

Designing the Electronic Interface for Qubit Control

van Dijk, J.P.G.

DOI

[10.4233/uuid:7abc3f2a-ed28-42d4-8ae4-f86426777884](https://doi.org/10.4233/uuid:7abc3f2a-ed28-42d4-8ae4-f86426777884)

Publication date

2021

Document Version

Final published version

Citation (APA)

van Dijk, J. P. G. (2021). *Designing the Electronic Interface for Qubit Control*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:7abc3f2a-ed28-42d4-8ae4-f86426777884>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

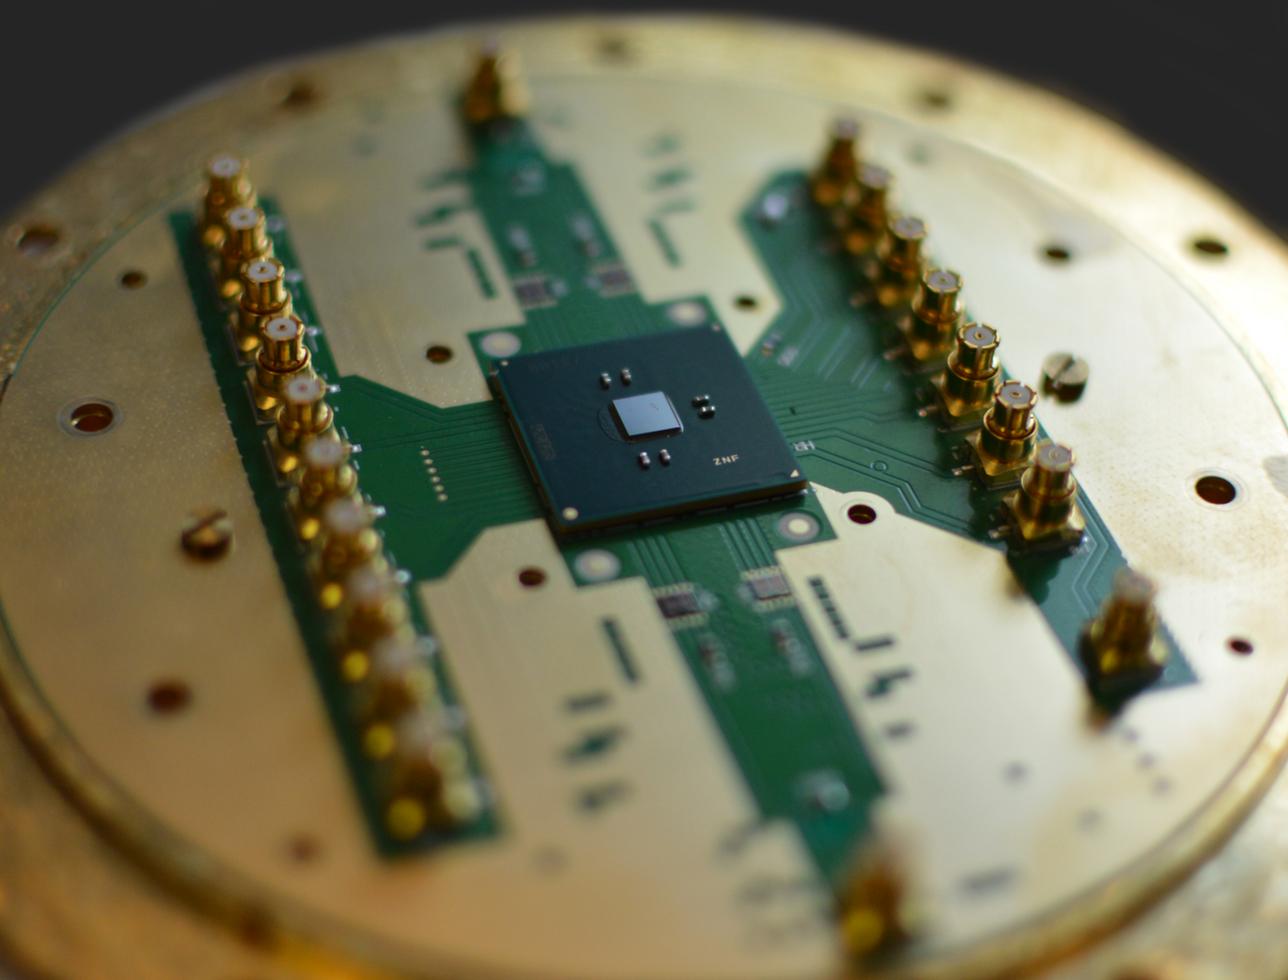
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Designing the Electronic Interface for Qubit Control

Jeroen van Dijk



Designing the Electronic Interface for Qubit Control

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus Prof. dr. ir. T.H.J.J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op dinsdag 5 oktober 2021 om 12:30 uur

door

Jeroen Petrus Gerardus VAN DIJK

Elektrotechnisch ingenieur,
Technische Universiteit Delft, Nederland,
geboren te Olst, Nederland.

Dit proefschrift is goedgekeurd door de

Promotor: Prof. dr. ir. E. Charbon
Copromotor: Dr. ir. F. Sebastiano

Samenstelling promotiecommissie:

Rector Magnificus, Voorzitter
Prof. dr. ir. E. Charbon, Technische Universiteit Delft, promotor
Dr. ir. F. Sebastiano, Technische Universiteit Delft, copromotor

Onafhankelijke leden:

Prof. dr. ir. L.C.N. de Vreede,
Technische Universiteit Delft
Prof. dr. A.S. Dzurak, University of New South Wales, Australië
Prof. dr. S.P. Voinigescu, University of Toronto, Canada
Dr. ir. J.C. Bardin, University of Massachusetts Amherst, U.S.
Prof. dr. K.A.A. Makinwa, Technische Universiteit Delft, reservelid

Overige leden:

Prof. dr. ir. A. Vladimirescu, University of California Berkeley, U.S.



QuTech



Keywords: cryogenic electronics, CMOS, FinFET,
quantum processor, qubit, electron spin

Printed by: IPSKAMP printing in the Netherlands

Cover by: Jeroen van Dijk, photo by Elizabeth Hatfield

Copyright © 2021 by J.P.G. van Dijk

ISBN 978-94-6421-504-5

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

*If you want to find the secrets of the universe, think in terms of energy,
frequency, and vibration.*

Nikola Tesla

Contents

Acronyms	xiii
Summary	xvii
Samenvatting	xxi
1 Introduction	1
1.1 Quantum computing.	1
1.2 Towards large-scale quantum computers	3
1.2.1 A cryogenic controller	4
1.3 Motivation and objectives	6
1.4 Challenges.	6
1.4.1 Classical controller functionality and its temperature of operation.	6
1.4.2 Cryogenic electronic technology	7
1.4.3 Lack of specifications for qubit controllers.	7
1.4.4 Unavailability of design tools	8
1.5 Organization of the dissertation	8
2 A Review of Quantum/Electrical Interfaces	11
2.1 Quantum Processor Requirements	12
2.1.1 Superconducting Qubits	13
2.1.2 Single-electron Semiconductor Qubits	14
2.1.3 Multiple-electron Semiconductor Qubits	16
2.1.4 Controller specifications.	17
2.2 Review of State-of-the-Art Controllers.	18
2.2.1 Superconducting Qubits	18
2.2.2 Semiconductor Qubits.	19
2.2.3 Controller Performance versus Signal Requirements.	22
2.3 Challenges in Scaling-up	22
2.4 Scalable Electronic Controller	23
2.4.1 Tailor-made room-temperature controllers.	23
2.4.2 Multiplexing solutions	24
2.4.3 Cryogenic Controllers	26
2.5 Conclusions	29
3 Deriving the specifications for the electrical control of quantum processors	31
3.1 A System-Level View of a Quantum Computer	32
3.1.1 The quantum processor	32
3.1.2 The classical electronic controller	34

3.2	Methods	35
3.3	Signal Specifications for Single-Qubit Operations	36
3.3.1	Fidelity of a single-qubit operation	36
3.3.2	Specifications for the idle operation and qubit frequency multiplexing	42
3.3.3	Case study for a single-qubit operation.	45
3.4	Signal Specifications for Two-Qubit Operations.	48
3.4.1	Fidelity of a two-qubit operation	48
3.4.2	Specifications for the idle operation	53
3.4.3	Case study for a two-qubit operation	53
3.5	Signal Specifications for Qubit Read-Out.	56
3.5.1	Fidelity for qubit read-out.	56
3.5.2	Case study for qubit read-out	61
3.6	Discussion.	63
3.7	Conclusion	65
4	A toolset for qubit/electronic co-simulation and co-design	67
4.1	Co-design Methodology	68
4.2	Toolset Implementation	69
4.2.1	Simulation of Quantum Physics	69
4.2.2	Hamiltonian Simulations	70
4.2.3	SPINE.	71
4.3	Design Examples	73
4.3.1	Optimization of Power Consumption	73
4.3.2	Finding Optimal Control Waveform	75
4.3.3	System-Level Verification	76
4.4	Future perspectives	76
4.5	Conclusion	78
5	System design of a qubit microwave generator	79
5.1	Requirements	79
5.1.1	System Specifications	80
5.1.2	Extending to Transmons	82
5.1.3	Summary	82
5.2	System Architecture	83
5.2.1	Analog/RF Section	83
5.2.2	Digital Signal Synthesis	84
5.2.3	Final Architecture	85
5.3	Circuit Specifications	86
5.3.1	Sample rate	87
5.3.2	Reconstruction filter	87
5.3.3	Digital blocks	87
5.3.4	Total digital system	93
5.3.5	Analog blocks	94
5.3.6	Power consumption estimate.	97

5.4	Application Example	98
5.4.1	Qubit Tune-Up	98
5.4.2	Multi-Qubit Simultaneous Excitation	100
5.5	Conclusion	100
6	The design of Horse Ridge	103
6.1	Challenges	103
6.2	System architecture & Specifications	106
6.3	Digital Circuit Design	108
6.4	Analog & RF Circuit Design	109
6.4.1	Reconstruction Filter	110
6.4.2	Digital-to-Analog Converter	111
6.4.3	Variable Gain Amplifier	113
6.4.4	Mixer	114
6.4.5	Output driver	115
6.4.6	Auxiliary circuits	117
6.5	Cryogenic Electrical Performance	120
6.5.1	Measurement setup	120
6.5.2	Electrical characterization	122
6.5.3	Comparison with state-of-the-art	125
6.6	Conclusion	127
7	Benchmarking Horse Ridge with Qubits	129
7.1	Setup	129
7.2	Single-Qubit Experiments	131
7.2.1	Rabi oscillation experiment	131
7.2.2	Ramsey-style experiment	132
7.2.3	Pulse shaping experiment	133
7.3	Single-Qubit Performance Benchmark	134
7.3.1	Two qubit Rabi oscillation experiment	134
7.3.2	AllXY and Quantum State Tomography experiment	134
7.3.3	Randomized benchmarking experiment	134
7.4	Two-Qubit Experiments	136
7.4.1	Coupled qubits resonance driving experiment	136
7.4.2	Two-qubit quantum algorithm experiment	136
7.5	Conclusion	139
8	Conclusion	141
8.1	Research overview	141
8.2	Main contributions	145
8.3	Future work	145
A	Derivations of Qubit Control Specifications	149
A.1	General Derivations	149
A.1.1	Numerical Simulations	150
A.1.2	Equivalent Noise Bandwidth	150

A.2	Derivations for Single-Qubit Operation	151
A.2.1	Inaccuracies	152
A.2.2	Quasi-static Noise	154
A.2.3	Noise Filtering	154
A.2.4	Jitter	157
A.2.5	Idle Gate	157
A.2.6	Frequency Multiplexing	159
A.3	Derivations for Two-Qubit Operation	160
A.3.1	Hamiltonian Eigenenergies	160
A.3.2	The C-Phase Gate	163
A.3.3	The Exchange Gate	164
A.3.4	Idle Gate	165
A.3.5	The noise sensitivity	166
A.4	Qubit Read-out	168
B	SPINE User Manual	171
B.1	Introduction	171
B.1.1	Simulation Platforms	171
B.2	SPINE	172
B.2.1	Hamiltonian Simulation	172
B.2.2	SPINE Simulator Core	173
B.2.3	Functions	173
B.3	Implemented Spin Systems	175
B.3.1	<code>system_spin</code>	177
B.3.2	<code>system_1_singlet_triplet</code>	178
B.3.3	<code>system_dispersive_readout</code>	178
B.4	Helper Functions	179
B.4.1	Plotting in MATLAB	179
B.4.2	Plotting in C++	180
B.5	Implemented Solvers	181
B.5.1	<code>solver_analytical_xz</code>	181
B.5.2	<code>solver_expm</code>	182
B.5.3	<code>solver_diagonalization</code>	182
B.5.4	<code>solver_taylor</code>	182
B.5.5	<code>solver_taylor_sparse_approx</code>	182
B.6	Examples	182
B.6.1	<code>example_1_spin</code>	183
B.6.2	<code>example_1_spin_rwa</code>	183
B.6.3	<code>example_2_spin_1_singlet</code>	183
B.6.4	<code>example_2_spin_2_singlet</code>	183
B.6.5	<code>example_2_spin_1_singlet_triplet</code> and <code>example_2_spin_2_singlet_triplet</code>	183
B.6.6	<code>example_n_spin_n_singlet</code>	183
B.6.7	<code>example_1_singlet_triplet</code>	184
B.6.8	<code>example_dispersive_readout</code>	184

B.7	Templates	184
B.7.1	MATLAB	184
B.7.2	C++	185
B.8	Verilog-A	186
C	Horse Ridge measurement setup	189
C.1	Horse Ridge Motherboard	189
C.2	Generic Measurement Setup	191
C.3	Dipstick Design	193
C.4	Fridge Integration	195
C.5	FPGA Functionality	198
	Acknowledgements	227
	Curriculum Vitæ	231
	List of Publications	233
	Chip Gallery	237

Acronyms

2DEG	Two-Dimensional Electron Gas
AC	Alternating Current
ADC	Analog-to-Digital Converter
APR	Automated Place-and-Route
APS	Arbitrary Pulse Sequencer
AWG	Arbitrary Waveform Generator
BGA	Ball Grid Array
BJT	Bipolar Junction Transistor
BPF	Band-Pass Filter
BW	Bandwidth
CML	Current-Mode Logic
CMOS	Complementary Metal Oxide Semiconductor
DAC	Digital-to-Analog Converter
DC	Direct Current
DDR	Double Data Rate
DDS	Direct Digital Synthesis/Synthesizer
DNL	Differential Non-Linearity
DPDT	Dual Pole Dual Throw
DRAG	Derivative Reduction by Adiabatic Gate
DRAM	Dynamic Random Access Memory
DSB	Double-Sideband
EDSR	Electric Dipole Spin Resonance
ENBW	Equivalent Noise Bandwidth
ENOB	Effective Number of Bits
ESR	Electron-Spin Resonance
FDMA	Frequency-Division Multiple Access
FET	Field-Effect Transistor
FPGA	Field-Programmable Gate Array
FTW	Frequency Tuning Word
GRAPE	Gradient Ascent Pulse Engineering
HD3	Third-order Harmonic Distortion
HEMT	High-Electron-Mobility Transistor
HPF	High-Pass Filter
I/O	Input/Output

IC	Integrated Circuit
IM2	Second-order Intermodulation Distortion
IM3	Third-order Intermodulation Distortion
INL	Integral Non-Linearity
IRR	Image Rejection Ratio
JFET	Junction Field Effect Transistor
JPA	Josephson Parametric Amplifier
LDO	Low Dropout
LNA	Low-Noise Amplifier
LO	Local Oscillator
LOR	LO Rejection
LPF	Low-Pass Filter
LSB	Lower Sideband
LSB	Least Significant Bit/Byte
LUT	Look-Up Table
MKL	Math Kernel Library
MOS	Metal Oxide Semiconductor
MOSFET	Metal Oxide Semiconductor Field-Effect Transistor
MRZ	Multiple-Return-to-Zero
MUX	Multiplexer
N-V	Nitrogen Vacancies
NCO	Numerically Controlled Oscillator
NMOS	N-Ch Metal Oxide Semiconductor
NMR	Nuclear Magnetic Resonance
PC	Personal Computer
PCB	Printed Circuit Board
PCIe	Peripheral Component Interconnect Express
PLL	Phase-Locked Loop
PM	Pulse Modulation
PMOS	P-Ch Metal Oxide Semiconductor
PSD	Power Spectral Density
QEC	Quantum Error Correction
QPC	Quantum Point Contact
RAM	Random Access Memory
RF	Radio Frequency
RMS	Root Mean Square
RSFQ	Rapid Single Flux Quantum
RWA	Rotating Wave Approximation
SAW	Surface Acoustic Wave

SEM	Scanning Electron Microscope
SET	Single-Electron Transistor
SFDR	Spurious-Free Dynamic Range
SFQ	Single Flux Quantum
SNR	Signal-to-Noise Ratio
SoC	System on Chip
SOS	Silicon-on-Sapphire
SP6T	Single Pole 6 Throw
SPI	Serial Peripheral Interface
SQNR	Signal-to-Quantization-Noise Ratio
SQUID	Superconducting Quantum Interference Device
SRAM	Static Random Access Memory
SSB	Single-Sideband
TDC	Time-to-Digital Converter
TDMA	Time-Division Multiple Access
TX	Transmitter
USB	Upper Sideband
USB	Universal Serial Bus
VGA	Variable Gain Amplifier
VLSI	Very Large Scale Integration
VSG	Vector Signal Generator
VSM	Vector Switch Matrix
WPD	Wilkinson Power Divider
XOR	Exclusive OR
ZOH	Zero-Order Hold

Summary

Quantum computers have gained widespread interest as they can potentially solve problems that are intractable even for today's supercomputers, such as the simulation of quantum systems as initially proposed by Feynman. To achieve such tremendous progress, a quantum computer relies on processing the information stored in quantum bits (qubits), the fundamental units of quantum computation, similar to the bits in a classical (i.e., non-quantum) processor. In general, a quantum processor comprising several qubits is connected through a quantum-to-classical interface consisting of classical electrical circuits for manipulating and reading out their quantum state. The accuracy of an operation on a qubit is characterized using the fidelity, which is 100 % in case of a perfect operation. In practice, a fidelity between 99 % and 99.9 % is typically achieved in today's experiments with less than 100 qubits, limited by non-idealities in the control signals and by the implementation of the physical qubit itself. Since such qubit fidelities are too low to execute relevant quantum algorithms, quantum error correction schemes have been developed which rely on a much larger number of qubits to correct for possible errors, resulting in the need for quantum processors with thousands or millions of qubits. Solid-state qubits promise the large-scale integration required to scale to millions of qubits, as they exploit the lithographic fabrication techniques borrowed from the semiconductor industry. However, state-of-the-art solid-state qubits, such as transmons and single-electron spin qubits, must be typically cooled to cryogenic temperatures in a dilution refrigerator to exhibit quantum behavior. Because of the limited number of qubits (< 100) in state-of-the-art solid-state quantum processors, the classical controller is currently implemented by general-purpose instruments or by tailor-made controllers operating at room temperature connected via several wires to the qubits in the cryogenic chamber. However, these systems must scale to support millions of qubits. While scaling up the underlying quantum processor is already exceptionally challenging, building the electronics required to interface such a large-scale processor is just as relevant and arduous.

To support the scaling to millions of qubits, this dissertation addresses the research question *'how to implement a scalable cryogenic electronic interface for quantum computers?'*. Its superior compactness and reliability, makes a tailor-made controller, operating at cryogenic temperatures physically close to the qubits, the sole practical alternative for large-scale quantum computers (Chapter 1). Whether the power consumption of a cryogenic controller can be compliant with the cooling capabilities of existing refrigerators and whether it is competitive with respect to the heat load of the wires and attenuators that are otherwise used, are among the questions that are further explored in this work. To be able to answer such questions, the design and experimental validation of a scalable cryogenic microwave

signal generator for the control of > 100 single-electron spin qubits and transmons is described in this dissertation.

While a general trend towards tailor-made room-temperature electronics optimized for size, power, and cost, has started, the greatest challenge left to overcome is the wiring complexity imposed by current quantum processors that require about one control signal per qubit, while operating at cryogenic temperatures. The growing number of cables present a heat load to the dilution fridge, incur latencies in the control signals, and reduce the reliability and scalability of the setup. In order to limit wiring complexity, two main approaches can be adopted: multiplexing control and readout signals, or moving the electronic interface closer to the qubits by operating it at cryogenic temperatures (Chapter 2). Both of these approaches have been adopted in this work, combining a CMOS controller operating around 4 K, and controlling a quantum processor at the same or lower temperature using multiplexed control signals.

Another challenge was the lack of clear specifications for the classical interface, as the electronic interface has been typically built from general-purpose components with good enough specifications. However, when designing a scalable controller operating at cryogenic temperatures, one cannot afford to significantly over-design as that would unnecessarily increase the controller's power consumption, thus limiting its scalability. In this dissertation, we methodically derive the effect of circuit non-idealities in the classical controller on the qubit fidelity for all possible operations, i.e., single-qubit gates, two-qubit gates, and readout (Chapter 3). Only with such a full set of specifications potential bottlenecks can be properly identified and optimized controllers be designed. Most interestingly, a case study targeting a 99.9% average gate fidelity confirms that setups using general-purpose instruments are over-designed for such a fidelity. Moreover, mapping the specifications onto already existing room-temperature CMOS circuits revealed that the specifications can be met with an expected power consumption in the order of 1 mW/qubit. Since cooling power in excess of 1 W is available in current refrigerators at temperatures as low as 4 K, that would be practical to target upcoming processors with up to few thousand of qubits.

As a final challenge, there is scarcity of verification methods and tools for the design flow of large-scale quantum computers spanning several technology domains. Specifically, no tool was available supporting both the simulation of classical electronics and quantum systems. To this end, SPINE (Spin Emulator) has been developed, encompassing a Hamiltonian solver that can be directly used in the industry-standard simulator for CMOS circuit design, for direct simulation of the quantum processor along with the electrical circuits (Chapter 4). A quantum/electronic interface co-design methodology is proposed that covers the entire flow from the definition of the target quantum processor till the system-level verification of the electronic interface with that quantum processor. Its effectiveness is furthermore demonstrated in the design of the scalable cryo-CMOS microwave signal generator presented in

this dissertation.

Based on the obtained specifications and the proposed design, the optimal architecture for the microwave signal generator supporting high-fidelity (99.99 %) multi-qubit (32 in a 2 GHz bandwidth) control was chosen (Chapter 5). Such a microwave signal generator is an integral part of the electronic interface, as it is used for performing single-qubit operations in a quantum processor based on single-electron spin-qubits or transmons. Frequency-multiplexing is assumed as starting point to deal with the wiring bottleneck. It is found that the most suitable architecture uses an I/Q-mixer-based single-sideband analog/RF front-end as it is expected to be most power efficient, while allowing high flexibility in the output frequency range, making this solution suitable for multiple qubit technologies. For the back-end, a numerically controlled oscillator (NCO)-based direct digital synthesis (DDS) system is found to be most suitable thanks to the relatively low hardware cost, and the ease with which coherent qubit operations can be ensured. The signal specifications for a 99.99 % fidelity, as obtained using the methods presented in this dissertation, have been translated into system-level requirements for the selected architecture.

Finally, the implementation of the proposed controller in 22-nm FinFET operating at 3 K is presented (Chapter 6). In addition to a bare DDS-based modulator, the proposed digital system includes a digital controller. This extra step is essential to minimize the interface to the room-temperature equipment, while the employed frequency-multiplexing minimizes the interface to the quantum processor. The digital controller can translate qubit gate operations into the microwave signals necessary for the execution of quantum algorithms, following a very basic instruction set for each qubit (with a maximum program length of 2048 instructions). Four controllers, each supporting 32 qubits, are implemented on a single die with an area of 16 mm^2 . Cryogenic measurements showed a power consumption of 1.7 mW/qubit for the analog circuitry, in line with previous expectations, but a large digital power consumption of 330 mW at 1 GHz. The large digital power consumption can be explained by the additionally integrated digital controller, which in the current generation lacks clock gating while many of the memory cells are implemented using relatively power-hungry flipflops, leaving plenty of room for future improvements. The entire 2 to 20 GHz output frequency band was verified to output sufficient power to drive qubits, while showing a typical spurious-free dynamic range (SFDR) $> 42 \text{ dB}$, 3rd-order intermodulation (IM3) $> 47 \text{ dB}$, and signal-to-noise ratio (SNR) of 48 dB in a 25 MHz bandwidth. These results are theoretically sufficient for the targeted 99.99 % average gate fidelity. Overall, this design, supporting a total of 128 qubits, demonstrates that cryogenic CMOS circuits can help solve the interconnect bottleneck between the quantum processor and its control electronics, thus enabling to scale up the number of qubits in quantum computers. The achieved power efficiency (12 mW/qubit) enabled by a digitally-intensive architecture and the frequency multiplexing allows for operating the chip at 3 K within the cooling capabilities of standard cryogenic refrigerators.

The overall performance of the quantum control chip in driving a real quantum processor with two single-electron spin-qubits is also demonstrated (Chapter 7). Experiments using off-resonance microwave bursts demonstrate the effectiveness of the controller's pulse shaping capabilities to suppress qubit rotations due to off-resonance bursts, as required in a frequency-multiplexed setup, an essential ingredient for this design to be considered scalable. The demonstration of both the Deutch-Josza quantum algorithm and experiments typically used during the bring-up of a quantum processor, such as the AllXY and Quantum State Tomography, show that the controller is flexible enough to be used both during multi-qubit quantum algorithm execution, as well as during the bring-up phase thanks to the versatile integrated digital controller. The performed randomized-benchmark experiment shows that the same gate fidelity ($\sim 99.7\%$) is obtained using the cryo-controller and the room-temperature setup using general-purpose equipment, with the infidelity limited by the qubit. While it can not be confirmed with an actual quantum processor that the presented controller supports the targeted 99.99% average gate fidelity, all measurements support the possibility of achieving this fidelity with a quantum processor with higher fidelity. Finally, the small footprint, power consumption, the ability to integrate multiple transmitters on one die, and operation at 3 K, demonstrate the scalability of the presented solution and highlight the promise of cryogenic controllers.

Generally, as in the setup used in this work, the quantum processor is limiting the achievable gate fidelity. However, great progress is being made in various quantum technologies, and the fidelity keeps improving. The control electronics could become the limiting factor for the fidelity, especially if the power consumption of the controller needs to be minimized. The work presented in this dissertation does however show that a low-power controller, operable at cryogenic temperatures, can be implemented while promising sufficient gate fidelity for quantum error correction schemes, thereby demonstrating one of the first steps towards scalable quantum computers.

Samenvatting

Kwantumcomputers hebben wijdverbreide belangstelling gekregen omdat ze mogelijk problemen kunnen oplossen die zelfs voor de hedendaagse supercomputers onoplosbaar zijn, zoals de simulatie van kwantumsystemen zoals oorspronkelijk voorgesteld door Feynman. Om zo'n enorme vooruitgang te boeken, vertrouwt een kwantumcomputer op het verwerken van de informatie die is opgeslagen in kwantumbits (qubits), de fundamentele eenheden in een kwantumberekening, vergelijkbaar met de bits in een klassieke (d.w.z. niet-quantum) processor. Over het algemeen is een kwantumprocessor met meerdere qubits verbonden via een kwantum-naar-klassieke interface die bestaat uit klassieke elektronische circuits voor het manipuleren en uitlezen van hun kwantumtoestand. De nauwkeurigheid van een bewerking op een qubit wordt gekarakteriseerd door de fidelity, die 100 % is in het geval van een perfecte bewerking. In de praktijk wordt een fidelity tussen de 99 % en 99.9 % doorgaans bereikt in de huidige experimenten met minder dan 100 qubits, beperkt door niet-idealiteiten in de stuursignalen en door de implementatie van de fysieke qubit zelf. Omdat dergelijke qubit fidelities te laag zijn om relevante kwantumalgoritmen uit te voeren, zijn er kwantumfoutcorrectieschema's ontwikkeld die afhankelijk zijn van een veel groter aantal qubits om mogelijke fouten te corrigeren, wat resulteert in de behoefte aan kwantumprocessors met duizenden of miljoenen qubits. Solid-state qubits beloven de grootschalige integratie die nodig is om op te schalen naar miljoenen qubits, omdat ze gebruikmaken van de lithografische fabricagetechnieken die zijn geleend van de halfgeleiderindustrie. State-of-the-art solid-state qubits, zoals transmons en single-electron spin-qubits, moeten echter typisch gekoeld worden tot cryogene temperaturen in een speciale koelkast (dilution refrigerator) om kwantumgedrag te vertonen. Vanwege het beperkte aantal qubits (< 100) in state-of-the-art solid-state kwantumprocessors, wordt de klassieke controller momenteel geïmplementeerd door instrumenten voor algemene doeleinden of door op maat gemaakte controllers die werken bij kamertemperatuur en verbonden zijn via verschillende draden naar de qubits in de cryogene kamer. Deze systemen moeten echter worden opgeschaald om miljoenen qubits te ondersteunen. Hoewel het opschalen van de onderliggende kwantumprocessor al een uitzonderlijke uitdaging is, is het bouwen van de elektronica die nodig is om een dergelijke grootschalige processor te besturen, net zo relevant en moeilijk.

Om de opschaling naar miljoenen qubits te ondersteunen, behandelt dit proefschrift de onderzoeksvraag *'hoe kan men een schaalbare cryogene elektronische interface voor kwantumcomputers implementeren?'*. Zijn superieure compactheid en betrouwbaarheid maakt een op maat gemaakte controller, die bij cryogene temperaturen, fysiek dicht bij de qubits, werkt het enige praktische alternatief voor grootschalige kwantumcomputers (Hoofdstuk 1). Of het stroomverbruik van een

cryogene controller compatibel kan zijn met de koelcapaciteiten van bestaande koelkasten en of het concurrerend is met betrekking tot de warmtebelasting van de draden en signaalverzwakkers die anders worden gebruikt, behoren tot de vragen die in dit werk verder worden onderzocht. Om dergelijke vragen te kunnen beantwoorden, wordt in dit proefschrift het ontwerp en de experimentele validatie van een schaalbare cryogene microgolfsignaalgenerator voor de besturing van > 100 single-electron spin-qubits en transmons beschreven.

Hoewel een algemene trend naar op maat gemaakte elektronica voor kamertemperatuur, die is geoptimaliseerd voor grootte, vermogen en kosten, is begonnen, is de grootste uitdaging die nog moet worden overwonnen de bedradingscomplexiteit die wordt opgelegd door de huidige kwantumprocessors die ongeveer één stuur-sigitaal per qubit nodig hebben, terwijl ze bij cryogene temperaturen werken. Het groeiende aantal kabels vormt een warmtebelasting voor de koelkast, veroorzaakt vertraging in de regelsignalen en vermindert de betrouwbaarheid en schaalbaarheid van de opstelling. Om de complexiteit van de bedrading te beperken, kunnen twee benaderingen worden gevolgd: het multiplexen van besturings- en uitleessignalen, of het dichterbij de qubits brengen van de elektronische interface door deze bij cryogene temperaturen te laten werken (Hoofdstuk 2). Beide benaderingen zijn in dit werk toegepast, waarbij een CMOS-controller, die rond 4 K werkt, wordt gecombineerd met een kwantumprocessor op dezelfde of lagere temperatuur en wordt bestuurd met behulp van gemultiplexte stuursignalen.

Een andere uitdaging was het ontbreken van duidelijke specificaties voor de klassieke interface, aangezien de elektronische interface doorgaans is opgebouwd met algemene componenten met voldoende specificaties. Bij het ontwerpen van een schaalbare controller die werkt bij cryogene temperaturen, kan men het zich echter niet veroorloven om aanzienlijk te overontwerpen, omdat dit het stroomverbruik van de controller onnodig zou verhogen, waardoor de schaalbaarheid ervan wordt beperkt. In dit proefschrift leiden we methodisch het effect van niet-idealiteiten van circuits in de klassieke controller op de qubit-fidelity af voor alle mogelijke operaties, d.w.z. single-qubit-poorten, twee-qubit-poorten en qubit uitlezing (Hoofdstuk 3). Alleen met zo'n volledige set specificaties kunnen potentiële knelpunten goed worden geïdentificeerd en kunnen geoptimaliseerde controllers worden ontworpen. Het meest interessante is dat een casestudy gericht op een 99.9% gemiddelde gate-fidelity bevestigt dat opstellingen die algemene instrumenten gebruiken, overontworpen zijn voor een dergelijke fidelity. Bovendien bleek uit het in kaart brengen van de specificaties op reeds bestaande CMOS-circuits bij kamertemperatuur dat aan de specificaties kan worden voldaan met een verwacht stroomverbruik in de orde grootte van 1 mW/qubit. Aangezien een koelvermogen van meer dan 1 W beschikbaar is in de huidige koelkasten bij temperaturen van slechts 4 K, zou dat praktisch zijn om toekomstige processors met maximaal enkele duizenden qubits te besturen.

Als laatste uitdaging is er een schaarste aan verificatiemethoden en hulpmiddelen voor het ontwerpen van grootschalige kwantumcomputers die verschillende tech-

nologiedomeinen bestrijken. Er was met name geen tool beschikbaar die zowel de simulatie van klassieke elektronica als kwantumsystemen ondersteunt. Zodoende is SPINE (Spin Emulator) ontwikkeld, met een Hamiltoniaan oplosser die direct kan worden gebruikt in de industriestandaard simulator voor CMOS-circuitontwerp, voor directe simulatie van de kwantumprocessor samen met de elektronische circuits (Hoofdstuk 4). Er wordt een co-design-methodologie voor kwantum/elektronische interfaces voorgesteld die de volledige ontwerpfase dekt, van de definitie van de beoogde kwantumprocessor tot de verificatie op systeemniveau van de elektronische interface met die kwantumprocessor. De effectiviteit ervan wordt verder aangetoond in het ontwerp van de schaalbare cryo-CMOS-microgolfsignaalgenerator die in dit proefschrift wordt gepresenteerd.

Op basis van de verkregen specificaties en het voorgestelde ontwerp is de optimale architectuur gekozen voor de microgolfsignaalgenerator die hoge fidelity (99.99 %) en meerdere qubits (32 in een 2 GHz bandbreedte) ondersteunt (Hoofdstuk 5). Een dergelijke microgolfsignaalgenerator is een integraal onderdeel van de elektronische interface, omdat deze wordt gebruikt voor het uitvoeren van single-qubit-bewerkingen in een kwantumprocessor op basis van single-electron spin-qubits of transmons. Frequentie-multiplexing wordt als uitgangspunt genomen om het bedradingsknelpunt aan te pakken. Het is gebleken dat de meest geschikte architectuur een op I/Q-mixer gebaseerde analoog/RF-front-end met enkele-zijband modulatie gebruikt, aangezien deze naar verwachting het meest energie-efficiënt is, terwijl een hoge flexibiliteit in het uitgangsfrequentiebereik mogelijk is, waardoor deze oplossing geschikt is voor meerdere qubit-technologieën. Voor de back-end blijkt een op een numeriek gestuurde oscillator (NCO)-gebaseerde directe digitale synthese (DDS) systeem het meest geschikt te zijn dankzij de relatief lage hardwarekosten en het gemak waarmee coherente qubit-bewerkingen kunnen worden gegarandeerd. De signaalspecificaties voor een 99.99 % fidelity, verkregen met behulp van de methoden die in dit proefschrift worden gepresenteerd, zijn vertaald naar systeemvereisten voor de geselecteerde architectuur.

Ten slotte wordt de implementatie van de voorgestelde controller in 22-nm Fin-FET, en werkend bij 3 K, gepresenteerd (Hoofdstuk 6). Naast een minimale, op DDS gebaseerde, modulator, omvat het voorgestelde digitale systeem een digitale controller. Deze extra stap is essentieel om de interface naar de apparatuur op kamertemperatuur te minimaliseren, terwijl de gebruikte frequentiemultiplexing de interface naar de kwantumprocessor minimaliseert. De digitale controller kan qubit-poortbewerkingen vertalen in de microgolfsignalen die nodig zijn voor de uitvoering van kwantumalgoritmen, volgens een zeer eenvoudige instructieset voor elke qubit (met een maximale programmalengte van 2048 instructies). Vier controllers, die elk 32 qubits ondersteunen, zijn geïmplementeerd op een enkele chip met een oppervlakte van 16 mm². Cryogene metingen toonden een stroomverbruik van 1.7 mW/qubit voor de analoge circuits, in lijn met eerdere verwachtingen, maar een groot digitaal stroomverbruik van 330 mW op 1 GHz. Het grote digitale stroomverbruik kan worden verklaard door de extra geïntegreerde digitale controller, die in de

huidige generatie geen klokpoorten heeft, terwijl veel van de geheugencellen zijn geïmplementeerd met relatief stroomverslindende flipflops, waardoor er voldoende ruimte is voor toekomstige verbeteringen. De volledige uitgangsfrequentieband van 2 tot 20 GHz is geverifieerd om voldoende vermogen te leveren om qubits aan te sturen, terwijl het een typisch spurious-free dynamic range (SFDR) > 42 dB vertoont, 3^e-orde intermodulatie (IM3) > 47 dB, en signaal-ruisverhouding (SNR) van 48 dB in een 25 MHz bandbreedte. Deze resultaten zijn theoretisch voldoende voor de beoogde 99.99 % gemiddelde gate-fidelity. Al met al toont dit ontwerp, dat in totaal 128 qubits ondersteunt, aan dat cryogene CMOS-circuits kunnen helpen bij het oplossen van het knelpunt tussen de kwantumprocessor en de besturingselektronica, waardoor het aantal qubits in kwantumcomputers kan worden vergroot. De bereikte energie-efficiëntie (12 mW/qubit), mogelijk gemaakt door een digitaal-intensieve architectuur en de frequentiemultiplexing, maakt het mogelijk om de chip op 3 K te laten werken binnen de koelcapaciteiten van standaard cryogene koelkasten.

De algehele prestatie van de chip bij het aansturen van een echte kwantumprocessor met twee single-electron spin-qubits wordt ook gedemonstreerd (Hoofdstuk 7). Experimenten met off-resonante microgolfbursts demonstreren de effectiviteit van de puls vormende mogelijkheden van de controller om qubit-rotaties als gevolg van off-resonante bursts te onderdrukken, zoals vereist in een frequentie-gemultiplexte opstelling, een essentieel ingrediënt voor dit ontwerp om als schaalbaar te worden beschouwd. De demonstratie van zowel het Deutch-Josza-kwantumalgoritme als experimenten die doorgaans worden gebruikt tijdens het opstarten van een kwantumprocessor, zoals de AllXY en Quantum State Tomography, tonen aan dat de controller flexibel genoeg is om zowel tijdens het multi-qubit-kwantumalgoritme te worden gebruikt, maar ook tijdens de opstartfase dankzij de veelzijdige geïntegreerde digitale controller. Het uitgevoerde randomized-benchmark-experiment laat zien dat dezelfde gate-fidelity ($\sim 99.7\%$) wordt verkregen met behulp van de cryo-controller en de opstelling op kamertemperatuur dat gebruik maakt van apparatuur voor algemeen gebruik, met de fidelity beperkt door de qubit. Hoewel met een echte kwantumprocessor niet kan worden bevestigd dat de gepresenteerde controller de beoogde 99.99 % gemiddelde gate-fidelity ondersteunt, ondersteunen alle metingen de mogelijkheid om deze fidelity te bereiken met een kwantumprocessor met een hogere fidelity. Ten slotte demonstreren de kleine footprint, het stroomverbruik, de mogelijkheid om meerdere zenders op één chip te integreren en de werking bij 3 K, de schaalbaarheid van de gepresenteerde oplossing en benadrukken ze de belofte van cryogene controllers.

Over het algemeen beperkt de kwantumprocessor de haalbare gate-fidelity, zoals in de opstelling die in dit werk wordt gebruikt. Er wordt echter grote vooruitgang geboekt in verschillende kwantumtechnologieën en de fidelity blijft verbeteren. De besturingselektronica kan de beperkende factor worden voor de fidelity, vooral als het stroomverbruik van de controller moet worden geminimaliseerd. Het werk gepresenteerd in dit proefschrift laat echter zien dat een low-power controller, die

werkt bij cryogene temperaturen, kan worden geïmplementeerd en tegelijkertijd voldoende gate-fidelity belooft voor kwantumfoutcorrectieschemas, waarmee een van de eerste stappen naar schaalbare kwantumcomputers is gedemonstreerd.

1

Introduction

1.1. Quantum computing

Quantum computers have gained widespread interest as they can potentially solve problems that are intractable even for today's supercomputers, such as the simulation of quantum systems as initially proposed by Feynman [1]. For instance, by enabling the efficient simulation of quantum systems, quantum computing could help in synthesizing a room-temperature superconductor, which would significantly reduce energy loss in power lines, power generators, and supercomputers [2]. It could also contribute to improving the efficiency of industrial processes, such as nitrogen fixation into fertilizers, which is currently achieved through the more-than-100-year-old Haber-Bosch process and uses up to 5% of the natural gas produced each year worldwide [3, 4]. Furthermore, quantum computers have the potential to solve classical problems with much higher efficiency, such as searching in large datasets using Grover's algorithm [5], factorizing large integers into prime numbers using Shor's algorithm [6], solving optimization problems [7], solving large linear systems, and even addressing privacy concerns by running algorithms without the operator being able to detect what is being computed [2]. The ability to solve currently intractable problems and to significantly accelerate certain computations represents a game-changer with the potential to revolutionize entire industries.

To achieve such tremendous progress, a quantum computer relies on processing the information stored in quantum bits (*qubits*), the fundamental units of quantum computation, equivalent to the bits in a classical (i.e., non-quantum) processor. Qubits can assume the value of one of the basis states $|0\rangle$ and $|1\rangle$, which are equivalent to the '0' and '1' in a classical computer. But, unlike classical bits, they can also exist in a *superposition* of $|0\rangle$ and $|1\rangle$, i.e. be simultaneously $|0\rangle$ and $|1\rangle$. The state of a single qubit is mathematically represented as:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle \tag{1.1}$$

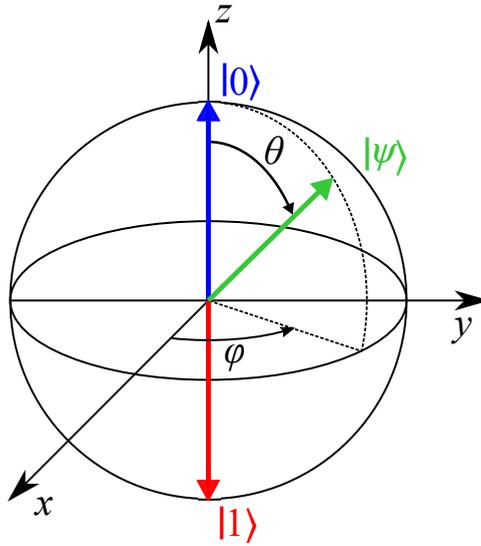


Figure 1.1: The qubit state $|\psi\rangle$ (green) is represented as a point on the surface of the Bloch sphere. The basis states $|0\rangle$ (blue) and $|1\rangle$ (red) are at the north and south pole of the Bloch sphere, respectively. A rotation by an angle θ of the state $|0\rangle$ to obtain the state $|\psi\rangle$ is an example of a single-qubit operation.

where α and β are complex coefficients restricted to $|\alpha|^2 + |\beta|^2 = 1$. All possible single qubit states can be represented by a point on the *Bloch sphere* (Fig. 1.1).

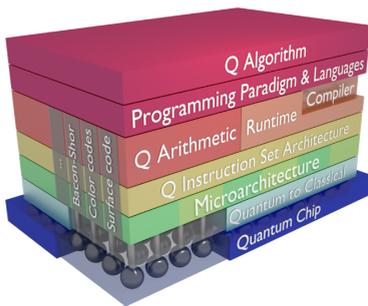
Quantum algorithms, like classical algorithms, are executed by performing operations on the qubits. These operations, or *gates*, fundamentally require *single-qubit operations*, that affect only the state of a single qubit (equivalent to e.g. a classical digital inverter), and *two-qubit operations* where the resulting quantum state depends on the state of two qubits (equivalent to e.g. a classical digital NAND or NOR gate).

