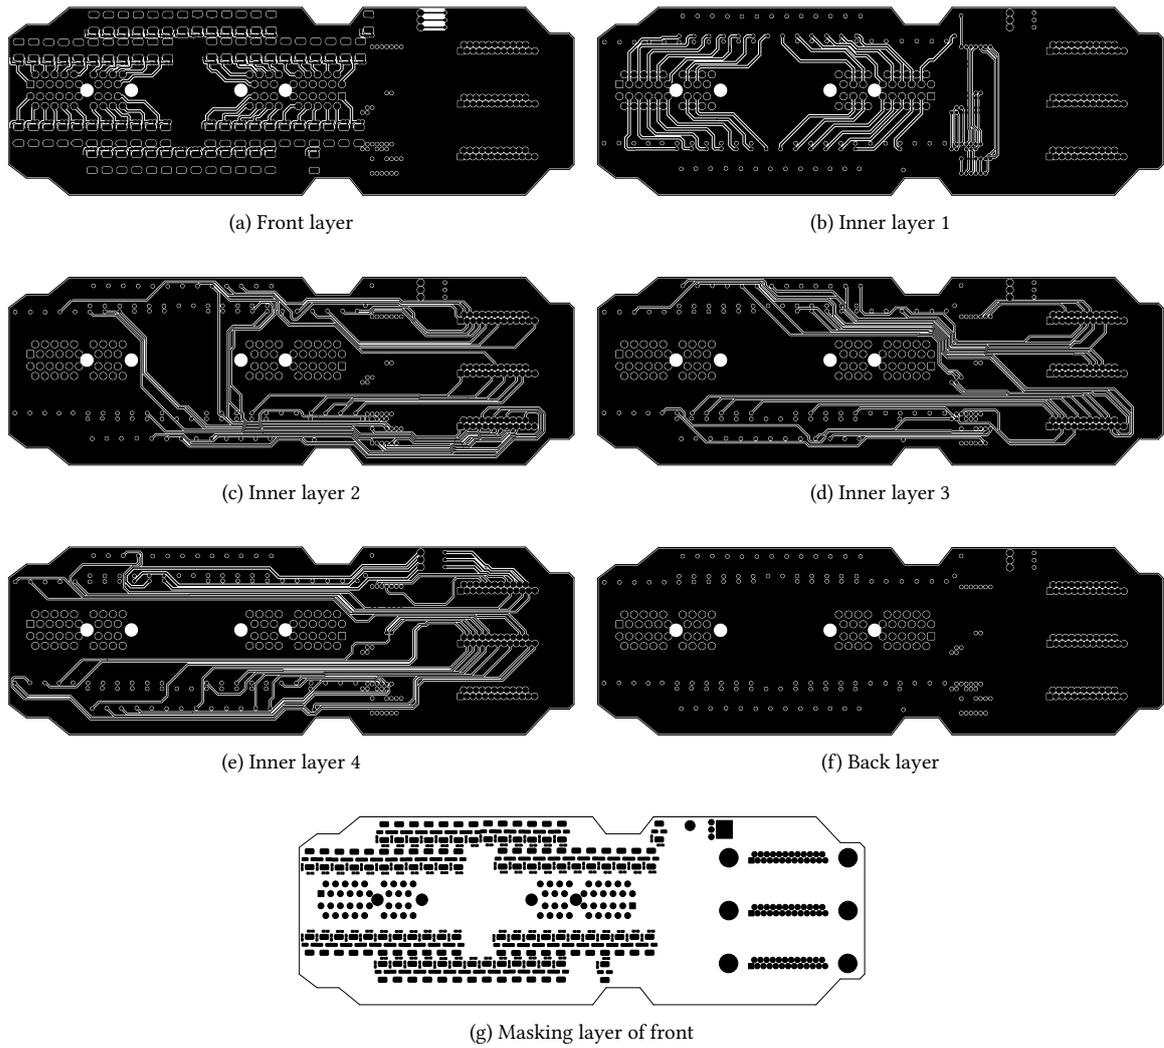


Figure B.12: Copper layers of the connector board.

C The SiMOS Spin Qubit

C.1 Fabrication Details of Device A and B.

Table C.1 lists the vertical gate stack (along the z-axis, top to substrate) for devices A and B.

Table C.1: Gate stack along the z-axis from the top surface down to the substrate. Layer heights indicated for device A and B. The microwave antenna is encapsulated in the passivation layer.

Description	Material	Device A	Device B
Passivation	SiO _x	250 nm	250 nm
Layer 3	poly Si	30 nm	30 nm
Insulator	SiO _x	5 nm	5 nm
Layer 2	poly Si	30 nm	30 nm
Insulator	SiO _x	5 nm	5 nm
Layer 1	poly Si	30 nm	30 nm
Insulator	SiO _x	8 nm	20 nm
Epitaxial	²⁸ Si	100 nm	100 nm
Wafer	nat. Si	775 μm	775 μm

C.2 Device A Charge Noise Spectroscopy

The analysis follows Sec. 4.2.3. First, the gate lever arm of the SET top gate was extracted from Coulomb-diamond spectroscopy by measuring the SET current as a function of V_{SD} and V_{ST} and fitting the diamond edges. Using the diamond height and width (Eq. (4.4)), the average lever arm for device A was found to be

$$\alpha = (5.6 \pm 0.1)\%$$

as shown in Fig. C.1.

With α determined, charge-noise spectroscopy was performed by parking the sensor on the flank of a Coulomb peak to obtain the local transconductance dI_{SET}/dV_{ST} , recording 100 s current time traces at $f_s = 100$ kHz, and estimating the power spectral density using Welch's method with a Hann window. The amplitude spectral density referred to the ST gate was then computed via Eq. (4.5). The procedure was repeated three times, yielding

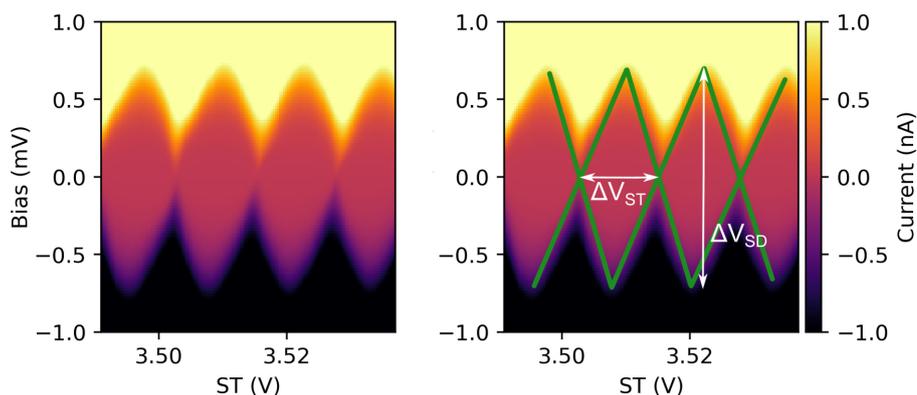


Figure C.1: Coulomb diamonds of the SET. **Left** SET current through the island formed under ST as a function of V_{SD} and V_{ST} . **Right** Same data with linear guides (green) along the diamond edges. The diamond height V_{SD} and width V_{ST} are indicated by white arrows. The lever arm α averaged over both diamonds leads to $\alpha = (5.6 \pm 0.1)\%$.

Table C.2: Charge noise spectroscopy of device A.

Measurement	Spectral noise density at 1 Hz	β
1	$(0.29 \pm 0.02) \mu\text{eV}\sqrt{\text{Hz}}$	-0.43 ± 0.04
2	$(0.35 \pm 0.02) \mu\text{eV}\sqrt{\text{Hz}}$	-0.45 ± 0.04
3	$(0.27 \pm 0.02) \mu\text{eV}\sqrt{\text{Hz}}$	-0.43 ± 0.04
average	$(0.30 \pm 0.02) \mu\text{eV}\sqrt{\text{Hz}}$	-0.44 ± 0.01

the 1 Hz noise amplitudes and power-law exponents listed in Tab. C.2. Averaging the three runs gives a 1 Hz spectral density of

$$(0.30 \pm 0.02) \mu\text{eV}/\sqrt{\text{Hz}}$$

and a frequency exponent $\beta = -0.44 \pm 0.01$ over the 0.1 Hz to 10 Hz fit band. A charge-insensitive working point was used to verify the instrument noise floor, as illustrated in Fig. C.2 d. An elevation of the setup noise is visible near $f = 1$ Hz.