Performing an operation on a single qubit is equivalent to performing a rotation of the qubit state in the Bloch sphere, as shown in Fig. 1.1. The accuracy of an actual qubit rotation can be measured by the *fidelity*, which is 100 % in case of a perfect operation. In practice, a fidelity of 99 %...99.9 % is more typical today, limited by non-idealities in the control signals and by the implementation of the physical qubit itself. In comparison, the bit error rate in typical digital data communication is below 10^{-13} . Since such qubit fidelities are too low to execute relevant quantum algorithms, *quantum error correction* schemes, such as surface codes [8], have been developed, which allow for the detection and correction of these errors. By encoding the state of a single *logical* qubit in many *physical* qubits, a logical error rate can theoretically be achieved that is low enough to execute practical quantum algorithms.

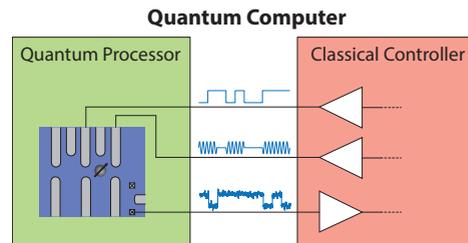
Figure 1.2a shows the generic architecture of a quantum computer [9]. The actual quantum processor, at the bottom of the stack, contains the qubits and it is connected to the quantum-to-classical interface consisting of classical (i.e. non-

quantum) electrical circuits for manipulating the qubits and reading out their quantum state [10]. The remaining upper layers, from the micro-architecture up to the algorithm layer, implement the quantum error correction and ensure proper algorithm execution by controlling the electronic hardware. From a physical perspective, the hardware components of a quantum computer can be fundamentally split into a quantum processor and a classical controller (Fig. 1.2b).

Building a functional quantum computer calls for innovations and developments in each of the functional layers mentioned above, both in the quantum hardware and in the classical hardware, hence representing a formidable scientific and engineering multidisciplinary challenge [11]. As detailed in the next sections, this dissertation addresses the development of the classical hardware.



(a) The layered architecture of a quantum computer (image reproduced from [9]).



(b) A quantum computer comprises a quantum processor and a classical (i.e. non-quantum) electronic controller that generates and processes the various signals going to and coming from the quantum processor.

Figure 1.2: Two representations of a quantum computer.

1.2. Towards large-scale quantum computers

The simplest non-trivial algorithms, such as quantum chemistry problems, require more than 100 logical qubits [12]. Even with the best known algorithms today, roughly 1000 physical qubits are required to implement a single logical qubit with sufficiently low error rate [8]. This requirement translates into the need for thousands or millions of physical qubits. Solid-state qubits promise the large-scale integration required to scale to millions of qubits, as they exploit the lithographic fabrication techniques borrowed from the semiconductor industry. State-of-the-art solid-state qubits, such as transmons and single-electron spin qubits, must be typically cooled to *cryogenic temperatures* in a dilution refrigerator (Fig. 1.3) to exhibit quantum behavior. Since state-of-the-art solid-state quantum processors comprise only a few qubits (< 100 qubits [13–15]), the classical controller is currently implemented by general-purpose instruments or tailor-made controllers operating at *room temperature* [16, 17] connected via several wires to the qubits in the cryogenic chamber. However, these systems must scale to support millions of qubits as required for practical quantum computing applications. While scaling up the underlying quantum processor is already exceptionally challenging, building the electronics required to interface such a large-scale processor is just as relevant and

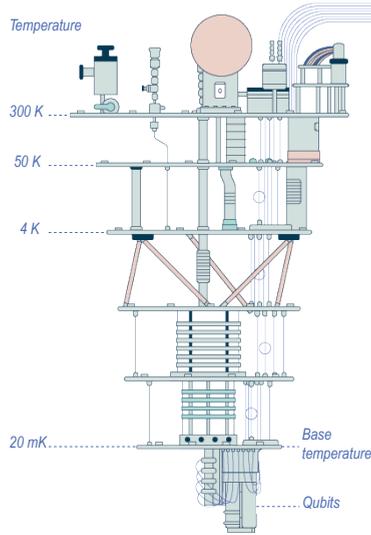


Figure 1.3: A schematic drawing of the cross-section of a dilution refrigerator showing the 4-K plate and the base temperature plate of 20 mK hosting the qubits (Courtesy: Dirma Janse).

arduous.

In the context of this dissertation, a scalable solution, key to achieving a practical quantum computer, refers to a system that can grow to the desired size under practical physical and economical constraints. Moreover, scalability implies that the addition of physical resources raises the performance of the system by a useful amount without introducing an excessively negative impact on the overall system [18].

In case of a large-scale quantum processors with millions of qubits, the implementation of the classical controller with general-purpose room-temperature instruments is impractical and offers limited scalability and reliability due to cost, size, and wiring complexity. A more practical and power-efficient approach could be to use tailor-made electronics optimized for this specific application in terms of power consumption, form factor, and cost [9, 10, 19–27]. A more drastic approach is to operate such tailor-made electronics at cryogenic temperatures, as discussed in the following.

1.2.1. A cryogenic controller

Several arguments can be given for moving the control electronics to cryogenic temperatures (e.g., 4 K) as proposed in Refs. [10, 19–24, 28], each with particular merit, as discussed in the following sections.

Thermal load of the wires and attenuators

The heat conducted through each wire from the warmer to the colder stages in a dilution refrigerator adds to its thermal load. However, a dilution refrigerator has

a very limited cooling power of only a few Watts at 4 K while requiring ~ 10 kW of power to achieve it [29]. While superconducting wires with negligible heat load can be used from the 4-K to the sub-1-K stage, the heat load can be up to ~ 1 mW/wire for coaxial cables from room temperature to the 4-K stage [30]. Thus, connecting each individual qubit with a wire to a room-temperature controller could potentially impose a large heat load on the refrigerator, even exceeding the fridge cooling power for a few hundred qubits.

Moreover, in order to achieve very low noise levels, as required to achieve a high qubit fidelity, today's setups typically take advantage of the reduced thermal noise at cryogenic temperatures by employing signal attenuators inside the dilution refrigerator. Such attenuators reduce the noise power coming from the room-temperature equipment, but require a larger signal power to be generated. Most of this power is dissipated in the attenuators in the dilution refrigerator, and, depending on the signal requirements, the power dissipated in an attenuator can be significantly higher (~ 10 mW) than the heat load of a wire. This further reduces the maximum number of wires with attenuators that can be placed inside the dilution refrigerator, and consequently the number of qubits that can be supported¹.

By using a cryogenic controller to generate the low-noise control signals, this power dissipation can be eliminated, and significantly less interconnect to room temperature will be required. However, the cryogenic controller will instead dissipate power itself, and will likely not solve the problem of heat load on the refrigerator.

Wire latency

Since long wires are required through the sizeable dilution refrigerator, the delay added by the wires to the round-trip time can easily be in the excess of 10 ns (2×1 m at $2/3$ of the speed of light) in case the classical controller is located at room temperature. Assuming typical qubit operation times around 50 ns [13, 15], the added latency of the wiring could potentially be significant in a system where the controller closes a feedback loop. An example of a quantum circuit requiring such a constant feedback loop is the quantum error correction cycle. In each cycle, qubits are read-out after four two-qubit operations, and afterward, a correcting qubit operation is eventually applied depending on the measurement outcome [8], resulting in a ~ 300 ns cycle time. The latency due to the long wires and the latency of the classical controller that processes the readout signals (Chapter 2) must therefore be minimized, e.g. by moving the controller to cryogenic temperatures.

Compactness & reliability

As superconducting wires with negligible thermal conductance can be used between the sub-1-K qubit temperature and the 4-K stage, it is easier to build complex wiring between these two stages than between the 4-K stage and room-temperature without exceeding the fridge cooling capability. Moreover, the interconnect between the base temperature and the 4-K stage can possibly be eliminated in the future by

¹Although the cooling power could improve in the future by, e.g., adopting custom-made refrigerators [31], the power consumption of the system is also expected to increase to serve an increasing number of qubits.

employing 'hot' qubits [32, 33], i.e. qubits that can operate above 1 K and eventually at the same temperature as the controller [30]. The resulting reduction in bulky interconnect when moving the controller to cryogenic temperatures reduces the complexity, size and cost of the system and, more importantly, increases the reliability. Furthermore, solid-state qubits can potentially be co-integrated with semiconductor-based control electronics, either on the same die or at least in the same package, thus altogether eliminating any off-chip interconnect and further enhancing the scalability [30]

1.3. Motivation and objectives

To support the scaling to millions of qubits, this dissertation will address the research question '*how to implement a scalable cryogenic electronic interface for quantum computers?*'. Its superior compactness and reliability, as discussed above, makes a cryogenic controller one of the most practical choices for large-scale quantum computers. Whether the power consumption of a cryogenic controller can be compliant with the cooling capabilities of existing refrigerators and whether it is competitive with respect to the heat load of the wires and attenuators that are otherwise used, are among the questions that are further explored in this work. To be able to answer such questions, in this dissertation, the design and experimental validation of a scalable cryogenic microwave signal generator for the control of > 100 single-electron spin qubits and transmons is described.

1.4. Challenges

Designing a scalable cryogenic electronic interface for quantum computers poses several challenges, as discussed in the following sections.

1.4.1. Classical controller functionality and its temperature of operation

The review of the state-of-the-art (Chapter 2) shows that several solid-state quantum technologies are viable that can support the scaling to a large number of qubits, each with different signal requirements and hence different implications for the design of the electronic interface. Furthermore, a review of the control approaches used today shows a consensus that tailor-made electronics are required to optimize the size and cost of the control system, and to minimize controller latency. However, proposals and implementations exist with different levels of integration in terms of functionality and operating at different temperatures, e.g., controllers with only multiplexing circuits operating at the qubit temperature versus full signal generators operating at 4 K. The lack of a general consensus in the research community makes it unclear what the optimal controller temperature and scaling strategy is.

1.4.2. Cryogenic electronic technology

Several technologies could be used for the design of electronics at cryogenic temperatures. Leveraging cryogenic operation, energy-efficient superconducting devices, such as Rapid Single Flux Quantum (RSFQ)s or Superconducting Quantum Interference Device (SQUID)s, could be used. On the other hand, leveraging, decades of advancements in the semiconductor industry, bipolar and field-effect transistors could also be used. However, it is known that silicon Bipolar Junction Transistor (BJT)s and germanium BJTs have significantly degraded performance at 4 K [34]. Therefore, today, Junction Field Effect Transistor (JFET)s, High-Electron-Mobility Transistor (HEMT)s, and compound semiconductor (InP, SiGe and GaAs) are mostly used for cryogenic operation below 4 K, showing favorable noise, gain and high-frequency behavior over regular Complementary Metal Oxide Semiconductor (CMOS) at cryogenic temperatures [35–37]. However, CMOS technology is currently the only option to allow the integration of billions of transistors on a single-chip, i.e. showing Very Large Scale Integration (VLSI), as required for a complex cryogenic controller capable of supporting millions of qubits, and can operate at least down to ~ 30 mK [38]. For these reasons, this work will focus on the development of cryogenic controllers for quantum processors using CMOS devices operating at cryogenic temperatures, generally referred to as *cryo-CMOS* [22].

While CMOS devices work at cryogenic temperatures, their behavior is affected and foundry models are no longer valid, making the design more challenging. For instance, the threshold voltage increases at cryogenic temperatures, which could give rise to headroom issues in an analog design [39]. While the increased threshold voltage would slow down digital circuits, the concurrent mobility increase at cryogenic temperatures keeps the speed of digital circuits mostly unaffected [39, 40]. Unfortunately, device matching degrades at cryogenic temperatures. This affects high-precision analog and mixed-signal circuits, such as Analog-to-Digital Converter (ADC)s and Digital-to-Analog Converter (DAC)s [41]. However, noise from active devices decreases at cryogenic temperatures, but not as much as expected from the lower operating temperature [42]. The deviation of those parameters, and many others, with respect to their room-temperature behavior create uncertainty on whether standard circuit architectures and design techniques can be straightforwardly applied to cryogenic CMOS circuits and on what changes must be considered.

1.4.3. Lack of specifications for qubit controllers

In order not to degrade qubit performance, the classical controller must provide high-accuracy low-noise control signals, and the read-out must be very sensitive and not induce any disturbance in the quantum processor, known as kickback, so as not to hinder detection of the weak signals from the quantum processor without altering the qubit states. However, a comprehensive analysis of the impact of controller inaccuracies on the qubit fidelity is missing, and as a result, currently implemented controllers use the best available electronics, as it is known to be good enough. The use of high-performance instruments today results in the fidelity of the quantum operations being limited by the quantum processor [43]. However, as the performance of the quantum processor improves, the classical controller

can become performance limiting. Consequently, it is crucial to understand how the controller impacts the performance of the whole quantum computer in order to derive the specifications required for designing a tailor-made controller (Chapter 3).

When operating the control electronics at cryogenic temperatures, the need for accurate specifications is even more severe. As the available cooling power is limited, the power dissipated by the electronics must be minimized by optimally allocating the available power across the various components of the classical controller. However, carrying out such optimization also demands a clear understanding of the impact of each component on the performance of the quantum computer.

1.4.4. Unavailability of design tools

Appropriate verification methods and tools are required to enable a reliable design flow for large-scale quantum computers spanning several technology domains (see Fig. 1.2a). Thus, a methodology and the related tools for the co-design of classical electronic and quantum systems are of paramount importance [24]. A co-simulation platform is the essential backbone of a co-design toolset. Electrical circuit simulators, such as SPICE and Spectre, are well-accepted industry standards and, equivalently, for the simulation of quantum systems, Hamiltonian solvers, such as QuTIP [44], are available. However, the actual interface between classical electronics and the quantum processor has mostly remained unexplored, and, to the best of the author's knowledge, no tool was available supporting both the simulation of classical electronics and quantum systems until SPINE was developed by the author (see Chapter 4).

1.5. Organization of the dissertation

This dissertation is organized as follows. Chapter 2 presents an overview of the different quantum technologies that promise scaling to a large number of qubits and their signal requirements. State-of-the-art electronic interfaces for quantum processors, both at room temperature and cryogenic temperatures, are described to create an understanding of the controller requirements and of the challenges in scaling such a controller. Chapter 3 presents a comprehensive analysis of the impact of the controller inaccuracies on the qubit fidelity to enable the design of a tailor-made controller optimized for power consumption. It also presents example specifications for high-fidelity qubit control, compares these requirements with what is achieved with general-purpose room-temperature controllers, and assesses the feasibility of a cryogenic controller meeting the same requirements. Chapter 4 proposes a co-design methodology, along with a toolset called SPINE (SPIN Emulator) for the co-simulation of the electrical circuit and a quantum processor based on single-electron spin qubits. By using this tool, designers can optimize their circuits while ensuring qubit performance, and exhaustive verification of the entire quantum computer can be performed. Chapter 5 presents an architecture for high-fidelity multi-qubit control, and translates the signal specifications into system-level requirements. Chapter 6 then discusses the implementation of this controller in cryo-CMOS and presents its electrical performance at cryogenic temperatures,

while Chapter 7 presents the results obtained in qubit control. A conclusion with an outlook to the future follows in Chapter 8.

2

A Review of Quantum/Electrical Interfaces

Due to the limited size of existing quantum processors, their electrical interfaces are not yet the main limitation in the performance of quantum computers. However, they could soon become a significant bottleneck, as both the performance and the complexity of quantum processors advance. This chapter addresses this concern by discussing the challenges in designing a scalable electronic interface for quantum processors. To that end, we discuss the requirements of such electronics when employed in combination with quantum processors fabricated in different qubit technologies (Section 2.1). Since we focus on the need for large-scale quantum computation, the scope will be limited to the qubit technologies that today promise large-scale integration, i.e., solid-state qubits requiring only purely electrical control, such as spin qubits and superconducting qubits. An overview of how state-of-the-art electronic controllers handle the requirements dictated by the quantum processor is given (Section 2.2). Then, the limitations in scaling up such state-of-the-art implementations of the electronic interface are analyzed (Section 2.3), and possible solutions to overcome those hurdles are reviewed (Section 2.4). Finally, the benefits offered by operating the electronic interface at cryogenic temperatures close to the low-temperature qubits are discussed. Although several significant challenges must still be faced by researchers in the field of cryogenic control for quantum processors, this chapter concludes that a cryogenic electronic interface appears the viable solution to enable large-scale quantum computers able to address world-changing computational problems (Section 2.5).

Parts of this chapter have been published in *Microprocessors and Microsystems* [27].

2.1. Quantum Processor Requirements

DiVincenzo *et al.* [45] defined five criteria that a quantum processor must satisfy to qualify for use in a quantum computer. In short, the quantum processor must be scalable and contain well-characterized qubits (I) that can be initialized (II) and can preserve their quantum state for a time¹ T_2^* sufficiently longer than the duration of a quantum operation (III). Finally, it must be possible to operate on the qubits with a universal set of quantum gates, i.e., both single-qubit and two-qubit operations, (IV), and to read out the state of a qubit (V).

In the following, a brief overview of the implementation of the read-out, the single-qubit and the two-qubit gates in different quantum technologies will be given. As the state of a qubit can be represented geometrically as a point on a sphere (Fig. 1.1) and single-qubit operations correspond to rotations along this sphere, these operations will also be referred to as rotations in the following², as commonly done in the literature. Specific attention will be devoted to assess the performance of each technology in terms of coherence time and to analyze their potentials and hurdles for scalability. The technologies under consideration are superconducting qubits (Fig. 2.1), specifically transmons, and spin qubits in semiconductors (Fig. 2.2), specifically single-electron spin-qubits (in quantum dots or bound to donors) and multiple-electron qubits, such as singlet-triplet qubits, exchange-only qubits, and hybrid qubits. Typical values and control methods, as used in current experiments, are reported, but these may shift in the future as the various quantum technologies develop.

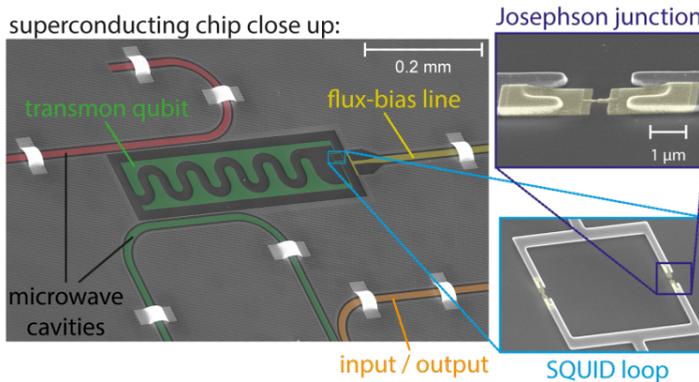


Figure 2.1: A typical transmon qubit, composed of a capacitor in parallel with a SQUID loop comprising two Josephson junctions. A flux-bias line to tune the qubit frequency and the microwave cavities that are used for control and read-out are also shown [46] (image reproduced from Ref. [47]).

¹The coherence time is the characteristic time constant of the process of a qubit losing its proper quantum behavior due to unwanted interactions with the environment.

²To satisfy DiVincenzo's criteria, it is in general sufficient to be able to rotate the qubit state around two axes, typically the X and Z axis.

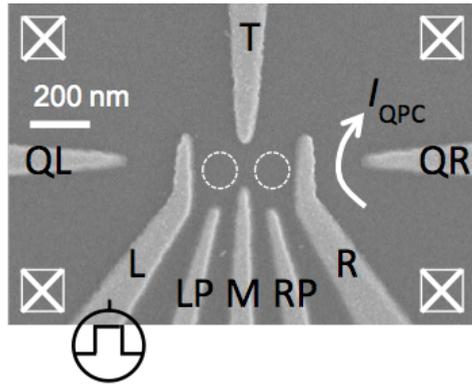


Figure 2.2: A typical semiconductor qubit: two quantum dots (dotted circles) are formed by a set of metallic electrodes depleting the Two-Dimensional Electron Gas (2DEG); electrons can be loaded from the reservoir connected by ohmic contacts (crossed squares); a Quantum Point Contact (QPC) for the read-out and its current (I_{QPC}) is present on the right side (image reproduced from Ref. [48]).

2.1.1. Superconducting Qubits

Various superconducting-qubit implementations have been proposed based on encoding the quantum information in the charge, phase, or flux state of a superconducting resonator. However, the research community has generally converged to the adoption of the *transmon*, a superconducting charge qubit with reduced sensitivity to charge noise, leading to improved coherence times [49–55], e.g., $T_2^* \sim 18.7 \mu\text{s}$ [55].

In a transmon, the qubit states are defined as the two lowest energy levels of a non-linear resonator implemented using Josephson junctions in a loop shunted by a capacitor (Fig. 2.1). A nearby current, called the flux bias, controls the magnetic flux coupled to the loop formed by the Josephson junctions, thus allowing for the qubit energy levels, and hence the qubit frequency, to be tuned externally. The flux-bias lines are typically implemented as transmission lines to allow for rapid changes of the qubit frequency.

Transmons can be fabricated as integrated circuits by depositing and patterning a superconducting film and integrating Josephson junctions on a solid-state substrate, typically made out of silicon or sapphire. Control and read-out are performed by coupling the qubit to microwave cavities on the same chip. Single-qubit rotations that flip the qubit state (X-rotation) can be obtained by applying a microwave pulse with a frequency matching the qubit resonance frequency, typically in the 4-6-GHz range [51, 54]. Microwave pulses with Gaussian envelopes and in/quadrature-phase (I/Q) modulation adopting the Derivative Reduction by Adiabatic Gate (DRAG) approach [56] are used to minimize leakage to higher energy states of the resonator. Typical parameters for the Gaussian envelope of the in-phase component are a 5-ns standard deviation and 20-ns total duration for a π -rotation [54]. As an alternative to a Z-rotation in software [57], a single-qubit rotation around the Z-axis can be obtained by temporarily detuning the qubit frequency by adapting the flux bias [50].

A two-qubit gate can be obtained between two qubits coupled to the same microwave cavity. The cavity is off-resonance with both qubits during regular operation, but, by detuning the qubit frequencies via their flux biases, the cavity can couple to the qubits by virtual photon exchange [50]. In this way, by applying proper pulses to the flux-bias lines of the two qubits, either a SWAP gate or a CPHASE gate can be implemented, depending on the qubit frequencies and whether the flux bias is pulsed suddenly (typically 12 ns [54]) or quickly but adiabatically (typically 40 ns [50, 51, 54]), respectively. Alternatively, a microwave-controlled two-qubit CPHASE gate can be engineered [58, 59].

The qubit state is read out by probing the frequency response of a microwave cavity coupled to the qubit. A frequency shift of the cavity resonance frequency dependent on the qubit state can be observed in the measurement of the transmission coefficient of the cavity. The microwave excitation used for the measurement has a typical frequency in the range of 7-8 GHz with a duration of $\sim 1 \mu\text{s}$ [54].

Superconducting qubits can be reliably fabricated nowadays, having led to quantum processors containing up to 53 qubits, for the first time demonstrating quantum supremacy [13]. However, the individual qubits are currently very large (in the order of $0.1\text{-}1 \text{ mm}^2$ per qubit), which will hinder further scaling to thousands and millions of qubits per quantum processor. Furthermore, superconducting qubits have to operate at extremely low temperatures (typically $< 100 \text{ mK}$). The interested reader is referred to Ref. [60] for more information on superconducting qubits.

2.1.2. Single-electron Semiconductor Qubits

Semiconductor qubits are significantly smaller than transmons, with a typical pitch between neighboring qubits in the order of 100 nm, therefore representing a promising candidate for large-scale quantum processors. As charge qubits in semiconductors show reduced coherence times due to their sensitivity to charge noise, most research is currently focused on spin qubits, in which the quantum state is encoded in the spin of one or more particles, such as nuclei and electrons, trapped in a specific location in a semiconductor solid-state substrate. The decoherence of spin qubits is typically limited by the interaction with nuclear spins of all the other atoms in the substrate, which in turn can be minimized by using isotopically purified silicon [61]. As a major advantage, semiconductor qubits can be fabricated using industry-standard CMOS fabrication processes [62], thus leveraging the very large-scale integration (VLSI) and the yield and the reproducibility of the semiconductor industry. More recently, Germanium, yielding faster transistors due to the high hole mobility, has also demonstrated appealing properties for semiconductor qubits [63]. An additional advantage of semiconductor qubits is that they can potentially be operated at higher temperatures of $\sim 1 \text{ K}$ [30, 32, 33]. However, the required small qubit pitch poses a challenge in routing all control lines individually.

The simplest spin qubit uses the spin of a single electron trapped in a quantum dot [15, 61, 64–66] with the qubit states encoded in the spin-up and spin-down states, as initially proposed by Loss and DiVincenzo [67]. A quantum dot hosting single electrons can be created by a set of gate electrodes locally forming an island in a 2DEG, as shown in Fig. 2.2. Coherence times of $T_2^* = 1 \mu\text{s}$ [15] and $T_2^* = 120 \mu\text{s}$

[61] can be achieved in natural silicon and isotopically purified silicon, respectively.

A large magnetic field in the order of one Tesla is required to obtain a sufficiently large energy difference between the spin-up and spin-down states. The energy splitting between those two states determines the resonant frequency of the qubit. While such a magnetic field can be relatively easily generated and applied, it becomes challenging when a field gradient is desired over many qubits on the same chip to create addressability [15, 65]. Single-qubit rotations across the X-axis can be obtained by applying a varying magnetic field interacting with the magnetic dipole of the electron. A current running in a transmission line nearby the qubit can generate such a field [61]. Alternatively, by applying a microwave excitation to a gate that is capacitively coupled to the quantum dot, the electron can be forced to oscillate in a magnetic-field gradient generated by a local magnet [65]. The required microwave frequency must match the energy difference between the two spin states, and is typically in the range of 13-40 GHz with gate times of 0.2-2 μ s for a π -rotation [15, 61, 65].

Tunnel barriers isolate electrons in neighboring quantum dots during regular operation. For a two-qubit operation, the voltages coupled to the quantum dots or the tunnel barriers can be adapted to allow interaction between two neighboring electrons, thereby implementing either a SWAP gate or a CPHASE gate depending on whether the pulses applied to the gates are sudden or adiabatic (analysis in Section 3.4). The first demonstrations used exchange couplings around 3-10 MHz [15, 66], leading to pulse durations \sim 100 ns.

Read-out is typically performed by converting information on the electron spin into the position of the electron and sensing its charge. For this spin-to-charge conversion, energy-selective tunneling to an electron reservoir [68] or Pauli-exclusion tunneling to another dot [69] can be used. To detect the charge, either a charge sensor, such as a QPC or Single-Electron Transistor (SET), or gate-based dispersive read-out [70] can be employed. In both cases, the read-out involves measuring an electrical impedance: for charge sensors, a variation of a few percents in a resistive impedance (\sim 25 k Ω for QPCs, \sim 100 k Ω for SETs); for gate-based read-out, a variation of \ll 1 fF on a capacitive impedance in the order of 1 pF.

Instead of trapping an electron in a quantum dot, an electron bound to a donor in the semiconductor substrate can be used, e.g., the excess electron provided by a phosphorous dopant atom in silicon [71–74]. While accurately placing the donor atoms poses a fabrication challenge, such a system has the advantage that, unlike quantum dots, all atoms are identical, thus allowing for reproducible behavior. Additionally, the atom nuclear spin can be used as a quantum memory for the long-term storage of a quantum state ($T_{2,\text{nuc}} = 30$ s has been demonstrated [74]). A typical donor-based system shows coherence times of $T_2^* = 55$ ns [71], and $T_2^* = 160$ μ s [74] in purified silicon. The operations on the donor-bound electron are implemented in the same way as for the electron trapped in a quantum dot. While the single-qubit operation times are similar, the qubit frequency in a donor system is typically slightly higher, i.e., in the range of 30-50 GHz [71, 72, 74]. Two-qubit gates have not yet been published since they would require the donors to be placed closer than 20 nm [30], which is currently hard to fabricate reproducibly, and could

hinder the scaling up of the number of qubits.

State-of-the-art quantum processors based on single-electron spin-qubits are currently limited to 2 qubits [15, 32]. The interested reader is referred to Ref. [75] for more information on single-electron semiconductor qubits.

2

2.1.3. Multiple-electron Semiconductor Qubits

The main disadvantage of the single-electron spin-qubits discussed above is that microwave signals are required for single-qubit operations and that these operations are generally much slower than the two-qubit operations based on the exchange interaction. As an alternative, it has been proposed to also use this much faster exchange interaction for single-qubit rotations, but at the cost of requiring multiple electrons to implement one qubit.

In the case of the singlet-triplet qubit, the singlet and triplet configurations of two electrons distributed over two quantum dots are used as the basis states of the qubit [76–82]. A coherence time of $T_2^* = 10$ ns has been demonstrated in a GaAs sample, which can be significantly boosted with, e.g., strong magnetic field gradients ($T_2^* = 700$ ns) [82] or pumping schemes to obtain dynamic nuclear polarization [78]. Single-qubit rotations that flip the qubit state (X-rotations) are then implemented by suddenly pulsing the voltage on the gates close to the quantum dots (similar to two-qubit gates for single-electron spin-qubits), with typical π -rotations of ~ 1 ns [77, 80]. Single-qubit rotations altering the phase of the qubit (Z-rotations) rely on an ever-present magnetic field gradient between the two quantum dots that continuously rotates the qubit. By timing the period between different operations, the desired rotation can be obtained. This rotation is generally slower, e.g., $10\times$ slower in Ref. [80]. Two-qubit gates can be implemented by capacitively coupling two pairs of quantum dots, each encoding a qubit, resulting in an operation time below $1 \mu\text{s}$ [80, 82]. The strength of this capacitive coupling is approximately proportional to the exchange interaction in both qubits. Read-out of the quantum state exploits the Pauli-exclusion principle to convert the quantum state encoded in the spin into a different charge occupation in the quantum dots, which is consecutively measured using a charge sensor.

A drawback of singlet-triplet qubits is that a magnetic field gradient is required between the two quantum dots. The design can be simplified by adding a third electron to represent a single qubit, as is the case for the exchange-only qubit and the hybrid qubit.

In the exchange-only qubit, three quantum dots in a row, each containing a single electron, are used to implement a single qubit [45, 83–87]. Single-qubit rotations around two axes at a 120° angle can be obtained by using the pairwise exchange interaction between the middle quantum dot and one of the outer quantum dots. In Ref. [85], exchange operations with a tunable duration from $0.01 \mu\text{s}$ to $1 \mu\text{s}$ have been measured around both axes, while obtaining a coherence time of $T_2^* = 2.3 \mu\text{s}$ in purified silicon. Proposals exist to implement two-qubit gates using only the exchange interactions between the quantum dots [45, 87].

In the hybrid qubit, the three electrons are placed in only two quantum dots, thereby simplifying the gate geometries and fabrication [48, 88–92]. Single-qubit

rotations around different axes are also obtained by suddenly pulsing the gates. In this way, π -rotations in less than 100 ps have been demonstrated, and coherence times of $T_2^* \sim 10$ ns have been achieved [48]. For the two-qubit gates, it has been shown that in theory fewer exchange-interactions are required between two double quantum dots than in the case of the exchange-only qubit [88].

To date, the presented types of multiple-electron semiconductor qubits have only been demonstrated in quantum processors containing a single qubit.

2.1.4. Controller specifications

Table 2.1: Summary of the typical signal specifications along with examples of the waveforms as required to implement the various operations in different qubit technologies.

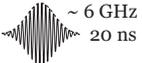
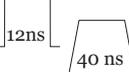
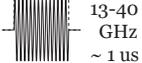
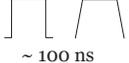
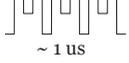
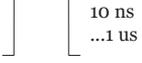
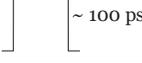
Technology	T_2^*	1-Qubit gate	2-Qubit gate	Qubit read-out	DC-Biasing
Superconducting qubits (Transmons)	18.7 μ s	 ~ 6 GHz 20 ns	 12 ns 40 ns	 7-8 GHz, ~ 1 μ s	flux-bias current
Single-electron spin qubits in a quantum dot	120 μ s	 13-40 GHz ~ 1 μ s	 ~ 100 ns		gate voltage
Single-electron spin qubits in a donor system	160 μ s	 30-50 GHz ~ 1 μ s	 ~ 100 ns		gate voltage
Singlet-triplet qubit	700 ns	 ~ 1 ns	 ~ 1 μ s		gate voltage
Exchange-only qubit	2.3 μ s	 10 ns $\dots 1$ μ s	Sequence of pulses between different quantum dots		gate voltage
Hybrid qubit	< 10 ns	 ~ 100 ps	Sequence of pulses between different quantum dots		gate voltage

Table 2.1 summarizes the electrical signals required for the different qubit technologies to drive the various gate operations and to be acquired for the read-out. In general, it is sufficient for the electronic controller to generate electrical voltage pulses in baseband, e.g., for single and two-qubit gates in singlet-triplet qubits, or at microwave frequency, e.g., for single-qubit gates for transmons. Likewise, the read-out electronics must demodulate pulses at baseband or microwave frequencies. Typically, the fast voltage or microwave pulses are applied directly to a gate, or to a 50- Ω load to prevent reflections in the interconnect. Additionally, but not shown in Table 2.1, the electronics must bias the quantum hardware appropriately, i.e., provide stable voltage bias for all gates in spin-based systems and provide accurate current for the flux bias for transmons. These gates typically present a small (< 1 pF) capacitive load to the controller, while the flux bias line is typically 50- Ω terminated.

On top of providing the required functionalities, the properties of the driving signals must be tightly controlled, both in terms of accuracy in the frequency, the amplitude, and the duration, and in terms of low electrical noise levels, in order not to degrade the inherent fidelity of the qubits. As the fidelity achieved in current qubit technologies is too low for algorithm execution, Quantum Error Correction (QEC) ([8]) can be used to improve the fidelity by encoding the quantum state on a larger number of physical qubits. QEC requires a minimum fidelity of the physical qubits and employs continuous cycles of qubit measurement with feedback control. Since feedback must be applied well within the qubit coherence timescale for QEC to be effective, fast qubit measurements and low controller latency are required. However, most state-of-the-art experiments do not apply QEC and feedback control, and, as a result, long measurement times (< 1 ms) are typically used to allow filtering of the noise and thereby increase the read-out fidelity in semiconductor qubits. For this reason, Table 2.1 does not report measurement times for semiconductor qubits, although QEC would ask for a read-out time in the order of 100 ns - 1 μ s. The control fidelity is inherently limited by non-idealities in the quantum processor but can be further limited by non-idealities in the applied control signals.

The effect of signal non-idealities can be either analytically evaluated (Chapter 3) or simulated (Chapter 4). Surprisingly, the electrical requirements for each qubit technology have not yet been systematically analyzed, resulting in quantum-computing experimentalists to overdesign their electronic setups, as described in the following section. In Chapter 3 a complete analysis of the electrical specification for single-electron spin qubits is presented, which shows a methodology that can, in principle, be adapted to any qubit technology.

2.2. Review of State-of-the-Art Controllers

In this section, we will show how state-of-the-art experimental setups generate the control signals and acquire the read-out signals, as required by existing quantum processors.

2.2.1. Superconducting Qubits

Figure 2.3 shows the experimental setup used in Ref. [54] to control a five-qubit transmon chip. A flux-bias line for each transmon is biased with a Direct Current (DC) current to set the frequency of the respective qubit.

For two-qubit gates, pulses on the flux lines are generated by an Arbitrary Waveform Generator (AWG) (here the 4-channel 1.2 GS/s Tektronix AWG5014) and added to the DC component via wideband 50- Ω bias-tees. For single-qubit gates, microwave signals need to be applied to the input port of the respective cavity. The microwave carrier from a Vector Signal Generator (VSG) (here the Agilent E8257D) can be I/Q modulated by an AWG to generate the required DRAG pulses, e.g., for qubit D_b in Fig. 2.3. Additionally, the AWG can also provide a Single-Sideband (SSB) modulation to drive a qubit at a frequency different from the carrier frequency provided by the VSG, as shown for all other qubits in Fig. 2.3.

The tones required for qubit read-out are generated in the same way, and all res-

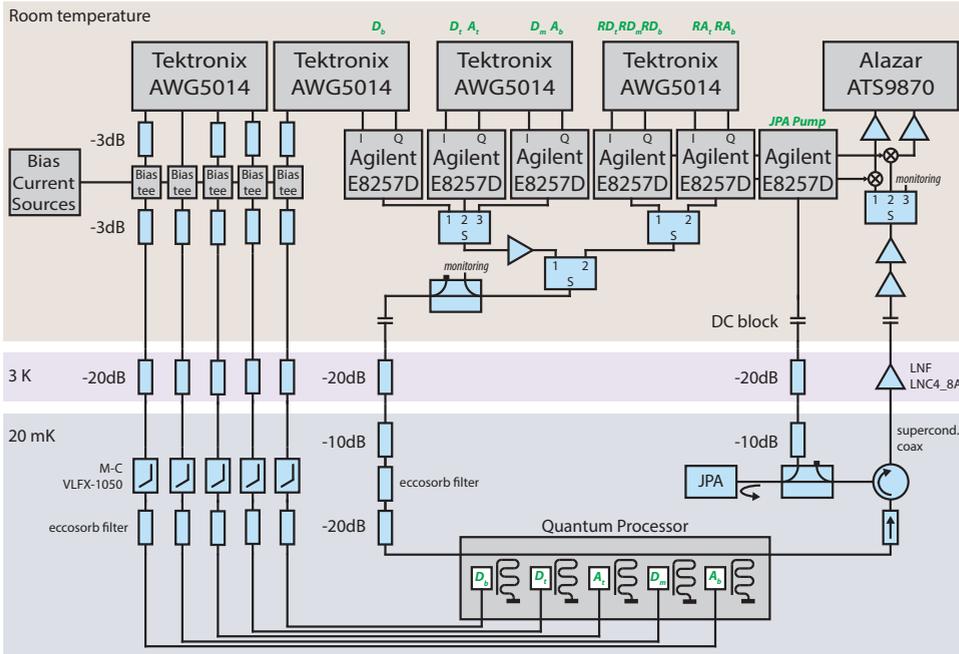


Figure 2.3: Example of the electrical interface typically used for superconducting transmon qubits [54].

ulting carriers are combined before entering the cryogenic refrigerator. However, for the read-out, the carrier generated by the VSG is also fed into the acquisition circuit to down-convert the signal received from the cavity. A Josephson Parametric Amplifier (JPA) [93] at base temperature and a cryogenic Low-Noise Amplifier (LNA) [94] are used to improve the Signal-to-Noise Ratio (SNR) to allow for shorter measurement times. The pump signal required by the JPA is provided by a VSG. Additional gain stages are used at room temperature before down-converting the signal and digitizing it using a high-speed ADC card. Additional demodulation and signal processing are performed on a Personal Computer (PC).

In addition to the components described above, several attenuators and filters are used at different temperature stages of the refrigerator both to minimize the heat conducted through the cables (heat conductivity is generally proportional to electrical conductivity for non-superconducting metals) to the cryogenic chamber and to reduce the noise (to take advantage of the significantly lower thermal noise at cryogenic temperatures). A more in-depth explanation can be found in 1.2.1.

2.2.2. Semiconductor Qubits

Due to the similarities in the different semiconductor qubit setups, we will only consider a setup for single-electron spin qubits in quantum dots and note where other implementations differ. Figure 2.4 shows the experimental setup used in Ref. [15] to control a two-qubit single-electron spin-qubit chip.

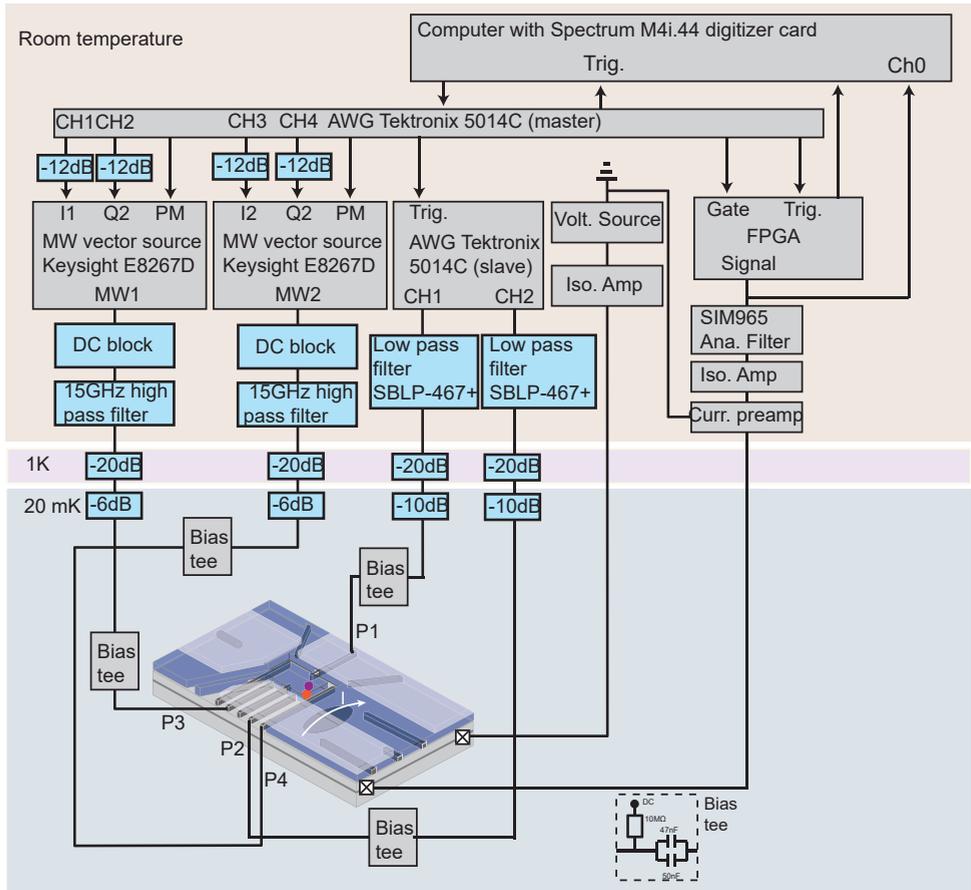


Figure 2.4: Example of the electrical interface typically used for single-electron spin qubits in quantum dots (image adapted from Ref. [15]).

Similar to superconducting qubits, DC biasing of the quantum processor's various gates is required. While not shown in Fig. 2.4, all gates, including the top gates and the ones for the SET charge sensor, are connected to the proper DC voltage generator.

The microwave signals required for the single-qubit gates are again generated by a VSG (the Keysight E8267D), which is modulated using an AWG (the Tektronix 5014C). Additional Pulse Modulation (PM) control is implemented to turn on the microwave carrier generation only when required to minimize signal leakage during idle periods.

The voltage pulses on the gates, as required for qubit initialization, read-out and two-qubit gates, are provided by an AWG connected to the gate via a low-pass filter and a bias-tee. The pulse generation of this AWG is controlled using an external trigger provided by the master AWG that ensures proper synchronization.

For multiple-electron spin qubits that leverage the high operation speeds offered by the exchange interaction, the speed of an AWG might be insufficient to provide pulses that are short enough or have sufficient timing resolution for accurate operations. In that case, dedicated pulse generators can be used, e.g., the Agilent 81134A pulse pattern generator [48].

For the read-out, the impedance of either an SET (as in Fig. 2.4) or a QPC is modulated by the qubit state and can be directly read by measuring the device current when biased at a fixed voltage (Fig. 2.4) [95, 96]. The applied signal conditioning consists of (cryogenic) amplification and filtering before the signal is digitized. Unfortunately, the bandwidth and hence the speed of this measurement is limited by capacitive parasitics due to the wiring connecting the quantum device to the amplifiers. Alternatively, Radio Frequency (RF)-reflectometry is also commonly used to increase the bandwidth of the charge-sensor impedance measurement [97, 98]. An example using this method is shown in Fig. 2.5 [99]. The nominal impedance of the charge sensor is adapted to $50\ \Omega$ by a matching network closely connected to the charge sensor. By sending an RF pulse and measuring the reflected power steered by a directional coupler, changes in the impedance are monitored. Instead of using a charge sensor, dispersive read-out of the quantum dot's capacitance can be used to detect changes in the dot occupation, which uses a setup similar to the RF-reflectometry read-out discussed above [70].

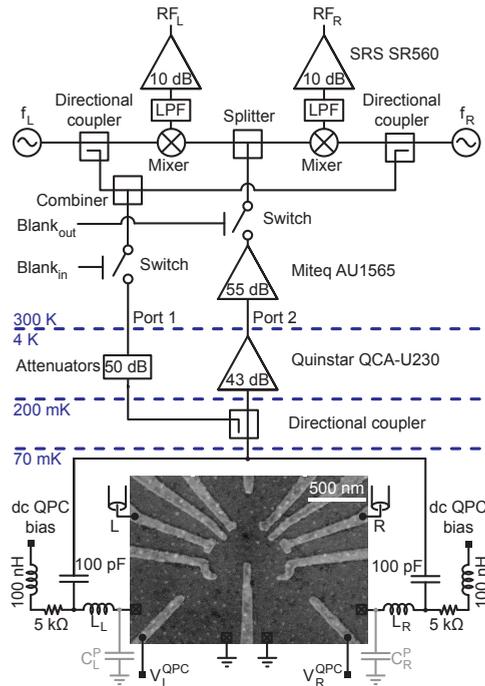


Figure 2.5: Example of a setup using RF-reflectometry of a QPC to sense the charge occupation of a triple quantum dot (image reproduced from Ref. [99]).

2.2.3. Controller Performance versus Signal Requirements

The use of general-purpose instruments, such as VSGs and AWGs, offers high flexibility, wide tuning ranges, and fast prototyping. Hence, it is often preferred in research laboratories working on qubit development, as the qubit requirements may be evolving during experimentation. However, to obtain accurate quantum operations as required to perform any quantum algorithm, the instrument specifications should be carefully evaluated. In Chapter 3, we will consider, as a representative example, the generation of the microwave pulses in the setup of Fig. 2.4, which uses the Tektronix 5014C AWG [100] and the Keysight E8267D VSG [101]. It is found that the performance of the analyzed setup is sufficient to drive single-qubit operations in an ideal quantum processor with a 99.9% fidelity. More in general, it can be concluded that typically adopted general-purpose instruments are not limiting the fidelity of state-of-the-art quantum computers, since solid-state semiconductor qubits with a single-qubit average gate fidelity exceeding 99.9% have been only recently demonstrated [43]. While this situation is acceptable for current developments focused on improving the performance of quantum processors, the substantial performance margin in the electronic interface may not be tolerated as the performance and the scale of quantum processors improve, as discussed in Section 2.3.

2.3. Challenges in Scaling-up

State-of-the-art controllers for quantum processors described in Section 2.2 are well suited for current qubit experiments but will show severe limitations when scaling up to the thousands of qubits as required for any practical quantum computer.

First, the general-purpose instruments used in current experiments are expensive, large, and power-hungry while only supporting a handful of qubits [15, 54]. Although this can be partially alleviated by sharing the hardware over multiple qubits, e.g., as in Fig. 2.3 (more details in Section 2.4.2), this can induce crosstalk effects between qubits, e.g., undesired phase shifts on unaddressed qubits [54, 102]. A more scalable solution is the adoption of tailor-made electronics ensuring the required specifications and optimized for size, power, and cost. Multiple tailor-made controllers that have already been employed will be further discussed in Section 2.4.1. Tailor-made controllers can also be optimized for speed, thus avoiding issues arising in feedback control due to general-purpose instruments generally being too slow. For instance, the Tektronix 5014C AWG (see Figs. 2.3 and 2.4) has a delay longer than 500 ns between the arrival of a trigger and the generation of the actual output signal [100].

Another limitation is that, at the moment, every qubit requires individual control lines connecting the quantum processor in the refrigerator to the instruments at room temperature, e.g., the flux-bias lines for transmons and the gates for spin qubits. Since the number of lines that can physically fit in a dilution refrigerator is limited, the wiring can become a severe bottleneck in scaling-up. Moreover, more connections, especially between stages at different cryogenic temperatures, cause the system to become more complex, more expensive, and less reliable. Furthermore, the heat conducted through each line from the warmer stages to the

colder ones, together with the power dissipated in the various signal attenuators, adds to the thermal load of the dilution fridge, thus increasing the requirements on its cooling power. Lastly, because of the size of the typical dilution fridge and because of the large number of bulky instruments, the cables are long enough to cause delays that cannot be neglected in every case. For instance, a 1-meter cable results in a round-trip time of ~ 10 ns, which can be comparable to the duration of the quantum operations (Table 2.1).

In order to limit the wiring complexity, two main approaches (or a combination of them) can be adopted [30, 103]: multiplexing the control and read-out signals over a reduced set of wires (Section 2.4.2), and moving the electronic interface closer to the qubits and operate it at cryogenic temperatures (Section 2.4.3).

2.4. Scalable Electronic Controller

2.4.1. Tailor-made room-temperature controllers

Recently, various tailor-made quantum-processor controllers have been developed [9, 25, 26, 104–106], as urged by the demand for highly integrated and flexible solutions for the rapidly evolving quantum processors. These solutions generally use a reconfigurable Field-Programmable Gate Array (FPGA) at their core to maximize flexibility. Low delay is of utmost importance for quantum algorithms requiring feedback control. Direct processing of the read-out and control signals in an FPGA has a significant advantage in terms of latency compared to any software solution, while keeping the flexibility of in-field programmability. Additionally, integrating the full controller around an FPGA eases the synchronization of the various components. A list of controllers and their capabilities is shown in Table 2.2. Moreover, commercial general-purpose instruments are evolving to fulfill the needs of quantum control, e.g., the Quantum Researchers Toolkit by Keysight offers a scalable modular control platform with a latency below 150 ns [17] and Zurich Instruments' series of HDAWGs a latency below 50 ns [16].

Table 2.2: Comparison of tailor-made room-temperature controllers.

Controller	Qin <i>et al.</i> [25]	Fu <i>et al.</i> [9]	Ryan <i>et al.</i> [26]
FPGA	Virtex-7	Cyclone V	Virtex-6
Host interface	Universal Serial Bus (USB) 2.0	USB	1 Gb Ethernet
Peripherals	2 \times Time-to-Digital Converter (TDC) 12 \times pulse generator	2 \times 200 MS/s 8-bit ADC 8 \times marker outputs	2 \times 1 GS/s 12-bit ADC 2 \times 1 GS/s 16-bit DAC
Slave modules with each:	1 DAC board 2 \times 1 GS/s 16-bit DAC	3 AWG modules 2 \times 200 MS/s 14-bit DAC	up to 9 Arbitrary Pulse Sequencer (APS) 2 modules 2 \times 1.2 GS/s 14-bit DAC 4 \times marker outputs
Latency	-	80 ns (AWG only)	428 ns (+ 110 ns cabling)

Controller	Salathé <i>et al.</i> [104]	Lin <i>et al.</i> [105]	Ofek <i>et al.</i> [106]
FPGA	Virtex-4	Xilinx	Virtex-6
Host interface		1 Gb Ethernet	Peripheral Component Interconnect Express (PCIe)
Peripherals	1 \times 100 MS/s 14-bit ADC	2 \times 2 GS/s 16-bit DAC	2 \times 1 GS/s ADC 2 \times 1 GS/s DAC
Slave modules with each:			
Latency	283 ns (+ 69 ns cabling)		200 ns

All electronic interfaces employ an FPGA as the central controller and include hardware for both the generation of signals and an ADC to digitize the read-out signals. Such hardware can be either tailor-made and assembled on separate slave boards or comprise general-purpose instruments as in the setups described in Section 2.2. The read-out signals are processed in the FPGA. In general, this includes signal demodulation, filtering, decimation, signal integration, and state discrimination. Special techniques can be used to reduce the latency, such as using frequency-selective kernels to combine the demodulation, filtering, and integration in a single processing stage [26], or using digital mixing and multiplier-less filters [104].

On the control side, additional functionality can be added on top of straightforward DACs, such as micro-operations [9]. The APS proposed in Ref. [26] provides additional Direct Digital Synthesis/Synthesizer (DDS) to implement SSB modulation, thus allowing multiple qubits to be controlled by a single VSG. Multiple Numerically Controlled Oscillator (NCO)s are implemented in the FPGAs to control the frequency offset and keep track of the qubit phase.

One common feature of the controllers in Table 2.2 is operation at room temperature, resulting in additional latency due to long wiring. Furthermore, these controllers still do not fully replace all required hardware for qubit interfacing, e.g., microwave signal generation is not included.

2.4.2. Multiplexing solutions

Multiplexing can help in reducing the number of wires but comes with additional constraints as summarized in this section for time-division and frequency-division multiplexing.

Time-division multiplexing In Time-Division Multiple Access (TDMA), the same line is alternately connected to different qubits through a switch. In the control architecture proposed in Ref. [19], calibrated control waveforms are applied to a single line ('prime-line') that is switched over multiple qubits depending on the addressing information supplied by another set of lines ('address-line'). The proposed switch matrix is controlled by an FPGA operating at 4 K and placed at the base-temperature stage of the dilution fridge (20 mK). Switches can be implemented with either GaAs HEMTs or a proposed capacitive switch. In addition to the basic on/off switching, phase and amplitude modulation can also be implemented in the switch matrix by tuning the impedance of each switch [19]. This allows for the generation of specific signals from a limited set of waveforms provided from room temperature, thus representing the simplest example of a cryogenic electronic controller. Similarly, a Vector Switch Matrix (VSM) operating at room temperature is proposed in Ref. [107] to switch single-qubit-gate control pulses over multiple transmons. Additional electronics included with each switch enables the tuning of both the phase and amplitude so that the signals to each qubit can be individually calibrated to take into account qubit fabrication inhomogeneity. In order to further reduce the number of lines, the address can be sent over a serial interface [103].

In Ref. [108], the addressing scheme is extended to a 2D array using word and bit lines as in classical memories. A crossbar addressing scheme, similar to the

use of word and bit lines, can theoretically also be directly used to control a large grid of qubits without the need for additional devices like switches, as proposed in Ref. [30, 109, 110]. While this is an effective method to reduce the amount of interconnects, it requires a certain level of homogeneity in the qubit grid to obtain sufficiently accurate operations. A proposal that attempts to tackle this problem is given in Ref. [110].

Time-division multiplexing may require the terminals of the quantum devices to remain floating for a given period. In order to keep such a terminal properly biased, as required, e.g., for gates in spin qubits, a sample-and-hold capacitor can be used to store the bias voltage at the gate and must be only periodically refreshed, similar to a Dynamic Random Access Memory (DRAM) cell (Fig. 2.6) [108, 111–114]. However, typically a relatively bulky capacitor in the order of 100 fF...10 pF is required to achieve low enough voltage fluctuations, limited by either charge quantization or sampling noise [30, 110].

While time-division multiplexing increases the scalability of the system, it strongly limits the amount of parallelism in the execution of quantum gates. As a result, not only the execution of the quantum algorithm becomes longer, but also the qubits will be idle for extended periods during which decoherence takes place, as discussed in detail in Ref. [103], thereby affecting qubit fidelity. Furthermore, such multiplexing often requires a switch for each terminal of the quantum device. For qubits demanding a small pitch, such as spin qubits requiring ~ 100 nm pitch between quantum dots, integrating even a single switching device, i.e., a transistor, per qubit will be extremely challenging, as the needed density for the switching devices is higher than what is currently offered by the most advanced CMOS technologies [115].

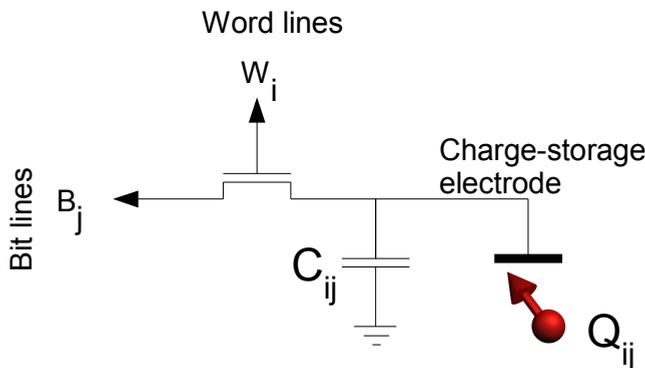


Figure 2.6: Gate biasing using a local storage element and time division multiplexed access via word and bit lines (image reproduced from Ref. [30]).

Frequency-division multiplexing In Frequency-Division Multiple Access (FDMA), different frequency bands on the same line are allocated for the simultaneous use

of different sub-blocks. As each superconducting transmon can be designed and tuned to a different resonance frequency at which it is sensitive (Section 2.1), frequency multiplexing can be directly used for control. In Ref. [116] a proposal is shown on how this can be leveraged to efficiently perform the surface code cycle in a scalable fashion. In Fig. 2.3, the read-out of different transmons also uses a different frequency band thanks to the use of different read-out resonators tuned to different frequencies and coupled to the same on-chip transmission line [117].

Similarly, for single-qubit gates, single-electron spin qubits are only sensitive to signals at a given frequency, which is determined by the externally applied magnetic field. Frequency differences among neighboring qubits, as required for FDMA, can be introduced using the Stark shift, i.e., by an electrical field provided with an extra confinement gate [61, 66], or by generating magnetic-field gradients by integrating on-chip micromagnets [15, 65, 118, 119]. For an RF-reflectometry-based read-out, the use of FDMA is straightforward, thanks to matching networks tuned to different frequencies [98], as shown in Fig. 2.5.

In order to limit the crosstalk between different channels, sufficient inter-channel frequency spacing is required. In the case of qubit control, the crosstalk causes the Alternating Current (AC)-Stark shift, i.e., the qubit frequency appears to shift in the presence of signal power at a nearby frequency. As a result, either complex pulse shaping is required [120], or the channels have to be time-multiplexed with the addition of a phase correction for all idling qubits after each operation [54]. Because of practical limitations on the available spectrum, and the need for sufficient channel spacing, a maximum number of channels of only 10-100 is expected for frequency-multiplexed readout [121]. Besides crosstalk, the application of frequency multiplexing in large-scale processors may be limited by the need for bulky frequency-selective components, such as waveguide resonators for transmons and matching networks in RF-reflectometry of spin qubits, although for the read-out this may be alleviated by crossbar approaches [30].

2.4.3. Cryogenic Controllers

In order to further reduce the wiring complexity and latency, placing the electronic controller close to the qubits, and hence operating it at cryogenic temperatures inside the dilution refrigerator, has been proposed [10, 20, 24, 30, 121–123]. Ideally, the electronics should operate directly at base temperature next to the quantum processor, as discussed in Section 1.2.1. However, as the cooling power of typical refrigerators is lower at colder temperatures [124], only limited functionality could be implemented at sub-K temperatures. Some researchers anticipate that the power dissipation of the electronics can be reduced to a level compatible with the cooling power of existing refrigerators at sub-K temperatures, thanks to the optimization of both the design and the microelectronic fabrication. Although Ref. [123] estimates that ~ 100 singlet-triplet qubits can be controlled by electronics fabricated in a commercial 65-nm CMOS process operating at 100 mK, most researchers advocate for the bulk of the electronic interface to operate around 4 K (Fig. 2.7), i.e., a temperature at which dilution refrigerators have significantly more cooling power (up to a few Watt). However, while a controller at 4 K solves the problems related

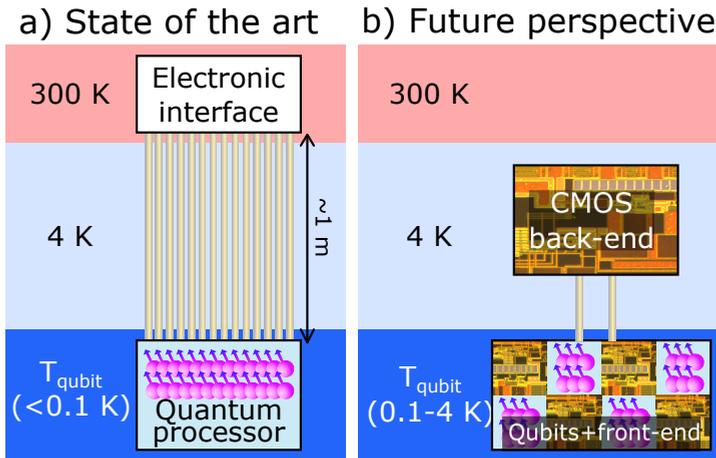


Figure 2.7: Compared to a state-of-the-art room-temperature electronic interface (left), implementing the controller inside the dilution refrigerator (right) can significantly reduce the wiring complexity. The wiring from the 4-K stage to the qubit stage can be further reduced using electronics with limited complexity to implement some form of multiplexing near the qubits.

to the wiring to room temperature, interfacing to a quantum processor operating at sub-K temperatures remains a challenge. Relatively high-power electronics could be placed close to the quantum processor if thermally-isolating interconnect (such as superconducting through-silicon vias [125]) could be employed to avoid heating the qubits. Using superconducting wires between the 4-K electronics and the quantum processor would drastically reduce the heat load on the sub-K stage [126, 127], but would not alleviate complexity and reliability issues, which could otherwise be circumvented by multiplexing at base temperature (Section 2.4.2). Alternatively, preliminary studies promise the feasibility of some qubit technologies, spin qubits more specifically, operating at higher temperatures [32, 33]. In that case, both the quantum processor and its electronic interface could be placed at the same temperature (1-4 K) [30].

Several electronic technologies have proven their functionality down to deep cryogenic temperatures, and can hence be used to implement the controller. These include JFET, HEMT, superconducting devices based on Josephson junctions (e.g., Single Flux Quantum (SFQ) logic [128]), compound semiconductors (e.g., GaAs) and CMOS transistors [36, 37]. CMOS has shown functionality down to 30 mK [22, 38], and it is the most advanced among these technologies. Thanks to the push of the semiconductor industry, CMOS offers an established design automation infrastructure and the possibility to integrate billions of transistors on a single chip, as would be required to interface millions of qubits. Consequently, it is generally considered the most promising choice for the integration of a cryogenic electronic interface for quantum processors.

Metal Oxide Semiconductor (MOS) transistors are fully functional at cryogenic temperatures, but their performance is different with respect to their standard tem-

perature range. Improvements appearing at cryogenic temperatures include increased mobility, and hence larger maximum current, higher subthreshold slope, lower leakage, and lower thermal noise [129, 130]. As a drawback, the threshold voltage increases, thus leading to less voltage headroom, and flicker noise performance and device matching degrade [10, 131]. As a major difference with respect to more mature CMOS technologies, advanced nanometer CMOS technologies are not affected by critical cryogenic non-idealities, such as current kink or hysteresis [129, 132]. However, accurate device modeling is required even for nanometer CMOS technologies to enable the design of complex optimized cryogenic CMOS (cryo-CMOS) circuits [22, 24, 129]. More information on the cryogenic behavior of a nanometer CMOS technology, and the impact on circuit design, can be found in Section 6.1.

Nevertheless, some commercial CMOS Integrated Circuit (IC)s are functional at cryogenic temperatures, well below their target temperature range. Most notably, some FPGAs remain fully functional down to 4 K with a marginal change in operating speed [40], and even DRAM seems functional down to 80 K [133]. The FPGAs could form the basis of a highly reconfigurable cryogenic control platform [20], similar to the current tailor-made controller at room temperature (Section 2.4.1). The ability to program the device in the field can prevent expensive and time-consuming warm-up/cool-down cycles of the dilution fridge during development. In Ref. [40], the flexibility of a cryogenic FPGA platform has been demonstrated with the implementation of an FPGA-based reconfigurable ADC.

While some commercial ICs can be directly used out of the box, the best performance and highest level of integration can only be achieved by custom cryo-CMOS designs. Several components required for the control and read-out of a quantum processor have been designed and optimized for operation at cryogenic temperatures. In Ref. [134], pulse generators for qubit control using either a mixed-signal or a fully-digital implementation have been integrated in a 500-nm Silicon-on-Sapphire (SOS) CMOS process. Both implementations can generate pulses with variable duration down to 10 ns. In Ref. [135], a 6-bit DAC designed in the same technology uses an analog calibration technique to overcome the increase in device mismatch at 4 K. A smaller 4-bit implementation demonstrates a fast rise time of 600 ps at 4.2 K [135]. A 40-nm CMOS digitally controlled oscillator for signal generation at cryogenic temperature has demonstrated state-of-the-art phase-noise performance at 4 K while generating frequencies (5.5-7 GHz) compatible with transmon control [10]. An improved oscillator with automatic calibration to achieve optimal phase noise performance at cryogenic temperatures is presented in [136]. Moreover, the first quantum controller operating at 3 K and integrated in 28-nm CMOS capable of modulating microwave signals for transmon control has been demonstrated [137]. For the read-out, an LNA fabricated in a 500-nm SOS CMOS process shows noise sufficiently low to measure the impedance of an SET in a measurement time of only ~ 520 ns [138]. Another LNA, integrated in a bulk 160-nm CMOS technology for the use in an RF-reflectometry read-out achieves a gain of 57 dB and bandwidth of 500 MHz at 4 K with an in-band noise figure of 0.29 dB (20 K noise temperature) [10, 139]. Parametric amplification using cryo-CMOS has been

shown in [140], and a cryogenic integrated circulator for qubit read-out is presented in [141]. Lastly, voltage references, as required in a fully-integrated cryogenic controller, have been demonstrated that work over the entire temperature range from 4 to 300 K [142, 143].

Besides these circuits, which are specifically designed for qubit control and read-out, various circuits operating at cryogenic temperatures exist for applications ranging from space missions to high-energy-physics but are out of the scope of this review. Furthermore, while cryo-CMOS has many advantages, other technologies show superior performance in terms of noise (e.g., SiGe) or power consumption (e.g., SFQs) and could be a better alternative in some specific applications, e.g., for the use in LNAs and FPGAs [128], respectively. The co-integration of a quantum processor and its classical controller based on a non-CMOS cryogenic technology has already been demonstrated for an SFQ-based controller [144].

2.5. Conclusions

While state-of-the-art setups are well capable of fulfilling the requirements set today by the different quantum technologies, they will be subject to severe limitations when quantum processors will scale up. By using tailor-made electronic controllers at room temperature, the cost, size, and power consumption can be reduced, while also enabling optimization for low-latency feedback. However, to obtain a truly scalable quantum computer, the interconnect complexity between the quantum processor and the electronic interface must be drastically reduced. While multiplexing techniques are useful in this sense, they still present several stringent drawbacks. For instance, in case of time-division multiplexing, operations can not be performed in parallel, and care must be taken to prevent crosstalk in frequency-multiplexed setups. On the contrary, the adoption of cryogenic electronics promises the deployment of truly scalable controllers. These cryogenic controllers can be implemented in standard CMOS technologies in order to exploit large-scale high-yield fabrication capabilities, thus allowing the ultimate vision of co-integrating the electronic interface with the qubits on a single chip or package, or, at least, to operate the electronics close to the quantum processor. Initial steps have already been taken in this direction, thus paving the way to large-scale cryo-CMOS electronic interfaces that will make the operation of future quantum computers addressing real world-changing problems possible. However, in order to realize a complex and power-efficient cryogenic electronic interface, such as Horse Ridge, a good understanding of the impact of controller inaccuracies on the qubit fidelity, and appropriate verification methods and tools are required.

3

Deriving the specifications for the electrical control of quantum processors

As the performance of quantum processors improves, non-idealities in the classical controller can become the performance bottleneck for the whole quantum computer. To prevent such limitation and to enable the design of tailor-made controllers, this chapter presents a systematic study of the impact of the classical electrical control signals on the qubit fidelity.

Analysis of the impact of the controller on the quantum computer's performance has been undertaken previously, but only for specific aspects of the control signals, i.e., the effect of microwave phase noise [145, 146], or for treating the effects of noise on qubit operations either theoretically [147–149] or experimentally [150, 151], e.g., to find sweet spots for the control [152–154]. To close this gap, the work presented here aims to provide a *comprehensive analysis* of the effect of non-ideal circuit blocks in the classical controller on the qubit fidelity for all possible operations, i.e., single-qubit gates, two-qubit gates, and read-out¹. This includes the effect of signal inaccuracies in the frequency, voltage, and time domain, and covers static and dynamic, systematic and random errors. Only with a full set of specifications can potential bottlenecks be identified, and tailor-made electronics be designed. Besides providing a general method for deriving the electronics specifications, the specifications resulting from a case study targeting a 99.9% average gate fidelity are

Parts of this chapter have been published in Physical Review Applied [102].

¹Initialization is assumed to be performed by relaxation or by read-out and will not be separately discussed.

mapped onto existing room-temperature ICs to assess the feasibility of a practical controller.

Although the proposed approach can be easily extended to any quantum technology, such as Nuclear Magnetic Resonance (NMR) [155, 156], ion traps [157, 158], superconducting qubits [159–161] or Nitrogen Vacancies (N-V) centers in diamond [162], we focus on the specific case of single-electron spin qubits, since this qubit technology offers promising prospects for large-scale quantum computing (Sections 2.1.2 and 2.4.3). In addition, while the approach can be generalized to other qubit types, the results obtained for the single-qubit gates can be generalized to any qubit system where single-qubit rotations are performed by applying a signal with a frequency matching the energy-level spacing between the $|0\rangle$ and $|1\rangle$ states. Qubit systems with such single-qubit gates are e.g., NMR [155, 156], ion traps [157, 158], N-V centers in diamond [162], and superconducting qubits [159–161]. Similarly, the results obtained for the two-qubit gates can be generalized to any qubit system that exploits the exchange gate.

This chapter is organized as follows: Section 3.1 describes the generic spin-qubit quantum processor used for the following analysis; Section 3.2 introduces the methods for deriving the fidelity for the various operations; in Sections 3.3 to 3.5 the electrical specifications required for single-qubit operations, two-qubit operations, and qubit read-out are derived, respectively. Detailed derivations along with additional closed-form expressions can be found in Appendix A. A discussion regarding the feasibility of these specifications using existing circuits follows in Section 3.6. Conclusions are drawn in Section 3.7.

3.1. A System-Level View of a Quantum Computer

3.1.1. The quantum processor

A single-electron spin qubit encodes the quantum state in the spin state of a single electron (Section 2.1.2). A generic model of a quantum processor based on single-electron spin qubits is shown in Fig. 3.1, which captures all of the properties relevant for the interaction with the controller. Moreover, although the figure illustrates a linear array of quantum dots, it can be extended to more complex geometries, such as a 2D grid of quantum dots, as shown in Refs. [30, 108, 109].

Quantum dots are formed using a set of gate electrodes that locally deplete a 2DEG on a semiconductor chip, e.g., a GaAs/AlGaAs heterojunction, a Si/SiGe heterojunction, or a Si-MOS structure [75, 163]. Due to the small size of the quantum dot, the charge states become discrete with an energy-level spacing related to the dot charging energy, thereby setting the required increase of the dot potential to add an electron to the dot. The dot potential, and thereby the number of electrons in the dot, is controlled by the plunger gate that capacitively couples to the quantum dot. Without loss of generality for the analysis of the electrical control signals, the following analysis assumes the availability of additional tunnel barrier gates which form tunnel barriers between neighboring dots by controlling the width of the depletion layer, thus allowing tunneling of electrons from and to the quantum dot in a tunable way. Figure 3.1 shows non-overlapping gates as common in early

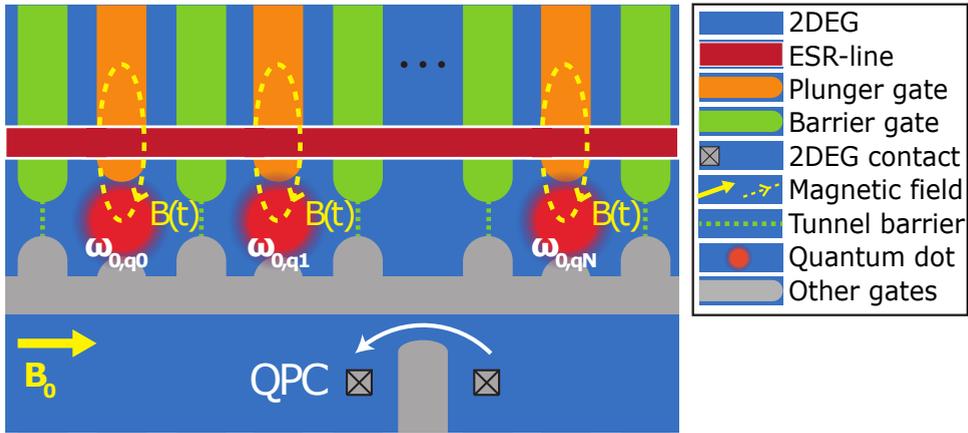


Figure 3.1: A generic model of a spin-qubit quantum processor comprising qubits encoded in the spin of electrons trapped in quantum dots and a charge sensor (e.g., a QPC). The blue background indicates the 2DEG where quantum dots, shown as red circles, are formed locally. Individual control over the dot potentials and the tunnel barriers is assumed, using plunger gates (orange) and barrier gates (green), respectively (other gates are shown in gray for completeness). Furthermore, each qubit can have a unique resonance frequency (Larmor frequency, $\omega_{0,qi}$).

integration schemes [65, 77, 119], while, in order to create better tunability and control, architectures now often include overlapping gates [61, 66, 164, 165]. The analysis in this chapter covers both approaches.

An external static magnetic field \mathbf{B}_0 induces an energy difference between electrons with spin up and spin down, called the Zeeman energy E_z . Because of the static magnetic field, when idle, the electron rotates around the Z-axis in the Bloch sphere with Larmor frequency $\omega_0 = \gamma_e \cdot |\mathbf{B}_0|$, where γ_e is the gyromagnetic ratio of the electron ($\gamma_e \approx 28 \text{ GHz/T}$ in silicon). As indicated in Fig. 3.1, each qubit can have a different Larmor frequency, which can be useful for two-qubit operations or for multiplexing single-qubit operations (Section 2.4.2) [15, 61, 66].

Single-qubit operations (Section 3.3) require the application of a varying magnetic field perpendicular to \mathbf{B}_0 and oscillating at the Larmor frequency. In the case of Electron-Spin Resonance (ESR), such a field is generated by a varying current in a nearby ESR-line [61, 66, 166]. Alternatively, the same effect can be obtained by applying a varying electric field to the electron in a spatial magnetic field gradient, as is the case for Electric Dipole Spin Resonance (EDSR) [62, 65, 118, 119]. In that case, the electric field variations are generated by a voltage on a nearby gate, e.g., through the plunger gate, without requiring an ESR-line. Although Fig. 3.1 shows an ESR line, the results of the analysis below are applicable to both ESR- or EDSR-based operations, as explained in the following section. Two-qubit operations (Section 3.4) and qubit read-out (Section 3.5) can be performed by pulsing the barrier and plunger gates.

The effect of crosstalk between different gates or the ESR-line is considered

negligible or compensated for in the classical controller, and is not further discussed here since it can be treated as a purely classical electrical effect.

Qubit read-out relies on a spin-to-charge conversion followed by the detection of the eventual electron movement [163], using either a gate dispersive read-out [70, 167] or an additional charge sensor. The latter is assumed in this chapter, as it is more straightforward and more commonly used. The extension to gate-based readout is left as future work. For such a charge sensor, different sensing techniques can be used, e.g., a QPC [68, 96, 97] or a SET [168]. As an example, Fig. 3.1 shows a QPC in close proximity to the quantum dots.

3.1.2. The classical electronic controller

The classical controller is responsible for generating the required electrical signals to bias and control each gate and, in case, the ESR line, and for reading the state of the charge sensor. The required electronics have been schematically summarized in Fig. 3.2. The electronic components in the figure are placeholders for the respective functionalities and are grouped by operation, i.e., single-qubit operation, two-qubit operation, and read-out. Thus, they do not necessarily correspond to a physical implementation as discussed in Section 2.2.

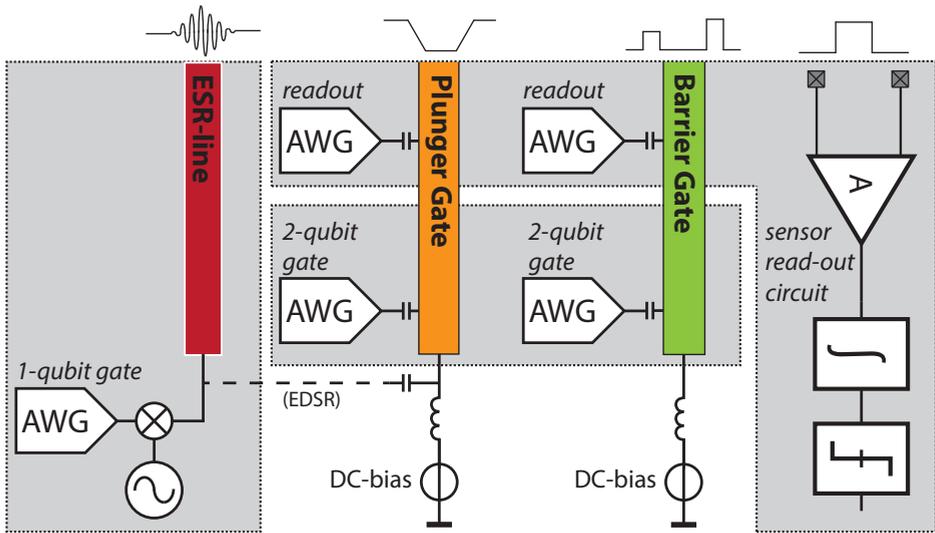


Figure 3.2: The classical control electronics required for each line type (ESR line, plunger gate, and barrier gate) of the quantum processor. AWGs are shown for the envelope and pulse generation.

When no operation is performed, each quantum dot must contain a single electron at the same dot potential, and the tunnel barriers must be tuned to ensure a negligible coupling between neighboring dots (Section 3.4.2). Such conditions are ensured by the use of DC-bias voltage generators, as shown in Fig. 3.2. An ideal LC bias-T is included in the figure to discriminate between purely DC biasing and AC excitations.

The oscillating magnetic field $B(t)$ required for single-qubit operations can be generated by an oscillating current $I(t)$, following the relation $B(t) = \alpha_I \cdot I(t)$ (in the case of ESR), or by an oscillating voltage $V(t)$, resulting in $B(t) = \alpha_V \cdot V(t)$ (in the case of EDSR). The conversion factors α_I and α_V depend on many factors, such as the exact geometry of the structures, and can be determined experimentally. Due to this abstraction, the results shown in Section 3.3 are valid both in the case of ESR and EDSR. The required microwave current or voltage is generated by modulating a carrier from a Local Oscillator (LO) with an envelope produced by an AWG. Although different hardware implementations are possible, this allows us, without loss of generality, to split the carrier specifications, i.e., the LO specifications, from the envelope specifications, i.e., the AWG specifications. In the case in which each qubit has a unique Larmor frequency, a single control line can be used to control multiple qubits independently via FDMA, i.e., frequency multiplexing, thus simplifying the wiring (Sections 2.4.2 and 3.3.2).

The voltage pulses required for the two-qubit gates and the read-out are generated by AWGs. Distinct AWGs are assumed for the two-qubit gates and the read-out since the specifications for such operations can be different.

Besides the presented control electronics, additional hardware is required to process the signal from the charge sensor. The required hardware depends on the employed read-out method, e.g., a direct measurement [96] or RF-reflectometry [70, 97]. As an example, a direct read-out, requiring a read-out amplifier, is shown in Fig. 3.2.

3.2. Methods

The evolution of the qubit state is evaluated by computing the system Hamiltonian (H_{sys}), which is a function of the electrical signals applied by the classical controller, and solving the related Schrödinger equation:

$$i\hbar \cdot \frac{\partial |\psi\rangle}{\partial t} = H_{sys} \cdot |\psi\rangle. \quad (3.1)$$

For static control signals, the Hamiltonian is time-independent and the unitary operation describing the evolution after a time T is trivially $U = e^{-i \cdot H_{sys} \cdot T}$ (assuming $\hbar = 1$).

For dynamic signals, such as for complex signal envelopes, the operation described by the time-varying Hamiltonian $H_{sys}(t)$ is approximated by the product of time-independent components, leading to:

$$U \approx \prod_{n=N}^0 e^{-i \cdot H_{sys}(n \cdot \Delta t) \cdot \Delta t}, \quad (3.2)$$

where Δt is the time step, which must be chosen to be small enough for the required accuracy of the approximation (more details in Section 4.2).

As a benchmark to evaluate how close U is to the ideal operation U_{ideal} as it would be performed by an ideal controller, the process fidelity is computed as

follows [169, 170]:

$$F = \frac{1}{n^2} \cdot \left| \text{Tr} \left[U_{\text{ideal}}^\dagger \cdot U \right] \right|^2, \quad (3.3)$$

for the n -dimensional complex Hilbert space ($n = 2$ for the single-qubit gate and $n = 4$ for the two-qubit gate).

In the case of random dynamic errors, the ensemble average of the fidelity over all realizations is evaluated, following Refs. [147, 148]. When treating random noise, the noise spectrum is relevant, as the operation can be affected differently by noise at different frequencies. The method presented in Refs. [147, 148] is used to evaluate the expected process fidelity, and is outlined in Appendix A.

Detailed derivations of the analytical results presented in this chapter are collected in Appendix A. The analytical results presented in the following are often the result of a series expansion of the fidelity truncated at the second order, as clearly stated in Appendix A. Since this dissertation is focused on high qubit fidelity ($> 90\%$), higher-order processes can be neglected. All presented results have been verified by numerical simulations of the Hamiltonian using an ad-hoc developed simulation tool named SPINE (Chapter 4). Unless otherwise stated, all angles (θ , ϕ etc.) are in radians.

3.3. Signal Specifications for Single-Qubit Operations

In this section, the fidelity of a single-qubit operation in the presence of non-idealities in the drive signal are first derived (Section 3.3.1). Next, the effects of non-idealities that affect the qubit when it is idle are considered (Section 3.3.2). Finally, example specifications for achieving a 99.9% fidelity for both an operation and idle period are given in a case-study (Section 3.3.3).

3.3.1. Fidelity of a single-qubit operation

As explained in Section 3.1, the qubit rotates around the Z-axis due to the applied external magnetic field. Using an LO tuned to a frequency equal to the qubit's Larmor frequency, the qubit phase can be tracked, and the qubit appears to be stationary in the reference frame of the LO. In this rotating frame, Z-rotations by an angle θ_Z can easily be obtained by instantaneously updating the LO's phase in software by an angle θ_Z [57, 155, 156]. For such a software-defined Z-rotation, only the accuracy of the phase update of the LO matters, which is limited by the finite resolution in the phase setting. A phase error $\Delta\phi = \Delta\theta_Z$ (the difference between the applied rotation angle and the intended rotation angle θ_Z) reduces the fidelity of the Z-rotation as follows:

$$F_{Z\phi} = 1 - \frac{1}{4} \cdot \Delta\phi^2. \quad (3.4)$$

In the remainder of this section, we will focus on rotations around the X/Y-axis that are obtained by applying a magnetic field $B(t)$ oscillating at the qubit Larmor frequency ω_0 and with a specific phase, which is generated by applying either a microwave current or a microwave voltage, as explained in Section 3.1.

Table 3.1: The fidelity of a single-qubit operation for inaccuracies and noise in the electrical control signals. In this table, θ is the intended rotation angle, ranging from $-\pi$ to π ; inaccuracies are denoted by a Δ ; noise Power Spectral Density (PSD)s by $S(\omega)$; qubit filter transfer functions by $H(\omega)$; and ω_{\min} denotes the lower limit of integration. Note that in the case of jitter, the full transfer function, with T_{clk} as the reference period, is given, as no Equivalent Noise Bandwidth (ENBW) is defined for a High-Pass Filter (HPF).

Inaccuracy				
Carrier	Frequency	$1 - \frac{1}{2} \cdot [1 - \cos(\theta)] \cdot \left(\frac{\Delta\omega_{\text{mw}}}{\omega_{\text{R}}}\right)^2$		
	Phase	$1 - \frac{1}{2} \cdot [1 - \cos(\theta)] \cdot \Delta\phi^2$		
Envelope	Amplitude	$1 - \frac{1}{4} \cdot \theta^2 \cdot \left(\frac{\Delta\omega_{\text{R}}}{\omega_{\text{R}}}\right)^2$		
	Duration	$1 - \frac{1}{4} \cdot \theta^2 \cdot \left(\frac{\Delta T}{T}\right)^2$		
Noise				
Carrier	Frequency	$1 - \frac{1}{\pi} \int_{\omega_{\min}}^{\infty} \frac{S_{\text{mw}}(\omega)}{\omega_{\text{R}}^2} \cdot H_{\text{mw}}(\omega) ^2 \cdot d\omega$		
	Additive noise	$1 - \frac{1}{\pi} \int_{\omega_{\min}}^{\infty} \frac{S_{\text{add}}(\omega - \omega_0)}{\omega_{\text{R}}^2} \cdot H_{\text{add}}(\omega) ^2 \cdot d\omega$		
Envelope	Amplitude	$1 - \frac{1}{\pi} \int_{\omega_{\min}}^{\infty} \frac{S_{\text{R}}(\omega)}{\omega_{\text{R}}^2} \cdot H_{\text{R}}(\omega) ^2 \cdot d\omega$		
	Duration	$1 - \frac{1}{\pi} \int_{\omega_{\min}}^{\infty} S_{\phi}(\omega) \cdot H_{\text{T}}(\omega) ^2 \cdot d\omega$		
Noise filter properties				
		Type	DC-gain $ H(0) ^2$	ENBW
Carrier	Frequency	Low-Pass Filter	$\frac{1}{2} \cdot [1 - \cos(\theta)]$	$\omega_{\text{R}} \cdot \frac{\pi \cdot \theta }{2 \cdot [1 - \cos(\theta)]}$
	Additive noise	LPF	$\frac{1}{4} \cdot \theta^2 + \frac{1}{2} \cdot [1 - \cos(\theta)]$	$\omega_{\text{R}} \cdot \frac{2 \cdot \pi \cdot \theta}{\theta^2 + 2 \cdot [1 - \cos(\theta)]}$
Envelope	Amplitude	LPF	$\frac{1}{4} \cdot \theta^2$	$\omega_{\text{R}} \cdot \frac{\pi}{ \theta }$
	Duration	HPF	$ H_{\text{T}}(\omega) ^2 = \frac{1}{8} \cdot \frac{\theta^2}{\pi^2} \cdot \frac{T_{\text{clk}}^2}{T^2} \cdot \sin^2\left(\frac{\omega}{2} \cdot T\right)$	

The Hamiltonian describing a single electron under microwave excitation in the laboratory frame is given by the following ($\hbar = 1$):

$$H_{\text{lab}} = -\omega_0 \cdot \frac{\sigma_z}{2} + \gamma_e \cdot B(t) \cdot \frac{\sigma_x}{2}, \quad (3.5)$$

where, here and in the following, σ_x , σ_y , and σ_z are the Pauli matrices. The microwave magnetic field can be described as $B(t) = 2/\gamma_e \cdot \omega_R(t) \cdot \cos(\omega_{\text{mw}} \cdot t + \phi)$. A constant amplitude ($\omega_R(t) = \omega_R$), i.e., a rectangular envelope, is considered unless stated otherwise. In the case of a rectangular envelope, ω_R is the Rabi frequency, i.e., the rotation speed for the single-qubit gate. Note that for more complex envelopes, the resulting specifications for the control electronics can differ, as the sensitivity to specific control parameters can be reduced when employing quantum optimum control, such as Gradient Ascent Pulse Engineering (GRAPE) [171].

For single-qubit rotations, the rotation axis is affected by the matching of the microwave frequency (ω_{mw}) to the Larmor frequency (ω_0) and by the phase of the microwave signal (ϕ), i.e., the carrier signal. The rotation angle ($\theta = \omega_R \cdot T$), on the other hand, is determined by the amplitude of the signal (ω_R) and the duration for which the microwave signal is applied (T), i.e., the signal envelope.

Table 3.1 summarizes the effect of noise and inaccuracy on the fidelity of single-qubit operations. Inaccuracies in the control signal are generally caused by finite instrument resolution and drift². In the case of envelope inaccuracies, the microwave amplitude ($\propto \omega_R$) and duration (T) of the signal together determine the rotation angle ($\theta = \omega_R \cdot T$). Hence, any error in either one leads to an under- or over-rotation, thereby reducing the fidelity. In the case of frequency inaccuracies, a better fidelity is achieved for larger Rabi frequencies, i.e., a larger microwave amplitude and a shorter pulse duration. However, a larger Rabi frequency ultimately requires a sufficiently larger Larmor frequency³, and it is harder to reach the same phase accuracy $\Delta\phi$ at higher LO frequencies.