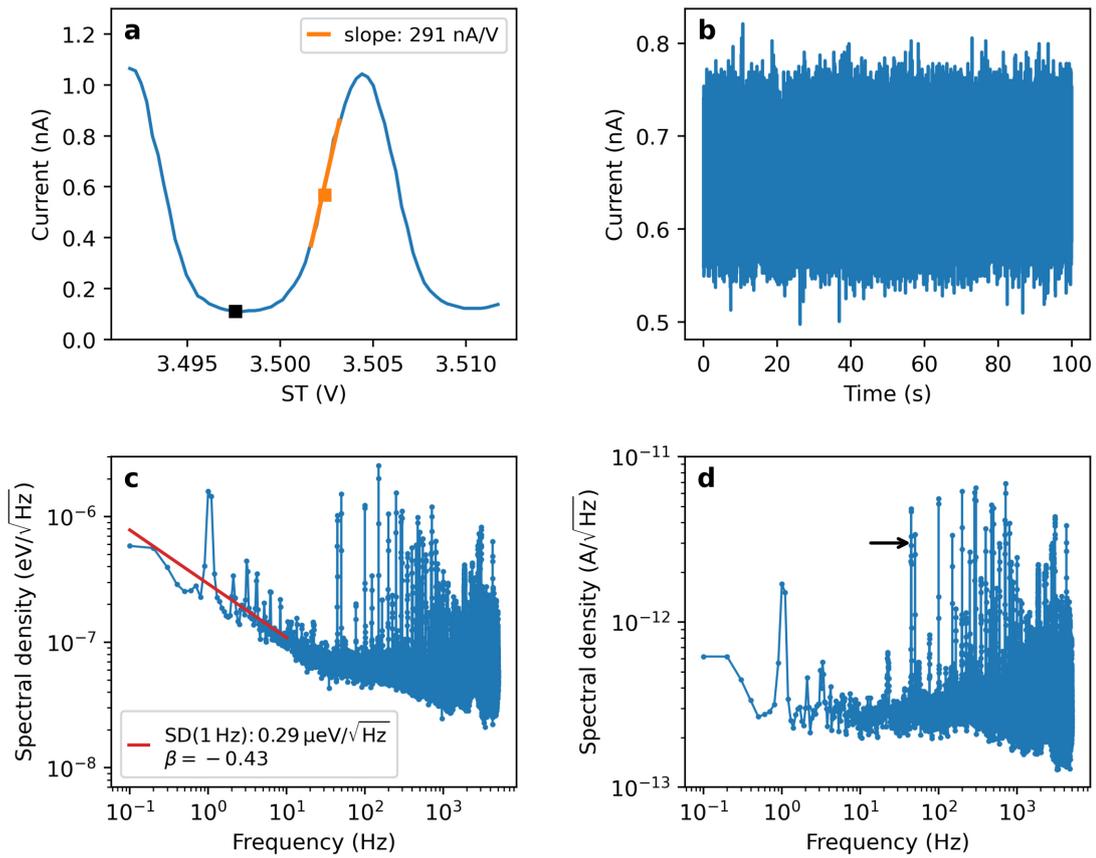


Figure C.2: Spectral noise density (SD) of the SET of device A. **a** Coulomb oscillations of the SET with two measurement positions indicated. The linear fit on the left peak flank (orange) sets $dI_{\text{SET}}/dV_{\text{ST}}$ used in **c**. The black dot marks a charge-noise-insensitive point used to estimate the setup noise floor in **d**. **b** Time trace of the SET current sampled at 100 kHz. **c** Spectral noise density of the trace in **b**, averaged with Welch's method (Hann window). The fit is performed over 0.1 Hz to 10 Hz. **d** Noise floor recorded at the charge-noise-insensitive point (black dot in **a**). The mains spur at 50 Hz is indicated (arrow). The IV-converter bandwidth of ~ 24 kHz appears as a roll-off of the white-noise floor above that frequency.

C.3 Device Working Voltages

Table C.3 lists representative DC gate biases used in this thesis. Device A was operated in single-dot mode, whereas device B was configured as a double dot. Absolute values vary slightly between cooldowns and tuning sessions.

Owing to the reduced SiO_x thickness in device B, gate-dot capacitances are larger and vertical confinement is stronger, which manifests as a smaller voltage offset required between the SET top gate (ST) and its barriers (BL/BR).

Table C.3: Gate voltages at the working points used for devices A and B.

Gate	Device A (V)	Device B (V)
ST	3.50	2.36
BL	0.22	-0.01
BR	0.16	0.11
C	-0.20	0.00
P1	0.78	1.17
B1	0.34	0.90
P2	3.00	0.69
B2	3.00	0.80
R	3.00	3.00

D Calculation Tools

D.1 Impedance Determination of a CPW Sandwich

This appendix provides a compact description on how to calculate the characteristic impedance Z_0 and the effective relative permittivity ϵ_{eff} of a coplanar waveguide (CPW) sandwiched between two ground planes (one above, one below). The metal thickness of the CPW is taken as $t \rightarrow 0$; in practice, a finite thickness slightly reduces Z_0 .

The calculation follows Ref. [250, pp. 16 ff.] and uses complete elliptic integrals as in Ref. [251]. The geometry and symbols are:

- Upper ground | dielectric (ϵ_2, h_2) | CPW (s, w) | dielectric (ϵ_1, h_1) | lower ground.
- s : centre conductor width; w : gap to ground (each side).
- h_1, ϵ_1 : thickness and relative permittivity of the lower dielectric.
- h_2, ϵ_2 : thickness and relative permittivity of the upper dielectric.

The (quasistatic) elliptic moduli for the two halfspaces are

$$k_1 = \frac{\tanh\left(\frac{\pi s}{4h_1}\right)}{\tanh\left(\frac{\pi(s+2w)}{4h_1}\right)}, \quad k_2 = \frac{\tanh\left(\frac{\pi s}{4h_2}\right)}{\tanh\left(\frac{\pi(s+2w)}{4h_2}\right)}, \quad k'_i = \sqrt{1 - k_i^2} \quad (i = 1, 2).$$

With $K(\cdot)$ the complete elliptic integral of the first kind (modulus notation), the perunit-length capacitances are

$$C' = 2\epsilon_0\epsilon_1 \frac{K(k_1)}{K(k'_1)} + 2\epsilon_0\epsilon_2 \frac{K(k_2)}{K(k'_2)}, \quad C'_{\text{air}} = 2\epsilon_0 \frac{K(k_1)}{K(k'_1)} + 2\epsilon_0 \frac{K(k_2)}{K(k'_2)}.$$

Hence

$$\epsilon_{\text{eff}} = \frac{C'}{C'_{\text{air}}}, \quad Z_0 = \frac{\sqrt{\epsilon_{\text{eff}}}}{c C'}.$$

Note on software conventions: `scipy.special.ellipk(m)` expects the parameter $m = k^2$, not the modulus k . This has to be wrapped correctly.

E Semiconductor Device Modelling with nextnano

The gate voltages used in the nextnano simulations of Ch. 7 are summarised in Tab. E.1.

Table E.1: Gate voltages used in the simulations. Values were derived from experimentally tuned devices and lightly adjusted for the numerical model.

Gate	Voltage (V)
ST	3.50
BL	0.18
BR	0.18
C	-0.20
P1	0.70
B1	0.47
P2	3.00
B2	3.00
R	3.00

List of Publications

- **Daniel Schroller**, Daniel Sitter, Thomas Koch, Viktor Adam, Noah Glaeser, Clément Godfrin, Stefan Kubicek, Julien Jussot, Roger Loo, Yosuke Shimura, Danny Wan, Yaorong Chen, Mario Ruben, Kristiaan De Greve, and Wolfgang Wernsdorfer. *Hybrid Quantum Systems: Coupling Single-Molecule Magnet Qudits to Industrial Silicon Spin Qubits*. Submitted to *Physical Review Applied*, October 2025; arXiv:2510.10110.
- Thomas Koch, Clément Godfrin, Viktor Adam, Julian Ferrero, **Daniel Schroller**, Noah Glaeser, Stefan Kubicek, Ruoyu Li, Roger Loo, Shana Massar, George Simion, Danny Wan, Kristiaan De Greve, and Wolfgang Wernsdorfer. *Industrial 300 mm-wafer-processed spin qubits in natural silicon/silicon–germanium*. *npj Quantum Information* 11, 59 (2025). doi:10.1038/s41534-025-01016-x.
- Julian Ferrero, Thomas Koch, Sonja Vogel, **Daniel Schroller**, Viktor Adam, Ran Xue, Inga Seidler, Lars R. Schreiber, Hendrik Bluhm, and Wolfgang Wernsdorfer. *Noise Reduction by Bias Cooling in Gated Si/Si_xGe_{1-x} Quantum Dots*. *Applied Physics Letters* 124, 204002 (2024). doi:10.1063/5.0206632.

Bibliography

We note that large language models were employed during the preparation of this manuscript for AI-assisted copyediting, including for example improvements to readability, grammar, and style. Their use was limited to refining human-written text generated by the author and did not involve autonomous content generation.

- [1] F. Arute et al. “Quantum supremacy using a programmable superconducting processor”. In: *Nature* 574.7779 (Oct. 2019), pp. 505–510. DOI: 10.1038/s41586-019-1666-5.
- [2] T. Tanttu et al. “Assessment of the errors of high-fidelity two-qubit gates in silicon quantum dots”. In: *Nature Physics* 20.11 (Nov. 2024), pp. 1804–1809. DOI: 10.1038/s41567-024-02614-w.
- [3] G. D. Fuchs et al. “Excited-state spin coherence of a single nitrogen–vacancy centre in diamond”. In: *Nature Physics* 6 (2010), pp. 668–672. DOI: 10.1038/nphys1716.
- [4] J. P. Dowling and G. J. Milburn. “Quantum technology: the second quantum revolution”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 361 (2003), pp. 1655–1674. DOI: 10.1098/rsta.2003.1227.
- [5] J. Preskill. “Quantum Computing in the NISQ era and beyond”. In: *Quantum* 2 (2018), p. 79. DOI: 10.22331/q-2018-08-06-79.
- [6] N. Gisin et al. “Quantum cryptography”. In: *Reviews of Modern Physics* 74.1 (2002), pp. 145–195. DOI: 10.1103/RevModPhys.74.145.
- [7] V. Giovannetti, S. Lloyd, and L. Maccone. “Quantum-Enhanced Measurements: Beating the Standard Quantum Limit”. In: *Science* 306.5700 (2004), pp. 1330–1336. DOI: 10.1126/science.1104149.
- [8] D. P. DiVincenzo. “The Physical Implementation of Quantum Computation”. In: *Fortschritte der Physik* 48.9-11 (2000), pp. 771–783. DOI: 10.1002/1521-3978(200009)48:9/11<771::AID-PROP771>3.0.CO;2-E.
- [9] Y. Wang et al. “Qudits and High-Dimensional Quantum Computing”. In: *Frontiers in Physics* 8 (2020), p. 589504. DOI: 10.3389/fphy.2020.589504.
- [10] A. G. Fowler et al. “Surface codes: Towards practical large-scale quantum computation”. In: *Physical Review A* 86.3 (2012), p. 032324. DOI: 10.1103/PhysRevA.86.032324.
- [11] E. T. Campbell, B. M. Terhal, and C. Vuillot. “Roads towards fault-tolerant universal quantum computation”. In: *Nature* 549.7671 (2017), pp. 172–179. DOI: 10.1038/nature23460.