Next, dynamic changes in the control signal can further limit the fidelity. In the event that such a change occurs on a timescale larger than the operation time, it can be considered a random static error. For a static but random error Δ for which $F = 1 - c \cdot \Delta^2$, the expected fidelity simply follows as $F = 1 - c \cdot \sigma^2$, if Δ follows a Gaussian distribution with standard deviation σ and zero mean (see Appendix A). Hence the equations for the inaccuracy as given in Table 3.1 apply.

In order to treat random noise with spectral content at frequencies higher than the operating rate, the method presented in Refs. [147, 148] is adopted to compute the expected operation fidelity as a function of the noise spectrum. The results for dephasing noise are reproduced from Refs. [145, 147, 148] and are repeated here for completeness, with additional analysis for different rotation angles. The fidelity due to the various noise sources is summarized in Table 3.1, where generally $|H(\omega)|^2$ is the intrinsic qubit filter function, implying that the qubit has a different sensitivity to noise at different frequencies. The amplitude responses of the intrinsic qubit filtering functions for frequency noise and amplitude noise are shown

²While it might be possible to calibrate for drift, this significantly increases the system's complexity.

³Assuming that the Rotating Wave Approximation (RWA) (discussed in Appendix A) should be satisfied.

in Figs. 3.3 to 3.5, with analytical formulas provided in Appendix A. These responses have a LPF characteristic, and their properties, the DC gain and ENBW, are summarized in Table 3.1 and highlighted in the plots as the brick wall approximation of the filter⁴.

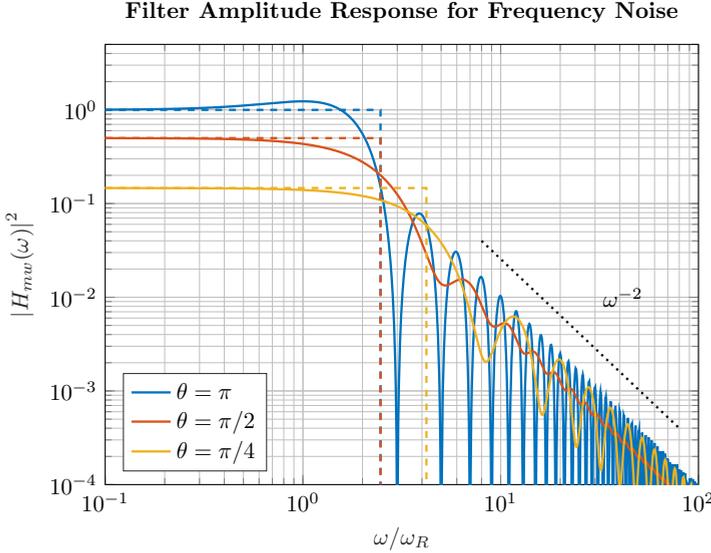


Figure 3.3: The amplitude response of the intrinsic qubit filter for frequency noise for various rotation angles θ . The brick wall approximations are shown with dashed lines.

Note that for frequency noise, the ENBW is proportional to the Rabi frequency, indicating that for faster operations, noise in a wider band affects the qubit. However, the lower limit of integration (ω_{\min}) is inversely proportional to the execution time of the quantum algorithm⁵ and is, therefore, also related to the operation time. In the case of white noise, a good approximation is obtained with $\omega_{\min} = 0$. Due to the factor $\frac{1}{\omega_R^2}$, it is advantageous to use the highest possible Rabi frequency. In the presence of flicker noise, the same conclusion holds, as then a higher ω_{\min} is desirable. Finally, the same discussion holds in the case of amplitude noise. However, a higher tolerance for amplitude noise at larger Rabi frequencies, i.e., larger amplitudes, simply means that the required signal-to-noise ratio ($\omega_R^2/\sigma_{\omega_R}^2$) in the qubit's band of sensitivity is fixed.

The microwave frequency noise ($S_{\omega}(\Delta\omega)$), as discussed before, is set by the phase noise ($S_{\phi}(\Delta\omega)$) of the LO generator and they can be related by $S_{\omega}(\Delta\omega) = \Delta\omega^2 \cdot S_{\phi}(\Delta\omega)$ at a frequency $\Delta\omega$ from the carrier ω_{mw} [145]. While the effect of

⁴The indicated ENBW is defined such that white noise integrated over this bandwidth results in the same noise power as white noise integrated over the actual filter transfer function, and can be used to simplify calculations in case the noise can be well approximated as white noise in the band of interest [172] (see Appendix A).

⁵Note that some quantum algorithms, such as dynamical decoupling sequences or error correction codes, can act as a high-pass filter for the noise, thereby setting ω_{\min} .

Filter Amplitude Response for Wideband Additive Noise

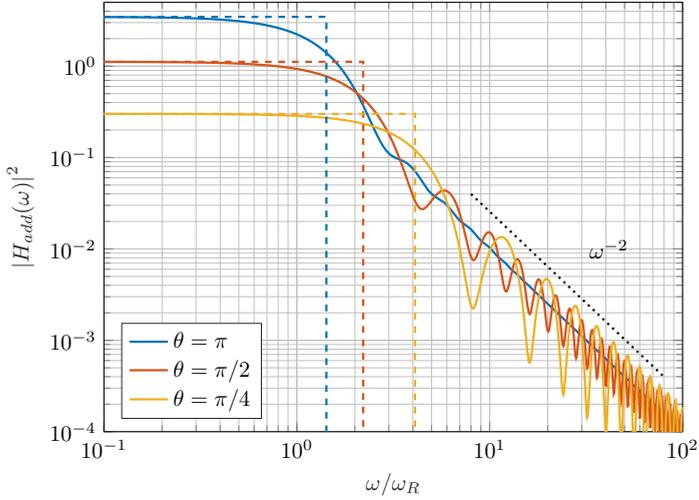


Figure 3.4: The amplitude response of the intrinsic qubit filter for wideband additive noise for various rotation angles θ . The brick wall approximations are shown with dashed lines.

Filter Amplitude Response for Amplitude Noise

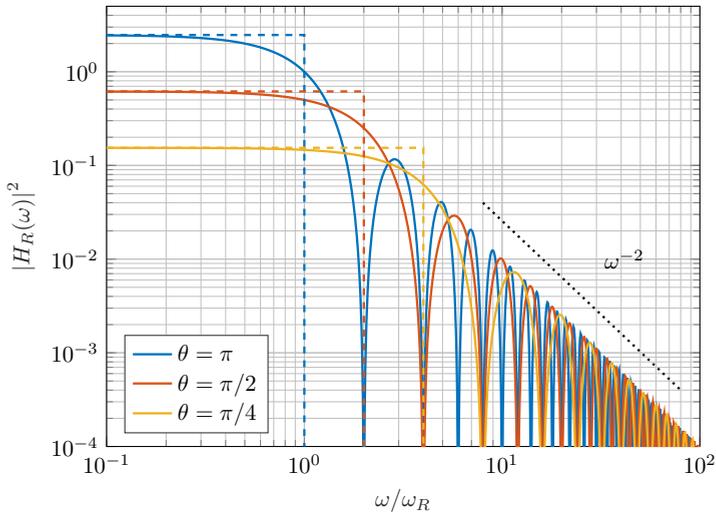


Figure 3.5: The amplitude response of the intrinsic qubit filter for amplitude noise for various rotation angles θ . The brick wall approximations are shown with dashed lines.

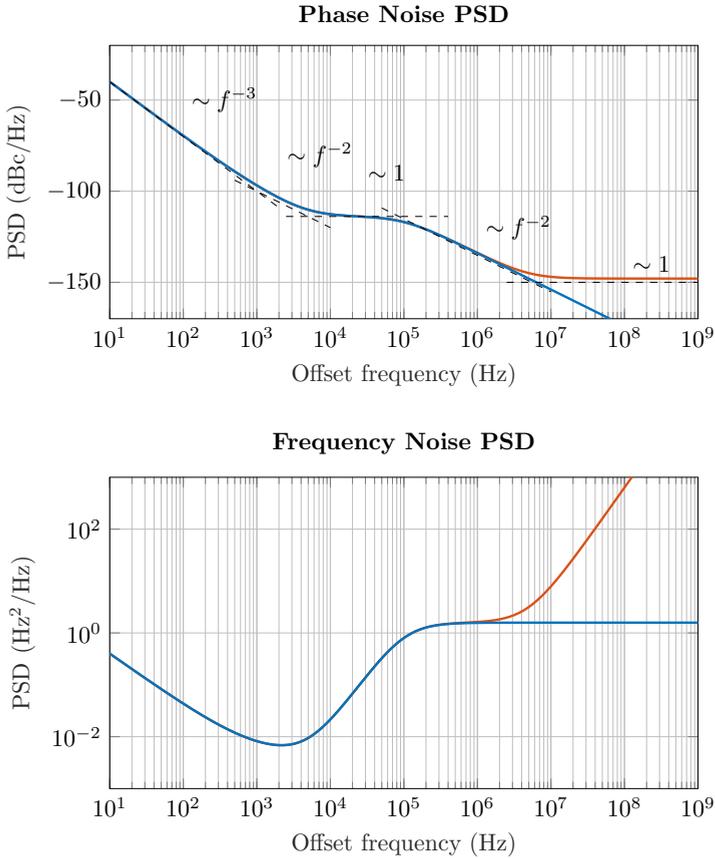


Figure 3.6: A typical plot of (a) the phase noise and (b) the resulting frequency noise PSD of a Phase-Locked Loop (PLL)-based frequency generator. The red line indicates the noise as measured by a phase noise analyzer, whereas the blue line indicates the part of the noise that is actually phase noise. At high offset frequencies, where the lines diverge, wideband additive noise shows up in the phase noise plot, giving rise to a noise floor of around -150 dBc/Hz in this example.

phase noise has already been extensively studied in Ref. [145], a more realistic oscillator noise model, including both phase noise and wideband additive noise, is adopted in this work. Derivation of the intrinsic qubit filtering function for each noise contribution leads to an improved estimation of the fidelity that deviates from Ref. [145], as elaborated in the following.

Consider, as an example, the typically reported plot for the phase noise of a PLL-based frequency generator, as shown in Fig. 3.6 (cf. Ref. [172]). Appended to this figure is a plot of the resulting frequency noise PSD. At low frequencies, the phase noise is typically limited by the flicker noise of the reference clock (the $\sim f^{-3}$ part). In the plot of the frequency noise PSD, this has a f^{-1} roll-off, making it important to maximize ω_{\min} , which could be resolved by using dynamical decoupling schemes, as they introduce an additional high-pass filtering [147, 148, 173, 174]. The part of the phase-noise plot highlighted in red may be a source of concern [145], as it results in a frequency noise increasing as f^2 that exactly cancels the roll-off of the intrinsic qubit filter (Fig. 3.3), thus resulting in a diverging integral for the fidelity (Table 3.1) if no additional band-pass filtering is applied. However, the noise highlighted in red, visible in the phase noise plot originates from thermal noise added to the microwave signal by, e.g., the output driver of the microwave signal generator [172, 175]. The additive noise, with generally a wide bandwidth, is more accurately modeled in the applied microwave magnetic field as follows:

$$B(t) = \frac{2\omega_R}{\gamma_e} \cdot \cos(\omega_{\text{mw}} \cdot t + \phi + \phi_n(t)) + B_{\text{add}}(t), \quad (3.6)$$

where $B_{\text{add}}(t)$ represents the additive noise with PSD $S_{\text{add}}(\omega)$. The actual phase noise $\phi_n(t)$, indicated by the blue line in Fig. 3.6, is clearly band-limited by the qubit filter function due to the absence of the f^2 factor. The PSD of this additive noise has the same frequency dependence as the PSD of the phase noise [176]. As a result, this noise is also filtered by the qubit, without the need for an external filter to limit the far-out phase noise, as suggested by prior art [145]. The fidelity of the qubit operation in the presence of this type of noise is given in Table 3.1.

Finally, the signal duration T is also subject to random variations, i.e., jitter. However, since the period cannot vary during the operation, the noise in the timing can be simply treated considering the period jitter as a quasi-static error. This period jitter is determined by the single-sideband phase noise $S_\phi(\omega)$ of the reference clock (period T_{clk}) used to set the duration [172, 177] (Table 3.1), following a HPF characteristic with the corner frequency set by the duration T .

3.3.2. Specifications for the idle operation and qubit frequency multiplexing

In a typical quantum algorithm, a qubit can be idle for a while, waiting for the operations on other qubits to finish, before being operated on, e.g., due to limitations in the hardware or data dependencies. This section discusses processes that cause the state of the qubit to degrade during an idle period lasting T_{nop} . The loss of the quantum state due to interactions with other qubits will be discussed in Section 3.4.2.

A qubit will perform an undesired Z-rotation (related to T_2^*) in the rotating frame if the microwave frequency is not matched to the qubit's Larmor frequency, even when the driving tone is not applied to the qubit. Evaluating the fidelity of an identity operation in the case of a frequency inaccuracy $\Delta\omega_{\text{mw}}$ leads to $F_{\text{nop,mw}} = 1 - \frac{1}{4} \cdot \Delta\omega_{\text{mw}}^2 \cdot T_{\text{nop}}^2$, which can easily be more stringent than the requirement due to a rotation (Table 3.1).

Besides Z-rotations, unintended X- and Y-rotations of the qubit (related to T_1) are possible in the case in which power is present at the qubit's Larmor frequency. In general, a tone could be present at the qubit frequency, e.g., due to signal leakage from the microwave source or non-linearities in the system leading to harmonic or inter-modulation tones. The presence of a spurious tone that would give a Rabi frequency of ω_{spur} will reduce the fidelity as follows: $F_{\text{nop,spur}} = 1 - \frac{1}{4} \cdot \omega_{\text{spur}}^2 \cdot T_{\text{nop}}^2$.

Besides a tone, residual thermal noise could be present on the driveline. Considering a noise signal with spectral density $S_{R_n}(\omega)$, the fidelity is: $F_{\text{nop,noise}} = 1 - \frac{1}{\pi} \int_0^\infty S_{R_n}(\omega) \cdot |H_n(\omega)|^2 \cdot d\omega$, where

$$|H_n(\omega)|^2 = 2 \cdot \frac{\sin^2\left(\frac{T_{\text{nop}}}{2} \cdot (\omega - \omega_0)\right)}{(\omega - \omega_0)^2}, \quad (3.7)$$

which indicates that the noise spectrum is filtered by a sinc-shaped band-pass filter centered around ω_0 , with the following brick wall approximation:

$$|H_n(\omega)|^2 \approx \begin{cases} T_{\text{nop}}^2/2 & |\omega - \omega_0| \leq \pi/T_{\text{nop}} \\ 0 & \text{elsewhere.} \end{cases} \quad (3.8)$$

Lastly, multiple qubits sharing the same control line, i.e., a single ESR line or control gates shorted together, can be controlled independently in the case in which each qubit has a unique Larmor frequency, using FDMA (Section 2.4.2). However, when rotating a qubit with Larmor frequency ω_0 by applying a microwave signal at frequency $\omega_{\text{mw}} = \omega_0$, any unaddressed qubit on the same line with Larmor frequency $\omega_{0,\text{other}} = \omega_0 + \omega_{0,\text{space}}$ will be affected. Similarly, even if not on the same driveline, another qubit could be unintentionally driven due to parasitic coupling such as capacitive or magnetic crosstalk.

An expression for the fidelity of the unaddressed qubit with respect to the ideal identity operation is reported in Appendix A for a microwave pulse with a rectangular envelope (Fig. 3.7a), and it is plotted in Fig. 3.7c, where we assume the same Rabi frequency ω_R for both qubits. As expected, driving the qubit with a larger amplitude (i.e., larger ω_R) results in a shorter pulse for a given rotation angle, thus leading to a wider pulse bandwidth and, consequently, to crosstalk extending to qubits that are further away in frequency.

Although the expectation may arise that reducing the pulse bandwidth by proper engineering of the pulse envelope can lead to lower crosstalk, Fig. 3.7d shows that also a Gaussian envelope (Fig. 3.7b) does not result in a much faster roll-off. As the figure shows, the fidelity can be limited by unintended Z-rotations of the

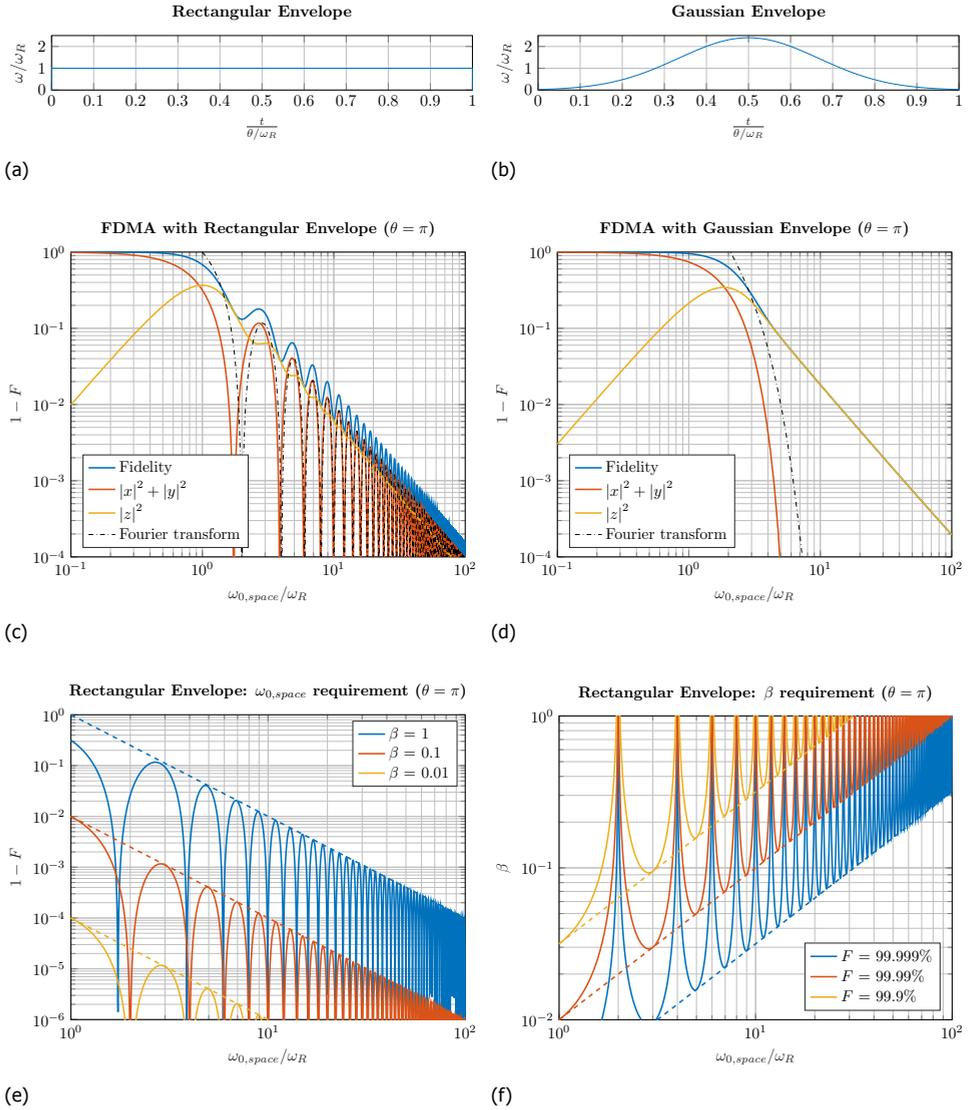


Figure 3.7: Qubit frequency multiplexing: the envelopes, achievable fidelity, and requirements in the case of a rectangular envelope. (a) The rectangular envelope under consideration. (b) The Gaussian envelope under consideration. (c) The infidelity of an identity operation (and the amount of X/Y rotation and Z rotation) on a qubit spaced at $\omega_{0,space}$ from the carrier for a rectangular envelope, along with the Fourier transform of the rectangular envelope. (d) The infidelity of an identity operation (and the amount of X/Y rotation and Z rotation) on a qubit spaced at $\omega_{0,space}$ from the carrier for a Gaussian envelope, obtained by numerical simulation, along with the Fourier transform of the Gaussian envelope. (e) The frequency spacing required to achieve a certain fidelity at given relative signal strength β , for a rectangular envelope. The upper bound (dashed lines) is given in Eq. (3.9). (f) The driving tone attenuation β required at a certain frequency spacing to achieve a given fidelity, for a rectangular envelope. The lower bound (dashed lines) is given in Eq. (3.10).

unaddressed qubit. However, by applying a simple correction for the Z-rotation, the fidelity of the identity operation on the unaddressed qubit improves to the following:

$$F_{\text{FDMA}} \approx 1 - \frac{\beta^2}{\alpha^2} \cdot \sin^2\left(\frac{\theta}{2}\alpha\right) \geq 1 - \frac{\beta^2}{\alpha^2}, \quad (3.9)$$

where $\alpha = \frac{\omega_{0,\text{space}}}{\omega_{\text{R}}}$ and $\beta = \frac{\omega_{\text{R,unaddressed}}}{\omega_{\text{R}}}$, and where, in general, the unaddressed qubit can have a different Rabi frequency ($\omega_{\text{R,unaddressed}}$) at the same microwave amplitude, e.g., due to a lower coupling to the drive signal. As expected, the fidelity given by Eq. (3.9) is approximately proportional to the spectrum of the envelope of the applied pulse (Fig. 3.7c). Consequently, reduction of the pulse bandwidth by proper engineering of the pulse envelope is an effective solution if the unintended Z-rotations are corrected, as shown in, e.g., Figs. 3.7b and 3.7d, where a Gaussian envelope is employed.

A certain minimum frequency separation is necessary to achieve a target fidelity, as shown in Fig. 3.7e for the rectangular envelope. The lower bound on the fidelity as given in Eq. (3.9) is plotted as well, as the notches in the graph move depending on θ . Similarly, if the coupling of the microwave drive is due to parasitic effects and is unwanted, a target fidelity for unaddressed qubits translates into a requirement in the driving tone attenuation (Fig. 3.7f):

$$\beta = \sqrt{1 - F_{\text{corr}}} \cdot \frac{\alpha}{\left|\sin\left(\frac{\theta}{2}\alpha\right)\right|} \geq \sqrt{1 - F_{\text{corr}}} \cdot \frac{\omega_{0,\text{space}}}{\omega_{\text{R}}}. \quad (3.10)$$

Finally, FDMA has the potential to perform single-qubit gates on several qubits at the same time, using a single driveline. In that case, it is not sufficient to apply a compensating Z-rotation afterward, on another qubit, if that qubit is also performing an operation. As the Z-rotation is obtained gradually when an off-resonance tone is applied, the driving tone applied to perform the operation should be altered to compensate for this Z-rotation (or AC Stark shift) during the operation. This requires proper engineering of all the microwave pulses that are applied simultaneously [120, 155, 178–180]. An example of such pulses is given in Section 5.4.2.

3.3.3. Case study for a single-qubit operation

With the information provided in Section 3.3, precise specifications for the control electronics can now be derived. Table 3.2 shows an example of how the total error budget can be allocated over the electronics specification to achieve a 99.9% fidelity for a π -rotation at a typical Rabi frequency of 1 MHz ($T_2^* = 120 \mu\text{s}$). Assuming an FDMA scheme, the same fidelity is targeted for preserving the state of the qubit when not operating on it for a time equal to the operation time ($T_{\text{nop}} = T$). The example considers the use of simple rectangular pulses, without any echo technique.

A Larmor frequency larger than 80 MHz would be sufficient not to get impaired by fast oscillating terms neglected by the RWA (see Appendix A). However, choosing $f_0 = 10 \text{ GHz}$ is more in line with values used in practice and allows for a large qubit frequency spacing. A frequency spacing of 1 GHz is selected, the same as

will be considered in the case study of two-qubit operations (Section 3.4.3). Such spacing is, however, approximately 10 times larger than required for minimizing the crosstalk due to FDMA (Section 3.3.2). The example also shows the effect of the frequency noise as expected from isotopically purified ^{28}Si (with only an 800 ppm concentration of ^{29}Si), highlighting that its contribution to the infidelity is negligible in this example.

The values provided for the microwave amplitude assume a qubit processor based on EDSR, where an amplitude of 2 mV (1.4 mV_{rms}) at the gate is required for a Rabi frequency of 1 MHz (close to the value reported in [65]). All specifications are valid at the gate so that wiring attenuation and filtering might need to be factored in to refer the specifications back to the electronics.

Following these specifications, the microwave envelope (amplitude and duration) can be generated by, e.g., an AWG with a sample rate of at least 150 MS/s, such that the sample time is less than 6.7 ns, resulting in a maximum inaccuracy of 3.3 ns. Furthermore, the AWG should have a resolution of 8 bits, such that at a full-scale swing of 4 mV, the quantization step is sufficiently low. An Effective Number of Bits (ENOB) of only 6.5 bits is required to meet the noise requirement and the specifications on the residual driving when not operating the qubit ('off-spur' in Table 3.2).

The LO used for the up-conversion requires a frequency resolution of approximately 20 kHz (for the inaccuracy). Assuming a -20 dB/dec slope of the phase noise, the single-side band phase noise at 1 MHz from the carrier, $\mathcal{L}(1\text{ MHz})$, needs to be below -106 dBc/Hz. Furthermore, the LO's phase inaccuracy needs to be below 0.64° .

Table 3.2: Example specifications for the control electronics for single-qubit operations. The noise levels provided assume a white spectrum for the amplitude and frequency noise (i.e., -20 dB/dec for the phase noise).

	Value	Infidelity contribution	
		to an operation	to idling
Frequency			
nominal	10 GHz	$0.64 \times 10^{-9(a)}$	
spacing	1 GHz		$1 \times 10^{-6(b)}$
inaccuracy	11 kHz	125×10^{-6}	308×10^{-6}
oscillator noise	11 kHz _{rms}	125×10^{-6}	308×10^{-6}
nuclear spin noise	1.9 kHz _{rms} ^(c)	3.6×10^{-6}	8.9×10^{-6}
wideband noise	12 μ V _{rms}	125×10^{-6}	
Phase			
inaccuracy	0.64°	125×10^{-6}	$31 \times 10^{-6(d)}$
Amplitude			
nominal	2 mV		
inaccuracy	14 μ V	125×10^{-6}	
noise	14 μ V _{rms}	125×10^{-6}	
off-spur	19 μ V ^(e)		217×10^{-6}
off-noise	10 μ V _{rms}		125×10^{-6}
Duration			
nominal	500 ns		
inaccuracy	3.6 ns	125×10^{-6}	
noise	3.6 ns _{rms}	125×10^{-6}	
		$F_{X,Y} = 99.9\%$	$F_I = 99.9\%$

Noise source	ENBW	Noise level
Frequency noise	2.5 MHz	$\mathcal{L}(1 \text{ MHz}) = -106 \text{ dBc/Hz}$
Wideband additive noise	2.9 MHz	$7.1 \text{ nV}/\sqrt{\text{Hz}}$
Amplitude noise	1.0 MHz	$14 \text{ nV}/\sqrt{\text{Hz}}$, SNR = -40 dB
Amplitude off-noise	2.0 MHz	$7.1 \text{ nV}/\sqrt{\text{Hz}}$

(a) Due to the RWA.

(b) Due to leakage in FDMA-setup using rectangular envelopes.

(c) From [61], $T_2^* = 120 \mu\text{s}$.

(d) FDMA Z-corrections limit the idling operation.

(e) Equivalent to -41 dBc .

3.4. Signal Specifications for Two-Qubit Operations

In this section, the fidelity of a two-qubit operation in the presence of non-idealities in the drive signal are first derived (Section 3.4.1). Next, effects of additional non-idealities that affect the qubits when they are idle are considered (Section 3.4.2). Finally, example specifications for achieving a 99.9% fidelity for both an operation and an idle period are given in a case-study (Section 3.4.3).

3.4.1. Fidelity of a two-qubit operation

As stated in Section 3.1, by default, the tunnel coupling between the qubits is negligible, and the qubits have the same potential, i.e., they are not detuned. By increasing the tunnel coupling and/or by detuning the qubits, the qubit interaction increases and a two-qubit gate can be obtained. In this system, by leveraging this exchange interaction, a two-qubit exchange gate and a C-phase gate can be implemented. With either of these gates and single-qubit operations, a universal set is obtained.

Higher energy levels need to be modeled in the Hamiltonian to describe the physical interactions required for the two-qubit gate. The analysis presented here is limited to the interaction between two neighboring qubits, A and B, and to the single-dot singlet states ($|0, 2\rangle$ represents the singlet state in the right dot and $|2, 0\rangle$ the singlet state in the left). In the basis $\Psi = [|\uparrow, \uparrow\rangle, |\uparrow, \downarrow\rangle, |\downarrow, \uparrow\rangle, |\downarrow, \downarrow\rangle, |0, 2\rangle, |2, 0\rangle]$, the Hamiltonian of a double quantum dot is given by the following ($\hbar = 1$) [66, 181, 182]:

$$H_{sys} = \begin{bmatrix} -\omega_0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{\delta\omega_0}{2} & 0 & 0 & t_0 & t_0 \\ 0 & 0 & -\frac{\delta\omega_0}{2} & 0 & -t_0 & -t_0 \\ 0 & 0 & 0 & \omega_0 & 0 & 0 \\ 0 & t_0 & -t_0 & 0 & U - \epsilon & 0 \\ 0 & t_0 & -t_0 & 0 & 0 & U + \epsilon \end{bmatrix}, \quad (3.11)$$

where $\omega_0 = (\omega_{0,A} + \omega_{0,B})/2$, $\delta\omega_0 = \omega_{0,B} - \omega_{0,A}$, and $\omega_{0,A}$ and $\omega_{0,B}$ are the Larmor frequencies of the two qubits. The charging energy (U) is assumed to be the same for both dots. The tunnel coupling between the quantum dots (t_0) has an exponential relation to the voltage on the barrier gate, and the detuning energy (ϵ) is controlled by the voltage difference on the plunger gates of the dots (V_d) via the lever arm $\alpha = \Delta\epsilon/\Delta V_d$.

An avoided crossing is observed in the energy level diagram for $|\epsilon| = U$ and $t_0 > 0$ (Fig. 3.8), which gives rise to eigenenergies that are different from the case of two isolated dots ($t_0 \sim 0$) for any detuning. This change of eigenenergy and the corresponding eigenstate forms the basis of the two-qubit operations. An investigation of the eigenenergies of the Hamiltonian in Eq. (3.11) reveals that the total change in eigenenergy equals the following ($\hbar = 1$):

$$\omega_{op} = 4 \cdot t_0^2 \cdot \frac{U}{U^2 - \epsilon^2}. \quad (3.12)$$

Note that the expression used here for ω_{op} derives directly from the Hamiltonian of

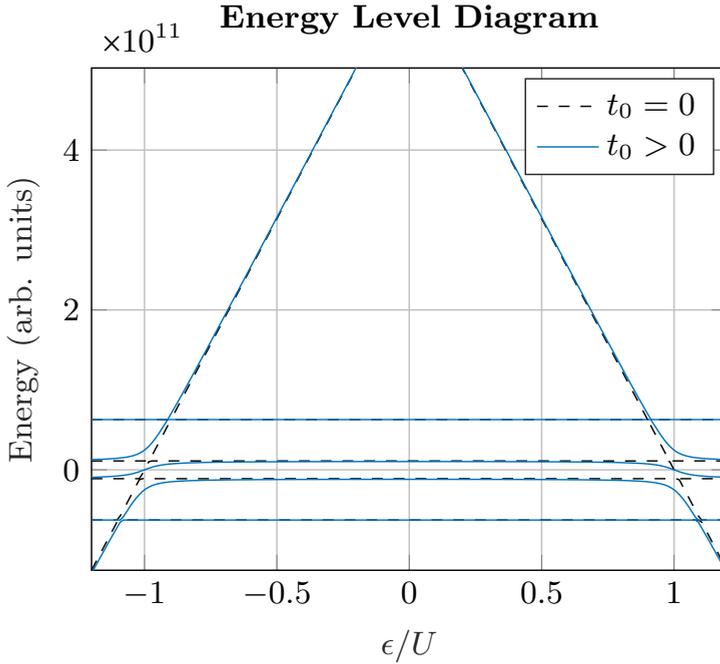


Figure 3.8: The energy level diagram of the two-qubit system. An avoided crossing is visible for $|\epsilon| = U$ when there is a finite tunnel coupling between the dots.

Eq. (3.11). However, experiments have reported ω_{op} as an exponential function of detuning [183].

As ω_{op} describes the amount of exchange interaction, it directly sets the speed of the two-qubit operation. A plot of ω_{op} versus the tunnel coupling and detuning is shown in Fig. 3.9. In order to perform the two-qubit operation, a control pulse must be applied to move the system away from the default point (negligible tunnel coupling and no detuning) to the desired operating point, where there is sufficient exchange interaction so that a two-qubit operation is performed. From Fig. 3.9, it is clear that a fast gate can be obtained at finite detuning, becoming faster closer to the avoided crossing, controlled by the detuning and/or the tunnel coupling. Alternatively, operation at no detuning (the charge symmetry point [152]) is possible, controlled by the tunnel coupling alone. Depending on whether the control parameter, the detuning and/or tunnel coupling, is changed adiabatically or diabatically, a C-phase or exchange gate, or a mixture of the two, is obtained.

The C-Phase Gate

In the case in which the control parameter changes slowly, i.e., adiabatically, the resulting operation, in the rotating frame, can be described by the following diagonal

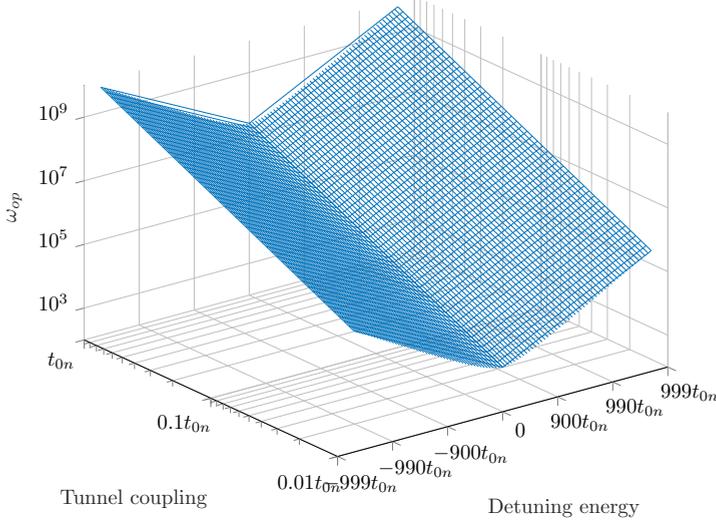
The Operation Speed ω_{op} ($U = 1000 t_{0n}$)


Figure 3.9: The two-qubit operation speed ω_{op} (Eq. (3.12)) versus the interdot tunnel coupling t_0 and detuning ϵ (charging energy $U = 1000t_{0n}$). A nominal tunnel coupling t_{0n} of 1 GHz is used.

matrix:

$$U_{cz}(t) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{-i\phi_{Z,A}} & 0 & 0 \\ 0 & 0 & e^{-i\phi_{Z,B}} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (3.13)$$

where $\phi_{Z,A}$ and $\phi_{Z,B}$ are the acquired phases in the rotating frame. Two additional Z-rotations with angles $\phi_{Z,A}$ and $\phi_{Z,B}$ can be applied to the right and left qubit, respectively, to obtain the C-phase gate with $\theta_{cz} = -(\phi_{Z,A} + \phi_{Z,B}) = \omega_{op} \cdot t$. These Z-rotations can easily be obtained by updating the software reference frame [155, 156], as described in Section 3.3.1. In the case in which $\theta_{cz} = \pi$, a controlled-Z operation is obtained. Interestingly, the total acquired phase ($\phi_{Z,A} + \phi_{Z,B}$) is independent of $\delta\omega_0$. However, when $\delta\omega_0 = \sqrt{2} \cdot t_0$, $\phi_{Z,A} = \phi_{Z,B}$ [66, 181, 182], whereas for $\delta\omega_0 = 0$, $\phi_{Z,A} = 0$.

The Exchange Gate

If, instead, the control parameter is changed rapidly, i.e., diabatically, and the Larmor frequency difference is negligible ($\delta\omega_0 \ll \omega_{op}$), the resulting operation, in the rotating frame, is as follows:

$$U_J(t) \approx \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1+e^{i\theta_J}}{2} & \frac{1-e^{i\theta_J}}{2} & 0 \\ 0 & \frac{1-e^{i\theta_J}}{2} & \frac{1+e^{i\theta_J}}{2} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (3.14)$$

where $\theta_J = \omega_{\text{op}} \cdot t$. In the case in which $\theta_J = \pi$, a SWAP operation is obtained. Note that since for an accurate exchange operation the Larmor frequency difference should be sufficiently small, the possibility of using FDMA for single-qubit operations (Section 3.3.2) is limited if the Larmor frequency cannot be dynamically changed.

Effect of non-idealities on the Exchange and C-phase Gate

From Eq. (3.11) it follows that the two-qubit operations are affected by the Larmor frequencies ($\omega_{0,A}, \omega_{0,B}$), the tunnel coupling (t_0), the charging energy (U), and the detuning (ϵ). Furthermore, the operation depends on the total duration (T) for which the two-qubit gate is active. The effect of errors, both static and dynamic, on the fully electrically controlled parameters (t_0 , ϵ and T) is analyzed in the subsequent section. Detailed derivations of the formulas can be found in Appendix A.

The resulting fidelity in the case of control signal inaccuracies is summarized in Table 3.3 for the exchange gate and the C-phase gate, both at zero detuning and finite detuning. For the exchange gate, we assume that no Larmor frequency difference between the qubits exists, since for such a gate $\delta\omega_0 \ll \omega_{\text{op}}$ is required, while for the C-phase gate, various scenarios are analyzed ($\delta\omega_0 = 0$, $\delta\omega_0 = \omega_{\text{op}}$ and $\delta\omega_0 = \sqrt{2} \cdot t_0$). In the table, T denotes the qubit gate operation time and inaccuracies are shown with the prefix Δ . The table provides values for different rotation angles (θ_{cz}, θ_J). In the case of the C-phase gate, additional Z-rotations might be required, which only have a minimal effect on the fidelity (Eq. (3.4)).

The error contributions have a quadratic relation with the infidelity except for detuning errors for $\epsilon = 0$, where a fourth-order dependence is found. This implies improved robustness to detuning errors when operating at the charge symmetry point ($\epsilon = 0$) [152, 153].

For low-frequency variations, i.e., those changing over a time-scale longer than the operation time, the same approach as for single-qubit operations holds and the expected fidelity follows the same equations as given in Table 3.3 when replacing the inaccuracy, such as $\Delta\epsilon$, with the standard deviation of the variation (assuming a Gaussian distribution). An exception is for detuning errors when operating at the charge symmetry point [152, 153], because of the fourth-order dependence. For a static but random error Δ for which $F = 1 - c \cdot \Delta^4$, the expected fidelity is $F = 1 - 3 \cdot c \cdot \sigma^4$, if Δ follows a Gaussian distribution with standard deviation σ and zero mean (see Appendix A). Consequently, there is a 3× higher sensitivity to noise than to static errors for the detuning.

For timing variations, only the total duration matters, and high-frequency noise is filtered as described in Table 3.1. Moreover, similar to the single-qubit gate, numerical simulations of the Hamiltonian have shown sensitivity to high-frequency noise ($> \omega_{\text{op}}$) only in a passband with a bandwidth that is inversely proportional to the operation duration, for both the electrically controlled detuning energy and tunnel coupling in the case of the two-qubit gate (see Appendix A). The quantum state is, however, also affected by the noise around the frequencies corresponding to the allowed energy transitions in a passband with similar bandwidth. Consequently, it is essential that the high-frequency noise components in the signals applied to the barrier gates and plunger gates are adequately filtered. However,

Table 3.3: The fidelity in the case of control signal inaccuracies for various two-qubit operations.

	Exchange $\delta\omega_0 = 0$	C-phase $\delta\omega_0 = 0$
Rotation $ \phi_{Z,B} $		θ_{cz}
Duration	$1 - \frac{3}{16} \cdot \theta_J^2 \cdot \left(\frac{\Delta T}{T}\right)^2$	$1 - \frac{3}{16} \cdot \theta_{cz}^2 \cdot \left(\frac{\Delta T}{T}\right)^2$
Tunnel coupling	$1 - \frac{3}{4} \cdot \theta_J^2 \cdot \left(\frac{\Delta t_0}{t_0}\right)^2$	$1 - \frac{3}{4} \cdot \theta_{cz}^2 \cdot \left(\frac{\Delta t_0}{t_0}\right)^2$
Detuning ($ \epsilon > 0$)	$1 - \frac{3}{4} \cdot \theta_J^2 \cdot \left(\frac{\frac{\epsilon}{U}}{1 - \left(\frac{\epsilon}{U}\right)^2}\right)^2 \cdot \left(\frac{\Delta\epsilon}{U}\right)^2$	$1 - \frac{3}{4} \cdot \theta_{cz}^2 \cdot \left(\frac{\frac{\epsilon}{U}}{1 - \left(\frac{\epsilon}{U}\right)^2}\right)^2 \cdot \left(\frac{\Delta\epsilon}{U}\right)^2$
Detuning ($\epsilon = 0$)	$1 - \frac{3}{16} \cdot \theta_J^2 \cdot \left(\frac{\Delta\epsilon}{U}\right)^4$	$1 - \frac{3}{16} \cdot \theta_{cz}^2 \cdot \left(\frac{\Delta\epsilon}{U}\right)^4$

	C-phase $\delta\omega_0 = \omega_{op}$	C-phase $\delta\omega_0 = \sqrt{2} \cdot t_0$
Rotation $ \phi_{Z,B} $	$\frac{\theta_{cz}}{\sqrt{2}}$	$\frac{\theta_{cz}}{2}$
Duration	$1 - \frac{7-4\sqrt{2}}{16} \cdot \theta_{cz}^2 \cdot \left(\frac{\Delta T}{T}\right)^2$	$1 - \frac{1}{16} \cdot \theta_{cz}^2 \cdot \left(\frac{\Delta T}{T}\right)^2$
Tunnel coupling	$1 - \frac{1}{2} \cdot \theta_{cz}^2 \cdot \left(\frac{\Delta t_0}{t_0}\right)^2$	$1 - \frac{1}{4} \cdot \theta_{cz}^2 \cdot \left(\frac{\Delta t_0}{t_0}\right)^2$
Detuning ($ \epsilon > 0$)	$1 - \frac{1}{2} \cdot \theta_{cz}^2 \cdot \left(\frac{\frac{\epsilon}{U}}{1 - \left(\frac{\epsilon}{U}\right)^2}\right)^2 \cdot \left(\frac{\Delta\epsilon}{U}\right)^2$	$1 - \frac{1}{4} \cdot \theta_{cz}^2 \cdot \left(\frac{\frac{\epsilon}{U}}{1 - \left(\frac{\epsilon}{U}\right)^2}\right)^2 \cdot \left(\frac{\Delta\epsilon}{U}\right)^2$
Detuning ($\epsilon = 0$)	$1 - \frac{1}{8} \cdot \theta_{cz}^2 \cdot \left(\frac{\Delta\epsilon}{U}\right)^4$	$1 - \frac{1}{16} \cdot \theta_{cz}^2 \cdot \left(\frac{\Delta\epsilon}{U}\right)^4$

since the exchange gate requires a diabatic change in the control parameter, only limited filtering can be applied. Closed-form analytical expressions for these effects have not been studied.

3.4.2. Specifications for the idle operation

Since, in practice, the tunnel coupling cannot be fully removed, the two-qubit operation is never completely turned off. The interaction strength can, however, be slowed down significantly, thus leading to a fidelity with respect to the ideal identity operation for the exchange and C-phase gates as follows:

$$F_{\text{Inop},J} = \begin{cases} 1 - \frac{3}{16} \cdot \omega_{\text{op,off}}^2 \cdot T_{\text{nop}}^2 & \delta\omega_0 = 0 \\ 1 - \frac{7-4\sqrt{2}}{16} \cdot \omega_{\text{op,off}}^2 \cdot T_{\text{nop}}^2 & \delta\omega_0 = \omega_{\text{op}} \\ 1 - \frac{1}{16} \cdot \omega_{\text{op,off}}^2 \cdot T_{\text{nop}}^2 & \delta\omega_0 = \sqrt{2} \cdot t_0 \end{cases}, \quad (3.15)$$

where $\omega_{\text{op,off}}$ is the reduced interaction strength during the time T_{nop} when no operation is applied.