- [12] C. Monroe and J. Kim. “Scaling the ion trap quantum processor”. In: *Science* 339.6124 (2013), pp. 1164–1169. DOI: 10.1126/science.1231298.
- [13] J. Benhelm et al. “Towards fault-tolerant quantum computing with trapped ions”. In: *Nature Physics* 4.6 (2008), pp. 463–466. DOI: 10.1038/nphys961.
- [14] P. Kok et al. “Linear optical quantum computing with photonic qubits”. In: *Reviews of Modern Physics* 79.1 (2007), pp. 135–174. DOI: 10.1103/RevModPhys.79.135.
- [15] E. Knill, R. Laflamme, and G. J. Milburn. “A scheme for efficient quantum computation with linear optics”. In: *Nature* 409.6816 (2001), pp. 46–52. DOI: 10.1038/35051009.
- [16] R. Barends et al. “Superconducting quantum circuits at the surface code threshold for fault tolerance”. In: *Nature* 508.7497 (2014), pp. 500–503. DOI: 10.1038/nature13171.
- [17] J. R. Petta et al. “Coherent Manipulation of Coupled Electron Spins in Semiconductor Quantum Dots”. In: *Science* 309.5744 (Sept. 2005), pp. 2180–2184. DOI: 10.1126/science.1116955.
- [18] C. Godfrin et al. “Operating Quantum States in Single Magnetic Molecules: Implementation of Grover’s Quantum Algorithm”. In: *Physical Review Letters* 119.18 (2017), p. 187702. DOI: 10.1103/PhysRevLett.119.187702.
- [19] E. Moreno-Pineda and W. Wernsdorfer. “Magnetic Molecules as Building Blocks for Quantum Technologies”. In: *Advanced Quantum Technologies* 7.24 (2024), p. 2300367. DOI: 10.1002/qute.202300367.
- [20] C. L. Degen, F. Reinhard, and P. Cappellaro. “Quantum sensing”. In: *Reviews of Modern Physics* 89.3 (2017), p. 035002. DOI: 10.1103/RevModPhys.89.035002.
- [21] M. J. Biercuk et al. “Ultrasensitive detection of force and displacement using trapped ions”. In: *Nature Nanotechnology* 5.9 (2010), pp. 646–650. DOI: 10.1038/nnano.2010.165.
- [22] P. Neumann et al. “High-Precision Nanoscale Temperature Sensing Using Single Defects in Diamond”. In: *Nano Letters* 13.6 (2013), pp. 2738–2742. DOI: 10.1021/nl401216y.
- [23] V. Giovannetti, S. Lloyd, and L. Maccone. “Advances in quantum metrology”. In: *Nature Photonics* 5.4 (2011), pp. 222–229. DOI: 10.1038/nphoton.2011.35.
- [24] J. M. Taylor et al. “High-sensitivity diamond magnetometer with nanoscale resolution”. In: *Nature Physics* 4.10 (2008), pp. 810–816. DOI: 10.1038/nphys1075.
- [25] L. Rondin et al. “Magnetometry with nitrogen-vacancy defects in diamond”. In: *Reports on Progress in Physics* 77.5 (2014), p. 056503. DOI: 10.1088/0034-4885/77/5/056503.
- [26] R. Kleiner et al. “Superconducting Quantum Interference Devices: State of the Art and Applications”. In: *Proceedings of the IEEE* 92.10 (2004), pp. 1534–1548. DOI: 10.1109/JPROC.2004.833655.

- [27] M. Hämmäläinen et al. “Magnetoencephalography theory, instrumentation, and applications to noninvasive studies of the working human brain”. In: *Reviews of Modern Physics* 65.2 (1993), pp. 413–497. DOI: 10.1103/RevModPhys.65.413.
- [28] D. Budker and M. Romalis. “Optical magnetometry”. In: *Nature Physics* 3.4 (2007), pp. 227–234. DOI: 10.1038/nphys566.
- [29] J. C. Allred et al. “High-Sensitivity Atomic Magnetometer Unaffected by Spin-Exchange Relaxation”. In: *Physical Review Letters* 89.13 (2002), p. 130801. DOI: 10.1103/PhysRevLett.89.130801.
- [30] S. M. Sze and K. K. Ng. *Physics of Semiconductor Devices*. 3rd ed. Hoboken, NJ, 2007. DOI: 10.1002/0470068329.
- [31] K. J. Kuhn. “Considerations for Ultimate CMOS Scaling”. In: *IEEE Transactions on Electron Devices* 59.7 (2012), pp. 1813–1828. DOI: 10.1109/TED.2012.2193129.
- [32] M. F. Gonzalez-Zalba et al. “Scaling silicon-based quantum computing using CMOS technology”. In: *Nature Electronics* 4 (2021), pp. 872–884. DOI: 10.1038/s41928-021-00681-y.
- [33] F. A. Zwanenburg et al. “Silicon quantum electronics”. In: *Reviews of Modern Physics* 85.3 (July 2013), pp. 961–1019. DOI: 10.1103/RevModPhys.85.961.
- [34] imec. *imec — the world’s leading independent nanoelectronics R&D hub*. <https://www.imec-int.com/en>. Company website. Leuven, Belgium, 2025.
- [35] B. Patra et al. “Cryo-CMOS Circuits and Systems for Quantum Computing Applications”. In: *IEEE Journal of Solid-State Circuits* 53.1 (2018), pp. 309–321. DOI: 10.1109/JSSC.2017.2737549.
- [36] F. Sebastiano et al. “Cryo-CMOS Interfaces for Large-Scale Quantum Computers”. In: *2020 IEEE International Electron Devices Meeting (IEDM)*. 2020, pp. 25.2.1–25.2.4. DOI: 10.1109/IEDM13553.2020.9372075.
- [37] F. Borsoi et al. “Shared control of a 16 semiconductor quantum dot crossbar array”. In: *Nature Nanotechnology* 19.1 (Jan. 2024), pp. 21–27. DOI: 10.1038/s41565-023-01491-3.
- [38] W. H. Lim et al. “A 2×2 Quantum Dot Array in Silicon with Fully Tunable Pairwise Interdot Coupling”. In: *Nano Letters* 25.26 (July 2025), pp. 10263–10269. DOI: 10.1021/acs.nanolett.4c06264.
- [39] J. Yoneda et al. “A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%”. In: *Nature Nanotechnology* 13.2 (Feb. 2018), pp. 102–106. DOI: 10.1038/s41565-017-0014-x.
- [40] D. M. Fleetwood et al. “Effects of oxide traps, interface traps, and “border traps” on metal–oxide–semiconductor devices”. In: *Journal of Applied Physics* 73.10 (May 1993), pp. 5058–5074. DOI: 10.1063/1.353777.
- [41] C. H. Yang et al. “Spin-valley lifetimes in a silicon quantum dot with tunable valley splitting”. In: *Nature Communications* 4 (June 2013), p. 2069. DOI: 10.1038/ncomms3069.

- [42] J. K. Gamble et al. “Valley splitting of single-electron Si MOS quantum dots”. In: *Applied Physics Letters* 109.25 (2016), p. 253101. DOI: 10.1063/1.4972514.
- [43] L. Bogani and W. Wernsdorfer. “Molecular spintronics using single-molecule magnets”. In: *Nature Materials* 7.3 (2008), pp. 179–186. DOI: 10.1038/nmat2133.
- [44] N. Ishikawa et al. “Lanthanide double-decker complexes functioning as magnets at the single-molecular level”. In: *Journal of the American Chemical Society* 125.29 (2003), pp. 8694–8695. DOI: 10.1021/ja029629n.
- [45] S. Thiele et al. “Electrical readout of individual nuclear spin trajectories in a single-molecule magnet spin transistor”. In: *Physical Review Letters* 111.3 (2013), p. 037203. DOI: 10.1103/PhysRevLett.111.037203.
- [46] S. Thiele et al. “Electrically driven nuclear spin resonance in single-molecule magnets”. In: *Science* 344.6188 (2014), pp. 1135–1138. DOI: 10.1126/science.1249802.
- [47] C. Godfrin et al. “Generalized Ramsey interferometry explored with a single nuclear spin qubit”. In: *npj Quantum Information* 4 (2018), p. 53. DOI: 10.1038/s41534-018-0101-3.
- [48] R. Vincent et al. “Electronic read-out of a single nuclear spin using a molecular spin transistor”. In: *Nature* 488.7411 (2012), pp. 357–360. DOI: 10.1038/nature11341.
- [49] M. Ganzhorn et al. “Strong spin–phonon coupling between a single-molecule magnet and a carbon nanotube nanoelectromechanical system”. In: *Nature Nanotechnology* 8.3 (2013), pp. 165–169. DOI: 10.1038/nnano.2012.258.
- [50] M. Urdampilleta et al. “Supramolecular spin valves”. In: *Nature Materials* 10.7 (2011), pp. 502–506. DOI: 10.1038/nmat3050.
- [51] G. LaRocca. *2025 State of the Industry Report: Investment and Innovation Amidst Global Challenges and Opportunities*. Semiconductor Industry Association (SIA) Blog. Blog post, 10 July 2025. July 2025. URL: <https://www.semiconductors.org/2025-state-of-the-industry-report-investment-and-innovation-amidst-global-challenges-and-opportunities/>.
- [52] B. Martinez and Y. Niquet. “Variability of Electron and Hole Spin Qubits Due to Interface Roughness and Charge Traps”. In: *Phys. Rev. Applied* 17.2 (Feb. 2022), p. 024022. DOI: 10.1103/PhysRevApplied.17.024022.
- [53] J. D. Cifuentes et al. “Impact of electrostatic crosstalk on spin qubits in dense CMOS quantum dot arrays”. In: *Physical Review B* 110 (Sept. 2024), p. 125414. DOI: 10.1103/PhysRevB.110.125414.
- [54] G. Burkard et al. “Semiconductor spin qubits”. In: *Reviews of Modern Physics* 95.2 (June 2023), p. 025003. DOI: 10.1103/RevModPhys.95.025003.
- [55] D. Loss and D. P. DiVincenzo. “Quantum computation with quantum dots”. In: *Phys. Rev. A* 57.1 (Jan. 1998), pp. 120–126. DOI: 10.1103/PhysRevA.57.120.
- [56] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. 10th Anniversary Edition. Cambridge: Cambridge University Press, 2010. DOI: 10.1017/CB09780511976667.