Following Eq. (3.12), the interaction strength can be reduced by lowering the tunnel coupling while not changing the detuning. A two-qubit operation performed at finite detuning could also be controlled using only the detuning. However, assuming that the interaction is considered to be off at no detuning, the operation might need to be performed at far detuning. As mentioned before, operating closer to the avoided crossing reduces the tolerance to inaccuracies and noise in the detuning (Table 3.3). Therefore, control over the tunnel coupling as a means to control the two-qubit operation can relax the electrical specifications.

3.4.3. Case study for a two-qubit operation

Specifications for the control electronics responsible for the two-qubit operation can now be derived using the results presented in Section 3.4. This example develops on the example given in Section 3.3.3 and, for instance, assumes that the same oscillator is used to keep the coherence with the qubits. Two examples will be given here, one at no detuning and one at finite detuning. Both examples focus on the C-phase gate operating at $\delta\omega_0 = \sqrt{2} \cdot t_0$. This choice for the Larmor frequency difference gives the most relaxed specifications for the control electronics, while at smaller $\delta\omega_0$ the specifications can be up to $\sqrt{3}$ times more demanding (see Table 3.3).

The Larmor frequency difference is chosen as 1 GHz to achieve a two-qubit operation speed of $\omega_{\text{op}} = 2$ MHz at no detuning, while maintaining $\delta\omega_0 = \sqrt{2} \cdot t_0$ (Eq. (3.12)). Example specifications for this operation are given in Table 3.4. To further increase the operating speed, an even higher qubit frequency spacing would be required or $\delta\omega_0 < \sqrt{2} \cdot t_0$. Alternatively, the operating speed can be enhanced to, e.g., 20 MHz, by operating the C-phase gate at finite detuning (Eq. (3.12)), as shown in another example (Table 3.4). Both examples target a fidelity of 99.9% for a C-phase gate with $\theta_{\text{cz}} = \pi$. The examples also indicate the specifications required for idling two qubits at 99.9% fidelity for a duration of 500 ns, the same as for the example in Section 3.3.3.

For the charging energy and tunnel coupling, typical values are chosen. As the relation to the gate voltage is device dependent, no values for the required electrical specifications are given. Note that in either example, the tunnel coupling only has to change by a factor of approximately 9 to turn the operation on or off. In the case of operation at finite detuning, this assumes that no detuning is applied when the operation is turned off.

The detuning energy is directly related to the voltage on the plunger gate via the lever arm, for which a typical value of $\alpha = 0.05$ eV/V is assumed [65]. When operating at finite detuning, the detuning energy is chosen at 95% of the charging energy. Even though higher operating speeds can be obtained by moving even closer to the avoided crossing, the electrical specifications become increasingly challenging. When operating at the charge symmetry point, very large detuning errors can be tolerated (at which point the approximations used to derive the expressions in Table 3.3 do not hold any more). When operating at moderate detuning, the error specification for the detuning is more than 100 times stricter. Moreover, as the operation at finite detuning is faster with the same tunnel coupling, the signal bandwidth must be larger, with a larger noise bandwidth. As a rough estimate, the ENBW has been chosen as 5 times the operating speed in both examples, which seems plausible as an adiabatic change is required (for the exchange gate, the situation might be worse). As a result, the maximum allowed noise spectral density, assuming white noise, is much lower. For the given example, this results in a difference of almost 5 orders of magnitude in the noise power spectral density, due to the difference in total tolerable integrated noise and ENBW.

In the example operation at finite detuning, the detuning control can be achieved by an AWG running at a sample rate of 1 GS/s for a maximum timing inaccuracy of 0.5 ns. Assuming that the AWG has to cover a voltage range of $-U \dots U$ (where U is the charging energy), it must have a 10-bit resolution to meet the accuracy specification of the detuning energy.

Table 3.4: Example specifications for the control electronics when operating a C-phase gate at no detuning (top) and at finite detuning (bottom). The PSD values provided assume a white spectrum with an ENBW of approximately 10 MHz ($\omega_{op} = 2$ MHz) when operating at no detuning and an ENBW of approximately 100 MHz ($\omega_{op} = 20$ MHz) at finite detuning. A nominal charging energy of 83 mV (4.1 meV, 1.0 THz) is assumed.

	Value	Infidelity contribution	
		to an operation	to idling
Frequency			
spacing	1 GHz		
inaccuracy	11 kHz	77×10^{-6}	308×10^{-6}
oscillator noise	11 kHz _{rms}	77×10^{-6}	308×10^{-6}
nuclear spin noise	1.9 kHz _{rms}	2.2×10^{-6}	8.9×10^{-6}
Duration			
nominal	250 ns		
error	5.3 ns	281×10^{-6}	
Detuning energy			
nominal	0 mV (0 μ eV, 0 GHz)		
error	12 mV (0.60 meV, 0.15 THz) $\sigma = 9.2$ mV _{rms} PSD = 2.9 μ V/ $\sqrt{\text{Hz}}$	281×10^{-6}	
Tunnel coupling			
nominal	0.71 GHz (2.9 μ eV)		
error	7.5 MHz (31 neV)	281×10^{-6}	
off-value	78 MHz (0.32 μ eV)		374×10^{-6}
		$F_{Cz} = 99.9\%$	$F_I = 99.9\%$
	Value	Infidelity contribution	
		to an operation	to idling
Frequency			
spacing	1 GHz		
inaccuracy	11 kHz _{rms}	0.8×10^{-6}	308×10^{-6}
oscillator noise	11 kHz _{rms}	0.8×10^{-6}	308×10^{-6}
nuclear spin noise	1.9 kHz _{rms}	0.02×10^{-6}	8.9×10^{-6}
Duration			
nominal	25 ns		
error	0.58 ns	333×10^{-6}	
Detuning energy			
nominal	78 mV (3.9 meV, 0.95 THz)		
error	0.10 mV (5.1 μ eV, 1.2 GHz) $\sigma = 0.10$ mV _{rms} PSD = 10 nV/ $\sqrt{\text{Hz}}$	333×10^{-6}	
Tunnel coupling			
nominal	0.71 GHz (2.9 μ eV)		
error	8.2 MHz (34 neV)	333×10^{-6}	
off-value	78 MHz (0.32 μ eV)		374×10^{-6}
		$F_{Cz} = 99.9\%$	$F_I = 99.9\%$

3.5. Signal Specifications for Qubit Read-Out

In this section, the fidelity of qubit read-out is first derived (Section 3.5.1), followed by a case-study presenting example specifications for achieving a 99.9% fidelity for the read-out (Section 3.5.2).

3.5.1. Fidelity for qubit read-out

For the read-out of the quantum state, the Pauli spin-blockade read-out [69] is analyzed, since it offers several advantages with respect to the other possible alternative, i.e., the Elzerman read-out [68]: no electron reservoir is required next to the quantum dot; and the Zeeman energy splitting does not have to be much higher than the thermal energy, thus enabling operation at higher temperatures and/or lower Larmor frequencies. As a drawback, the Pauli spin-blockade read-out involves two quantum dots, where the measurement involves discrimination between the singlet and triplet states.

Even though relaxation, which is quantified by the relaxation time T_1 , is an important limiting factor in qubit read-out, its effect is not considered in the following analysis as all gates are assumed to be performed in a time significantly smaller than T_1 . Furthermore, the following analysis assumes that the spin-dependent charge state resulting from a Pauli spin-blockade read-out is measured using a charge sensor. As a result, the read-out fidelity is determined by various factors:

- P_{charge} : the probability that the spin-state is correctly projected to the charge state.
- P_{sense} : the probability that the charge sensor correctly detects the charge state.
- P_{detect} : the probability that the read-out circuit correctly discriminates the signal of the charge sensor.

The overall read-out fidelity is then:

$$F_{\text{read}} \approx P_{\text{charge}} \cdot P_{\text{sense}} \cdot P_{\text{detect}}. \quad (3.16)$$

The probability P_{sense} is limited by, e.g., interference on one of the charge-sensor bias gates and charge noise in the substrate. As this depends highly on the type of sensor employed and the sensor integration, this error contribution will not be discussed further.

The quantum-dot control electronics limit P_{charge} , while the read-out electronics limit P_{detect} , as discussed next.

Specifications for the electronics controlling the spin-to-charge conversion

For the analysis of the charge transfer in the Pauli spin-blockade read-out, the Hamiltonian of Eq. (3.11) is extended with the lowest-energy triplet states (either due to the valley splitting or the orbital energy splitting). Those states are spaced by a singlet-triplet energy splitting E_{ST} from the singlet energy level (for the Hamiltonian, see Eq. (A.118)). A plot of the energy of the stationary states versus the

detuning near the avoided crossing is shown in Fig. 3.10. For the following discussion, only the $|\downarrow, \downarrow\rangle$ and $|\downarrow, \uparrow\rangle$ states, highlighted in Fig. 3.10 with dashed lines, need to be considered.

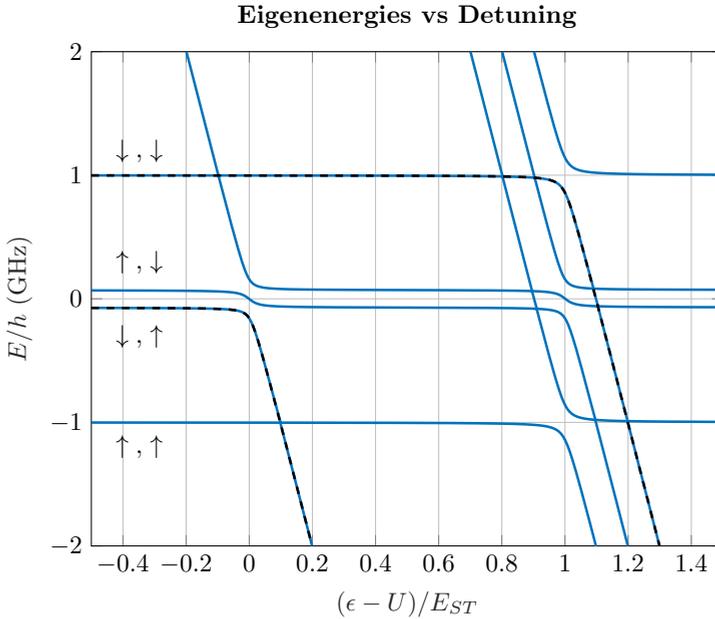


Figure 3.10: The energy of the stationary states versus the detuning near the avoided crossing. The black dashed lines indicate the states where the left qubit was initially in the $|\downarrow\rangle$ state.

The Pauli spin-blockade read-out relies on E_{ST} for the discrimination of the single-dot singlet configuration from three possible single-dot triplet configurations. The state of the right qubit, in a pair of neighboring qubits, can be measured as follows. The left qubit is initialized in the $|\downarrow\rangle$ state. By detuning adiabatically to a point between the singlet avoided crossing and the triplet avoided crossing (with $t_0 > 0$), only the $|\downarrow, \uparrow\rangle$ -state (at $\epsilon = 0$) becomes a singlet, and both electrons will move into the same dot. This charge movement can be measured using a charge sensor. Based on the measurement result, it is then clear whether the qubits are in the singlet or one of the three triplet configurations. This scenario is analyzed here.

When starting from the $|\downarrow, \downarrow\rangle$ -state, there is a small probability $P(\text{transfer} | |\downarrow, \downarrow\rangle)$ that both electrons will end up in the same dot. Similarly, starting from the $|\downarrow, \uparrow\rangle$ -state, there is a small probability $P(\text{no transfer} | |\downarrow, \uparrow\rangle)$ that no charge will transfer. The probability of a correct spin-to-charge conversion can be defined as follows:

$$P_{\text{charge}} = 1 - P(\text{transfer} | |\downarrow, \downarrow\rangle) - P(\text{no transfer} | |\downarrow, \uparrow\rangle). \quad (3.17)$$

The analysis is again simplified by assuming an ideal adiabatic change in the detuning energy. The results presented in this section are obtained from numerical simulations of the Hamiltonian.

Simulations show that the highest P_{charge} is obtained by detuning to $\epsilon = U + \frac{E_{\text{ST}}}{2}$, i.e., equidistant between the singlet and triplet avoided crossings, as shown in Fig. 3.11. The shape of the probability versus $\frac{\epsilon - U}{E_{\text{ST}}}$ plot is independent of the Larmor frequency, assuming that $\omega_0 \ll E_{\text{ST}}$. Although the shape remains the same, the obtainable maximum P_{charge} scales with the tunnel coupling and the singlet-triplet energy splitting, as can be seen in Fig. 3.12. From this figure, an upper bound for the tunnel coupling can be found, which must be maintained even with errors caused by limitations in the control electronics.

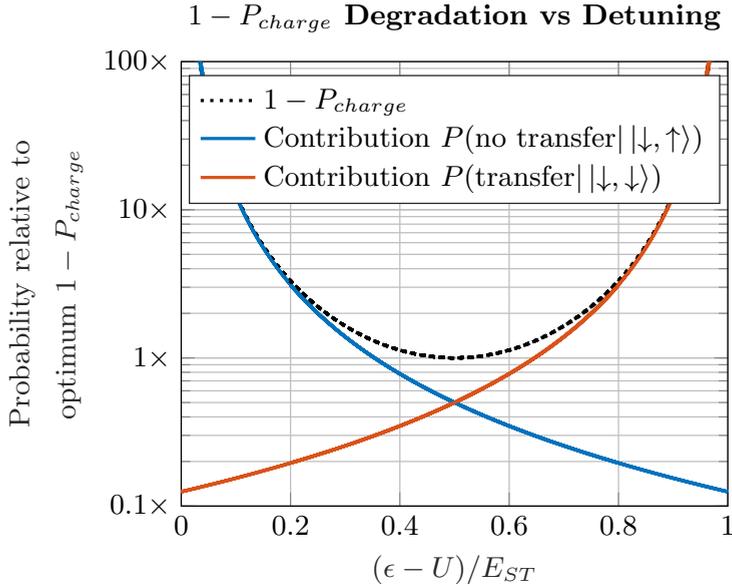


Figure 3.11: The error probability $1 - P_{\text{charge}}$, along with the individual error contributors, at various points of detuning as simulated for various tunnel rates, Larmor frequencies, charging energies, and singlet-triplet energy splittings (each varied over a decade; the resulting plots are overlapping). The obtained probabilities are plotted relative to the optimum, i.e., the lowest error probability at $\epsilon = U + \frac{E_{\text{ST}}}{2}$, thereby showing the degradation when moving away from the optimum detuning.

Even though P_{charge} is highly influenced by the achievable tunnel couplings and singlet-triplet energy splittings in the system (Fig. 3.12), the detuning value has a minor influence (provided that there is a sufficient singlet-triplet energy splitting), since $1 - P_{\text{charge}}$ is relatively flat around its minimum, as shown in Fig. 3.11. For instance, for a twofold increase in $1 - P_{\text{charge}}$, the detuning must stay in the range $\frac{\epsilon - U}{E_{\text{ST}}} \approx 0.5 \pm 0.235$. We can then conclude that a large singlet-triplet splitting is desired to limit the influence of the control electronics on the read-out.

Specifications for the electronics processing the read-out signal

In this section, we will consider a direct read-out. A model of a typical read-out chain analyzed here is shown in Fig. 3.13. For simplicity, the sensor is modeled as a

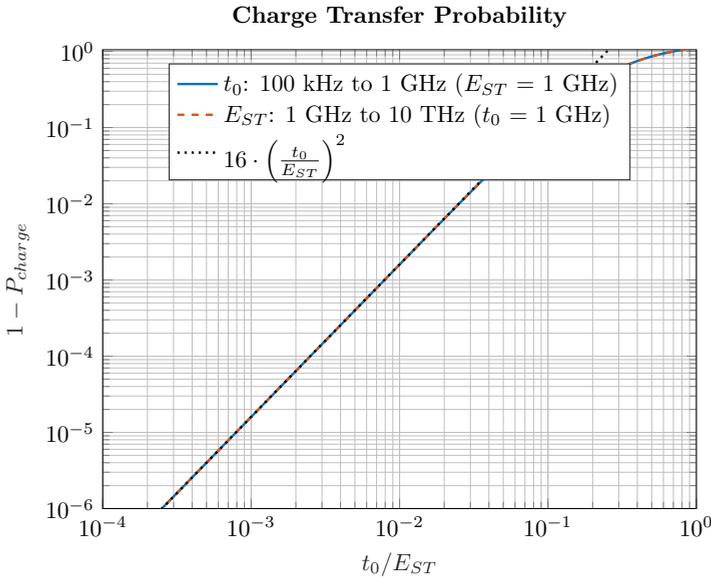


Figure 3.12: The simulated probability $1 - P_{\text{charge}}$ versus the singlet-triplet splitting, normalized to the tunnel coupling at $\epsilon = U + \frac{E_{ST}}{2}$, while sweeping either the tunnel coupling or the singlet-triplet energy splitting.

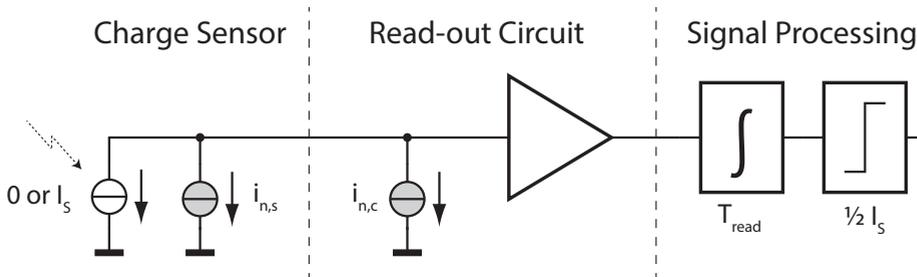


Figure 3.13: A model of the analyzed read-out chain, showing the sensor, the read-out electronics, and the required signal processing for the measurement discrimination. Additional sources modeling the noise are shown in gray.

current source with a value of either 0 or I_s , depending on the sensed charge. The read-out fidelity is limited by the noise introduced by the sensor and by the read-out circuit, indicated in Fig. 3.13 as $i_{n,s}$ and $i_{n,c}$, respectively. Assuming the typical matched-filter detection [184], i.e., integrating the signal current for a duration T_{read} and comparing the result to a threshold, the probability of a correct measurement under the presence of Gaussian-distributed noise is given by the following (Fig. 3.14):

$$P_{\text{detect}} = \frac{1 + \operatorname{erf}\left(\sqrt{\frac{\text{SNR}}{8}}\right)}{2}, \quad (3.18)$$

with

$$\text{SNR} = \frac{I_s^2}{\int_0^\infty S_i(f) \cdot \left(\frac{\sin(\pi \cdot f \cdot T_{\text{read}})}{\pi \cdot f}\right)^2 df}, \quad (3.19)$$

where $S_i(f)$ is the PSD of the total noise $i_n = i_{n,s} + i_{n,c}$. When the noise is white, this simplifies to the following:

$$\text{SNR} = \frac{I_s^2}{S_i \cdot \text{ENBW}}, \quad (3.20)$$

with effective noise bandwidth $\text{ENBW} = 1/(2 \cdot T_{\text{read}})$.

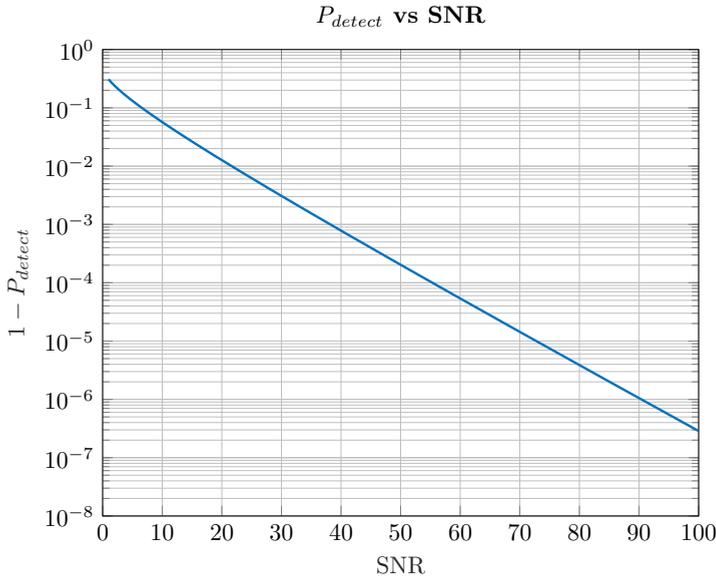


Figure 3.14: A plot of $1 - P_{\text{detect}}$ versus SNR (linear scale) in the case of Gaussian-distributed noise.

3.5.2. Case study for qubit read-out

The example specifications presented in this section build on those presented in Sections 3.3.3 and 3.4.3 and hence assume the same charging energy and lever arm for the detuning control. For the singlet-triplet energy splitting, a typical value of $E_{ST} = 50 \mu\text{eV}$ is used. As a result, the optimum detuning is at 83.2 mV. The resulting specifications are summarized in Table 3.5 and assume equal contributions from P_{charge} , P_{sense} , and P_{detect} (Eq. (3.16)).

Following Fig. 3.12, to achieve the required P_{charge} , the tunnel coupling must be even lower than is required to turn off the two-qubit operation (see Section 3.4.3), thereby extending the required tunnel coupling tuning range to approximately 18 \times . In the example of Table 3.5, the detuning control can be achieved by an AWG with a low sample rate, as the detuning must change adiabatically, and the read-out generally takes a relatively long time. Assuming that the AWG has to cover a voltage range from 0 to $2U$, it must have a 9-bit resolution to meet the accuracy specification of the detuning energy (Fig. 3.11). As a result, the same circuitry as used for the two-qubit operation (Section 3.4.3) could potentially be used.

In this example, P_{detect} assumes a direct read-out of a QPC, following the numbers provided in [96] ($I_s = 400 \text{ pA}$ and $i_{n,s} = 57 \text{ fA}/\sqrt{\text{Hz}}$). Assuming that the read-out circuit is designed to contribute about half the noise compared to the shot noise limit of the QPC ($i_{n,c} \approx i_{n,s}/2$), an integration time of at least $T_{\text{read}} = 0.6 \mu\text{s}$ is required to achieve an SNR of 46 for a P_{detect} of 99.967% (for $F = 99.9\%$, Eq. (3.16)). For such short read-out times, the assumption of white noise is valid, and the effects of qubit relaxation (T_1) could be negligible.

Table 3.5: Example specifications for the control electronics. The PSD values provided assume a white spectrum with an ENBW of approximately 1 MHz for the detuning control and a measurement time of $T_{\text{read}} = 0.6 \mu\text{s}$. A nominal charging energy of 82.7 mV (4.1 meV, 1.0 THz) is assumed, and a singlet-triplet energy splitting of 1.0 mV (50 μeV , 12 GHz).

	Value	Infidelity contribution to the read-out
Detuning energy		
nominal	83.2 mV (4.2 meV, 1.0 THz)	
error	0.24 mV (12 μeV , 2.8 GHz) $\sigma = 0.24 \text{ mV}_{\text{rms}}$ PSD = 0.24 $\mu\text{V}/\sqrt{\text{Hz}}$	167×10^{-6}
Tunnel coupling		
nominal	39 MHz (0.16 μeV)	167×10^{-6}
		$P_{\text{charge}} = 99.967 \%$
Charge sensor		
		333×10^{-6}
		$P_{\text{sense}} = 99.967 \%$
Quantum Point Contact		
signal	400 pA	
noise	53 pA _{rms} , PSD = 57 fA/ $\sqrt{\text{Hz}}$	222×10^{-6}
Readout Circuit		
input-referred noise	26 pA _{rms} , PSD = 28 fA/ $\sqrt{\text{Hz}}$	111×10^{-6}
		$P_{\text{detect}} = 99.967 \%$
$P_{\text{charge}} \cdot P_{\text{sense}} \cdot P_{\text{detect}}$		$F = 99.9 \%$

3.6. Discussion

Case studies targeting a 99.9% average gate fidelity have been presented in Sections 3.3.3, 3.4.3 and 3.5.2. This target fidelity is particularly relevant since the minimum error rate to reach fault-tolerant quantum computing is around 99% for a complete error correction cycle, thus requiring a single-operation fidelity above 99.9% for a typical cycle of ten operations [8]. Reaching a 99.9% fidelity for all operations is currently an ambitious goal for all qubit platforms. However, our model can be directly applied to analyze specifications for any given fidelity.

The derived electronic specifications can now be compared to the performance achieved by state-of-the-art electronics. To this end, Table 3.6 summarizes the performance of commonly used arbitrary waveform generators and microwave vector sources [100, 101, 185–187].

Table 3.6: Specifications of commonly used AWGs and VSGs.

	Tektronix AWG5014C [100]	Keysight M9330A [185]	Tabor WX1282C [186]
Sample rate	1.2 GS/s	1.25 GS/s	1.25 GS/s
Resolution	14 bit	15 bit	14 bit
Jitter	5.0 ps _{rms}	-	-
Output noise	-(^a)	-150 dBc/Hz	-
Spurious-Free Dynamic Range (SFDR)	< -56 dBc	< -65 dBc	< -44 dBc
		Agilent E8267D [101]	R&S SMW200A [187]
Max. Output Frequency		44 GHz	40 GHz
Frequency Resolution		1 mHz	1 mHz
Phase Noise (100 kHz)		< -100 dBc/Hz	< -100 dBc/Hz
Wideband noise (10 dBm)		< -141 dBc/Hz	< -134 dBc/Hz

(^a) Amplitude resolution: 1 mV.

For the generation of the detuning control for two-qubit gates and read-out and the microwave envelope for single-qubit gates, an AWG is required. We compare the specifications of the Tektronix 5014C, as used in, e.g., [15, 65, 119], with the specifications derived in the case study. This AWG achieves a sample rate of 1.2 GS/s with 14-bit resolution, thereby providing enough resolution for the amplitude and duration of the microwave envelope (150 MS/s and 8 bits). The worst-case spurious-free dynamic range of -56 dBc is well below the required -41 dBc. The specified random jitter of 5.0 ps_{rms} is well below the required value of 3.6 ns_{rms}. Finally, the output noise level is not clearly specified but can be assumed to be not much larger than the amplitude resolution of 1 mV. In the case in which the AWG's output is attenuated by approximately 40 dB, this also meets the specifications. This

AWG can also be used for detuning control in two-qubit gates. The sample rate is high enough to meet the required timing resolution (> 1 GS/s), and the resolution is sufficiently high to reach the detuning requirements (a worst-case 0.10 mV for a < 100 mV pulse) with a 20 dB attenuator. As the specifications for Pauli-spin blockade read-out are found to be more relaxed, the same AWG again suffices.

Finally, for the generation of the microwave carrier for single-qubit gates, some setups use the Agilent microwave VSG E8267D [15, 43, 61, 65, 66, 119], which has a frequency resolution well below the requirements (1 MHz versus approximately 20 kHz). The single-sided phase noise is also well below the required -106 dBc/Hz at a 1 MHz offset from the carrier (at the worst point, the E8267D achieves better than -100 dBc/Hz at a 100 kHz offset). The broadband noise is specified as 63 nV/ $\sqrt{\text{Hz}}$ (-141 dBc/Hz at 10 dBm output power) and therefore at least 20 dB attenuation is required to meet the specification of 7.1 nV/ $\sqrt{\text{Hz}}$.

It can be concluded that typically adopted instruments are capable of supporting a 99.9% fidelity. However, for the currently used instrumentation, the specifications on the amplitude noise and wideband additive noise are the most stringent and consequently require the use of attenuators to reduce the noise reaching the quantum devices. Moreover, these instruments are bulky, consume several Watts of power, and cannot directly be operated at cryogenic temperatures, therefore hindering scalability.

Fully integrated CMOS circuits operating at cryogenic temperatures can be adopted to tackle this problem [10, 19–24]. In order to assess the feasibility of such a solution, the power consumption of the required circuit blocks will be estimated by using room-temperature CMOS circuits as a reference. This is valid since cryogenic CMOS circuits are expected to show significantly less noise for the same power budget, as shown in refs [10, 23]. As a result, the estimates given here likely overestimate the required power consumption. Furthermore, we assume a $50\text{-}\Omega$ load for each circuit, which may not be the case for a fully integrated controller.

The core component determining the specifications of an AWG is its DAC. The 10 -bit 500 MS/s DAC presented in [188] meets the specifications for the microwave envelope generation at a power consumption of 24 mW. For detuning control, the DAC specifications are stricter but can be met by the 12 -bit 1.6 GS/s DAC presented in [189], with a power consumption of 40 mW. Although for the tunnel barrier the specifications will depend highly on the gate structure, a similar DAC is assumed to be sufficient.

The core component of a microwave carrier generator, the PLL, is also available as a CMOS circuit operating over the required frequency range ($9.2 - 12.7$ GHz) at a power consumption of around 13 mW [190]. Its phase noise performance is slightly worse than required. However, with operation at cryogenic temperatures, the thermal noise contribution is expected to improve, although perhaps not linearly with temperature [42].

In a linear qubit array, one DAC is required for the barrier gate and one for the plunger gate for each qubit. This leads to an estimated power of 80 mW per qubit. For the microwave signals, the envelope DAC and PLL together consume approximately 40 mW. Without any form of multiplexing, this indicates a power

consumption of 120 mW/qubit. For a state-of-the-art dilution refrigerator with a cooling power of a few watts at 4 K, this suggests a maximum of a few tens of qubits when operating the classical controller at 4 K.

However, the power consumption of the DACs controlling the barrier gates and plunger gates could be highly reduced if it is not being limited to a 50- Ω system. To get a sufficient signal swing, in [189], a 16-mA current is delivered to a 50- Ω load, thereby setting a lower bound to the power consumption. A much lower current would be required for a higher impedance, or even for a lower swing as acceptable in this application, in which power consumption is ultimately limited by the speed or noise requirement. Furthermore, the same fast DAC can be used to generate frequency multiplexed microwave envelopes. With a sample rate of 1.6 GS/s, a bandwidth of roughly 640 MHz is available [189]. This can be used to drive 64 qubits with a Rabi frequency of 1 MHz spaced by 10 MHz using, e.g., a Gaussian envelope (Fig. 3.7d). The combined power of the fast DAC and PLL, i.e., 53 mW, is then shared over 64 qubits, thus resulting in a power consumption below 1 mW/qubit. For the read-out, on the other hand, cryogenic CMOS circuits have already been proposed that can achieve a power consumption < 1 mW/qubit [10, 23].

In summary, a cryogenic CMOS controller for a large-scale quantum processor appears to be feasible for a target fidelity of 99.9%. However, for minimum power consumption, the trade-offs in the electronics design must be systematically investigated. The analysis proposed in this chapter provides the foundations for such optimization and will help electronics designers to build a functional controller.

3.7. Conclusion

In this chapter, the effect of non-idealities in the classical controller for a quantum processor have been analyzed. A comprehensive approach is proposed, by covering the effect of both static and dynamic errors on all quantum operations, i.e., single-qubit gates, two-qubit gates, and read-out. The presented approach can be used to analyze the performance of a quantum processor in any qubit technology.

The results of this analysis allow to quantify the impact of the controller on the performance of the quantum computer as a whole. This is required to ensure that the controller does not become the performance bottleneck as the qubit performance keeps improving. Moreover, a full set of electrical specifications targeting a given qubit fidelity can be derived by applying the presented results. The availability of these specifications enables the design of the next-generation controllers tailor-made for a target quantum processor and optimized for performance, power, cost, and size, so as to improve the scalability of the quantum computer.

As future controllers will need to operate physically close to the quantum processor, i.e., at cryogenic temperatures where the cooling power is limited, the power optimization of the controller will be critical in enabling large-scale quantum computing. With the results obtained in this chapter, the trade-offs between qubit fidelity and power spent in the controller can be analyzed, representing the foundation for such a power optimization. A practical example of this optimization aimed at the design of a cryo-CMOS microwave driver for single-qubit operations will be

described in Chapter 5.

4

A toolset for qubit/electronic co-simulation and co-design

The generic architecture of a large-scale quantum computer was introduced in Chapter 1 and Fig. 1.2a, and comprises the actual quantum processor, the quantum-to-classical interface, and the upper layers that ensure proper algorithm execution by controlling the electronic hardware [9]. For such large-scale quantum computers spanning several technology domains, appropriate verification methods and tools are required to enable a reliable design flow. Thus, a methodology and related tools for the co-design of the classical electronic interface and the quantum devices are of paramount importance [24], and their introduction is urged by the fast pace at which larger quantum processors are currently being developed [13, 14].

Previous works focused on the upper layers of the quantum computer stack, from the quantum algorithm layer to the quantum execution layers [193], and even down to the micro-architecture [9]. For the layers below, electrical circuit simulators, such as SPICE and Spectre, are well-accepted industry standards. Equivalently, for the simulation of quantum systems, Hamiltonian solvers, such as QuTiP [44], are available. However, the actual interface between classical electronics and the quantum processor has mostly remained unexplored, and, to the best of the authors' knowledge, no tool is available supporting the simulation of classical electronics and quantum systems together.

In this chapter, a co-design methodology is proposed, along with a toolset called SPINE (SPIN Emulator) for the co-simulation of classical electronic circuits and a quantum processor based on single-electron spin qubits (details in Appendix B). Using this tool, circuit designs can be optimized while ensuring qubit performance,

Parts of this chapter have been published at DATE [191] and IWASI [192].

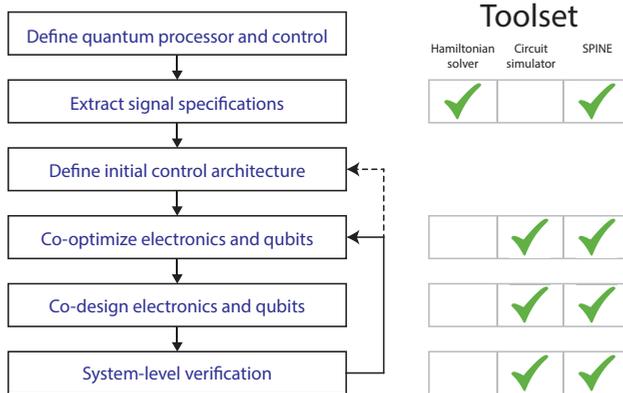


Figure 4.1: Outline of the proposed classical electronic/quantum system co-design methodology along with the tools used in every step.

and exhaustive verification of qubit operations can be performed, although restricted to quantum processors with limited complexity.

The chapter is organized as follows: Section 4.1 introduces the proposed co-design methodology; Section 4.2 discusses the simulation of a quantum system and the implementation of the proposed toolset; design examples are given in Section 4.3; possible future developments are discussed in Section 4.4 and conclusions are drawn in Section 4.5.

4.1. Co-design Methodology

The proposed methodology for the co-design of the classical electronic interface and of the quantum system is summarized in Fig. 4.1. Every step of this process is aided by the introduced toolset, as exemplified in Section 4.3.

The design starts with an underlying qubit technology (e.g. spin qubits) and control methods (e.g. operations driven by the exchange interaction or microwave signals). Next, specifications for the control and read-out signals need to be extracted using qubit simulations such that the desired performance of the quantum processor is obtained (Section 4.3.1). Based on these results, trade-offs between qubit performance and performance of the control electronics, e.g., power and area, can be identified. With this information, an initial architecture can be defined, and an error budget for the different circuit blocks can be drafted. For the chosen control architecture, further co-optimization of the electronics and the qubits can be performed, e.g., optimizing the number of qubits that can be frequency multiplexed over a single electronic channel (Section 4.3.2). Finally, the classical electronics and the quantum processor can be fully designed, and a system-level verification of the final design can be performed (Section 4.3.3).

4.2. Toolset Implementation

The SPINE toolset focuses on qubits in solid-state quantum dots, where the quantum information is encoded in the spin state of a single electron (Appendix B). However, the toolset can be directly extended to other qubit technologies.

4.2.1. Simulation of Quantum Physics

Before diving into the implementation of SPINE, the simulation of quantum systems is reviewed. The quantum state $|\psi\rangle$ (Chapter 1) is evolved by multiplication with a unitary matrix U representing the quantum operation:

$$|\psi_{i+1}\rangle = U \cdot |\psi_i\rangle, \quad (4.1)$$

where ψ_i is the quantum state after operation number i .

In the case of quantum algorithm simulators, U represents one ideal quantum gate, i.e., a simple operation on the qubit, as required for the execution of the quantum algorithm. To find the link between the desired quantum operation (described by U) and the physical behavior of the actual quantum processor, the evolution of the quantum state is found by solving the Schrödinger equation given a Hamiltonian H describing the physical system:

$$i\hbar \cdot \frac{\partial|\psi\rangle}{\partial t} = H \cdot |\psi\rangle, \quad (4.2)$$

where \hbar is Planck's constant. In general, the Hamiltonian is time-dependent and depends on classical signals applied to the quantum processor. For instance, for an isolated single-electron spin qubit, under the excitation of a microwave current $i_{\text{mw}}(t)$ (ESR, Chapter 3):

$$H = \frac{\hbar}{2} \cdot \begin{bmatrix} -\omega_0 & \alpha \cdot i_{\text{mw}}(t) \\ \alpha \cdot i_{\text{mw}}(t) & \omega_0 \end{bmatrix}, \quad (4.3)$$

where α is a constant coefficient that can be determined experimentally, and ω_0 is the precession frequency, i.e., the rotation speed of the electron spin. Finding the exact solution of the Schrödinger equation for an arbitrary current $i_{\text{mw}}(t)$ is not trivial, and numerical simulations are used instead. For every simulation time step, the Hamiltonian parameters are considered to be piecewise constant (Fig. 4.2), and a solution to the Schrödinger equation can be found:

$$U(t_i) = e^{-i/\hbar \cdot H(t_i) \cdot dt}, \quad (4.4)$$

where $U(t_i)$ is the quantum operation for the Hamiltonian $H(t_i)$ valid at time step t_i for a duration dt . For the time step dt , an oversampling of the signal by a factor of 10 has been found to give accurate results for this system.

The overall quantum operation is then found by combining the results from all N time steps:

$$U = \prod_{n=N}^0 U(t_n) \quad (4.5)$$

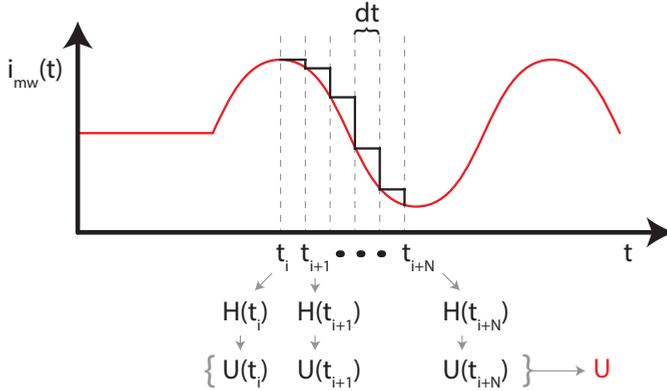


Figure 4.2: All signals are considered piecewise constant in a numerical simulation of the quantum physical system.

For a physical system, the states $|0\rangle$ and $|1\rangle$ represent the energy levels of the system, e.g., in case of a single-electron spin qubit, the energy level of a spin-down electron ($|\downarrow\rangle = |1\rangle$) or a spin-up electron ($|\uparrow\rangle = |0\rangle$). More energy levels could be required to describe all physical effects. For instance, the interaction between 2 spin qubits is mediated by a higher energy level, i.e., the singlet state [66]. The quantum state is updated to incorporate these extra energy levels:

$$|\psi\rangle = \alpha_{\uparrow\uparrow} \cdot |\uparrow\uparrow\rangle + \alpha_{\uparrow\downarrow} \cdot |\uparrow\downarrow\rangle + \alpha_{\downarrow\uparrow} \cdot |\downarrow\uparrow\rangle + \alpha_{\downarrow\downarrow} \cdot |\downarrow\downarrow\rangle + \alpha_{20} \cdot |S_{20}\rangle + \alpha_{02} \cdot |S_{02}\rangle \quad (4.6)$$

where $|S_{20}\rangle$ and $|S_{02}\rangle$ represent the two possible singlet states and α_i the probability amplitudes, restricted to $|\alpha_{\uparrow\uparrow}|^2 + |\alpha_{\uparrow\downarrow}|^2 + |\alpha_{\downarrow\uparrow}|^2 + |\alpha_{\downarrow\downarrow}|^2 + |\alpha_{20}|^2 + |\alpha_{02}|^2 = 1$. In this case, for the quantum algorithm, the basis used in the quantum computation is $|00\rangle = |\uparrow\uparrow\rangle$, $|01\rangle = |\uparrow\downarrow\rangle$, $|10\rangle = |\downarrow\uparrow\rangle$, $|11\rangle = |\downarrow\downarrow\rangle$. Instead of 4-dimensional vectors and matrices, as would be minimally required for 2 qubits, a simulation of the full 6-dimensional quantum state vectors and a 6×6 Hamiltonian is required (Eq. (4.6)).

Thus, it is clear that to simulate a system with more energy levels, or more qubits, the size of the vectors and matrices rapidly grows. Together with the many time steps required for an accurate simulation (Eq. (4.5)), this highlights the challenge of accurately simulating quantum physics.

4.2.2. Hamiltonian Simulations

Following the proposed co-design methodology, quantum simulations are required at various steps that do not necessarily entail the co-simulation of the electronics and quantum physics. For such simulations, a generic Hamiltonian simulator based on Eqs. (4.1), (4.4) and (4.5) can be used, which is also integrated with SPINE (Section 4.2.3).

In this dissertation, Hamiltonian simulations are implemented both in MATLAB and C++. The optimized C++ implementation (with and without multi-threading) uses the Intel[®] Math Kernel Library (MKL). To benchmark these Hamiltonian sim-

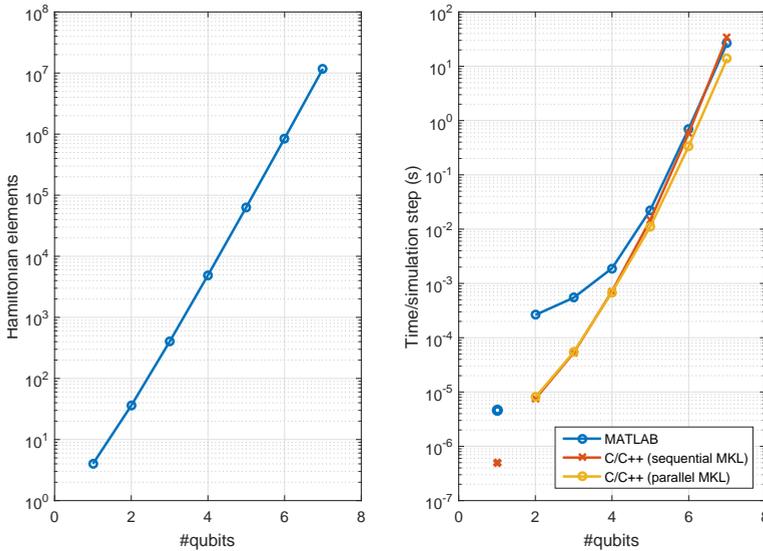


Figure 4.3: The performance of the Hamiltonian simulators in MATLAB and C++. For C++, both implementations with and without multi-threading have been tested.

ulators, an N -qubit system with singlet states included (N electrons in N quantum dots) has been simulated, with a finite tunnel coupling between each pair of quantum dots. The results obtained on an Intel[®] Core[™] i7-4700HQ with 8.0 GB Double Data Rate (DDR)3 Random Access Memory (RAM) are shown in Fig. 4.3¹. It can be seen that the computation time scales with the size of the Hamiltonian under simulation. For small qubit systems, the C++ implementation is up to ~ 35 times faster, and for larger systems almost twice as fast². On our system, a 7-qubit Hamiltonian is the limit, showing a peak memory usage of 2.5 GB in MATLAB and 1.0 GB for the C++ implementation. However, as every qubit operation requires thousands of simulation steps (Section 4.3), simulations may need to be limited in practice to even fewer qubits to avoid excessive simulation times.

4.2.3. SPINE

For the co-simulation of the classical electronics and the quantum system, SPINE was developed. As advanced electrical circuit simulators use quasi-static time-domain solvers, they provide a favorable environment for the inclusion of a time-discrete Hamiltonian simulation. Using Cadence[®] as a framework, the quantum physical system is included in the electrical simulation as a module that takes as input the control signals for the quantum system and outputs the quantum operation U (Figs. 4.4 and 4.5), from which the state probabilities or fidelity can be computed.

¹The results shown for a single qubit system use the analytical solution available for Eq. (4.4).

²MATLAB uses multi-threading.

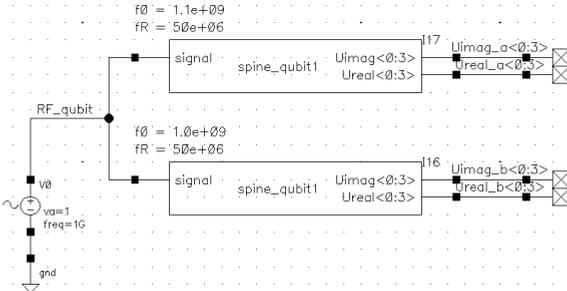


Figure 4.4: Two modules, each emulating one single-electron spin qubit, have been instantiated as a verilog-A module in the electrical circuit simulator; the two qubits are uncoupled and cannot be entangled.

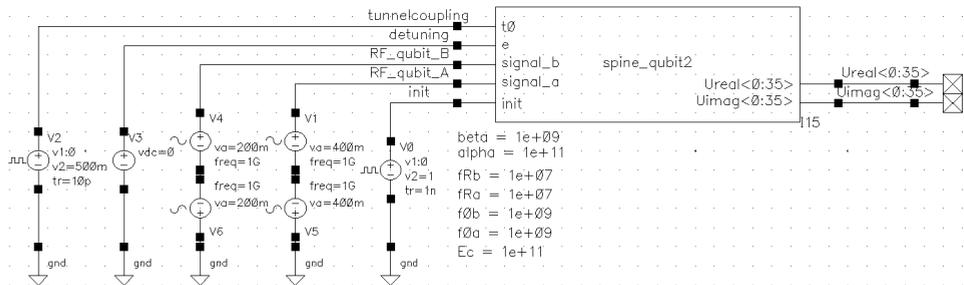


Figure 4.5: A system of two coupled single-electron spin qubits is included as a verilog-A module in the electrical circuit simulator.

Since many circuit simulators support Verilog-A, this language is employed for the implementation of our module. The quantum state is updated during every time step of the electrical simulation, following Eqs. (4.4) and (4.5). Different modules have been written for the different Hamiltonians, as required for a different number of qubits or energy levels. In SPINE, only modules emulating either one single-electron spin qubit (Fig. 4.4) or a system of two coupled single-electron spin qubits (Fig. 4.5) are currently available, but this can be easily expanded in the future.

The inputs to the Verilog-A blocks are the electrical signals applied to the quantum processor (Figs. 4.4 and 4.5; `signal`, `signal_a`, `signal_b` are the microwave signals for single qubit operations, `e` is the detuning of the quantum dots and `t0` the tunnel coupling between two neighboring dots) and `init` to reset the operation to the identity matrix. The resulting complex operation U , which can be used to calculate the operation fidelity, is available at the output with separated real (`Ureal<>`) and imaginary (`Uimag<>`) parts in row-major order. Parameters of the physical system are set as a module parameter when instantiating the module in the circuit schematic (Figs. 4.4 and 4.5; `Ec` is the charging energy of the quantum dot, `f0` the spin precession frequency, `fR` the rotation frequency at 1-V RF-signal, `beta` the tunnel coupling at 1 V and `alpha` the detuning energy at 1 V).

Due to the inclusion of the Hamiltonian simulation in the circuit simulation, the time-step control of a transient simulation is managed by the circuit simulator, which relaxes the time step when tolerable for simulating the electrical circuit with the desired accuracy. This speeds up both the circuit simulation and the Hamiltonian simulation while also recording and taking into account in the quantum simulation events occurring on a very short timescale, such as glitches. On top of this, a maximum time step is set by the Verilog-A module to ensure accurate simulation of the quantum physics by preventing the circuit simulator from relaxing the time step too much.

4.3. Design Examples

4.3.1. Optimization of Power Consumption

As mentioned before, a major concern in scaling quantum computers is the power consumption required by the control electronics. A reduction of the power consumption can be obtained at the cost of the quality of the signal being generated for qubit control, or more errors during qubit read-out. Quantum simulations considering signal non-idealities are required in order to assess the minimum signal quality that can ensure a tolerable error rate in the quantum processor.

As an example, we will consider the control signal required to perform a single-qubit operation for single-electron spin qubits (Hamiltonian in Eq. (4.3)). The microwave control signal $i_{\text{mw}}(t)$ can display several signal non-idealities that reduce the qubit operation fidelity (Chapter 3). Once the effect of each of these errors on the qubit fidelity can be quantified, a larger error budget can be allocated to the sources of error that require more power in the electrical circuit to be mitigated, thereby optimizing the total power consumption while ensuring a certain qubit fidelity.

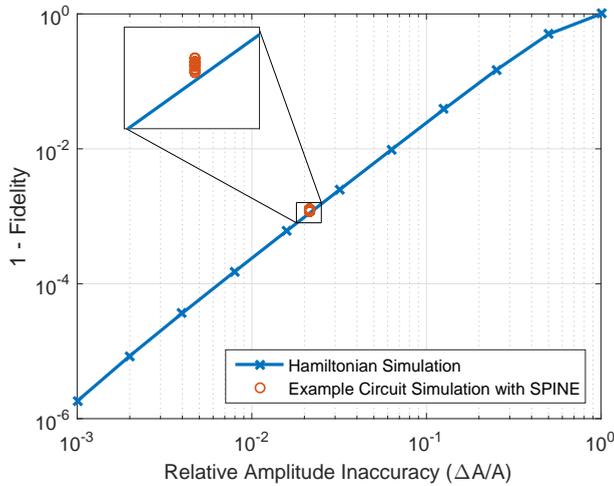


Figure 4.6: The simulated fidelity for an inaccuracy in the microwave amplitude.

As an example, we will focus on deriving the effect of inaccuracies in the microwave amplitude. First, simulations of the quantum system have been performed for control signals with different amounts of amplitude errors ΔA on an otherwise ideal signal with amplitude A , as shown in Fig. 4.6. From this plot, the required signal specifications can be derived when the tolerable qubit infidelity is known. For specific cases, these requirements can also be derived analytically (Chapter 3) [145].

Based on the signal specifications found using these initial simulations and on proper budgeting of the various errors, control circuits can be designed meeting the desired specifications and can be validated using a co-simulation of the electronics and the physical system. As an example, Fig. 4.7a shows an output driver circuit

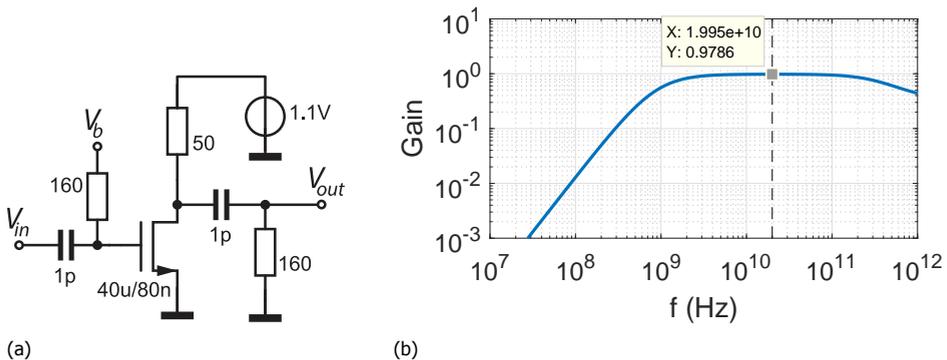


Figure 4.7: (a) An example output driver circuit. (b) The simulated gain versus frequency.

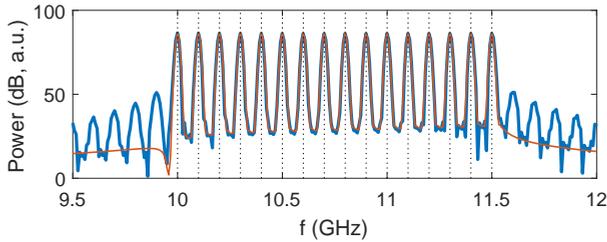


Figure 4.8: The signal applied to the qubit, shown in the frequency domain (red: the unoptimized signal; blue: the optimized signal; black dashed lines: the qubit frequencies).

designed in a standard CMOS technology that can be used to apply the microwave signal to the qubit processor. Due to circuit non-idealities, this circuit has a gain lower than unity (Fig. 4.7b), resulting in an amplitude inaccuracy and a reduced control fidelity. The fidelity obtained using a co-simulation of this electrical circuit in Cadence equipped with SPINE is plotted in Fig. 4.6, which is in good agreement with the results obtained from the Hamiltonian simulation. To illustrate the strength of this co-simulation setup, several simulations have been run with transient noise enabled, and the resulting fidelity of each simulation run is plotted. As expected, the noise further degrades the fidelity, but as it is a dynamic error, each operation has a slightly different fidelity. With this setup, the effect of actual device noise, which consists of different noise types (e.g. flicker noise and thermal noise), can easily be verified.

4.3.2. Finding Optimal Control Waveform

FDMA of qubits can be used to reduce the wiring to the quantum processor, as discussed in Section 2.4.2. In FDMA, each qubit can be tuned to respond to a different microwave frequency, allowing each qubit to be independently addressed using a single driveline. However, a qubit is also sensitive to energy at a frequency close-by its individual frequency, thus, ultimately, the required bandwidth of the control signal is set by how closely the qubits can be spaced in frequency. As covering a larger bandwidth requires more power in the control electronics, an optimization of the signal bandwidth is required.

By optimizing the signal applied to the qubits using a Hamiltonian simulator, this crosstalk effect can be minimized and partially compensated. Following the procedure outlined in [120] for only 2 qubits, as an example, a more elaborate signal optimization has been applied to the simultaneous X-rotation of 16 otherwise uncoupled qubits in 120 ns, starting from a Gaussian envelope. The obtained results are summarized in Fig. 4.8. For the unoptimized signal, the fidelity can be as low as 98 %, whereas this improves to > 99.95 % for all qubits when using the optimized signal, therefore allowing much closer spacing of the qubits, and thus a lower signal bandwidth.

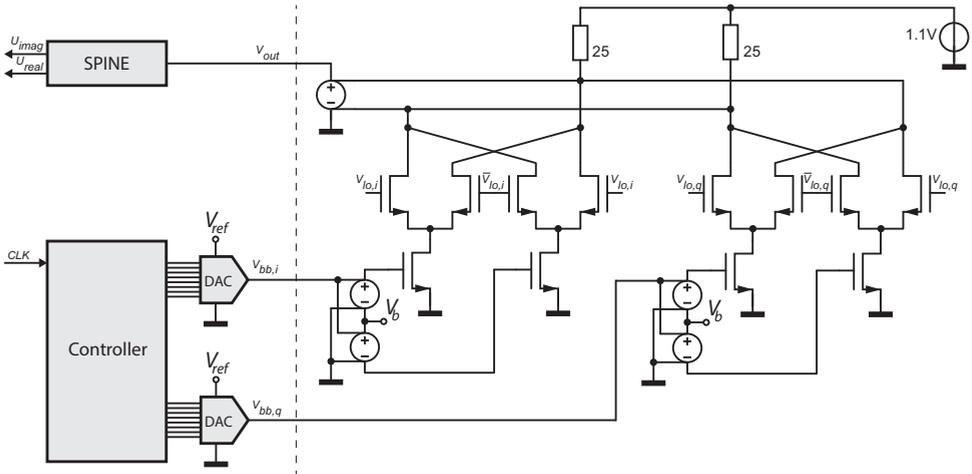


Figure 4.9: Schematic of a full system containing a high-level description of the controller, Verilog models for the DACs, an analog mixer circuit, and finally, SPINE.

4.3.3. System-Level Verification

To show the power of the proposed toolset for the verification of the quantum computer, a full system containing a high-level description of the quantum computer's controller, Verilog-A models of the DACs and an analog mixer circuit has been integrated together with SPINE (Fig. 4.9). The performance was verified by simulating a small quantum algorithm consisting of 4 gates, see Fig. 4.10. It can be seen that in response to the controller, the DACs generate the required in-phase and quadrature-phase signals for the mixer, and the analog circuit performs the required upconversion. In response to this signal, the qubit performs the expected rotations, finally achieving a 99.98% chance of success ($P = |U_{11}|^2$).

4.4. Future perspectives

A first step has been made towards the co-simulation of classical electronics and quantum devices. The proposed toolset is used throughout all the proposed co-design methodology and is, in fact, used in the design of a complex cryogenic integrated circuit for the control of spin qubits (Chapter 5). The toolset is used to evaluate the effect of signal non-idealities, both dynamic, e.g. noise, and static, i.e. inaccuracies, in the various control signals required for performing operations on single or multiple qubits.

However, there are still some opportunities for the design automation community to enhance this toolset. For instance, tremendous speed-ups are required to enable the simulation of larger qubit systems, e.g., to facilitate the inclusion of higher levels of the quantum computer stack for system-level verification. Next, modules can be added for the simulation of quantum systems in other technologies (examples in Chapter 2), and the modules can be enhanced to include quantum processor readout, e.g. with changing electrical properties depending on the qubit state,

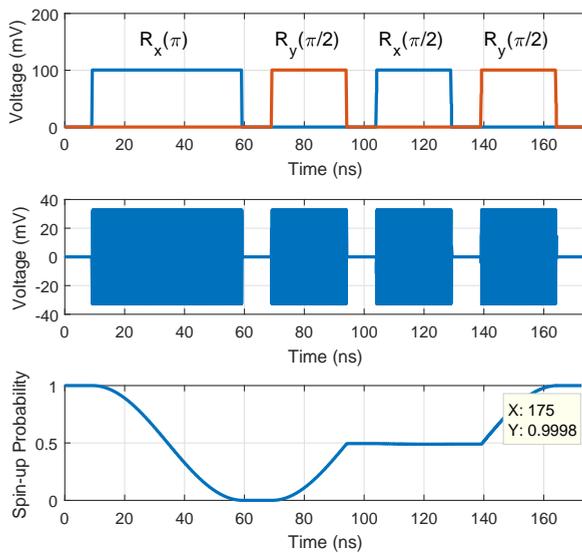


Figure 4.10: Result of the full system simulation, from top to bottom: the voltage at the output of the DACs, the I- and Q-signals, driving the mixer, along with a description of the quantum gate; the voltage at the output of the mixer, driving the qubit; and finally the qubit spin-up probability assuming the qubit was initialized to spin-up (and is not subject to decoherence and relaxation).

which can be the basis for designing readout circuitry.

4.5. Conclusion

A tailor-made electronic interface for quantum processors is required to enable the scaling of a quantum computer to the size required for any practical application. Co-designing the classical electronics and the quantum processor is essential to obtain a complete system that meets the required performance under practical constraints, such as cost, size, power, and reliability. In this chapter, a co-design methodology and the related toolset that meets such demand have been proposed, and it has been demonstrated, via practical examples, how such methodology can be applied to co-optimize, co-design, and verify the classical electronic/quantum system. This represents the first fundamental step towards a full co-design platform. The true effectiveness of the proposed approach is verified through the design of a complex cryo-CMOS integrated circuit for the control of spin qubits presented in Chapter 5.

5

System design of a qubit microwave generator

In order to design custom electronics for the control of large scale quantum processors based on solid-state qubits, such as transmons and spin qubits, circuit specifications need to be estimated/simulated to produce a power-efficient design. This is especially indispensable for circuits operating at cryogenic temperatures, due to limited cooling power of the dilution refrigerator.

In this chapter, we address this by proposing the systematic design of the electronic controller for single-qubit operations, employing frequency multiplexing to reduce interconnects and power consumption. The architecture and specifications of a power-efficient qubit control system are presented, to achieve a single-qubit gate fidelity up to 99.99 % by complying with the signal specifications for qubit control, outlined in Chapter 3 and verified using SPINE (Chapter 4).

In the following, Section 5.1 presents the requirements for the qubit control system. Section 5.2 discusses the trade-offs between possible generator architectures and describes the chosen system architecture. In Section 5.3, the specifications for the different architectural sub-blocks are determined to assess the feasibility of the system. Finally, Section 5.4 demonstrates the flexibility of the design for various applications, and a conclusion follows in Section 5.5.

5.1. Requirements

The main focus of this work is on single-electron spin qubits, since they are very promising both in terms of scaling opportunities and co-integration with CMOS electronics [30], and a co-simulation platform is readily available (Chapter 4). Since the

Parts of this chapter have been published in IEEE Transactions on Circuits and Systems I: Regular Papers [194].

control signals required by spin qubits and transmons are very similar, we will also describe the minor changes required to ensure compatibility with transmons.

5.1.1. System Specifications

In a single-electron spin qubit, information is encoded in the spin of a single electron hosted in a quantum dot. Microwave pulses are required to perform single-qubit gates (Section 2.1.2)¹. Importantly, the phase of the microwave signal needs to remain coherent with the phase of the qubit, which implies keeping a coherent phase for the whole duration of the quantum algorithm, even over different pulses, as changing the relative phase results in a rotation along a different axis. The qubit frequency is typically 12-40 GHz with a microwave-pulse duration in the order of 1 μ s, and, to achieve a typical Rabi frequency of 1 MHz, a power of ~ -45 dBm is usually required at the quantum processor (Section 2.1.2). For future systems, it is desirable to operate at lower qubit frequencies to ease electronic design and higher Rabi frequencies to be less affected by decoherence [32]. Hence, the system presented here will be designed for an output frequency range of 5-20 GHz, and Rabi frequencies in the range of 1-10 MHz, with a maximum rotation angle of π along the X- or Y-axis. This sets the nominal duration of a π -rotation to 50-500 ns. The required output power for spin qubits ranges then from -45 dBm to -25 dBm for the selected Rabi frequency range. However, as attenuators (e.g., with 6 dB loss) are typically employed before the qubits to reduce the heat injected into the quantum processor, and the sensitivity of the qubit can easily vary by $\pm 50\%$, the required output power range is extended to -48 dBm-- -16 dBm (50 mV_r).

As discussed in Section 2.4.2, frequency multiplexing can reduce the amount of interconnects and potentially make the control electronics more area and power efficient. However, pulse shaping techniques must be used to minimize spectral leakage, generally known as cross-talk, to other qubits, and unintended Z-rotations caused by the AC-Stark shift must be corrected (Section 3.3.2). Current experiments on single-electron spin-qubits typically do not use frequency multiplexing, and hence, rectangular envelopes for the microwave pulses are allowed [15, 61]. However, for our system, more complex pulse shaping, e.g. Gaussian envelopes, are necessary to support frequency multiplexing (Section 3.3.2). Moreover, for flexibility, it is desirable to program any envelope, with support of I/Q-modulation for the benefit of having X- and Y-rotations.

Since modern CMOS processes allow processing of extremely wide bandwidths, we aim at the maximum feasible bandwidth to maximize the number of qubits that can be served. Figure 5.1 shows the number of qubits that can be multiplexed in a 1 GHz bandwidth for different microwave pulse envelopes when assuming uniformly distributed qubit frequencies, a π -rotation at the maximum supported Rabi frequency of 10 MHz, and Z-corrections to compensate for the AC-Stark shift (derived from the results presented in Chapter 3). Less than 5 qubits can be served at a 99.9% fidelity with a rectangular envelope. By employing Gaussian pulses, this

¹Two-qubit operations, qubit initialization and qubit readout typically require unmodulated pulses to be applied to the quantum processor (see Section 2.1.2), and are here assumed to be generated by other control electronics and are therefore outside the scope of this chapter.

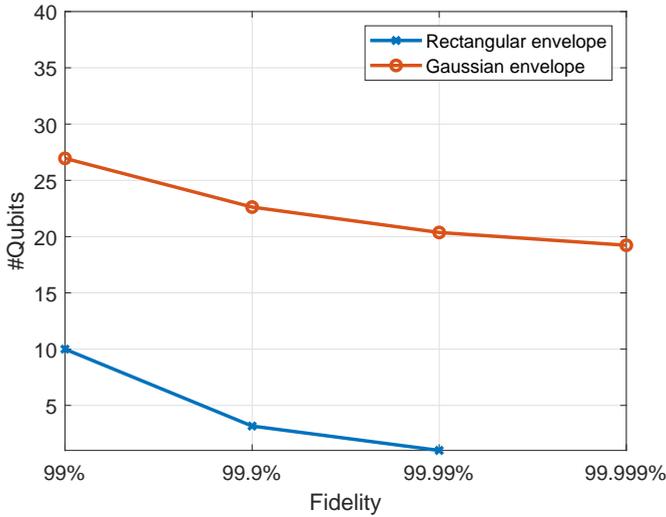


Figure 5.1: The number of qubits that can be allocated in a 1 GHz band when driving with different envelopes and the required Z-correction, each performing a π -rotation in 50 ns.

can be significantly improved, resulting in ~ 40 qubits operating at a 99.99 % fidelity in a 2 GHz bandwidth. Equivalently, to support 40 qubits while using rectangular envelopes, 8 systems each supporting 5 qubits over an individual output are required. While generating Gaussian envelopes comes with some added complexity, this is believed to be more power efficient and scalable than having 8 systems with individual drive lines. For the system discussed here, 32 qubits are targeted in a 2 GHz bandwidth, since this number allows for easy binary addressing of the qubits ($32 = 2^5$), and for a theoretical fidelity $> 99.999\%$.

Even though frequency multiplexing allows for operating on multiple qubits simultaneously, the system will be optimized assuming sequential execution of the operations on the different qubits, as more complicated measures than a simple Z-correction are required when operating on multiple frequency-multiplexed qubits concurrently [120]. However, as a scalable solution is desired, the chosen system architecture should support the simultaneous excitation of multiple qubits.

The fidelity of single-qubit operations is typically above 99 % for single-electron spin-qubits [43]. For fault-tolerant quantum computing, a minimum qubit fidelity, typically around 99.9 %, is required when using error-correction techniques [8]. In order not to limit the performance of the whole quantum computer, the proposed electronic interface targets a fidelity of 99.99 % for a π -rotation performed on a spin-qubit, when taking into account only the errors due to the electronic interface and assuming a perfect qubit. Considering frequency multiplexing, the system will be designed such that both the addressed qubit achieves a 99.99 % fidelity for the targeted π -rotation (which generally gives the lowest fidelity, see Section 3.3) while the idle qubits reach a 99.99 % for the identity, or idle, operation.

Table 5.1: The requirements of the multi-qubit control system.

Qubit technology	Focus on single-electron spin-qubits, support transmons
Qubit frequency range	5 GHz to 20 GHz
Qubit Rabi frequency range	1 MHz to 10 MHz
Output power range	−60 dBm to −16 dBm
FDMA, bandwidth	2 GHz
FDMA, number of qubits	32
FDMA, parallel operations	Supported
Operation, maximum angle	π
Operation, duration	Specifications guaranteed from 50 ns to 500 ns
Operation, modulation	I/Q-modulation with any envelope (nominally Gaussian for FDMA, rectangular otherwise)
Fidelity, addressed qubit	99.99 % for a π -rotation on a spin-qubit
Fidelity, idle qubit	99.99 % for identity operation
Power consumption	Minimum

5.1.2. Extending to Transmons

The control of transmons is very similar to spin qubits, but there are a few key differences that could affect the system specifications (Chapter 2). The qubit frequency is typically around 6 GHz for transmons, and microwave pulses as short as 20 ns are used with a signal power of < -60 dBm. Hence, the duration and output power specifications are extended to include this. Additionally, pulse shaping (DRAG) is typically used to minimize spectral leakage to higher energy levels of the same qubit. This specific pulse requires I/Q modulation, which is already supported to allow X- and Y-rotations seamlessly. Finally, as state-of-the-art transmons typically achieve fidelities not better than 99.99 % [51], the control system should still not limit the achievable fidelity.

5.1.3. Summary

A summary of the discussed specifications is given in Table 5.1.

Following the methods presented in Chapter 3, preliminary signal specifications can be estimated for performing a π -rotation on the addressed spin qubit with either a Rabi frequency of 1 or 10 MHz and a rectangular envelope, see Table 5.2. Equal error contributions are assumed, and the value given for the amplitude inaccuracy assumes a peak amplitude of 50 mV_p , which corresponds to the maximum required output power. These preliminary specifications will be used to assess the feasibility of different proposed architectures. Most notably, a high SFDR is required, as spurious tones could interfere with the idle qubits in a frequency multiplexing scheme.

Table 5.2: Example specifications for achieving a 99.99 % fidelity for a π -rotation

<i>Rabi frequency:</i>	1.0	10	MHz
Addressed qubit:			
Phase imbalance	0.20	0.20	°
Frequency inaccuracy	3.5	35	kHz
Frequency noise	3.5	35	kHz _{rms}
Duration inaccuracy	1.1	0.11	ns
Timing jitter	1.1	0.11	ns _{rms}
Amplitude inaccuracy	0.011	0.11	mV
Amplitude noise	50	50	dB SNR
Wideband additive noise	5.6	18	nV/ $\sqrt{\text{Hz}}$
Idle qubit:			
SFDR	44	44	dB

5.2. System Architecture

Based on the signal requirements for qubit control (Table 5.2), the feasibility of several architectures is discussed, and the chosen architecture is presented in this section.

5.2.1. Analog/RF Section

The most straightforward design of the analog/RF front-end consists of a DAC operating at 40 GS/s (Fig. 5.2(a)), but such a design has a very high power consumption due to its large data bandwidth [195]. Instead, a more power efficient Multiple-Return-to-Zero (MRZ) DAC exploiting higher Nyquist zones capable of synthesizing frequencies up to 20 GHz can be used [196] (Fig. 5.2(b)). In such a design, the output frequency band (centered around $N \cdot f_s$), is directly coupled to the bandwidth, reducing the flexibility as desired when targeting various qubit technologies. Moreover, the output spectrum is corrupted by DAC replicas, which would require a high-order filter to remove. To overcome these concerns, several low-speed DACs along with I/Q mixers (Fig. 5.2(c)), can be used to generate envelopes at distinct frequencies [197]. Such a solution is power/area inefficient for multi-qubit control, as an LO signal per qubit is required, which additionally comes with its own set of design challenges such as frequency pulling. Instead, a DAC at 4 GS/s and a single mixer can be used for controlling multiple qubits from a single RF cable (Fig. 5.2(d)) in a Double-Sideband (DSB) modulation scheme. Moving instead to SSB modulation, the same bandwidth can be achieved with half the DAC sampling frequency. Such SSB modulation can be achieved either by filtering (Fig. 5.2(e)), or by using a modulator with I/Q DAC and mixer (Fig. 5.2(f)). The latter can more easily achieve an Image Rejection Ratio (IRR) > 44 dB for output frequencies close to the carrier (as required by the SFDR specification), by removing the need for filters with a very high order. Note that the architecture in Fig. 5.2(f) still requires a reconstruction filter as alias frequencies may fall in the upconverted 2 GHz output band when the signal bandwidth is comparable to the carrier frequency, as will be

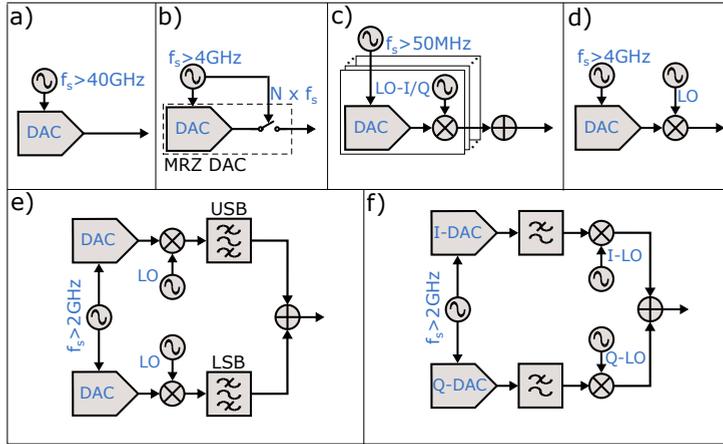


Figure 5.2: Possible transmitter architectures: (a) Very high-speed DAC, (b) MRZ DAC, (c) low-speed DAC with mixer, (d) high-speed DAC with a mixer (e) SSB modulation with bandpass filters for the sidebands, (f) High-speed DAC with reconstruction filter and I/Q mixer.

shown in Section 5.3 (Fig. 5.5). This prevents the use of a more efficient mixing DAC [198] that lacks a reconstruction filter.

5.2.2. Digital Signal Synthesis

A digital back-end is required to generate multiple SSB-modulated tones with such a front-end design. This work assumes the availability of a reprogrammable on-chip memory that is used to store calibrated waveforms for each of the desired qubit rotations. A qubit algorithm is then executed by playing the various stored waveforms in the desired order. The most straightforward approach is to store all possible combinations of qubit instructions in an Static Random Access Memory (SRAM) (Fig. 5.3(a)). The required memory of such an SRAM can be estimated as $SRAM_{\text{mem}} = N \times f_s \times t_{\text{pulse}} \times m^n$, where N and f_s are the number of bits and the sampling frequency of the DAC, respectively, t_{pulse} is the pulse duration, m the number of possible instructions per qubit and n the number of qubits. Assuming an 8-bit DAC operating at 2.5GS/s to address 32 qubits and a maximum pulse duration of 500 ns, it would require an impractically large memory of 3.7×10^{19} bits, considering merely 3 instructions per qubit. Moreover, since qubits require coherent control, intermittent sequential operations on any qubit demand keeping track of the phase of all qubits. Consequently, an individual reference clock would be required for each qubit.

To reduce the required memory, an alternative approach is to store only the amplitude information in the SRAM, which can modulate the amplitude of a sinusoidal waveform with a programmable phase (Fig. 5.3(b)). Following the same scenario, less than 1-Mb SRAM is required (scaling as $m \times n$ instead of m^n). To update the phase for each qubit and ensure coherent control, sine and cosine waveforms scaled by appropriate coefficients can be combined to generate the required phase

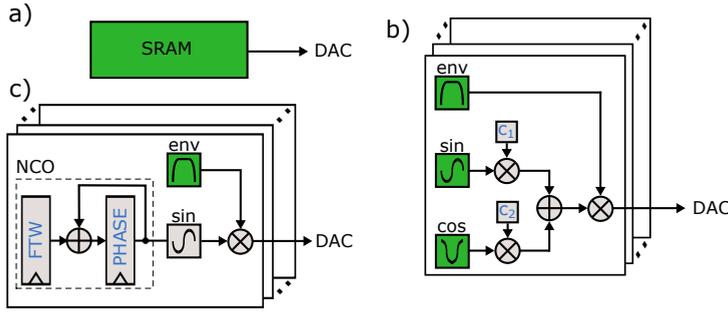


Figure 5.3: Possible backends: (a) SRAM for all possible instructions, (b) Reduced memory for on-chip modulation, (c) NCO-based modulation. The green blocks are programmable memories.

offset. However, this adds an overhead of 2 multipliers per qubit running at the full sampling speed. A more power-efficient approach is to use an NCO, i.e. a phase accumulator running at f_s , for each qubit to generate both the required frequency and phase offset [199], with the added advantage that the NCO can keep track of the phase of individual qubits, thus allowing coherent operation [26]. In such a system, the phase accumulator’s step size, Frequency Tuning Word (FTW), sets the NCO’s frequency ($f_{out} = FTW \times f_s / 2^N$, where N is the number of bits in the phase accumulator). The phase output of the NCO is fed into a sine Look-Up Table (LUT) to generate the corresponding sinewave value, which is then multiplied with an envelope (stored in the SRAM) to obtain the necessary modulated signal, as shown in Fig. 5.3(c).

5.2.3. Final Architecture

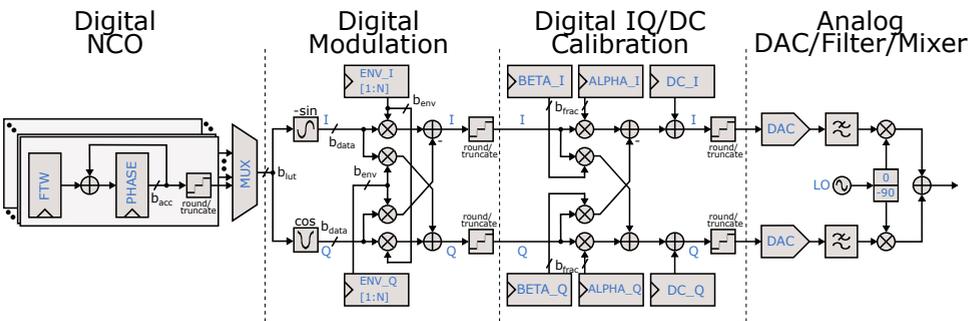


Figure 5.4: Block diagram of the proposed controller.

Considering the above-mentioned trade-offs, a digitally intensive architecture based on DDS with digital modulation, as shown in Fig. 5.4, has been selected. Such architecture benefits from the scaling advantage of advanced CMOS technology nodes in terms of speed and power efficiency and offers the flexibility and robustness of digital signal processing.

Multiple NCOs (one per qubit) are used to keep track of the phase evolution of the qubits. However, the NCO outputs are time-multiplexed to allow operation on one qubit at a time to reduce system complexity, as mentioned in Section 5.1. The multiplexed output is fed into LUTs to generate the sinusoidal signals, which are then modulated by the envelope memory (ENV_I, ENV_Q) for various gate operations and pulse shaping [57, 200], providing flexibility in qubit control.

Because of the stringent IRR requirement of 44 dB (originating from SFDR) corresponding to a maximum phase and gain imbalance of 0.3° and 0.1 dB, respectively, an I/Q digital correction network is required to compensate for analog I/Q mismatch. Moreover, a DC offset correction is added to cancel the LO feed-through to the output.

Finally, the same transmitter as in Fig. 5.2(f), comprising I/Q DACs, reconstruction filter, and an I/Q mixer, translate its digital input to the RF band. The only required analog input is then a quadrature LO signal to drive the mixer.

5

5.3. Circuit Specifications

When increasing the signal's dynamic range, the rate at which the power consumption increases is much lower in a digital circuit than in its analog counterpart, especially in nanometer CMOS technologies [201]. Therefore, the error budget for the digital section is set an order of magnitude tighter than the target fidelity, i.e., it is set to a 99.999 % fidelity, so as to contribute negligibly to the target fidelity of the controller.

To this purpose, a MATLAB simulation model of the entire system is developed, comprising an accurate representation of the digital section (including quantization and rounding effects), an ideal model of every analog block, and a model of the 32-qubit quantum processor (32 individual uncoupled qubits, Chapter 4).

The following calculations are based on a rectangular envelope, while the simulations consider both a rectangular and Gaussian envelope. Moreover, as the specifications are typically stricter when operating at a Rabi frequency of 10 MHz, this will be the default assumption, unless otherwise specified. Besides that, the lowest output frequency band of 5-7 GHz will be used in the simulations as this band suffers more from sampling replicas. When simulating the idle qubits, any Z-error is ignored, as these can be corrected in software [155].

The design strategy is as follows. First, the sample rate is chosen (Section 5.3.1), which then allows for the selection of an appropriate reconstruction filter (Section 5.3.2). Next, the effects of a limited bit length in each digital block on the targeted and idle qubits are individually simulated while keeping the other blocks ideal, i.e., not quantized (Section 5.3.3). The results of this sensitivity analysis are used to select the number of bits required in each block to achieve the targeted fidelity. The final digital system, including all non-idealities, which are simultaneously accounted for, is simulated in a final verification step (Section 5.3.4). Finally, in Section 5.3.5, the specifications of analog blocks can be readily derived from the requirements in Table 5.2.

5.3.1. Sample rate

Due to the chosen I/Q-modulation architecture, there is individual control over the upper and lower sidebands of the upconverted signal. For the required 1-GHz sideband (for a 2-GHz bandwidth), it is sufficient to run the DACs at a sample rate of 2 GS/s to fulfill the Nyquist criteria. However, considering the inherent Zero-Order Hold (ZOH) operation of DACs, the -3-dB bandwidth of a DAC is roughly 40 % of the sample rate. Hence, in this design, a sample rate of 2.5 GHz is chosen, thus resulting in a timing resolution of 400 ps for the microwave envelopes². The shortest operation of 20 ns is then supported (50 points), while the longest operation (500 ns) sets a minimum memory of, e.g., 160 kS, assuming four instructions for each of the 32 qubits.

5.3.2. Reconstruction filter

The lowest qubit frequency of 5 GHz is achieved using a carrier frequency of 6 GHz and the 1 GHz sideband. A sketch of the output spectrum for this condition is shown in Fig. 5.5. The negative frequencies are shown for clarity to illustrate that negative sampling replicas fold back to positive frequencies and eventually fall back close to an in-band qubit. Since the ZOH suppression of the replicas corresponds to a worst-case SFDR of 21 dB, an additional attenuation of at least 33 dB is required at 11 GHz to achieve an SFDR better than 54 dB, as required for 99.999 % fidelity (10 dB more than Table 5.2 for a 10× smaller error). Since a second-order filter is at least required, a 2nd order Chebyshev-I with 3-dB passband ripple and a 1.8-GHz corner frequency was chosen. The combination of the ZOH and reconstruction filter provides an SFDR better than 58 dB in all cases, resulting in a simulated fidelity of the idle qubit of > 99.9996 %.

In addition, the chosen filter improves the in-band flatness to 0.14 dB over the full 2-GHz data band. While this is not a strict requirement, this removes the need to predistort the envelopes. As a result, a qubit driven at 5.1 GHz with a rectangular envelope can achieve a fidelity of 99.999 95 % without any predistortion in an otherwise ideal system. In comparison, a 3rd-order Butterworth filter with a 1.7-GHz corner frequency has an in-band flatness of 2.6 dB, which results in a fidelity of only 99.998 % for a non-predistorted rectangular envelope. This is an important result, as it shows that, with proper design, one does not have to pre-distort the envelope to achieve the intended performance.

5.3.3. Digital blocks

Number of NCO accumulator bits

The number of bits in the accumulator register b_{acc} (see Fig. 5.4) sets the frequency resolution f_{res} of the numerically controlled oscillator according to [199]:

$$f_{\text{res}} = \frac{f_{\text{clk}}}{2^{b_{\text{acc}}}}. \quad (5.1)$$

²The specified duration inaccuracy cannot be guaranteed with a 400-ps timestep. However, as the total rotation angle is set by both the duration and amplitude, such an under/over-rotation error can be corrected by using the amplitude instead (Chapter 3).

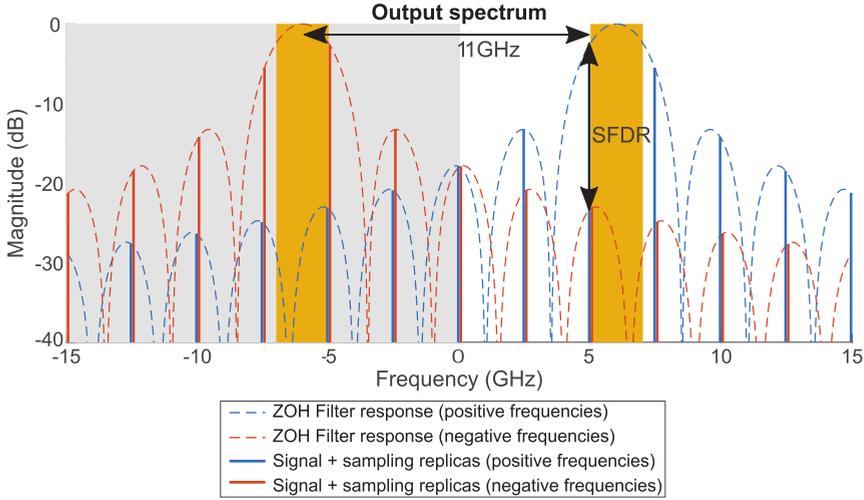


Figure 5.5: Output spectrum assuming an output tone around 5 GHz, a carrier frequency of 6 GHz, a sample rate of 2.5 GHz for an ideal sampler, and no reconstruction filter.

This results in a maximum frequency error $\Delta f = f_{\text{res}}/2$, which results in a theoretical infidelity of

$$1 - F = \left(\frac{\Delta f}{f_R}\right)^2 = \left(\frac{1}{2^{b_{\text{acc}}+1}} \frac{f_{\text{clk}}}{f_R}\right)^2, \quad (5.2)$$

when performing a π -rotation using a rectangular envelope (Chapter 3). This result, along with the simulated fidelity in the case of both a rectangular and Gaussian envelope is shown in Fig. 5.6. In the simulation, the target qubit frequency is chosen such that the frequency error is maximized. As the Gaussian envelope has a longer duration, a larger frequency error is accumulated. At least 16 accumulator bits are required to achieve a 99.999 % fidelity³.

Number of LUT entries

For a more efficient design, the minimum number of entries ($2^{b_{\text{lut}}}$) should be used in the sine/cosine lookup table. However, as this requires the number of bits out of the accumulator (b_{acc}) to be reduced to the number of LUT address bits (b_{lut}), a periodic error would appear, and, as a result, the spectrum will show spurious tones. While the spectrum depends on the generated frequency (see Fig. 5.7a), the spurs are associated with a limited *SFDR* equal to [199]:

$$SFDR = 6 b_{\text{lut}} \text{ dB}. \quad (5.3)$$

³At very small frequency errors, the simulated infidelity deviates from the expected infidelity, as the practical reconstruction filter limits the achievable fidelity.

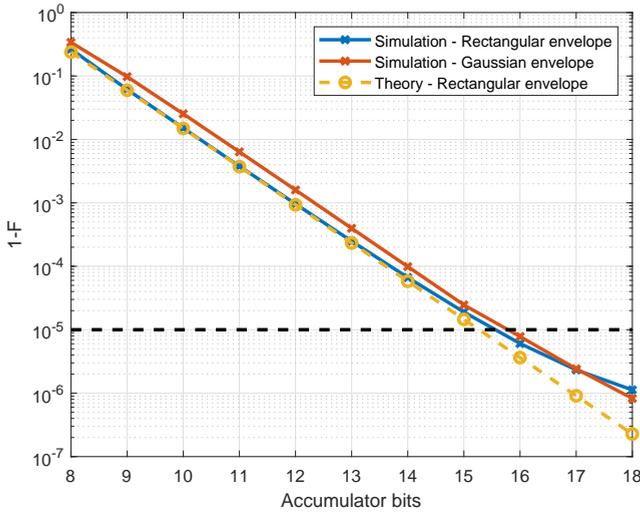


Figure 5.6: Infidelity of a π -rotation as a function of the NCO accumulator number of bits. Eq. (5.2), valid for rectangular envelopes, is plotted as the theoretically expected fidelity.

As such a spurious tone can be at the frequency of an idle qubit, its infidelity is expected to increase to (Chapter 3):

$$1 - F = \frac{\theta^2}{4} \cdot 10^{-SFDR/10} \approx \frac{\theta^2}{4^{b_{lut}+1}}. \quad (5.4)$$

The above theoretical bound is compared to simulations in Fig. 5.8. As the effects of Gaussian and rectangular envelopes are similar, only the results of the Gaussian envelope are presented. Different target frequencies have been simulated, and, in each condition, an idle qubit is considered at the frequency of the largest spur. For the accumulator output bit reduction, both truncation and rounding are considered. Eq. (5.4) well predicts the fidelity of the idle qubit, both for rounding and truncation. In the case of rounding, the idle-qubit fidelity requires at least 9 bits for a 99.999 % fidelity. In the case of truncation, the targeted qubit is affected more, and at least 10 bits are required. When targeting a certain fidelity, the required b_{lut} is one bit less when rounding the accumulator output. Note that, saving 1 bit, is significant as it halves the number of entries required in the LUT.