- [57] W. Pauli. “Über den Zusammenhang des Abschlusses der Elektronengruppen im Atom mit der Komplexstruktur der Spektren”. In: *Zeitschrift für Physik* 31 (1925), pp. 765–783. DOI: 10.1007/BF02980631.
- [58] P. Steinacker et al. “Industry-compatible silicon spin-qubit unit cells exceeding 99% fidelity”. In: *Nature* 646.8083 (Oct. 2025), pp. 81–87. DOI: 10.1038/s41586-025-09531-9.
- [59] S. M. Sze and K. K. Ng. *Physics of Semiconductor Devices*. 3rd ed. Wiley, 2006.
- [60] K. M. Itoh and H. Watanabe. “Isotope engineering of silicon and diamond for quantum computing and sensing applications”. In: *MRS Communications* 4.4 (2014), pp. 143–157. DOI: 10.1557/mrc.2014.32.
- [61] A. M. Tyryshkin et al. “Electron spin coherence exceeding seconds in high-purity silicon”. In: *Nature Materials* 11.2 (2012), pp. 143–147. DOI: 10.1038/nmat3182.
- [62] A. M. J. Zwerver et al. “Qubits made by advanced semiconductor manufacturing”. In: *Nature Electronics* 5.3 (2022), pp. 184–190. DOI: 10.1038/s41928-022-00727-9.
- [63] C. Kittel. *Introduction to Solid State Physics*. 8th ed. Hoboken, NJ: Wiley, 2005.
- [64] M. Cardona and M. L. W. Thewalt. “Temperature Dependence of the Energy Gap of Semiconductors in the Low-Temperature Limit”. In: *Physical Review Letters* 92.19 (2004), p. 196403. DOI: 10.1103/PhysRevLett.92.196403.
- [65] A. Pályi and G. Burkard. “Hyperfine-induced valley mixing and the spin-valley blockade in carbon-based quantum dots”. In: *Physical Review B* 80.20 (Nov. 2009), p. 201404. DOI: 10.1103/PhysRevB.80.201404.
- [66] F. Schäffler. “High-mobility Si and Ge structures”. In: *Semiconductor Science and Technology* 12.12 (Dec. 1997), pp. 1515–1549. DOI: 10.1088/0268-1242/12/12/001.
- [67] D. Culcer, X. Hu, and S. Das Sarma. “Interface roughness, valley-orbit coupling, and valley manipulation in quantum dots”. In: *Physical Review B* 82.20 (Nov. 2010), p. 205315. DOI: 10.1103/PhysRevB.82.205315.
- [68] J. K. Gamble et al. “Disorder-induced valley-orbit hybrid states in Si quantum dots”. In: *Physical Review B* 88.3 (July 2013), p. 035310. DOI: 10.1103/PhysRevB.88.035310.
- [69] A. Hollmann et al. “Large, Tunable Valley Splitting and Single-Spin Relaxation Mechanisms in a Si/Si_xGe_{1-x} Quantum Dot”. In: *Physical Review Applied* 13.3 (Mar. 2020), p. 034068. DOI: 10.1103/PhysRevApplied.13.034068.
- [70] P. Huang and X. Hu. “Spin relaxation in a Si quantum dot due to spin-valley mixing”. In: *Physical Review B* 90.23 (Dec. 2014), p. 235315. DOI: 10.1103/PhysRevB.90.235315.
- [71] W. H. Lim et al. “Spin filling of valley-orbit states in a silicon quantum dot”. In: *Nanotechnology* 22.33 (July 2011), p. 335704. DOI: 10.1088/0957-4484/22/33/335704.
- [72] M. Lodari et al. “Valley Splitting in Silicon from the Interference Pattern of Quantum Oscillations”. In: *Physical Review Letters* 128.17 (Apr. 2022), p. 176603. DOI: 10.1103/PhysRevLett.128.176603.

- [73] H. Bluhm et al. *Electrons in Solids: Mesoscopics, Photonics, Quantum Computing, Correlations, Topology*. Graduate Texts in Condensed Matter. Berlin; Boston: De Gruyter, 2019. DOI: 10.1515/9783110438321.
- [74] L. M. K. Vandersypen and I. L. Chuang. “NMR techniques for quantum control and computation”. In: *Rev. Mod. Phys.* 76.4 (Jan. 2005), pp. 1037–1069. DOI: 10.1103/RevModPhys.76.1037.
- [75] A. Caneschi et al. “Alternating current susceptibility, high field magnetization, and millimeter band EPR evidence for a ground $S = 10$ state in $[\text{Mn}_{12}\text{O}_{12}(\text{CH}_3\text{COO})_{16}(\text{H}_2\text{O})_4] \cdot 2\text{CH}_3\text{COOH} \cdot 4\text{H}_2\text{O}$ ”. In: *J. Am. Chem. Soc.* 113.15 (1991), pp. 5873–5874. DOI: 10.1021/ja00015a057.
- [76] R. Sessoli et al. “High-Spin Molecules: $[\text{Mn}_{12}\text{O}_{12}(\text{O}_2\text{CR})_{16}(\text{H}_2\text{O})_4]$ ”. In: *J. Am. Chem. Soc.* 115.5 (Mar. 1993), pp. 1804–1816. DOI: 10.1021/ja00058a027.
- [77] J. R. Friedman et al. “Macroscopic Measurement of Resonant Magnetization Tunneling in High-Spin Molecules”. In: *Phys. Rev. Lett.* 76.20 (May 1996), pp. 3830–3833. DOI: 10.1103/PhysRevLett.76.3830.
- [78] L. Thomas et al. “Macroscopic quantum tunnelling of magnetization in a single crystal of nanomagnets”. In: *Nature* 383.6596 (Sept. 1996), pp. 145–147. DOI: 10.1038/383145a0.
- [79] M. Atzori and R. Sessoli. “The Second Quantum Revolution: Role and Challenges of Molecular Chemistry”. In: *Journal of the American Chemical Society* 141.29 (2019), pp. 11339–11352. DOI: 10.1021/jacs.9b00984.
- [80] W. Wernsdorfer and M. Ruben. “Synthetic Hilbert Space Engineering of Molecular Qudits: Isotopologue Chemistry”. In: *Advanced Materials* 31.26 (2019), p. 1806687. DOI: 10.1002/adma.201806687.
- [81] D. Janković et al. “Noisy qudit vs multiple qubits: Conditions on gate efficiency for enhancing fidelity”. In: *npj Quantum Information* 10 (2024), p. 59. DOI: 10.1038/s41534-024-00829-6.
- [82] S. Lumetti et al. “Single-molecule devices with graphene electrodes”. In: *Dalton Transactions* 45 (2016), pp. 16570–16574. DOI: 10.1039/C6DT02445A.
- [83] G. Taran. “Quantum Effects in Molecular Magnets”. PhD thesis. Karlsruhe, Germany: Karlsruhe Institute of Technology, 2020.
- [84] N. Ishikawa et al. “Determination of Ligand-Field Parameters and f-Electronic Structures of Double-Decker Bis(phthalocyaninato)lanthanide Complexes”. In: *Inorganic Chemistry* 42.7 (Apr. 2003), pp. 2440–2446. DOI: 10.1021/ic026295u.
- [85] K. W. H. Stevens. “Matrix Elements and Operator Equivalents Connected with the Magnetic Properties of Rare Earth Ions”. In: *Proc. Phys. Soc. A* 65.3 (Mar. 1952), pp. 209–215. DOI: 10.1088/0370-1298/65/3/308.
- [86] N. Ishikawa, M. Sugita, and W. Wernsdorfer. “Quantum tunneling of magnetization in lanthanide single-molecule magnets: bis(phthalocyaninato)terbium and bis(phthalocyaninato)dysprosium anions”. In: *Angewandte Chemie International Edition* 44.19 (2005), pp. 2931–2935. DOI: 10.1002/anie.200462072.

- [87] E. Moreno-Pineda et al. “Molecular spin qubits for quantum algorithms”. In: *Chemical Society Reviews* 47.2 (2018), pp. 501–513. DOI: 10.1039/C5CS00933B.
- [88] C. Zener. “Non-adiabatic Crossing of Energy Levels”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 137.833 (1932), pp. 696–702. DOI: 10.1098/rspa.1932.0165.
- [89] L. D. Landau. “Zur Theorie der Energieübertragung. II”. German. In: *Physikalische Zeitschrift der Sowjetunion* 2 (1932), pp. 46–51.
- [90] P. L. Scott and C. D. Jeffries. “Spin-Lattice Relaxation in Some Rare-Earth Salts at Helium Temperatures; Observation of the Phonon Bottleneck”. In: *Phys. Rev.* 127.1 (July 1962), pp. 32–51. DOI: 10.1103/PhysRev.127.32.
- [91] K. N. Shrivastava. “Theory of Spin–Lattice Relaxation”. In: *physica status solidi (b)* 117.2 (June 1983), pp. 437–458. DOI: 10.1002/pssb.2221170202.
- [92] F. Delgado and J. Fernández-Rossier. “Spin decoherence of magnetic atoms on surfaces”. In: *Progress in Surface Science* 92.1 (Feb. 2017), pp. 40–82. DOI: 10.1016/j.progsurf.2016.12.001.
- [93] M. Dailey and C. Besson. “Selective crystallization of four bis(phthalocyaninato)lanthanoid(III) polymorphs”. In: *CrystEngComm* 23.40 (Sept. 2021), pp. 7151–7161. DOI: 10.1039/D1CE00936B.
- [94] S. Klyatskaya, A. Eichhöfer, and W. Wernsdorfer. “X-ray Crystallographic Analysis of a Tailor-Made Bis(phthalocyaninato)-Tb^{III} Single-Molecule Magnet as a Fundamental Unit for Supramolecular Spintronic Devices”. In: *Eur. J. Inorg. Chem.* 2014.25 (Sept. 2014), pp. 4179–4185. DOI: 10.1002/ejic.201402421.
- [95] C. Godfrin et al. “Electrical Read-Out of a Single Spin Using an Exchange-Coupled Quantum Dot”. In: *ACS Nano* 11.4 (Apr. 2017), pp. 3984–3989. DOI: 10.1021/acsnano.7b00451.
- [96] J. D. Jackson. *Classical Electrodynamics*. 3rd ed. Wiley, 1999.
- [97] G. Taran et al. “Direct determination of high-order transverse ligand field parameters via μ SQUID-EPR in a Et₄N[¹⁶⁰GdPc₂] SMM”. In: *Nature Communications* 14.1 (June 2023), p. 3361. DOI: 10.1038/s41467-023-39003-5.
- [98] R. Hanson et al. “Spins in few-electron quantum dots”. In: *Reviews of Modern Physics* 79.4 (Oct. 2007), pp. 1217–1265. DOI: 10.1103/RevModPhys.79.1217.
- [99] G. Scappucci et al. “Crystalline materials for quantum computing: Semiconductor heterostructures and topological insulators exemplars”. In: *MRS Bulletin* 46.7 (2021), pp. 596–606. DOI: 10.1557/s43577-021-00147-8.
- [100] E. Vahapoglu et al. “Coherent control of electron spin qubits in silicon using a global field”. In: *npj Quantum Information* 8 (2022), p. 126. DOI: 10.1038/s41534-022-00645-w.
- [101] M. Veldhorst et al. “An addressable quantum dot qubit with fault-tolerant control-fidelity”. In: *Nature Nanotechnology* 9.12 (Dec. 2014), pp. 981–985. DOI: 10.1038/nnano.2014.216.

- [102] L. F. Peña et al. “Modeling Si/SiGe quantum dot variability induced by interface disorder reconstructed from multiperspective microscopy”. In: *npj Quantum Information* 10 (2024), p. 33. DOI: 10.1038/s41534-024-00827-8.
- [103] T. Koch et al. “Industrial 300 mm wafer processed spin qubits in natural silicon/silicon–germanium”. In: *npj Quantum Information* 11.59 (2025). DOI: 10.1038/s41534-025-01016-x.
- [104] F. van Riggelen-Doelman et al. “Coherent spin qubit shuttling through germanium quantum dots”. In: *Nature Communications* 15.1 (2024), p. 5716. DOI: 10.1038/s41467-024-49358-y.
- [105] N. W. Hendrickx et al. “Fast two-qubit logic with holes in germanium”. In: *Nature* 577.7791 (Jan. 2020), pp. 487–491. DOI: 10.1038/s41586-019-1919-3.
- [106] L. Mauro et al. “Hole spin qubits in unstrained Germanium layers”. In: *npj Quantum Information* 11 (Oct. 2025), p. 167. DOI: 10.1038/s41534-025-01108-8.
- [107] B. Martinez, S. de Franceschi, and Y.-M. Niquet. “Mitigating variability in epitaxial-heterostructure-based spin-qubit devices by optimizing gate layout”. In: *Physical Review Applied* 22 (Aug. 2024), p. 024030. DOI: 10.1103/PhysRevApplied.22.024030.
- [108] R. Xue et al. “Si/SiGe QuBus for single electron information-processing devices with memory and micron-scale connectivity function”. In: *Nature Communications* 15.1 (Mar. 2024), p. 2296. DOI: 10.1038/s41467-024-46519-x.
- [109] S. D. Ha et al. “Two-Dimensional Si Spin Qubit Arrays with Multilevel Interconnects”. In: *PRX Quantum* 6.3 (Aug. 2025), p. 030327. DOI: 10.1103/sgn1-1t2d.
- [110] S. Neyens et al. “Probing single electrons across 300-mm spin qubit wafers”. In: *Nature* 629 (May 2024), pp. 80–85. DOI: 10.1038/s41586-024-07275-6.
- [111] M. D. Smet et al. “High-fidelity single-spin shuttling in silicon”. In: *Nature Nanotechnology* (2025). Published online 9 June 2025. DOI: 10.1038/s41565-025-01920-5.
- [112] T. Struck et al. “Spin-EPR-pair separation by conveyor-mode single electron shuttling in Si/SiGe”. In: *Nature Communications* 15 (2024). DOI: 10.1038/s41467-024-45583-7.
- [113] A. Abragam. *The Principles of Nuclear Magnetism*. International Series of Monographs on Physics. Oxford: Clarendon Press, 1961.
- [114] F. Pobell. “The ^3He – ^4He Dilution Refrigerator”. In: *Matter and Methods at Low Temperatures*. Berlin, Heidelberg: Springer, 2007, pp. 239–276. DOI: 10.1007/978-3-540-46360-3_7.
- [115] G. Frossati. “Obtaining ultralow temperatures by dilution of ^3He into ^4He ”. In: *Journal de Physique Colloques* 39.C6 (1978), pp. C6-1578–C6-1589. DOI: 10.1051/jphyscol:19786604.
- [116] J. Ferrero. “Development of a Cryogenic Platform with Fast-Turnaround In-Situ Characterization of Semiconductor Circuits for Qubit Applications”. Supervisors: Prof. Dr. Wolfgang Wernsdorfer, Prof. Dr. Hendrik Bluhm. PhD thesis. Karlsruhe, Germany: Karlsruhe Institute of Technology (KIT), Faculty of Physics, Mar. 2024.

-
- [117] J. Kamioka et al. “Evaluation of a physically defined silicon quantum dot for design of matching circuit for RF reflectometry charge sensing”. In: *AIP Advances* 13.3 (2023), p. 035219. DOI: 10.1063/5.0141092.
- [118] Y.-Y. Liu et al. “Radio-Frequency Reflectometry in Silicon-Based Quantum Dots”. In: *Physical Review Applied* 16.1 (July 2021), p. 014057. DOI: 10.1103/PhysRevApplied.16.014057.
- [119] V. Adam. “Coherent Two-Qubit Control in Industrial Si/SiGe Spin Qubits”. Supervisors: Prof. Dr. Wolfgang Wernsdorfer, Prof. Dr. Sebastian Kempf. PhD thesis. Karlsruhe, Germany: Karlsruhe Institute of Technology (KIT), Faculty of Physics, July 2025.
- [120] qkitgroup and contributors. *Qkit – a quantum measurement suite in Python*. <https://github.com/qkitgroup/qkit>. 2025.
- [121] D. M. Pozar. *Microwave Engineering*. 4th ed. Hoboken, NJ, USA: John Wiley & Sons, 2011.
- [122] Marki Microwave. *IQ, IR, and SSB Mixer Primer*. <https://markimicrowave.com/technical-resources/white-papers/iq-ir-ssb-mixer-primer/>. White Paper. 2020.
- [123] P. Horowitz and W. Hill. *The Art of Electronics*. 3rd. Cambridge University Press, 2015.
- [124] P. Scarlino et al. “Spin-Relaxation Anisotropy in a GaAs Quantum Dot”. In: *Phys. Rev. Lett.* 115.10 (2015), p. 106802. DOI: 10.1103/PhysRevLett.113.256802.
- [125] Jäger Computergesteuerte Messtechnik GmbH. *ADwin Pro II System Manual*. Product manual for the ADwin Pro II real-time system. Jäger Computergesteuerte Messtechnik GmbH. Lorsch, Germany, 2024. URL: <https://www.adwin.de/de/download/cdrom.html>.
- [126] S. Fuhrmann. “Measurement Setup for the Characterization of Semiconductor Qubits”. Reviewer: Prof. Dr. Wolfgang Wernsdorfer; Second reviewer: Dr. Philip Willke; Advisor: Julian Ferrero. Master’s Thesis. Karlsruhe, Germany: Physikalisches Institut (PHI), Karlsruhe Institute of Technology (KIT), Nov. 2021.
- [127] D. Schroller. *ADwin_Pro2_V3.py: Python driver for ADwin Pro II and associated ADbasic files in ADwin_Pro/ADbasic_files*. https://github.com/ThomasK-AGWW/Semicon_qkit/tree/master/qkit/drivers. Accessed: 2025-08-22. 2025.
- [128] J. B. Johnson. “Thermal Agitation of Electricity in Conductors”. In: *Physical Review* 32.1 (1928), pp. 97–109. DOI: 10.1103/PhysRev.32.97.
- [129] H. Nyquist. “Thermal Agitation of Electric Charge in Conductors”. In: *Physical Review* 32.1 (1928), pp. 110–113. DOI: 10.1103/PhysRev.32.110.
- [130] V. Adam. *Private communication*. Karlsruhe Institute of Technology (KIT), group of Prof. Wolfgang Wernsdorfer. 2025.
- [131] S. G. J. Philips et al. “Universal control of a six-qubit quantum processor in silicon”. In: *Nature* 609 (2022), pp. 919–924. DOI: 10.1038/s41586-022-05117-w.