Number of LUT data bits

A finite number of data bits in the sine/cosine lookup table (b_{data}) results in a quantization error. Generally, such a quantization error can be modeled as white noise spread over the full Nyquist bandwidth $f_s/2$ with associated Signal-to-Quantization-Noise Ratio (SQNR) of [202]:

$$SQNR = 4^{b_{data}} \cdot \frac{3}{2}. \quad (5.5)$$

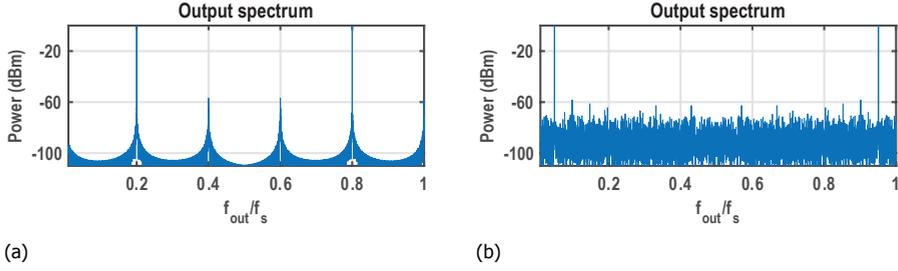


Figure 5.7: Depending on the choice of f_{out} , the spectrum will either show (a) spurious tones when $f_{out} = f_s/N$ (with an integer N , e.g. $N = 5$ in the plot) due to the repetitive behavior of errors, or (b) a white spectrum when f_{out} is not an integer sub-multiple of f_s as the periodic behavior of the errors is disturbed.

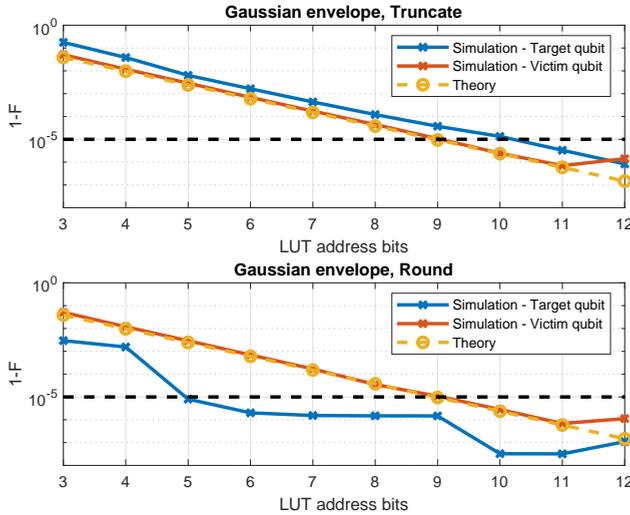


Figure 5.8: The simulated infidelity as a function of the number of LUT entries in case of a Gaussian envelope. The top and bottom plots assume truncation and rounding of the accumulator output, respectively. Multiple plotted lines correspond to the different offset frequencies of 450, 495, 499.5, and 500 MHz, while the theoretically expected fidelity refers to Eq. (5.4).

Since the qubit is only sensitive to noise in a bandwidth $ENBW = f_R \cdot \frac{\pi}{\theta}$ due to the intrinsic noise filtering of the qubit (see Section 3.3, also for a general treatment of a qubit's response to noise), the expected infidelity for the driven qubit is given by:

$$1 - F = \frac{\theta^2}{4} \cdot \frac{1}{SQNR} \cdot \frac{ENBW}{BW} = \frac{\pi\theta}{3} \cdot \frac{f_R}{f_s} \cdot \frac{1}{4^{b_{data}}}. \quad (5.6)$$

This quantization noise affects both the targeted and idle qubits. For certain output frequencies, however, quantization noise is more tonal (similar to Fig. 5.7a), and the spur could be at the frequency of an idle qubit. To capture these different

cases, again, different offset frequencies are used when determining the number of LUT entries, and a victim qubit is simulated at the frequency of the highest spur. In Fig. 5.9, only the simulated fidelity for an offset frequency of 450 MHz is shown for clarity, as the spectrum shows many spurious tones resulting in significant tones affecting the qubit more than expected from the white-noise model. The simulations with the various offset frequencies show that at least 8 data bits are required for a 99.999 % fidelity.

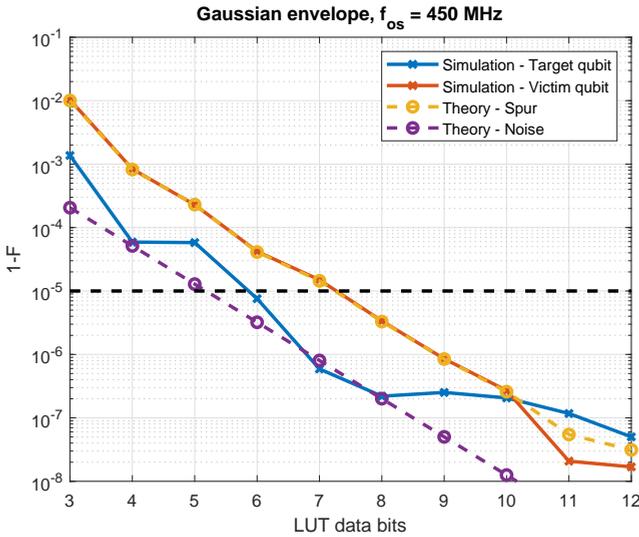


Figure 5.9: The simulated infidelity when reducing the number of data bits in the LUT when using a Gaussian envelope for an offset frequency of 450 MHz. The theoretically expected fidelity due to the spur and noise are given by Eqs. (5.4) and (5.6), respectively.

Number of envelope bits

A limited number of bits used for the envelope in the I/Q-modulation (b_{env} , signed) causes an error in the pulse amplitude. For a rectangular envelope, the maximum amplitude inaccuracy is

$$\frac{\Delta A}{A} = \frac{1}{2^{b_{\text{env}}}}, \quad (5.7)$$

leading to an infidelity of (Chapter 3):

$$1 - F = \frac{\theta^2}{4} \cdot \left(\frac{\Delta A}{A} \right)^2 = \frac{\theta^2}{4^{b_{\text{env}}+1}}. \quad (5.8)$$

While the amplitude could be different for a rectangular envelope due to quantization noise, the shape of the envelope is unaffected. This is not the case for e.g., a Gaussian envelope, where quantization leads to distortion of the envelope, affecting the signal spectrum.

A simulation is performed by setting the qubit properties such that the ideal driving amplitude for a rectangular envelope is in-between two quantization levels. Although the effect of the quantization noise on another qubit may be relevant for a Gaussian envelope, and it is hence simulated as well, the results in Fig. 5.10 indicate that such an effect is negligible. The simulated fidelity follows the prediction of Eq. (5.8), resulting in a minimum of 9 bits for a fidelity of 99.999 %.

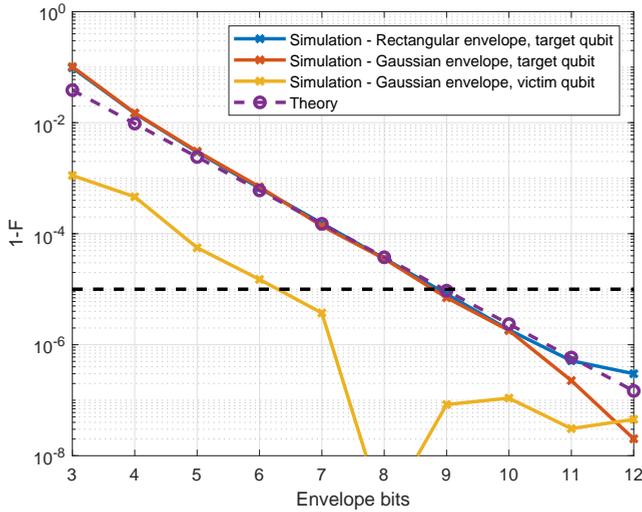


Figure 5.10: The simulated infidelity versus the number of envelope bits when using a rectangular or Gaussian envelope for an offset frequency of 500 MHz.

Number of bits in the correction network

The tolerable phase imbalance (ϕ) and gain imbalance (ϵ) follow from the required image rejection ratio [203]:

$$IRR \approx \frac{4}{\epsilon^2 + \phi^2}. \quad (5.9)$$

For an SFDR of 54 dB, to achieve a fidelity of 99.999 % due to the image spur (see Section 3.3.2), the maximum gain imbalance and the maximum phase imbalance are 0.4 % (0.035 dB) and 0.32°, respectively.

A correction network is added to compensate for inaccuracies in the analog blocks (see Fig. 5.4). In this correction network, the coefficients α_I , α_Q , β_I and β_Q are unsigned fractions of b_{frac} bits⁴.

A gain imbalance can be compensated for by lowering either α_I or α_Q , and for a maximum error of 0.4 %, at least 7 bits are required ($\frac{\Delta A}{A} = \frac{1}{2^{b_{\text{frac}}+1}}$). However, since the relation is non-linear for phase imbalance, both the α and β coefficients need

⁴In case different corrections are required at different frequencies, these coefficients could be selected based on the selected NCO.

to be adapted. As it is difficult to predict the worst-case scenario, a system-level simulation is performed where any phase imbalance from -25° to 25° is introduced and subsequently corrected using a finite number of bits. The situation of the worst-case IRR is further considered when simulating the system along with the quantum processor. The results of this simulation, when using a Gaussian envelope, are shown in Fig. 5.11.

It can be clearly seen that the fidelity of the qubit at the image frequency equals the fidelity as expected from the spur power (Eq. (5.4)). Besides the victim qubit, the targeted qubit seems affected in the same way. From this simulation, it follows that at least 9 fractional bits in the fixed-point number are required to achieve a fidelity of 99.999 %.

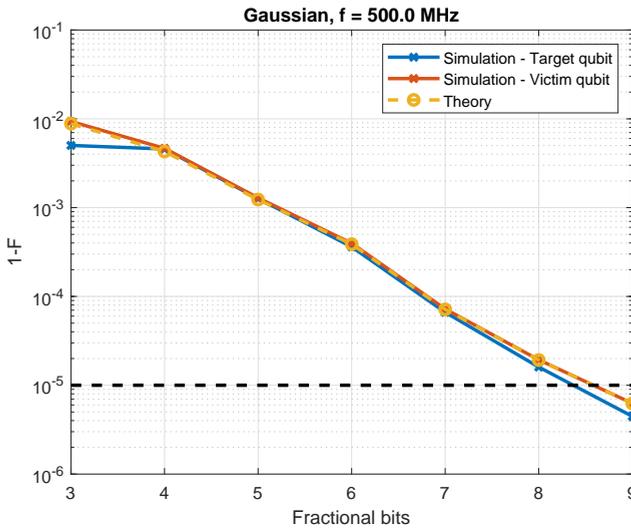


Figure 5.11: The simulated infidelity when reducing the number of fractional bits in the fixed-point number in the I/Q-correction network. Each simulation is performed at the worst-case phase imbalance and uses a Gaussian envelope to drive a qubit at an offset frequency of 500 MHz. The victim qubit is placed at the image frequency of -500 MHz and its fidelity is estimated from the simulated SFDR following Eq. (5.4).

5.3.4. Total digital system

To summarize, for a 99.999 % fidelity, it was found that at least a 16-bit accumulator is required, of which the 9 most-significant bits, after rounding, are used to index the LUT holding 8-bit values. Moreover, both the envelope and I/Q-correction network require a 9-bit resolution. In the sensitivity analysis, only part of the digital datapath under investigation was quantized, and hence all multiplier *outputs* were not quantized. As an initial estimate for the entire digital system, these minimum specifications were used, and all multiplier outputs were truncated to 9 bits, as at least 9 bits were found necessary for the envelope. Reducing the number of multiplier output bits is critical to save power and to find the minimum number of bits

required for the DAC.

A full system simulation was done, where, on each qubit, a Gaussian-shaped microwave pulse was applied to perform a π -rotation at a 10 MHz Rabi frequency. The operating frequencies of the 32 qubits are evenly spaced over the available 2 GHz band. Furthermore, the system was again simulated with the worst DC and I/Q errors. The fidelity of the performed rotation is recorded, as well as the fidelity of all unaddressed qubits, including an additional one placed at the highest spectral spur.

The fidelity of the resulting system was limited to $\sim 99.996\%$ by the unaddressed qubit at the image frequency when truncating the multiplier outputs. After implementing rounding in the multipliers of the I/Q correction network, to reduce the effect of quantization, the fidelity improved to $\sim 99.998\%$, limited by the unaddressed qubit at the highest spectral spur. When increasing the number of LUT entry bits (b_{lut}) by 1, we are at the edge of achieving the desired fidelity. The result of this simulation is shown in Fig. 5.12. Finally, the number of accumulator bits (b_{acc}) is increased to 19 to ensure the required frequency accuracy when operating at the lowest Rabi frequency of 1 MHz. A summary of the specifications is given in Table 5.3.

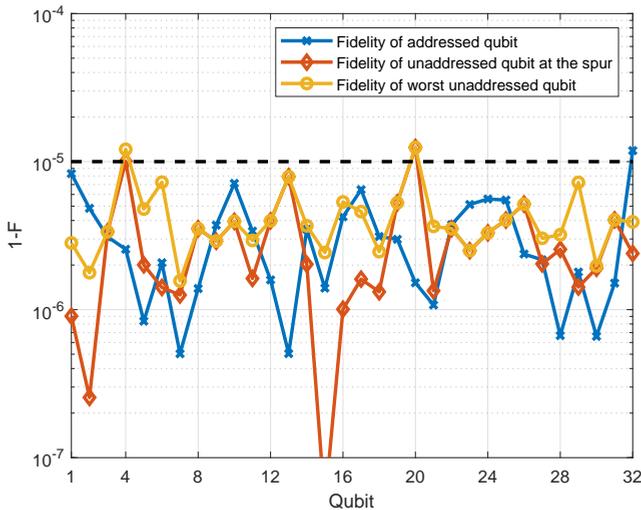


Figure 5.12: The simulated infidelity of the digital system with specifications in Table 5.3.

5.3.5. Analog blocks

The coarse specifications for a 99.99% fidelity (100 ppm infidelity) in Table 5.2 assume an equal contribution from the different errors (~ 10 ppm each), with the previously discussed digital system contributing another ~ 10 ppm to the infidelity. While the assumption of equal error contribution is useful for drafting initial specifications, the trade-offs between these specifications are analyzed in this section

Table 5.3: Specifications for the digital system.

Number of accumulator bits	19 bits
Number of LUT entries	10 bits (after rounding)
Number of LUT data bits	8 bits
Number of envelope bits	9 bits
Result of envelope multiplication	9 bits (after truncation)
Number of bits in the correction network	9 bits
Result of I/Q correction multiplication	9 bits (after rounding)
Envelope memory	~ 100 kS

in order to budget the different errors for feasibility.

As long as the digital clock frequency and analog gain are stable enough, the pulse amplitude, generated frequency, I/Q phase imbalance, and duration can be guaranteed by the digital section. Following Table 5.2, a variable gain of 44 dB with a stability of 0.22 % is required from the analog circuit. The frequency accuracy of 3.5 kHz (for a 1-MHz Rabi frequency and a 20-GHz output) requires a 0.18 ppm frequency stability. Such stability can be achieved by a crystal oscillator [204], and easily satisfies the required duration accuracy of $0.11 \text{ ns}/50 \text{ ns} = 0.22 \%$. Hence, the duration inaccuracy will hardly contribute to the infidelity.

Assuming that the same frequency generator is used to derive the clock and the LO, the tolerable frequency noise (σ_f) can be translated to the required clock jitter (σ_t) as

$$\sigma_t = \frac{1}{2\pi} \frac{1}{f_0} \sqrt{\frac{f_b}{f_a} \frac{\sigma_f}{f_b}}, \quad (5.10)$$

where f_0 is the clock frequency and a phase noise profile of a narrowband PLL with $\sim 1/f^2$ over the frequency range of interest from f_a to f_b ($f_a \ll f_b$) is assumed. A qubit is only sensitive to noise in a bandwidth of $f_b = f_R \cdot \frac{\pi^2}{4}$ for a π -rotation at a Rabi frequency of f_R (Chapter 3). For the case of a 1-MHz Rabi oscillation ($\sigma_f = 3.5 \text{ kHz}_{\text{rms}}$) and a 2.5-GHz clock, this requires an absolute jitter of $\sigma_t < 0.9 \text{ ps}$ ($f_a = f_b/100$ for a total duration of ~ 100 quantum operations). Consequently, the timing jitter requirement of $0.11 \text{ ns}_{\text{rms}}$ is well satisfied, and this error source will hardly contribute to the infidelity. Achieving such a frequency noise is, however, not trivial; assuming the same phase noise profile, a single-sideband phase noise of -116 dBc/Hz is required at a 1 MHz offset from the carrier.

As the maximum output swing of -16 dBm (50 mV_p) can be directly generated by the DAC (and lower output powers by attenuation at the output), no gain is assumed in the following stages; thus, each stage contributes equally to the noise⁵ and distortion. As a representation of those blocks, a single-stage CMOS class-A resistive-loaded common-source amplifier, that can serve as the $50\text{-}\Omega$ output driver⁶, is analyzed in the following.

⁵The DAC quantization noise is already accounted for in the digital specifications.

⁶The same noise analysis is also valid for, e.g., a current-steering DAC.

The maximum Root Mean Square (RMS) output voltage of such an amplifier is given by

$$V_{\text{out}} = \frac{I_d \cdot R_L}{\sqrt{2}}, \quad (5.11)$$

and the RMS output noise voltage by

$$v_n = \sqrt{4 \cdot k_B \cdot T \cdot \gamma \cdot g_m \cdot BW \cdot R_L}, \quad (5.12)$$

where I_d is the bias current, R_L is the load resistance, k_B is Boltzmann's constant, T the temperature, $\gamma \sim 2$ is the excess noise factor for sub-micron devices, and g_m the device transconductance. The SNR follows as

$$SNR = \frac{V_{\text{out}}^2}{v_n^2} = \frac{I_d}{8 \cdot k_B \cdot T \cdot \gamma \cdot \left(\frac{g_m}{I_d}\right) \cdot BW}. \quad (5.13)$$

Assuming $T = 300\text{K}$, a transistor overdrive voltage where $\left(\frac{g_m}{I_d}\right) = 10\text{V}^{-1}$, and the bandwidth for which the qubit is sensitive to amplitude noise $BW = f_R$ (for a π -rotation, Chapter 3), it is found that a bias current $I_d > 0.66\ \mu\text{A}$ is required to achieve the 50-dB SNR requirement with a 10-MHz Rabi frequency. Note that this is easily satisfied as $I_d > 1\ \text{mA}$ is required to obtain the desired output voltage swing over a 50- Ω load.

Assuming a CMOS single-ended amplifier, and an ideal square law device, the 2nd-order distortion is given by [203]

$$HD2 = \frac{1}{4} \cdot \frac{V_{\text{in}}}{V_{\text{gs}} - V_T}, \quad (5.14)$$

where V_{in} is the input voltage, V_{gs} is the gate-to-source voltage, and V_T is the device threshold voltage. Given the requirement of $HD2 < -44\ \text{dB}$, and assuming no gain ($V_{\text{in,max}} = 50\ \text{mV}_p$), an unrealistic overdrive voltage $V_{\text{gs}} - V_T > 2\ \text{V}$ is required. In order not to be limited by HD2, a differential circuit topology can be considered with a 3rd-order distortion of [203]

$$HD3 = \frac{1}{18} \left[\frac{V_{\text{out,p}}}{(V_{\text{gs}} - V_T) \cdot g_m \cdot R_L} \right]^2, \quad (5.15)$$

where $V_{\text{out,p}}$ is the peak amplitude. Achieving $HD3 < -44\ \text{dB}$ requires an overdrive $V_{\text{gs}} - V_T > 0.15\ \text{V}$ (assuming the gain $g_m \cdot R_L = 1$). For a device in saturation, $V_{\text{gs}} - V_T = 2 \cdot \left(\frac{g_m}{I_d}\right)^{-1}$. With a g_m of $\frac{1}{50\ \Omega}$, a bias current larger than 1.5 mA is required⁷. Finally, an SFDR $< -44\ \text{dB}$ requires a DAC with an ENOB of 7.

To summarize, the proposed design requirements are specified in Table 5.4. Of these requirements, the reference clock stability and LO frequency noise requirements appear most stringent. As the duration accuracy, timing jitter, and amplitude

⁷In case of simultaneous excitation of multiple qubits (Section 5.4.2), an $IM3 = 3\ HD3 < -44\ \text{dB}$ is required to meet the SFDR specification and, consequently, a bias current larger than 2.6 mA.

Table 5.4: Specifications for the analog system.

Clock sample rate (f_s)	2.5 GHz
Reference clock stability	0.18 ppm
Clock jitter	0.9 ps _{rms}
LO frequency range	6-19 GHz
LO phase noise	-116 dBc/Hz at 1 MHz offset
DAC Resolution	9 bit
DAC ENOB	7 bit
Reconstruction filter	2 nd -order Chebyshev-I at 1.8 GHz
Analog gain stability	0.22 %
Analog gain control	44 dB
SNR	50 dB ($I_d > 0.66 \mu\text{A}$)
HD3	-44 dB ($I_d > 1.5 \text{ mA}$)

noise specifications of Table 5.2 are most easily satisfied, their error contribution can be reduced to relax the specifications on the more stringent ones to save power. However, the SFDR, as specified in Table 5.2, is not part of any error budgeting as it is the only error source considered affecting idle qubits, and hence this specification cannot be relaxed.

5.3.6. Power consumption estimate

While an accurate estimation of the power consumption requires knowledge of the exact digital and analog circuit implementation, in this section, an estimate is given based on the previously found specifications and implementation examples found in the literature.

A direct digital synthesizer with similar specifications (9-bit amplitude, 2-GHz clock, and 55-dB SFDR), has been implemented in 55-nm CMOS while consuming 25 mW in the 32-bit NCO and 37 mW in the phase-to-amplitude conversion [205]. Considering our system with 32 19-bit NCOs operating at 2.5 GHz and a single phase-to-amplitude conversion block, a power consumption of 640 mW is expected. Similarly, in 65-nm CMOS, a 10-bit multiplier operating at 2.5 GHz consumes 14 mW [206], and hence an additional 112 mW is expected in our digital modulation and I/Q correction network (8 multipliers), bringing the total digital power consumption to ~ 750 mW. Based on the study presented in [207], a power consumption of ~ 160 mW is expected in a 22-nm CMOS node, with 80 % of the power consumed in the NCOs, i.e., 4 mW per NCO. For a design with a single NCO, instead of 32, a total power consumption of 36 mW is expected.

As found from the analog specifications, a single-transistor bias current of 1.5 mA is required to meet the linearity requirement if the entire circuit consists of a single stage. However, a more realistic implementation consists of at least 2 stages contributing to the distortion, e.g., current-steering DACs driving a 50- Ω passive reconstruction filter which in turn drives a double-balanced I/Q mixer driving the 50- Ω output load. Considering a 2-stage implementation, a single-transistor bias current

of $1.5 \text{ mA} \cdot \sqrt{2} = 2.1 \text{ mA}$ is required⁸. As there are 2 stages, each differential, with I/Q, a total current of at least 17 mA is required (17 mW with a 1-V supply).

As about 36 mW is expected for the digital section in case of a single NCO and further reduction in digital power is promised when going to a more advanced CMOS node, the power consumption is well-balanced between the analog and digital section, with another 4 mW required for every NCO, i.e., qubit, that is added.

5.4. Application Example

Compared to state-of-the-art controllers based on general-purpose instruments or tailor-made controllers employing FPGAs (see Sections 2.2 and 2.4.1), the presented solution offers the highest number of frequency multiplexed control channels and is maximally tailored to the quantum processor requirements allowing for a reduced power consumption. Implementing the proposed controller as a CMOS System on Chip (SoC) will reduce its form factor, potentially enabling operating this power-efficient controller physically close to the qubits. The advantages of such a digital-intensive microwave signal generator can be observed by considering application examples for qubit control, as illustrated in this section.

5

5.4.1. Qubit Tune-Up

Besides the intended application of performing single-qubit operations, the control architecture can, for example, be used to tune-up the qubit processor. Part of this tune-up protocol is to find the qubit resonance frequency. The adiabatic fast passage technique uses a frequency chirp to sweep the microwave frequency across the spin resonance frequencies of multiple qubits in an FDMA setup, thereby smoothly rotating all spins whose resonant frequencies lie within the range [208]. Generating such a chirp pulse using the system architecture presented in this chapter can be readily implemented using the following waveform for the in-phase (I) and quadrature-phase (Q) part of the envelope:

$$\phi[n] = 2\pi \cdot \frac{f_{\max} - f_{\min}}{f_s} \cdot \sum_{i=0}^{n-1} \left(\frac{i}{N} - \frac{1}{2} \right) \quad (5.16)$$

$$I[n] = A \cdot \sin(\phi[n]) \quad (5.17)$$

$$Q[n] = A \cdot \cos(\phi[n]), \quad (5.18)$$

for envelope samples $n = 1$ to N , resulting in a chirp from frequency f_{\min} to f_{\max} using N samples (total chirp time $T_{\text{chirp}} = N/f_s$) and amplitude A .

As an example, a binary search for the qubit frequency is shown in Fig. 5.13 using the designed system. Multiple frequency chirps are used over a frequency band that is halved every cycle of the search, narrowing down on the actual qubit frequency indicated by the black dashed line. Only if the frequency chirp contains the frequency corresponding to the qubit frequency, the qubit will perform a rotation. In the first chirp, the frequency is swept over the Lower Sideband (LSB)

⁸Each stage requires a 6 dB stricter HD3 of -50 dB [203].

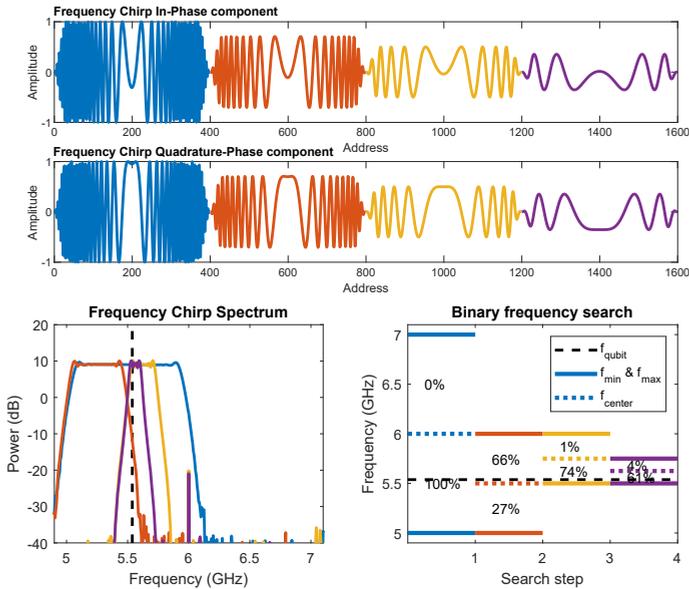


Figure 5.13: Binary search for the qubit resonance frequency using multiple frequency chirps. Top: the in-phase and quadrature-phase envelopes used in the different cycles of the search. Bottom-left: the resulting spectra of the frequency chirps. Bottom-right: the start/stop/center frequency in every cycle of the search protocol, along with the simulated probability of finding the qubit in the excited state after performing the chirp, which determines the next chirp frequency band.

from 5 to 6 GHz, and it is observed whether the qubit rotates or not. In case the qubit rotates, the qubit resonance frequency is in the LSB, and the search continues there; otherwise, the search continues in the Upper Sideband (USB). In order to keep the power spectral density the same when the frequency band is halved, the signal amplitude is gradually reduced with each step.

In the presented example (Section 5.4.1), a linear frequency sweep is implemented. However, any other profile can be implemented as well, which could be more efficient in determining the qubit resonance frequency [208]. Thanks to the high-speed DACs and digital back-end that allows modulation over the full data bandwidth, such frequency chirp can be easily implemented in the presented system.

Besides the qubit resonance frequency, the required pulse duration and amplitude should be determined during tune-up to calibrate the rotation angle. This is typically done by performing a Rabi oscillation where either the pulse duration or amplitude is incremented in small steps, and the resulting rotation angle is measured by fitting the measured Rabi oscillation to a cosine. Due to the option to program any pulse envelope, the pulse duration and/or amplitude can easily be varied to perform such a Rabi oscillation and finalize the calibration of the qubit operation.

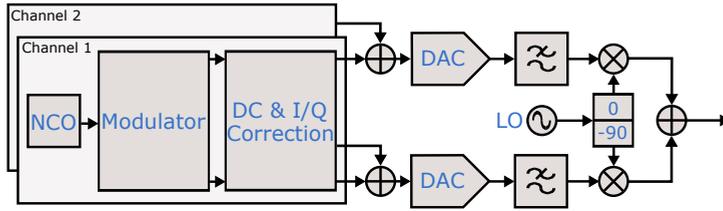


Figure 5.14: Simplified block diagram of the system for the case of 2 simultaneous excitation channels.

5.4.2. Multi-Qubit Simultaneous Excitation

As stated previously, the system is optimized by assuming sequential execution of operations on different qubits. However, the chosen system architecture supports the excitation of multiple qubits simultaneously when having a digital modulator and correction network for each channel (Fig. 5.14). The required DAC resolution increases to:

$$b_{\text{dac}} = 9 + \log_2(N_{\text{ch}}), \quad (5.19)$$

where N_{ch} is the number of simultaneous channels. For the following example, the simulation model is adapted to allow for the simultaneous excitation of 2 qubits with a 10-bit DAC.

When simultaneously exciting 2 qubits using standard Gaussian envelopes (Fig. 5.16, amplitude modulation only), the fidelity is limited when the qubits are close in frequency due to the AC-Stark shift [120]⁹ (Fig. 5.15). This effect shifts the resonance frequency of the qubit when an off-resonance pulse is applied. To account for this frequency shift and to compensate for its effect, phase modulation must be added besides the Gaussian amplitude modulation [120], as shown in Fig. 5.16 (top). The resulting in-phase and quadrature-phase components that are used for the digital modulation are shown in Fig. 5.16 (bottom). With these compensated Gaussian envelopes, a high fidelity can be achieved for 2 qubits spaced closely in frequency while being driven simultaneously in the presented control system (Fig. 5.15).

Thanks to the digital-intensive back-end that allows individual I/Q modulation for each channel, simultaneous excitation of multiple qubits is easily implemented in the presented system.

5.5. Conclusion

Deriving the system specifications of the classical electronic controller for qubits and determining the optimal error budget are crucial in designing power-efficient circuits. To meet these specifications, design trade-offs between several system architectures have been compared in this chapter, resulting in the proposal of an efficient architecture exploiting frequency multiplexing for multi-qubit control. Co-simulation of the proposed electronic system and the qubits was used to assess the effect of non-idealities of each circuit block on qubit fidelity. Based on such

⁹Moreover, additional spurious tones can be present due to intermodulation distortion affecting the performance of other qubits.

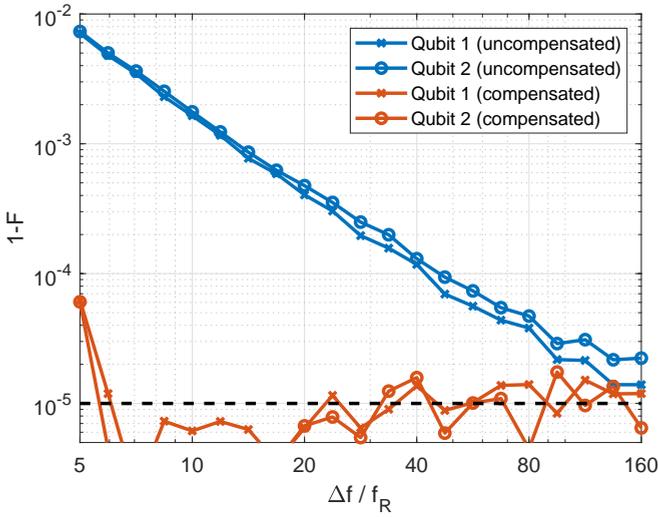


Figure 5.15: Fidelity for the simultaneous excitation of 2 qubits spaced by a frequency Δf when using uncompensated and compensated Gaussian envelopes.

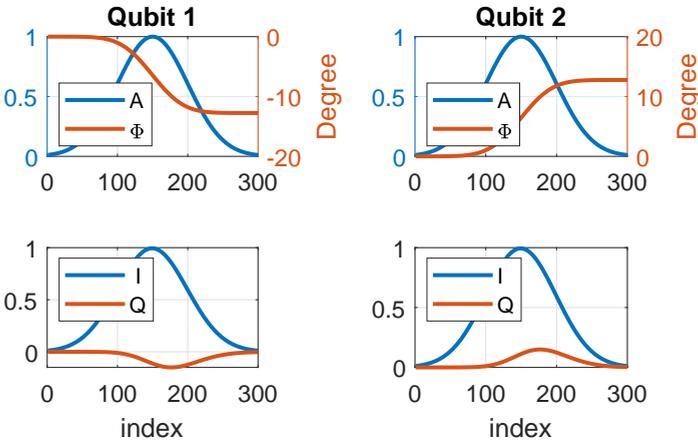


Figure 5.16: The compensated Gaussian envelope for the simultaneous excitation of 2 qubits spaced by a frequency $\Delta f = 5f_R$. Top: amplitude and phase components; bottom: the in-phase and quadrature-phase components.

analysis, the design specifications of each block have been determined to achieve the required gate fidelity while optimizing power consumption. Finally, the effectiveness and flexibility of such a system have been shown by demonstrating relevant practical applications, such as qubit tune-up and simultaneous qubit excitation. As a result of the proposed design methodology, we have obtained the blueprint for a power-efficient integrated electronic controller to realize single-qubit operations for practical large-scale quantum computers.

6

The design of Horse Ridge

Cryogenic circuits must be developed to achieve the goal of co-integrating qubits and control electronics on the same die or package, ultimately operating at the same temperature. A cryo-CMOS pulse modulator for the control of a single transmon qubit has previously been demonstrated [137]. As a first step towards a more scalable cryogenic electronic interface for a large-scale quantum processor, this chapter demonstrates a single-chip cryo-CMOS controller (operating at 3 K) optimized for controlling 128 qubits (operating at 20 mK). The controller minimizes the interface to room-temperature equipment, thanks to the use of multiple cryogenic controllers employing frequency multiplexing that share a single interface to room temperature, as shown in Fig. 6.1.

In this chapter, the focus is on the design of this controller, named Horse Ridge, in Intel 22 nm FinField-Effect Transistor (FET) (22FFL) technology and its electrical characterization. Measurement results of Horse Ridge with qubits are presented in Chapter 7. In Section 6.1 the challenges in designing a scalable cryogenic controller are discussed. Section 6.2 describes the system-level architecture of the controller, and is followed by an in-depth discussion on the design of the digital (Section 6.3) and analog & RF (Section 6.4) circuitry. Measurements describing the electrical performance of the controller while operating at cryogenic temperatures are presented in Section 6.5 and conclusions follow in Section 6.6.

6.1. Challenges

Designing a cryogenic controller for large-scale quantum computing comes with several challenges.

As extensively discussed in Chapter 3, qubits require highly accurate and low-noise microwave control signals to ensure high-fidelity single-qubit operations (cf. Table 3.2).

Parts of this chapter have been published at ISSCC [209] and in IEEE Journal of Solid-State Circuits [210].

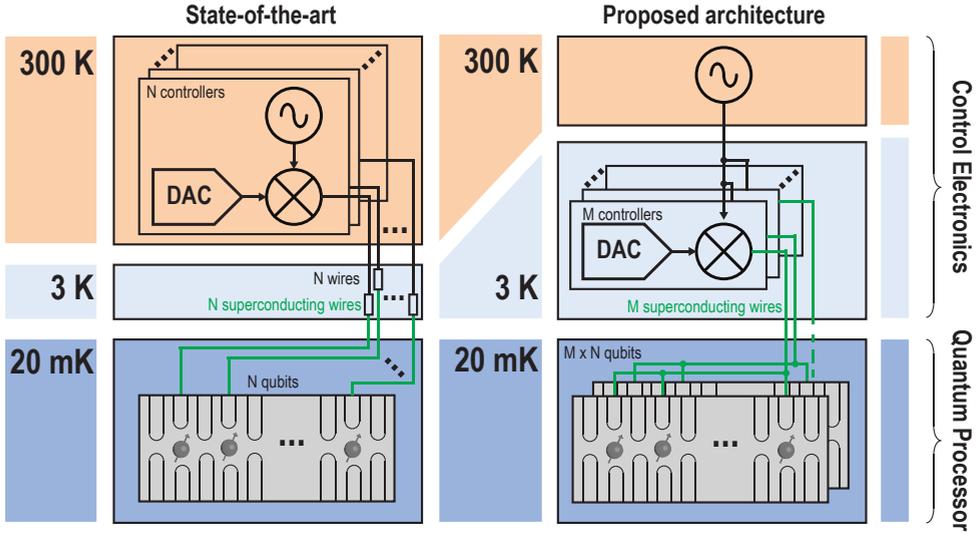


Figure 6.1: Current state-of-the-art controller and the proposed cryogenic controller employing frequency multiplexing and a shared interface to room-temperature equipment.

6

Besides that, accurate control of the phase of the microwave signal with respect to the qubit's phase is essential to perform coherent qubit operations, i.e. rotations around a well-controlled axis, over the entire duration of the quantum algorithm.

Furthermore, as discussed in Section 1.2.1, the cooling power available at cryogenic temperatures in typically employed dilution refrigerators is strictly limited, thus complicating the integration of a large number of high-performance microwave signal generators. In this work, the focus is on the design of a controller operating at 3 K, because of the higher available cooling power. This does not restrict a future co-integration with qubits at the same temperature as the electronics, since 'hot' qubits operating at temperatures above 1 K have recently been demonstrated and are likely to evolve further in the next few years [32, 211, 212].

While cryo-CMOS circuits have been shown to operate down to 30 mK [38], the device characteristics are different at cryogenic temperatures, and no mature models were available at the time of design to accurately predict the behavior of passive and active devices at cryogenic temperatures. Consequently, the circuits need to be designed for robustness against these variations and additional tuning circuitry is required. For instance, a higher threshold voltage is expected for CMOS transistors at cryogenic temperatures, limiting the stacking of transistors in analog circuits [39], unless a supply voltage higher than the nominal is adopted, probably at the cost of reduced reliability, based on the limited research available on cryo-CMOS reliability [213]. As an example, Fig. 6.2 demonstrates the expected change in g_m/I_d -efficiency and linearity of an N-Ch Metal Oxide Semiconductor (NMOS) transistor over temperature. The intrinsic gain is enhanced at cryogenic temperatures, in line with what is reported for other CMOS technologies [39]. It has been shown that device matching degrades at cryogenic temperatures [214]. This directly impacts

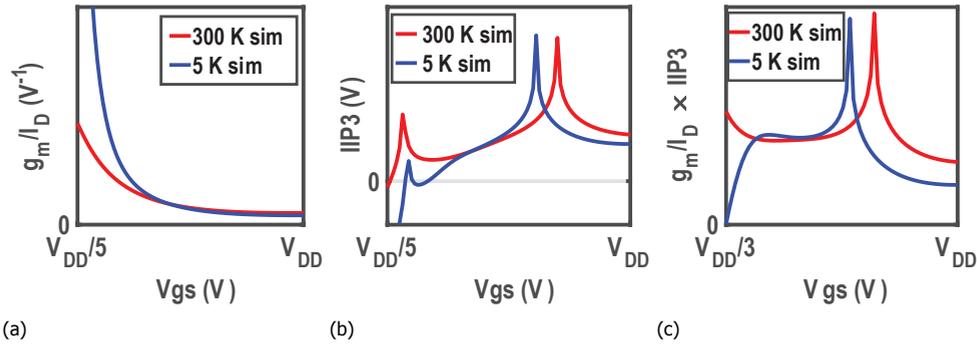


Figure 6.2: Simulated (a) g_m/I_D , (b) IIP3 & (c) $g_m/I_D \times IIP3$ at 300 K and 5 K for the adopted 22-nm FinFET CMOS technology (because of confidentiality reasons the scale has been omitted). The 5 K simulation models were developed from preliminary device characterization performed at 5 K instead of 3 K due to limitations in the probe station temperature control, and are expected to be valid from 3 K to 20 K, as demonstrated in [41] for another CMOS node.

the linearity of ADCs and DACs, and leads to increased offset in differential amplifiers. On the contrary, carrier mobility increases, offering higher driving currents [39], and thermal noise is lower, potentially allowing a lower power consumption. However, the noise power spectral density does not scale linearly with temperature and is only expected to be approximately 10 \times lower at 3 K as compared to 300 K [42]. Some devices are not strongly affected by the cryogenic operation, e.g. the thin film resistors used in this work show negligible change at 5 K compared to 300 K. The capacitance of metal-oxide-metal capacitors and the inductance of on-chip inductors are expected to slightly change at cryogenic temperatures (due to dielectric change and skin-depth change resp.), while the inductor quality factor can double (due to conductivity change and lower substrate loss) [215].

Finally, relocating the controller physically closer to the qubits is advantageous for scaling only if a limited number of control lines from room temperature are required. Hence, all or part of the quantum algorithm execution controller needs to be co-integrated at cryogenic temperatures. To ensure a power-efficient design for such a complex SoC with algorithm capabilities, this design leverages the use of qubit FDMA (Section 2.4.2) [15] to obtain a power-efficient multi-qubit controller. However, employing FDMA introduces several additional challenges. Firstly, the required data bandwidth scales with the number of qubits and the qubit operation speed, ultimately requiring a data bandwidth in the order of 1 GHz. To pack more qubits in the available frequency spectrum, pulse shaping needs to be applied to optimize the spectral content of the microwave pulses. Moreover, a high SFDR is required to ensure that no power is delivered to the qubits that are not addressed at a given time, and a mechanism should be incorporated to efficiently track the phase of all qubits to ensure coherent operations. In addition, phase-corrections must be applied to all qubits nearby in frequency after every operation to compensate for the AC-Stark shift in a frequency multiplexing scheme [155].

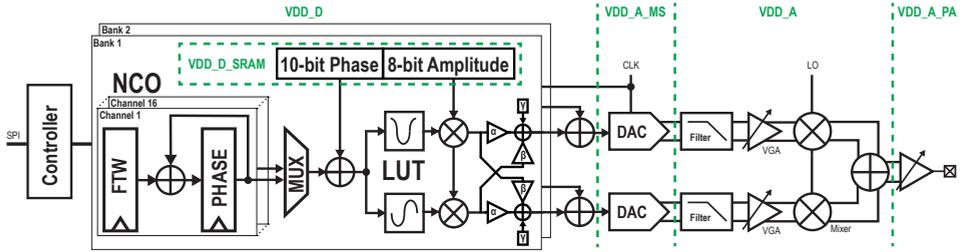


Figure 6.3: Architecture of the controller employing direct digital synthesis in the digital back-end and single-sideband modulation in the analog front-end. The power supply domains are indicated in green, and use the following nominal supply voltages at 3 K unless otherwise stated: $VDD_D = 0.7V$, $VDD_D_SRAM = 0.7V$, $VDD_A_MS = 1.0V$, $VDD_A = 1.0V$, $VDD_A_PA = 1.0V$.

6.2. System architecture & Specifications

The proposed controller contains 4 transmitters, each designed following the methodology presented in Chapter 5 to achieve a 99.99% fidelity for controlling 32 frequency-multiplexed single-electron spin qubits or transmons in the frequency range from 5 GHz to 20 GHz. In order to drive X- or Y-rotations on a single qubit¹ in a frequency-multiplexing setup, power needs to be applied only at the frequency of the qubit that needs to perform an operation, with the amount of power setting the speed of the operation, or Rabi frequency (f_R). To avoid addressing the wrong qubit in a frequency multiplexing scheme, pulse shaping and sufficient qubit frequency spacing are required. It can be shown that a 2 GHz output data bandwidth is required to achieve the target fidelity for 32 frequency multiplexed qubits with an operation speed up to $f_R = 10$ MHz, as a channel spacing larger than $5f_R$ is required (Section 3.3.2).

As discussed in Chapter 5, the architecture employing SSB modulation in the analog front-end and DDS in the digital back-end, as shown in Fig. 6.3, is the most suitable for multiple reasons. Such a system requires only a single LO, setting the desired output frequency band, while supporting multiple qubits at different frequencies in the 2 GHz output band, thanks to the DDS containing NCOs to keep track of the phase of each qubit. To realize a highly-linear RF transmitters with low output power, operating up to 20 GHz, MRZ DAC topologies are typically preferred over a standard RF DAC architecture [196]. However, in comparison, a generic SSB architecture offers the widest data bandwidth, as required to control many qubits, and more flexibility in the choice of the output frequency band independently from the data bandwidth so as to efficiently address different qubit types.