- [132] Optiprint AG. *Leiterplatten Technologie und Produktion*. <https://optiprint.ch/>. Optiprint AG develops and manufactures high-quality PCB solutions in Switzerland. 2025.
- [133] A. Crippa et al. “Gate-reflectometry dispersive readout and coherent control of a spin qubit in silicon”. In: *Nature Communications* 10 (2019), p. 2776. DOI: 10.1038/s41467-019-10848-z.
- [134] I. Wolff. *Coplanar Microwave Integrated Circuits*. See p. 40. Hoboken, NJ: John Wiley & Sons, 2006.
- [135] S. Hageman. “Via spacing on high-performance PCBs”. In: *EDN* (Feb. 2013). URL: <https://www.edn.com/via-spacing-on-high-performance-pcbs/>.
- [136] E. J. Connors, J. Nelson, and J. M. Nichol. “Rapid High-Fidelity Spin-State Readout in Si/Si-Ge Quantum Dots via rf Reflectometry”. In: *Physical Review Applied* 13.2 (2020), p. 024019. DOI: 10.1103/PhysRevApplied.13.024019.
- [137] F. Vigneau et al. “Probing quantum devices with radio-frequency reflectometry”. In: *APL Physics Reviews* 10.2 (2023), p. 021305. DOI: 10.1063/5.0088229.
- [138] T. A. Fulton and G. J. Dolan. “Observation of Single-Electron Charging Effects in Small Tunnel Junctions”. In: *Physical Review Letters* 59.1 (July 1987), pp. 109–112. DOI: 10.1103/PhysRevLett.59.109.
- [139] C. W. J. Beenakker. “Theory of Coulomb-blockade oscillations in the conductance of a quantum dot”. In: *Physical Review B* 44.4 (July 1991), pp. 1646–1656. DOI: 10.1103/PhysRevB.44.1646.
- [140] L. P. Kouwenhoven et al. “Electron transport in quantum dots”. In: *Mesoscopic Electron Transport*. Ed. by L. L. Sohn, L. P. Kouwenhoven, and G. Schön. Vol. 345. NATO ASI Series E: Applied Sciences. Dordrecht: Kluwer Academic Publishers, 1997, pp. 105–214. DOI: 10.1007/978-94-015-8839-3_4.
- [141] H. van Houten, C. W. J. Beenakker, and A. A. M. Staring. “Coulomb-Blockade Oscillations in Semiconductor Nanostructures”. In: *Single Charge Tunneling*. Ed. by H. Grabert and M. H. Devoret. Vol. 294. NATO ASI Series. Boston, MA: Springer, 1992, pp. 167–216. DOI: 10.1007/978-1-4757-2166-9_5.
- [142] A. Morello et al. “Single-shot readout of an electron spin in silicon”. In: *Nature* 467.7316 (Oct. 2010), pp. 687–691. DOI: 10.1038/nature09392.
- [143] L. Petit et al. “Spin Lifetime and Charge Noise in Hot Silicon Quantum Dot Qubits”. In: *Physical Review Letters* 121 (2018), p. 076801. DOI: 10.1103/PhysRevLett.121.076801.
- [144] M. J. Kirton and M. J. Uren. “Noise in solid-state microstructures: A new perspective on individual defects, interface states and low-frequency (1/f) noise”. In: *Advances in Physics* 38.4 (1989), pp. 367–468. DOI: 10.1080/00018738900101122.
- [145] B. M. Freeman, J. S. Schoenfield, and H. Jiang. “Comparison of low frequency charge noise in identically patterned Si/SiO₂ and Si/SiGe quantum dots”. In: *Applied Physics Letters* 108.25 (June 2016), p. 253108. DOI: 10.1063/1.4954700.

- [146] K. W. Chan et al. “Assessment of a Silicon Quantum Dot Spin Qubit Environment via Noise Spectroscopy”. In: *Physical Review Applied* 10.4 (Oct. 2018), p. 044017. DOI: 10.1103/PhysRevApplied.10.044017.
- [147] C. E. Shannon. “Communication in the Presence of Noise”. In: *Proceedings of the IRE* 37.1 (Jan. 1949), pp. 10–21. DOI: 10.1109/JRPROC.1949.232969.
- [148] P. D. Welch. “The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms”. In: *IEEE Transactions on Audio and Electroacoustics* AU-15.2 (June 1967), pp. 70–73. DOI: 10.1109/TAU.1967.1161901.
- [149] SciPy Developers. *scipy.signal.welch — Welch’s method*. SciPy. 2025. URL: <https://docs.scipy.org/doc/scipy/reference/generated/scipy.signal.welch.html> (visited on 09/30/2025).
- [150] P. Dutta and P. M. Horn. “Low-frequency fluctuations in solids: 1/f noise”. In: *Reviews of Modern Physics* 53.3 (July 1981), pp. 497–516. DOI: 10.1103/RevModPhys.53.497.
- [151] A. Elsayed et al. “Low charge noise quantum dots with industrial CMOS manufacturing”. In: *npj Quantum Information* (2024). DOI: 10.1038/s41534-024-00864-3.
- [152] B. Paquelet Wuetz et al. “Reducing charge noise in quantum dots by using thin silicon quantum wells”. In: *Nature Communications* 14 (Mar. 2023), p. 1385. DOI: 10.1038/s41467-023-36951-w.
- [153] D. Maradan et al. “GaAs Quantum Dot Thermometry Using Direct Transport and Charge Sensing”. In: *Journal of Low Temperature Physics* 175.5–6 (Apr. 2014), pp. 784–798. DOI: 10.1007/s10909-014-1169-6.
- [154] T. Koch. “Fidelity Benchmarking of Industrial Silicon/Silicon-Germanium Qubits”. Doctoral dissertation; oral defence on 13 December 2024. Referees: Prof. Dr. Wolfgang Wernsdorfer; Prof. Dr. Kristiaan De Greve. PhD thesis. Karlsruhe, Germany: Karlsruhe Institute of Technology (KIT), KIT-Fakultät für Physik, Dec. 2024.
- [155] S. Gasparinetti et al. “Nongalvanic thermometry for ultracold two-dimensional electron domains”. In: *Applied Physics Letters* 100.26 (June 2012), p. 253502. DOI: 10.1063/1.4729388.
- [156] A. Rossi, T. Ferrus, and D. A. Williams. “Electron temperature in electrically isolated Si double quantum dots”. In: *Applied Physics Letters* 100.13 (Mar. 2012), p. 133503. DOI: 10.1063/1.3697832.
- [157] J. M. Elzerman et al. “Single-shot read-out of an individual electron spin in a quantum dot”. In: *Nature* 430 (July 2004), pp. 431–435. DOI: 10.1038/nature02693.
- [158] V. N. Golovach, A. Khaetskii, and D. Loss. “Phonon-Induced Decay of the Electron Spin in Quantum Dots”. In: *Physical Review Letters* 93.1 (July 2004), p. 016601. DOI: 10.1103/PhysRevLett.93.016601.
- [159] M. Garwood and L. DelaBarre. “The Return of the Frequency Sweep: Designing Adiabatic Pulses for Contemporary NMR”. In: *Journal of Magnetic Resonance* 153.2 (Aug. 2001), pp. 155–177. DOI: 10.1006/jmre.2001.2340.

- [160] K. Li, D. C. Spierings, and A. M. Steinberg. “Efficient adiabatic rapid passage in the presence of noise”. In: *Physical Review A* 108.1 (July 2023), p. 012615. DOI: 10.1103/PhysRevA.108.012615.
- [161] F. Bloch. “Nuclear Induction”. In: *Physical Review* 70.7–8 (Oct. 1946), pp. 460–474. DOI: 10.1103/PhysRev.70.460.
- [162] C. B. Harris and R. J. Hoover. “Optically Detected Adiabatic Inversion in Phosphorescent Triplet States and the Measurement of Intramolecular Energy Transfer Processes”. In: *The Journal of Chemical Physics* 56.5 (1972), pp. 2199–2206. DOI: 10.1063/1.1677520.
- [163] A. Doll et al. “Adiabatic and fast passage ultra-wideband inversion in pulsed EPR”. In: *Journal of Magnetic Resonance* 230 (2013), pp. 27–39. DOI: 10.1016/j.jmr.2013.01.002.
- [164] H. Wu et al. “Geometric phase gates with adiabatic control in electron spin resonance”. In: *Physical Review A* 87.3 (2013), p. 032326. DOI: 10.1103/PhysRevA.87.032326.
- [165] J. Zhang et al. “Experimental Implementation of Assisted Quantum Adiabatic Passage in a Single Spin”. In: *Physical Review Letters* 110.24 (2013), p. 240501. DOI: 10.1103/PhysRevLett.110.240501.
- [166] I. Niemeyer et al. “Broadband excitation by chirped pulses: application to single electron spins in diamond”. In: *New Journal of Physics* 15.3 (2013), p. 033027. DOI: 10.1088/1367-2630/15/3/033027.
- [167] A. Laucht et al. “High-fidelity adiabatic inversion of a ^{31}P electron spin qubit in natural silicon”. In: *Applied Physics Letters* 104.9 (2014), p. 092115. DOI: 10.1063/1.4867905.
- [168] X.-F. Liu et al. “Accelerated Adiabatic Passage of a Single Electron Spin Qubit in Quantum Dots”. In: *Physical Review Letters* 132.2 (2024), p. 027002. DOI: 10.1103/PhysRevLett.132.027002.
- [169] N. Chanda, P. Patnaik, and R. Bhattacharyya. “Optimal population transfer using adiabatic rapid passage in the presence of drive-induced dissipation”. In: *Physical Review A* 107.6 (June 2023), p. 063708. DOI: 10.1103/PhysRevA.107.063708.
- [170] P. Krantz et al. “A quantum engineer’s guide to superconducting qubits”. In: *Applied Physics Reviews* 6.2 (2019), p. 021318. DOI: 10.1063/1.5089550.
- [171] L. Allen and J. H. Eberly. *Optical Resonance and Two-Level Atoms*. Dover ed. New York, NY: Dover Publications, 1987.
- [172] G. Ithier et al. “Decoherence in a superconducting quantum bit circuit”. In: *Phys. Rev. B* 72 (2005), p. 134519. DOI: 10.1103/PhysRevB.72.134519.
- [173] A. Farolfi et al. “Design and characterization of a compact magnetic shield for ultracold atomic gas experiments”. In: *Review of Scientific Instruments* 90.11 (Nov. 2019), p. 115114. DOI: 10.1063/1.5119915.

- [174] T. Ruster et al. “A long-lived Zeeman trapped-ion qubit”. In: *Appl. Phys. B* 122 (Sept. 2016), p. 254. DOI: 10.1007/s00340-016-6527-4.
- [175] E. Kawakami et al. “Electrical control of a long-lived spin qubit in a Si/SiGe quantum dot”. In: *Nature Nanotechnology* 9.9 (Sept. 2014), pp. 666–670. DOI: 10.1038/nnano.2014.153.
- [176] I. I. Rabi. “Space Quantization in a Gyating Magnetic Field”. In: *Phys. Rev.* 51 (1937), pp. 652–654. DOI: 10.1103/PhysRev.51.652.
- [177] N. F. Ramsey. “A Molecular Beam Resonance Method with Separated Oscillating Fields”. In: *Physical Review* 78 (1950), pp. 695–699. DOI: 10.1103/PhysRev.78.695.
- [178] O. E. Dial et al. “Charge Noise Spectroscopy Using Coherent Exchange Oscillations in a Singlet-Triplet Qubit”. In: *Physical Review Letters* 110.14 (Apr. 2013), p. 146804. DOI: 10.1103/PhysRevLett.110.146804.
- [179] E. L. Hahn. “Spin Echoes”. In: *Physical Review* 80.4 (Nov. 1950), pp. 580–594. DOI: 10.1103/PhysRev.80.580.
- [180] J. Bylander et al. “Noise spectroscopy through dynamical decoupling with a superconducting flux qubit”. In: *Nature Physics* 7.7 (July 2011), pp. 565–570. DOI: 10.1038/nphys1994.
- [181] R. Zhao et al. “Single-spin qubits in isotopically enriched silicon at low magnetic field”. In: *Nature Communications* 10 (2019), p. 5500. DOI: 10.1038/s41467-019-13416-7.
- [182] R. de Sousa and S. Das Sarma. “Theory of nuclear-induced spectral diffusion: Spin decoherence of phosphorus donors in Si and GaAs quantum dots”. In: *Phys. Rev. B* 68.11 (Sept. 2003), p. 115322. DOI: 10.1103/PhysRevB.68.115322.
- [183] W. M. Witzel et al. “Electron Spin Decoherence in Isotope-Enriched Silicon”. In: *Phys. Rev. Lett.* 105.18 (Oct. 2010), p. 187602. DOI: 10.1103/PhysRevLett.105.187602.
- [184] F. G. Kondev et al. “The NUBASE2020 evaluation of nuclear physics properties”. In: *Chinese Physics C* 45.3 (2021), p. 030001. DOI: 10.1088/1674-1137/abddae.
- [185] D. Schroller et al. *Hybrid Quantum Systems: Coupling Single-Molecule Magnet Qudits with Industrial Silicon Spin Qubits*. Oct. 2025. DOI: 10.48550/arXiv.2510.10110. arXiv: 2510.10110 [cond-mat.mes-hall].
- [186] M. A. Fogarty et al. “Integrated silicon qubit platform with single-spin addressability, exchange control and single-shot singlet-triplet readout”. In: *Nature Communications* 9.1 (2018), p. 4370. DOI: 10.1038/s41467-018-06039-x.
- [187] N. Gläser. *KIT / PHI / Nanospin / SemiSpin / Denoise (Python library)*. <https://gitlab.kit.edu/kit/phi/nanospin/semispin/denoise>. GitLab repository. Accessed: 2025-10-24. 2025.

- [188] N. Gläser. “Enhancement of Molecular Spin Sensitivity and Blip Detection Efficiency in Silicon Spin Qubits”. Reviewer: Prof. Dr. Wolfgang Wernsdorfer; Second Reviewer: TT-Prof. Dr. Philip Willke; Advisors: M. Sc. Viktor Adam, M. Sc. Thomas Koch, M. Sc. Daniel Schroller. Master’s Thesis. Physikalisches Institut (PHI), Karlsruhe Institute of Technology (KIT), June 2024.
- [189] G. Casella and R. L. Berger. *Statistical Inference*. 2nd ed. Pacific Grove, CA: Duxbury, 2002.
- [190] M. Perfetti et al. “Molecular Order in Buried Layers of TbPc₂ Single-Molecule Magnets Detected by Torque Magnetometry”. In: *Advanced Materials* 28.32 (2016), pp. 6946–6951. DOI: 10.1002/adma.201600791.
- [191] A. Nabiałek et al. “The Influence of Magnetic History on the Stability of Critical State and the Dynamics of Flux Jumps in Conventional NbTi Superconductor”. In: *Acta Physica Polonica A* 118.2 (2010), p. 343. DOI: 10.12693/APHYSPOLA.118.343.
- [192] W. Wernsdorfer and R. Sessoli. “Quantum Phase Interference and Parity Effects in Magnetic Molecular Clusters”. In: *Science* 284.5411 (1999), pp. 133–135. DOI: 10.1126/science.284.5411.133.
- [193] E. Moreno-Pineda, S. Paul, and A. Sunil. *μSQUID precharacterization of [TbPc₂] crystals*. Private communication. 2025.
- [194] O. Edholm and C. Blomberg. “Stretched exponentials and barrier distributions”. In: *Chemical Physics* 252.1-2 (2000), pp. 221–225. DOI: 10.1016/S0301-0104(99)00349-3.
- [195] D. N. Woodruff, R. E. P. Winpenny, and R. A. Layfield. “Lanthanide Single-Molecule Magnets”. In: *Chemical Reviews* 113.7 (2013), pp. 5110–5148. DOI: 10.1021/cr400018q.
- [196] K. Ono et al. “Current Rectification by Pauli Exclusion in a Weakly Coupled Double Quantum Dot System”. In: *Science* 297.5585 (2002), pp. 1313–1317. DOI: 10.1126/science.1070958.
- [197] Y. Ban et al. “Pauli spin blockade at room temperature in double-quantum-dot tunneling through individual deep dopants in silicon”. In: *Communications Physics* 8 (2025), p. 293. DOI: 10.1038/s42005-025-02177-z.
- [198] J. K. Perron, J. Stewart M. D., and N. M. Zimmerman. “A new regime of Pauli-spin blockade”. In: *Journal of Applied Physics* 119.13 (2016), p. 134301. DOI: 10.1063/1.4945393.
- [199] S. Foletti et al. “Universal quantum control of two-electron spin quantum bits using dynamic nuclear polarization”. In: *Nature Physics* 5 (2009), pp. 903–908. DOI: 10.1038/nphys1424.
- [200] X. Wu et al. “Two-axis control of a singlet–triplet qubit with an integrated micromagnet”. In: *Proceedings of the National Academy of Sciences* 111.33 (2014), pp. 11938–11942. DOI: 10.1073/pnas.1412230111.
- [201] X. Zhang et al. “Universal control of four singlet–triplet qubits”. In: *Nature Nanotechnology* 20 (2024), pp. 209–215. DOI: 10.1038/s41565-024-01817-9.

-
- [202] G. Burkard, D. Loss, and D. P. DiVincenzo. “Coupled quantum dots as quantum gates”. In: *Physical Review B* 59.3 (1999), pp. 2070–2078. DOI: 10.1103/PhysRevB.59.2070.
- [203] J. R. Johansson, P. D. Nation, and F. Nori. “QuTiP 2: A Python framework for the dynamics of open quantum systems”. In: *Computer Physics Communications* 184 (2013), pp. 1234–1240. DOI: 10.1016/j.cpc.2012.11.019.
- [204] M. D. Shulman et al. “Demonstration of entanglement of electrostatically coupled singlet–triplet qubits”. In: *Science* 336.6078 (2012), pp. 202–205. DOI: 10.1126/science.1217692.
- [205] T. Meunier, V. E. Calado, and L. M. K. Vandersypen. “Efficient controlled-phase gate for single-spin qubits in quantum dots”. In: *Physical Review B* 83.12 (2011), p. 121403. DOI: 10.1103/PhysRevB.83.121403.
- [206] M. Veldhorst et al. “A two-qubit logic gate in silicon”. In: *Nature* 526 (2015), pp. 410–414. DOI: 10.1038/nature15263.
- [207] L. R. Schreiber et al. “Coupling artificial molecular spin states by photon-assisted tunnelling”. In: *Nature Communications* 2 (2011), p. 556. DOI: 10.1038/ncomms1561.
- [208] R. Ferdous et al. “Interface-induced spin-orbit interaction in silicon quantum dots and prospects for scalability”. In: *Physical Review B* 97.24 (2018), p. 241401. DOI: 10.1103/PhysRevB.97.241401.
- [209] T. Tanttu et al. “Controlling Spin-Orbit Interactions in Silicon Quantum Dots Using Magnetic Field Direction”. In: *Physical Review X* 9.2 (2019), p. 021028. DOI: 10.1103/PhysRevX.9.021028.
- [210] R. M. Jock et al. “A silicon metal-oxide-semiconductor electron spin-orbit qubit”. In: *Nature Communications* 9 (2018), p. 1768. DOI: 10.1038/s41467-018-04200-0.
- [211] M. D. Reed et al. “Reduced sensitivity to charge noise in semiconductor spin qubits via symmetric operation”. In: *Physical Review Letters* 116.11 (2016), p. 110402. DOI: 10.1103/PhysRevLett.116.110402.
- [212] S. N. Shevchenko, S. Ashhab, and F. Nori. “Landau–Zener–Stückelberg interferometry”. In: *Physics Reports* 492.1 (2010), pp. 1–30. DOI: 10.1016/j.physrep.2010.03.002.
- [213] E. A. Laird et al. “Effect of exchange interaction on spin dephasing in a double quantum dot”. In: *Physical Review Letters* 97.5 (2006), p. 056801. DOI: 10.1103/PhysRevLett.97.056801.
- [214] P. Huang, N. M. Zimmerman, and G. W. Bryant. “Spin decoherence in a two-qubit CPHASE gate: the critical role of tunneling noise”. In: *npj Quantum Information* 4 (2018), p. 62. DOI: 10.1038/s41534-018-0112-0.
- [215] S. F. Neyens et al. “The critical role of substrate disorder in valley splitting in Si quantum wells”. In: *Applied Physics Letters* 112.24 (2018), p. 243107. DOI: 10.1063/1.5033447.

- [216] J. D. Cifuentes et al. “Bounds to electron spin qubit variability for scalable CMOS architectures”. In: *Nature Communications* 15.1 (2024), p. 4299. DOI: 10.1038/s41467-024-48557-x.
- [217] C. H. Yang et al. “Orbital and valley state spectra of a few-electron silicon quantum dot”. In: *Physical Review B* 86.11 (2012), p. 115319. DOI: 10.1103/PhysRevB.86.115319.
- [218] A. E. Seedhouse et al. “Pauli blockade in silicon quantum dots with spin-orbit control”. In: *PRX Quantum* 2.1 (2021), p. 010303. DOI: 10.1103/PRXQuantum.2.010303.
- [219] Y. Song et al. “Coherence of a field gradient driven singlet–triplet qubit coupled to multielectron spin states in $^{28}\text{Si}/\text{SiGe}$ ”. In: *npj Quantum Information* 10 (2024), p. 77. DOI: 10.1038/s41534-024-00869-y.
- [220] C. Barthel et al. “Relaxation and readout visibility of a singlet–triplet qubit in an Overhauser field gradient”. In: *Physical Review B* 85.3 (2012), p. 035306. DOI: 10.1103/PhysRevB.85.035306.
- [221] J. C. C. Hwang, C. H. Yang, J. C. C. Hwang, et al. “Impact of g-factors and valleys on spin qubits in a silicon metal-oxide-semiconductor double quantum dot”. In: *Physical Review B* 96 (2017), p. 045302. DOI: 10.1103/PhysRevB.96.045302.
- [222] C. Barthel et al. “Rapid single-shot measurement of a singlet–triplet qubit”. In: *Physical Review Letters* 103.16 (2009), p. 160503. DOI: 10.1103/PhysRevLett.103.160503.
- [223] M. T. Mađzik et al. “Controllable freezing of the nuclear spin bath in a single-atom spin qubit”. In: *Science Advances* 6.30 (2020), eaba3442. DOI: 10.1126/sciadv.aba3442.
- [224] E. Paladino et al. “ $1/f$ noise: Implications for solid-state quantum information”. In: *Reviews of Modern Physics* 86.2 (2014), pp. 361–418. DOI: 10.1103/RevModPhys.86.361.
- [225] K. Takeda et al. “Rapid single-shot parity spin readout in a silicon double quantum dot with fidelity exceeding 99%”. In: *npj Quantum Information* 10 (2024), p. 22. DOI: 10.1038/s41534-024-00813-0.
- [226] J. R. Prance et al. “Single-shot measurement of triplet–singlet relaxation in a Si/SiGe double quantum dot”. In: *Physical Review Letters* 108.4 (2012), p. 046808. DOI: 10.1103/PhysRevLett.108.046808.
- [227] nextnano GmbH. *nextnano — Simulation Software for Electronic and Optoelectronic Semiconductor Nanodevices*. <https://nextnano.com/>. 2025.
- [228] nextnano GmbH. *nextnano+ Documentation*. <https://www.nextnano.com/docu/nextnanoplus/index.html>. Online documentation for the nextnano+ simulator. 2025. (Visited on 09/05/2025).
- [229] Y. Okada and Y. Tokumaru. “Precise determination of lattice parameter and thermal expansion coefficient of silicon between 0 and 1000 °C”. In: *Journal of Applied Physics* 56.2 (1984), pp. 314–320. DOI: 10.1063/1.333965.

- [230] G. Bastard. “Superlattice band structure in the envelope-function approximation”. In: *Physical Review B* 24.10 (Nov. 1981), pp. 5693–5697. DOI: 10.1103/PhysRevB.24.5693.
- [231] A. Trellakis, T. Andlauer, and P. Vogl, eds. *Large-Scale Scientific Computing: 5th International Conference, LSSC 2005, Sozopol, Bulgaria, June 6–10, 2005, Revised Papers*. Vol. 3743. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2006. DOI: 10.1007/11666806.
- [232] nextnano GmbH. *Charge densities (nextnano3 models)*. https://www.nextnano.com/documentation/tools/nextnano3/models/charge_densities.html. Online documentation. 2025. (Visited on 09/08/2025).
- [233] C. T. Kelley. *Iterative Methods for Linear and Nonlinear Equations*. Vol. 16. Frontiers in Applied Mathematics. Philadelphia: SIAM, 1995.
- [234] Netlib. *BLAS (Basic Linear Algebra Subprograms)*. <https://netlib.org/blas/>. Accessed 2025-09-05. 2025.
- [235] Netlib. *LAPACK — Linear Algebra PACKage*. <https://netlib.org/lapack/>. Accessed 2025-09-05. 2025.
- [236] N. Ries. *Magnetic Particle Sensing Using SiMOS Quantum Dots: Simulation-Driven Gatestack Optimization with nextnano*. Bachelor’s Thesis. First examiner: Prof. Dr. Wolfgang Wernsdorfer; Second examiner: Dr. Christoph Sürgers; Advisor: M.Sc. Daniel Schroller; Work period: 01.03.2024–01.09.2024. Karlsruhe, Germany, Sept. 2024.
- [237] R. Acharya et al. “Highly ^{28}Si enriched silicon by localised focused ion beam implantation”. In: *Communications Materials* 5 (2024), p. 57. DOI: 10.1038/s43246-024-00498-0.
- [238] N. P. de Leon et al. “Materials challenges and opportunities for quantum computing hardware”. In: *Science* 372.6539 (2021), eabb2823. DOI: 10.1126/science.abb2823.
- [239] S. J. Pearton, R. J. Shul, and F. Ren. “A Review of Dry Etching of GaN and Related Materials”. In: *MRS Internet Journal of Nitride Semiconductor Research* 5.1 (2000), e11. DOI: 10.1557/S1092578300000119.
- [240] A. Wolff. *Focused ion beams: An overview of the technology and its capabilities*. Wiley Analytical Science. URL: <https://analyticalscience.wiley.com/content/article-do/focused-ion-beams-overview-technology-and-its-capabilities> (visited on 09/12/2025).
- [241] B. Bhuvu and S. Kerns. “Plasma Process-induced Damage”. In: *Encyclopedia of Materials: Science and Technology*. Ed. by K. H. J. Buschow et al. Elsevier, 2001, pp. 7030–7034. DOI: 10.1016/B0-08-043152-6/01247-X.
- [242] S. J. Fonash. “An Overview of Dry Etching Damage and Contamination Effects”. In: *Journal of The Electrochemical Society* 137.12 (1990), pp. 3885–3892. DOI: 10.1149/1.2086322.

- [243] H.-C. Cheng. “Wet Etching”. In: *Handbook of Visual Display Technology*. Ed. by J. Chen, W. Cranton, and M. Fihn. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 861–870. DOI: 10.1007/978-3-540-79567-4_59.
- [244] Sigma-Aldrich. *Buffered Oxide Etch (BOE) Safety Data Sheet*. <https://aggiefab.tamu.edu/wp-content/uploads/2021/04/BOE-MSDS.pdf>. Hazard statements include: “Fatal in contact with skin”; “Causes severe skin burns and eye damage.” 2021.
- [245] K. R. Williams, K. Gupta, and M. Wasilik. “Etch Rates for Micromachining Processing—Part II”. In: *Journal of Microelectromechanical Systems* 12.6 (2003), pp. 761–778. DOI: 10.1109/JMEMS.2003.820936.
- [246] S. Verhaverbeke et al. “The Etching Mechanisms of SiO₂ in Hydrofluoric Acid”. In: *Journal of The Electrochemical Society* 141.10 (1994), pp. 2852–2857. DOI: 10.1149/1.2059243.
- [247] H. Proksche, G. Nagorsen, and D. Ross. “The Influence of NH₄F on the Etch Rates of Undoped SiO₂ in Buffered Oxide Etch”. In: *Journal of The Electrochemical Society* 139.2 (1992), pp. 521–524. DOI: 10.1149/1.2069249.
- [248] M. Neyens et al. “Wet Etchant Diffusion through Photoresist during Gate Oxide Patterning”. In: *ECS Transactions* 69.8 (2015), pp. 177–183. DOI: 10.1149/06908.0177ECST.
- [249] Optiprint AG. *Optiprint Technical Information, Version 4.0*. Technical datasheet. Optiprint AG. Berneck, Switzerland, 2024. URL: <https://optiprint.ch/default-wAssets/docs/technische-informationen/Optiprint-Technical-Information-4.0.pdf>.
- [250] H. Muzaffar and K. S. Williams. “Evaluation of Complete Elliptic Integrals of the First Kind at Singular Moduli”. In: *Mathematics of Computation* 54.189 (1990), pp. 303–313. DOI: 10.2307/2008791.
- [251] R. N. Simons. *Coplanar Waveguide Circuits, Components, and Systems*. Wiley Series in Microwave and Optical Engineering. New York, NY, USA: John Wiley & Sons, 2001.

Acknowledgments

I'm deeply grateful to everyone who made the past years not only possible but genuinely fun. The atmosphere at the institute was always welcoming and cheerful. Everyone was unfailingly helpful.

First and foremost, I thank Prof. Wolfgang Wernsdorfer, my “Doktorvater”, for the opportunity to work on this project in his group. His love for science and cryostats is inspiring, especially the mechanical craftsmanship on display in his tool-filled hall. The freedom and trust he gives his students made all the difference.

I also thank Prof. David Hunger for being the second reviewer of this thesis.

Special thanks go to the other members of “Team Semiconductor”: Julian Ferrero, Thomas Koch, and Viktor Adam. We became great colleagues, sometimes rivals, and, above all, friends. Thank you for the many internal meetings, sharp discussions, and fresh ideas. A big shout-out to my students, Daniel Sitter and Niklas Ries, whose support made this project both stronger and more enjoyable. Thank you to Sonja Vogel for bringing so much cheerfulness to the group.

Clément Godfrin, who really should be called the “second Doktorvater” of this thesis, gave immense input and motivation throughout. Thank you for the excellent supervision and for building up the collaboration with imec.

In the office, many thanks to Michael Schulze, Luca Kosche, Richard Wulff, Safa Ahmed, Appu Sunil, and Stefan Pfleging. You made it a pleasure to come in day after day. Our overboarding coffee-gear experiments are still something to remember.

I want to thank Christoph Sürgers for technical support at the experiment, and the whole AG Wernsdorfer. I'm grateful for the group BBQs, Christmas parties, coffee-room cakes, and Thursday restaurant lunch breaks. Will Thai Friday ever return?

For electrical support, many thanks to Jannis Ret, Kevin Zych, and Patrick Rust; to Michael Meyer & team at the precision machining shop; and to Michael Pfirrmann for the endless liquid-helium supply that kept the experiments running.

I also thank Marcel Schrodin, Alexander Zilz, and Eugen Stach from the Qinu team.

I'm grateful to the Karlsruhe School of Quantum Matter (KSQM) for organising excellent lab visits, workshops, and summer schools.

Finally, to my family and to Christine: thank you for your support through all of this. Without you, none of it would have been possible.