The specifications for the digital back-end and analog front-end to achieve the desired fidelity of 99.99% are summarized in Table 6.1. The 44-dB output power range is dictated by the support of different qubit types, the desired operating speed from 1 MHz to 10 MHz, and the expected variability between different quantum processors. In order to also support transmons, the targeted output frequency range was extended to include 5 GHz to 9 GHz, a range typically used in transmon-based

¹Z-rotations are obtained by updating the tracked phase.

quantum processors (see Chapter 2). The output power has also been made controllable over a vast range as transmons generally require a lower output power. Finally, baseband polar modulation was added to support the generation of complex envelopes, such as DRAG pulses, which are typically employed in high-fidelity transmon control [56]. The wide output power range, output frequency range, and the support for polar modulation of the envelopes ensure the compatibility with both spin qubits and transmons.

Table 6.1: The requirements of the multi-qubit control system, and the specifications for achieving a 99.99% fidelity for a π -rotation.

System	Qubit technology	Single-electron spin-qubits Transmons
	Qubit frequency range	5 GHz to 20 GHz
	Data bandwidth	2 GHz
	Number of qubits	32
	Parallel operations	Max. 2 simultaneously
	Qubit Rabi frequency range	1 MHz to 10 MHz
	Output power range	-60 dBm to -16 dBm
	Analog gain control	44 dB
	Pulse duration	Specifications guaranteed from 50 ns to 500 ns
	Pulse modulation	Polar modulation with any envelope
	Number of RF-outputs	1 for 32 qubits
	Number of high-speed inputs	4 for SPI at ~1 kb/s
	Power consumption	minimize
Digital	Clock frequency	2.5 GHz
	acrshortsnr	54 dB
	SFDR	54 dB
	NCO frequency	22 bit resolution
	LUT address	10 bits (after rounding)
	LUT data	8 bit resolution
	Envelope amplitude	8 bit resolution
Analog	Envelope phase	10 bit resolution
	Correction network (α, β, γ)	9 bit resolution
	Reference clock stability	0.18 ppm
	Clock jitter	0.9 ps _{rms}
	LO frequency	6 GHz to 14 GHz
	LO phase noise	-116 dBc/Hz at 1 MHz offset
	Phase imbalance	0.20°
	DAC Resolution	10 bit
	Reconstruction filter	> 33 dB attenuation > 11 GHz with < 1 dB inband ripple < 1 GHz
	Analog front-end SNR	50 dB in 10 MHz bandwidth
	Analog front-end Third-order Harmonic Distortion (HD3)	-44 dB

The digital back-end is designed for the 10× higher fidelity of 99.999%, in order to relax the analog specifications, and to achieve a proper balance between the expected analog and digital power consumption as shown in Section 5.3.6. Consequently, the number of bits in the data path are optimized to obtain an SFDR and SNR of 54 dB, as required to achieve this higher fidelity. The number of bits in the NCO is chosen to ensure a frequency inaccuracy lower than the frequency noise of state-of-the-art qubits, i.e. 1.9 kHz_{rms}, determined by the nuclear spin noise in isotopically purified silicon [61]. As shown in Section 5.3, with an electronics/quantum

co-simulation of the system considering only the finite number of bits in the digital circuitry, the desired fidelity is obtained for all qubits.

The specifications for the analog front-end are derived following the results presented in Section 3.3 under the assumption of equal error contributions from the different error sources. A reconstruction filter is required to sufficiently attenuate the DAC replicas that can fall in-band after upconversion. Additionally, sufficient in-band flatness should be achieved in order to ease qubit control by removing the need for pre-distorting the microwave pulses.

6.3. Digital Circuit Design

The digital back-end comprises a controller for algorithm execution and for memory management, and a digital signal-generation unit. The signal-generation unit employs a DDS, but, unlike the quadrature modulation as presented in 5.2, polar modulation is adopted to reduce the power consumption by saving two multipliers and an adder. The coefficients used in the I/Q calibration network are selected based on the active qubit channel, i.e. based on the output frequency band, to compensate for the frequency-dependent phase and gain imbalances in the analog circuit. The entire DDS block is replicated to allow for the simultaneous excitation of 2 qubits (Fig. 6.3).

As the two signal-generation units require an input data rate of $(8 \text{ bit} + 10 \text{ bit}) \cdot 2.5 \text{ GHz} \cdot 2 = 90 \text{ Gb/s}$, a quantum algorithm execution controller has been integrated, comprising an envelope memory containing the desired pulse envelopes, an instruction table for each qubit referencing the envelopes, and an instruction list containing the sequence of instructions to be executed (Fig. 6.4). Since a pulse of 500 ns, or 1250 samples at 2.5 GHz, is required for the lowest operating speed of 1 MHz and the largest rotation angle of π , 2560 samples are available per qubit in the envelope memory (40960 samples shared over 16 qubits). The envelopes can be efficiently reused for rotations around different axes, as the axis, and the respective phase shift, are defined in the instruction table, which has 8 entries per qubit to define the instructions. This is expected to be sufficient, as a typical instruction set will contain a limited set of rotations, e.g. a π , $\pi/2$, and $\pi/4$, around the X and Y axis. Moreover, the controller automatically performs the qubit Z-rotations required to compensate for the AC-Stark shift in a frequency multiplexing scheme [120], by applying a phase shift defined from a programmable Z-correction table to all NCOs after each generated pulse. Thanks to this level of digital integration, the external data rate is lowered to $\approx 1 \text{ kb/s}$ using the instruction list during the quantum algorithm execution.

The external interface consists of an Serial Peripheral Interface (SPI) interface for programming the various internal memories and triggering the start of the algorithm execution, and a dedicated 150 Mb/s shift register to quickly trigger the execution of a single quantum instruction, as often as every $\approx 75 \text{ ns}$. As an alternative to the execution of the pre-programmed instruction list (see Fig. 6.4), this operation mode allows for fast feedback and conditional branching in the quantum algorithm execution.

Since a cryogenic model of the standard-cell library was not available at the

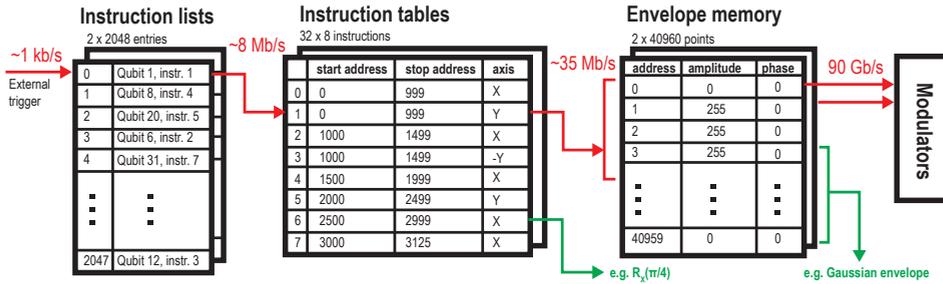


Figure 6.4: Memory organization of the integrated controller comprising instruction lists, instruction tables and envelope memories to gradually reduce the data rate, while finally providing the required 10-bit phase and 8-bit amplitude information to the modulator (Fig. 6.3).

time of design to close timing for synthesis and Automated Place-and-Route (APR), derating factors were implemented to extrapolate the timing behavior at 3 K from the room-temperature models of the standard cell library. The derating factor for the delay of sequential gates is extracted by comparing the simulated oscillation frequency of a 9-inverter ring oscillator at room temperature (using the standard foundry device models) and at 5 K (using a preliminary cryogenic DC device model). Similarly, the derating factors for set-up and hold times are extracted by transistor-level simulations of a standard D Flip Flop. A common derating factor of ~ 1.3 is determined for all cases, implying about 30% reduction in gate delay and set-up/hold times. Using such derating factor for gate delays for synthesis and APR results in effective timing slacks at 5 K for both min and max delay, i.e. timing margins for hold and set-up violations, equal or greater than the values targeted for room temperature. Interconnect delay should also be scaled accordingly when using room temperature models to predict 5 K behavior. From transistor measurement de-embedding data, it is evident that interconnect capacitance does not change significantly at 5 K, while resistance is reduced by about 50%. A $0.5\times$ derating factor is therefore used for the room temperature extracted resistances during APR to model 5 K interconnect delays.

The SoC is implemented as a digital-on-top system with 4 transmitters sharing one common Input/Output (I/O) block. Timing is resolved at 2.5 GHz, with the SRAMs for the envelope memory operating at 1.25 GHz with $2\times$ time-interleaving. The SRAM supply voltage can be controlled independently to ensure correct operation in the presence of an increased threshold voltage at cryogenic temperatures. Standard digital-circuit design-optimization techniques, such as pipelining and time interleaving, along with the aforementioned derating, were used to resolve timing at 2.5 GHz.

6.4. Analog & RF Circuit Design

A current-mode design is adopted for the analog baseband, as the baseband circuitry requires a fairly high bandwidth and linearity (> 44 dB) and the RF mixer (see Section 6.4.4) requires an input current, thereby preventing transconductance

non-linearity. The baseband circuitry comprises a current-steering DAC, a current-mode gm-C reconstruction filter, and a current-mirror-based Variable Gain Amplifier (VGA) feeding the mixer. The filter is discussed first, as it sets the required baseband signal swing to achieve the desired dynamic range while obtaining the lowest power consumption.

6.4.1. Reconstruction Filter

A 2nd-order Chebyshev-I filter with 1.8 GHz cut-off frequency is chosen as it meets the stopband requirement, while its peaking results in the required improved inband flatness near the end of the passband by compensating for the DAC zero-order-hold filter response². A passive implementation of such a filter, as desirable for low noise, distortion and power consumption, would require a prohibitively large inductor of a few nH, limiting future scaling of the controller. Instead, an active current-mode gm-C filter implementation (structure in Fig. 6.5a) is considered [216]. Due to the cross-coupled transistor pair, the impedance at the output is effectively negative and the equivalent single-sided circuit of Fig. 6.5b is obtained, from which the transfer function $H(s)$ and input impedance $Z_{in}(s)$ follow as

$$H(s) = \frac{1}{1 + 2 \frac{C_B}{g_m} s + 4 \frac{C_A C_B}{g_m^2} s^2} \quad (6.1)$$

$$Z_{in}(s) = H(s) \cdot 2 \frac{C_B}{g_m^2} s, \quad (6.2)$$

assuming the same transconductance g_m for all transistors. The transfer function and input impedance of the designed filter are plotted in Fig. 6.5. The ratios C_A/g_m and C_B/g_m set the transfer function and are therefore fixed by the desired filter response.

The linearity of such a filter is limited by the 3rd-order distortion in the transconductance of the transistor, which leads to a non-linear modulation of the capacitor voltage resulting in non-linear components in the capacitor current, and hence in the output current. Therefore, the transistors are biased at an overdrive $V_{gt,opt}$ corresponding to the first peak in the IIP3 plot (Fig. 6.2b) at both 3 K and 300 K as guaranteed by the tunable filter bias current. A high intrinsic gain is also obtained, ensuring an accurate filter transfer function.

For a given linearity, and hence a fixed overdrive $V_{gt,opt}$, the maximum signal current swing scales proportionally to the bias current I_{bias} . This assumes that the filter components are scaled appropriately to maintain the filter transfer function, i.e. by scaling $C_{A,B} \propto I_{bias}$ and the transistor width $\propto I_{bias}$ so that $g_m \propto I_{bias}$. The current noise of the filter is dominated by the bias current sources, which scales $\propto \sqrt{I_{bias}}$. Consequently, the dynamic range of the filter increases by 3 dB when doubling the bias current (I_{bias}), and the minimum bias current to achieve the required dynamic range can be found. Moreover, the required bias current is expected to be $\sim 10\times$ lower at 3 K than at 300 K, as the transistor linearity is not

²The residual phase response does not limit the fidelity

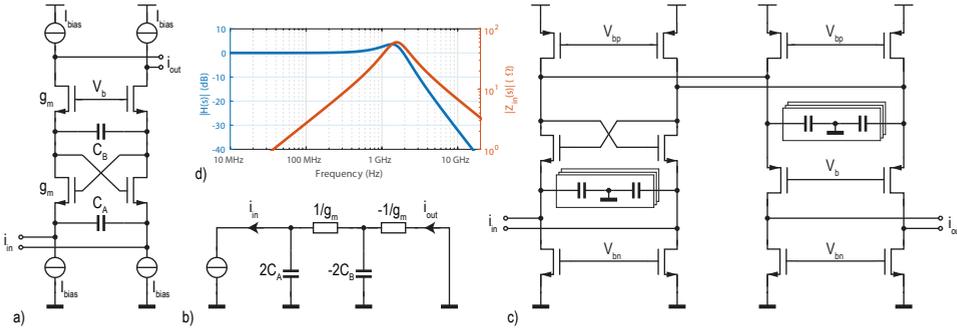


Figure 6.5: a) The structure of the 2nd-order active current-mode gm-C filter implementation, b) its equivalent small-signal single-sided circuit, where the capacitance is effectively negated due to the cross-coupled transistors, c) schematic of the final folded circuit with tunable capacitors, and d) the transfer function [Eq. (6.1)] and input impedance of the filter [Eq. (6.2)].

expected to change significantly over temperature, while the thermal noise power is expected to be $\sim 10\times$ lower at 3 K compared to 300 K. As it is impractical to design the circuit to work over a decade in bias current change, the minimum bias current is chosen for achieving the desired dynamic range only at 3 K, and the resulting bias and signal current is used over the entire temperature range from 3 to 300 K, with a lower expected dynamic range at 300 K³.

Due to the peaking of the filter transfer function, and the DAC sampling replica in the 2nd Nyquist zone, the peak signal swing is about $1.67\times$ higher than the amplitude of the fundamental near the end of the band, requiring a larger bias current for the same linearity. Moreover, as the structure of Fig. 6.5a requires the stacking of four transistors, and as the threshold voltage is expected to increase at 3 K, the structure is folded, resulting in a $4\times$ higher power consumption. The final circuit is shown in Fig. 6.5c. The capacitors are tunable (C_A from 50% to 125% and C_B from 75% to 200% of their nominal value, respectively) to account for the DAC output capacitance and changes in the transfer characteristic at cryogenic temperatures, as the transistor transconductance is expected to increase at 3 K (see Fig. 6.2a). The optimal differential input current of the filter to achieve the required dynamic range at 3 K is $125\mu\text{A}_p$, and is used at both 3 K and 300 K as guaranteed by the on-chip bias current generator (Section 6.4.6). The single-ended input impedance [Eq. (6.2)] peaks to a worst-case 60Ω around the corner frequency at 300 K (Fig. 6.5).

6.4.2. Digital-to-Analog Converter

From the system specifications and the filter design, it follows that a 10-bit current-steering DAC is required, with a unit current of $125\mu\text{A}/2^{10} = 122\text{nA}$ from a P-Ch Metal Oxide Semiconductor (PMOS) current source. Due to the significant overdrive voltage to reduce the effect of threshold-voltage mismatch and noise, a low $g_m/I_d \sim$

³Along with 3 K operation, the controller has been designed to operate at 300 K for convenient circuit validation and debugging.

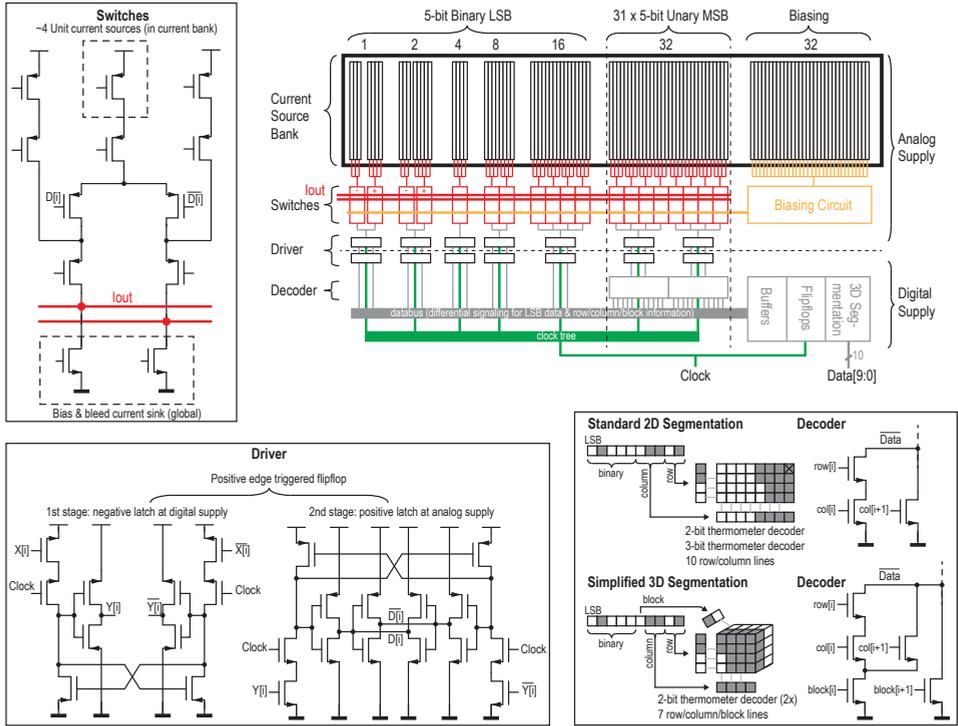


Figure 6.6: Organization of the 10-bit current-steering DAC (top right, with only 1 of the 31 unary cells shown. The layout of the current source bank is not accurately represented in this schematic drawing.), schematic of the switch circuit with bleed currents and cascodes, schematic of the switch driver, and diagram explaining the simplified 3D segmentation extending the 2D row/column decoder with a 3rd dimensions named 'block' (only the single-ended pull-down network is shown of the fully-differential push-pull CMOS decoder).

$5V^{-1}$ is expected at 3K. Moreover, assuming, a typical device noise excess factor $\gamma \sim 2$ for short-channel devices and a pessimistic junction temperature $T = 30\text{K}$ when operating at 3K, the integrated noise in a 10MHz bandwidth is:

$$i_{n,rms} = \sqrt{4kT\gamma \left(\frac{g_m}{I_d}\right) I_d BW \cdot NEF \cdot N} \approx 6.4\text{nA}_{rms} \quad (6.3)$$

for the total DAC with $N = 2^{10}$ current sources, and a circuit noise excess factor $NEF = 2$ to account for the noise from the bias current sources at the DAC output (Fig. 6.6, top left). This corresponds to an expected peak SNR of 77 dB for a single tone, making the noise contributed by the DAC negligible.

The DAC is segmented in 5-bit unary and 5-bit binary sections as a trade-off between Differential Non-Linearity (DNL) and decoder complexity, see Fig. 6.6. A unit current source matching of 0.5% is targeted to achieve a 99.7% yield for 0.5 Least Significant Bit/Byte (LSB) Integral Non-Linearity (INL) [217]. To account

for the expected increase in mismatch at cryogenic temperatures [41], the area of the current sources is doubled. A Monte Carlo simulation shows about 3 dB loss in SFDR due to current source mismatch at 3 K, achieving ~ 56 dB SFDR for a single tone (half the DAC swing). Another important source of distortion is the code-dependent output impedance, leading to [218]:

$$HD3 = \left(\frac{|Z_L|N}{4|Z_o|} \right)^2 \quad (6.4)$$

where Z_o is the output impedance of the unit current source, and Z_L is the load impedance, i.e., the input impedance of the filter ($Z_{L,\max} = 60 \Omega$). Note that scaling the filter to obtain the desired dynamic range does not affect the achievable HD3 of the DAC, since, while Z_L decreases for a larger filter, a larger DAC unit current is required, lowering Z_o accordingly, leaving the ratio $\frac{|Z_L|}{|Z_o|}$ constant. For the stricter two-tone IM3 requirement of 56 dB, $|Z_o| > 290 \text{ k}\Omega$, equivalent to 0.55 fF at 1 GHz, is required. Consequently, bleed currents are used to lower the effective switching impedance [218], see Fig. 6.6⁴.

Due to the very small DAC unit current and constraints in transistor size, there are settling issues. This is resolved by switching the combined current of 3, 4 or 5 current sources using a single switching pair with current bleeding and cascoding. The currents for the two least significant bits are obtained by subtraction of larger currents (i.e. 4 LSB - 3 LSB = 1 LSB and 5 LSB - 3 LSB = 2 LSB) at the output. As the switches are still implemented using minimum size devices, the switch glitch energy is minimized, and a single switch driver can drive up to 4 switches. The switch driver consists of two latches, with the last stage supplied from the analog supply, and with a back-to-back inverter at the output for improved symmetrical switching. For the thermometer decoder, the standard row-column decoder [218] has been extended to a 3D row-column-block decoder, as it only requires trivial 2-bit thermometer decoders, and reduces the number of lines routed differentially (for minimum crosstalk) to the switch drivers (Fig. 6.6). Although the 3D decoder is slightly slower due to the increased number of stacked transistors, it is not a limiting factor for the required sample rate in the adopted technology.

6.4.3. Variable Gain Amplifier

The variable gain amplifier is implemented as a tunable current mirror. An additional output branch feeding a buffer is added to monitor the baseband output signal ($I_{\text{out,test}}$ in Fig. 6.7). The filter output current is $\sim 15\times$ smaller than required by the mixer to generate the maximum required output voltage. Hence, the circuit in Fig. 6.7 is used to provide a 4-bit tunable gain up to $15\times$. As the filter bias current is much higher than the mixer bias current, part of it is sunk at the filter output, while maintaining sufficient VGA linearity, and the residual excess bias current is removed at the VGA output (i.e. the mixer bleed current in Section 6.4.4). Both of these current-bleeding sources are tunable to ensure optimal performance at 3 K.

⁴As the bleed current is about the same as the unit current, the SNR is degraded, but still significantly higher than required.

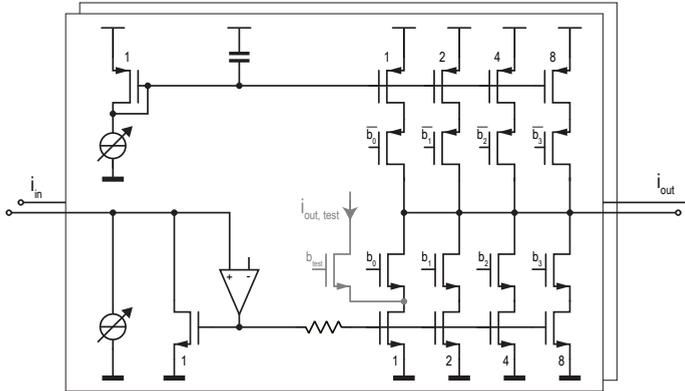


Figure 6.7: The schematic of the variable gain amplifier, based on a current mirror [219].

Due to the reduced bias current in the VGA, and the significantly large output transistor, achieving the required linearity over the full bandwidth is difficult, but it is ensured by adding a single-stage amplifier (PMOS differential pair with current-mirror load) that increases the loop gain and delivers the non-linear current required on the large mirror gate capacitance (as a result of the non-linear gate voltage in response to a linear input current). Finally, to reduce the LO signal leaking back, a $500\ \Omega$ thin-film resistor is added in the current mirror, providing first-order filtering.

6

6.4.4. Mixer

Figure 6.8a shows the schematic of the dual-frequency-band mixer. The fundamental and third harmonic output current (obtained by hard switching) of the mixer are used to extend the operating frequency range compared to a traditional double-balanced Gilbert cell I/Q mixer. Cascode devices are added at the output of the switching devices to steer the output current into either a resistive load for the lower frequency band (O_L), i.e. 2-15 GHz, or an inter-stage matching transformer for the higher frequency band (O_H), i.e. 15-20 GHz.

Lower frequency band

The required bandwidth (15 GHz), and the parasitic capacitance of the output driver, mixer, and their interconnection sets a maximum limit on the load resistance (R_L , e.g., $\leq 70\ \Omega$). The gain of the VGA is chosen such that its output current (I_{sw}) gives the desired mixer output voltage swing ($V_{out,mixer}$):

$$V_{out,mixer} = \sqrt{2} \times 2 \times \frac{2}{\pi} \times I_{sw} \times R_L. \quad (6.5)$$

Current bleeding is implemented in the mixer to remove excess bias current from the VGA output, without sacrificing the required linearity [220]. This tackles the voltage-headroom issue due to stacking of 4 transistors and a resistor, and allows for smaller switching devices, resulting in a lower load capacitance presented to the LO driver and hence a lower power consumption in the LO driver.

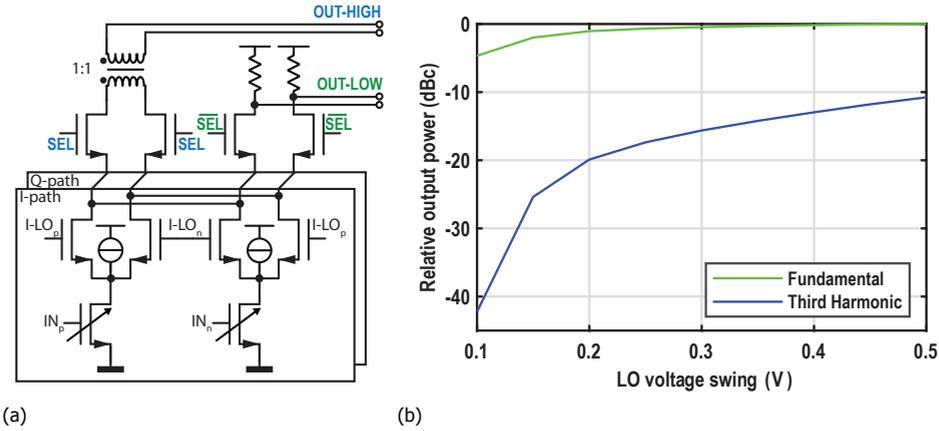


Figure 6.8: (a) Double balanced Gilbert cell I/Q active mixer [219], with two output frequency bands, (b) Schematic level simulation of fundamental and third harmonic output current from the mixer versus LO swing at 300 K.

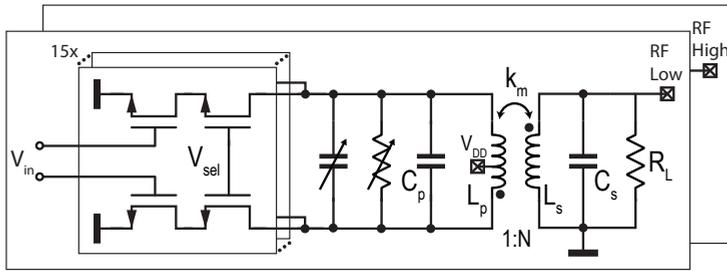


Figure 6.9: Schematic of the output driver used in both the lower and higher frequency band.

Higher frequency band

The output current of the mixer at $3 \times f_{LO} - f_{BB}$ versus the LO swing is shown in Fig. 6.8b. The LO swing is chosen to be 300 mV since a further increase in swing does not significantly improve the conversion gain at the cost of higher power consumption in the LO driver. Note that the third harmonic output current is 15 dB lower than the fundamental at 300 mV LO swing, which is compensated by amplification in the following stages. A narrow-band tuned inter-stage matching network is designed to amplify the third harmonic while attenuating the fundamental tone.

6.4.5. Output driver

Figure 6.9 shows the schematic of the output driver consisting of a class-A amplifier with an output matching network. This design is used for both lower- and higher-frequency-band outputs with different device sizing and matching networks.

The specification of the output driver is to deliver -16 dBm output power (P_{out}) to a 50Ω load, with 50 dB SFDR setting an OIP3 requirement of 9 dBm. Since, $V_{OIP3} = V_{IP3} \times g_m/I_D \times I_D \times R_L$, and both V_{IP3} and g_m/I_D are determined by the

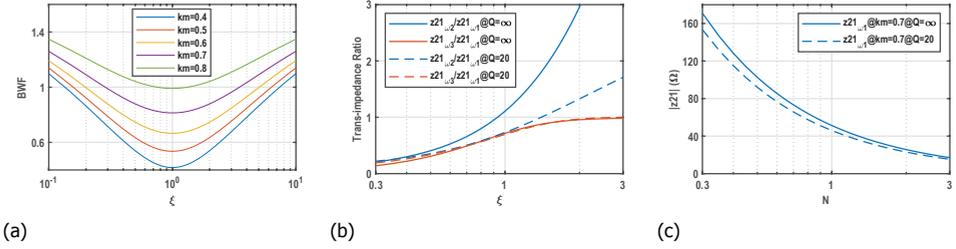


Figure 6.10: Optimization of matching networks: (a) Trans-impedance ratio, (b) BWF, (c) Trans-impedance vs N .

intrinsic device characteristics, the maximum point of this product ($V_{IIP3} \times g_m/I_D$) at 3K is chosen to obtain the required linearity at the lowest power consumption, while taking into account voltage headroom and signal swing, as shown in Fig. 6.2c. Thus, an overdrive voltage $V_{gt} = 0.25$ V has been chosen, leading to $V_{IIP3} = 0.63$ V and $g_m/I_D = 8$. Consequently, the maximum input swing ($V_{in,max}$) to obtain an IM3 of 50 dB can be calculated as

$$V_{in,max} = \frac{V_{IIP3}}{10^{IM3/40}} = \frac{0.63 \text{ V}}{10^{50/40}} = 35 \text{ mV} \quad (6.6)$$

The output matching network can be analyzed as a trans-impedance (Z_{21}) network to convert the drain-current swing of the driver transistor to the required voltage swing at the output. The pole ($\omega_{1,2}$) and minimum (ω_3) frequencies of a matching network can be derived from the maxima and minima of Z_{21} , respectively, given as

$$\omega_{1,2} = \sqrt{\frac{1 + \xi \pm \sqrt{1 - 2\xi + 4k_m^2\xi + \xi^2}}{2C_sL_s - 2C_s k_m^2 L_s}} \quad (6.7)$$

$$\omega_3 = \sqrt{\frac{1 + \xi + \sqrt{1 + 14\xi - 12k_m^2\xi + \xi^2}}{6C_sL_s - 6C_s k_m^2 L_s}}$$

where L_s is the secondary inductance, C_s is the secondary capacitance, $\xi = \frac{L_s C_s}{L_p C_p}$, L_p is the primary inductance, C_p is the primary capacitance and k_m is coupling factor. For a lossless matching network, a flat transfer function (the transimpedance at these poles should be equalized, i.e. $|Z_{21}(\omega_2)/Z_{21}(\omega_1)| = 1$) is obtained for $\xi=1$ [221]. However, the bandwidth factor ($BWF = \frac{\omega_1 - \omega_2}{\omega_3}$) is minimum at $\xi = 1$, as shown in Fig. 6.10a. Hence, to obtain a flat transfer function ($\xi \sim 1$) and high BWF, one has to maximize k_m , which is ultimately limited by the physical realization of the transformer. To further increase the BWF for the maximum attainable k_m , ξ should be increased at the cost of flatness in the transfer function as shown in Fig. 6.10b. However, the flatness can be restored by lowering the quality factor at the cost of passive efficiency.

Figure 6.10c shows the dependency of $|Z_{21}|$ on N ($N = \sqrt{\frac{L_s}{L_p}}$). A higher $|Z_{21}|$ or lower N increases the equivalent resistance seen by the driver transistors. Hence, a relatively lower current swing can produce the same output voltage swing. This in turn would demand lower DC bias current and improve the efficiency. Moreover, this leads to smaller transistors and consequently, higher bandwidth of the mixer due to lower input capacitance presented by the output driver. Since ξ is already fixed by the flatness and BWF , minimizing N would require maximizing C_s and minimizing C_p , as $N = \sqrt{\xi \frac{C_p}{C_s}}$. The minimum value of C_p is determined by the parasitic capacitance of the output driver, while the optimum C_s can be obtained from the value of loaded quality factor of the secondary side ($Q_L = R_L C_s \omega$) that maximizes the passive efficiency of the matching network at a given frequency [222]. Finally, $N = 0.8$ is obtained.

An increase in the quality factor (Q) of a transformer by a factor of ~ 2 expected at cryogenic temperatures, due to lower substrate losses and a higher metal conductivity [215], can affect the flatness of the transfer function. The transfer function can shift towards higher frequencies due to a reduction in effective inductance and capacitance of the transformer at cryogenic temperatures [215]. To compensate for these variations that are not easily predictable, capacitor and resistor tuning networks were implemented at the windings of all matching networks.

To maintain a better efficiency at lower output voltage swing, a gain control of 24 dB is achieved by selectively switching 15 unit cells, each consisting of a class-A amplifier and cascode transistor. To further improve the power efficiency, the supply voltage of the driver is lowered without significant impact on linearity, since the required output voltage swing is significantly lower than the supply voltage.

6.4.6. Auxiliary circuits

An overview of the main circuit blocks is given in Fig. 6.11. An LO driver, a clock receiver and a constant- g_m bias circuits are also implemented in each transmitter (Transmitter (TX)). Four transmitters are integrated into a single chip to increase the number of qubits that can be controlled, and to allow for the simultaneous control of 4 qubits at the same frequency through individual transmitter outputs.

LO driver

An LO driver with 20 dB voltage gain and 15 GHz bandwidth is designed to deliver the required voltage swing to the mixer, while incorporating single-ended to differential conversion. On-chip co-planar waveguide transmission lines are used to connect the input of the LO driver to the I/O bumps. This allows to reduce phase and gain imbalance by allowing the LO driver output to be abutted to mixer switches. Figure 6.12a shows the schematic of the LO driver. The first stage serves as an active balun converting a single-ended signal to a differential signal while providing wideband input-impedance matching [223]. For proper operation, the input matching is achieved by adjusting M_1 gate bias such that $1/g_{m,M_1} = 50 \Omega$ and by setting the gain of the common-gate (CG) path $g_{m,M_1} \cdot R_{CG}$ equal to the gain of the common-source (CS) path $g_{m,M_2} \cdot R_{CS}$.

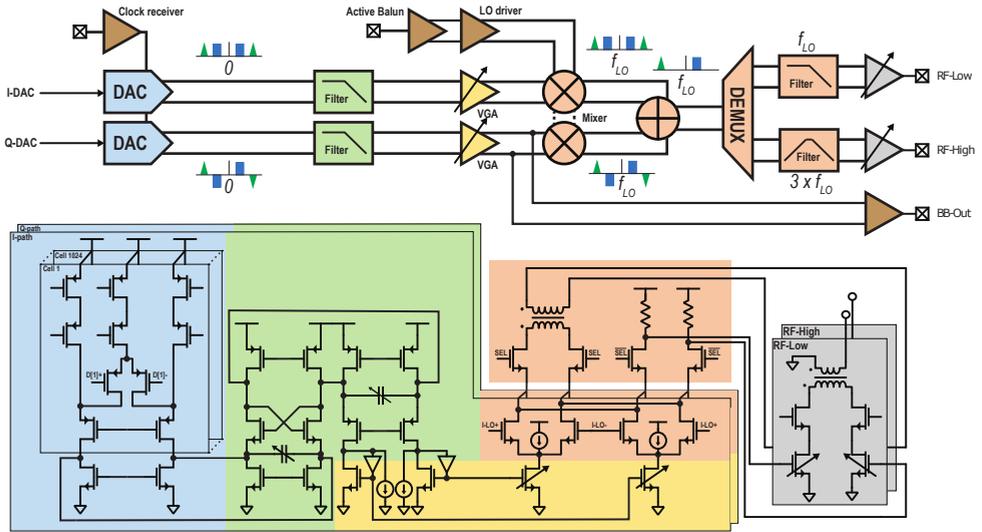


Figure 6.11: Block diagram and transistor level schematic (bias circuitry not shown).

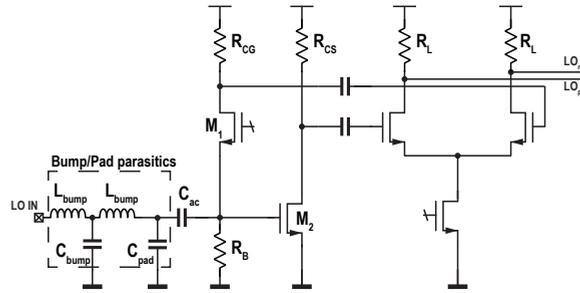
6

The required gain of $5\times$ at 15 GHz sets the required gain-bandwidth to be 75 GHz. For the active balun to directly drive the mixer switches, a load capacitance $C_L = 40$ fF (due to parasitic capacitance of mixer switches, M1/M2 devices, and routing traces) limits the maximum load resistance to $180\ \Omega$ and consequently, the gain to 3.6. Hence, to achieve the required gain-bandwidth, a high-speed differential Current-Mode Logic (CML) amplifier stage is cascaded to the first stage.

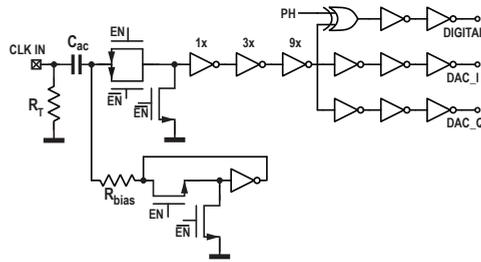
The required phase noise specification of -116 dBc/Hz at 1 MHz offset from the carrier is achieved over the entire frequency range with a power consumption of 7 mW for both I&Q branches.

Clock receiver

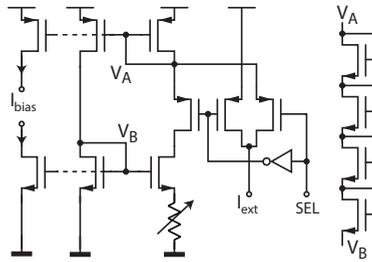
A clock-receiver circuit powered from the digital supply provides the rail-to-rail-swing clock signals for the DAC and the digital blocks. All supplies are substantially decoupled on-chip to reduce the supply noise feedthrough between different circuit blocks. To share a single external clock signal between all 4 transmitters, each transmitter is AC coupled with an input termination of $200\ \Omega$ (R_T) to present an equivalent input impedance of $50\ \Omega$. A self-biased inverter with power-down option and a transmission gate are employed to individually switch off the clock receiver in each transmitter while preventing feed-through during the off state. A half-period time shift can be introduced between the clock fed to the digital circuits (DIGITAL) with respect to the DAC (DAC_I, DAC_Q), enabled by a digitally controlled on-chip register PH via an Exclusive OR (XOR) gate, as shown in Fig. 6.12b. This can address any potential data timing issue at the digital/DAC interface due to layout mismatch and changes in digital propagation delay at 3 K. A fan-out of 3 is maintained at each stage to obtain the required jitter.



(a)



(b)



(c)

Figure 6.12: Auxiliary circuits: (a) LO driver (b) Clock receiver (c) Constant- g_m bias circuit [219].

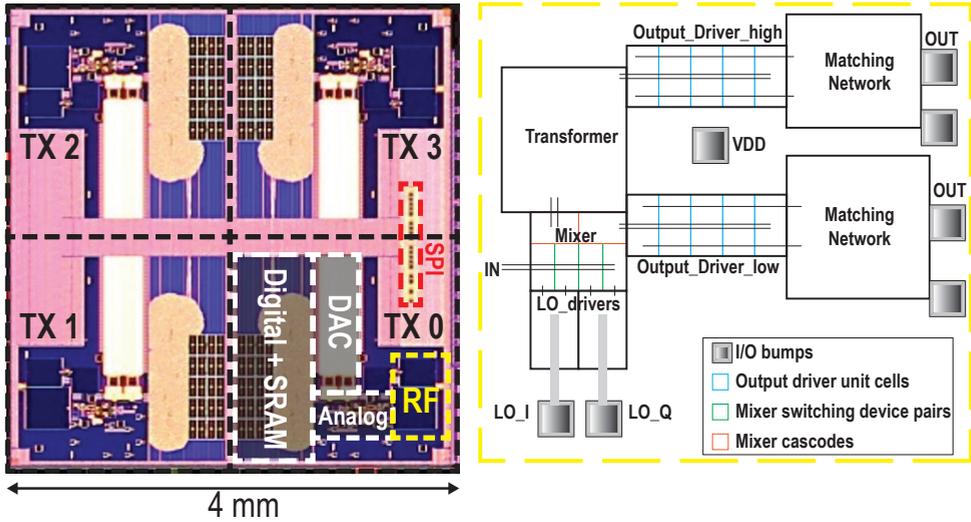


Figure 6.13: Chip micrograph and RF circuits top level layout.

6

Bias circuit

The bias currents are generated by a standard constant- g_m circuit, see Fig. 6.12c. The desired $g_m = 1/R$ is set by a tunable resistor, that allows the output bias current, I_{bias} , to be adjusted over a range of 50% to 200% relative the nominal value at 300 K, to ensure the same signal current at 3 K and 300 K, while accounting for changes in the device transconductance at 3 K. A stack of 4 diodes, between nodes V_A and V_B , is used to start-up the bias circuitry. An externally applied bias current, I_{ext} , can also be selected, and used to start-up the circuit if the stack of diodes is not sufficiently strong due to the increased threshold voltage at cryogenic temperatures.

6.5. Cryogenic Electrical Performance

Figure 6.13 shows the micrograph of the chip fabricated in Intel 22 nm FinFET (22FFL) technology [224]. The transmitter architecture shown in Fig. 6.3 is replicated 4 times (TX0...TX3) with each instance occupying an area of 4 mm^2 with a single shared SPI controller on the die.

6.5.1. Measurement setup

The chip is placed on the 3 K plate of a dilution refrigerator. Dual-pole-dual-throw (Dual Pole Dual Throw (DPDT)) microwave switches are used in the fridge to select the chip or the room-temperature signal generator on one side, and the qubit device or the room-temperature spectrum analyzer on the other side (see Fig. 6.14). This enables proper characterization of the chip performance and the comparison of the qubit control by the room-temperature equipment and the designed chip. An FPGA is used as the master to synchronize the chip with the other instruments used for

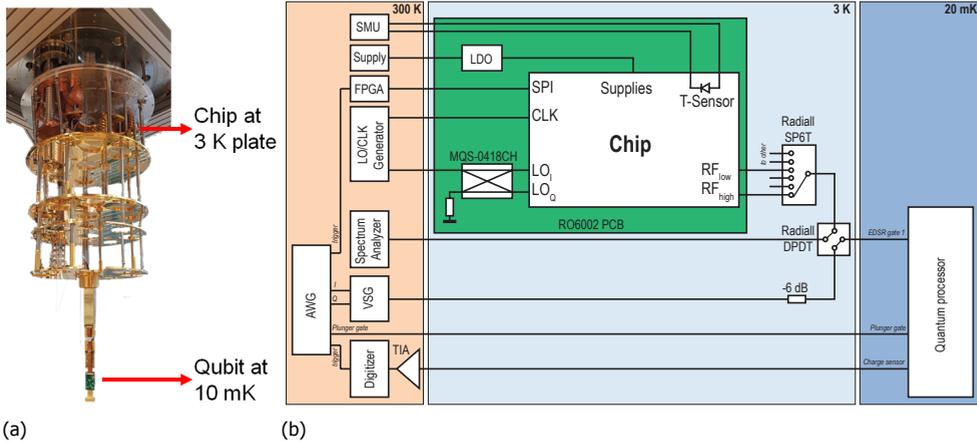


Figure 6.14: Measurement setup: (a) placement in dilution refrigerator (b) schematic representation. The employed Low Dropout (LDO) is custom designed using discrete components (AD8086 opamp with TSM2314 Metal Oxide Semiconductor Field-Effect Transistor (MOSFET)). The Band-Pass Filter (BPF) used in the measurement presented in Fig. 7.2 is placed directly at the RF_{low} output, before the Single Pole 6 Throw (SP6T) switch. A Niobium Titanium (NbTi) coax cable, without attenuators, is used between the 3 K and 20 mK stage.

qubit readout and initialization (more details in Section 7.1).

The die is flip-chip bonded to a Ball Grid Array (BGA)-324 package with impedance-matched traces and on-package discrete capacitors for supply decoupling. A 6-layer Printed Circuit Board (PCB) is designed to route the RF signals on the top layer with RT/duroid 6002 microwave substrate and DC signals on the bottom layers with FR4 dielectric. The solder-mask areas on the top and bottom layers are minimized to allow better heat transfer. To reduce the number of cables between the room-temperature LO generator and the chip inside the dilution refrigerator, each LO line is shared between two transmitters. A custom-designed Wilkinson power divider (Wilkinson Power Divider (WPD)) on the PCB with discrete wire-bonded quadrature hybrids were used to generate the required LO signals for the transmitters, as shown in Fig. 6.15. All the above-mentioned components were individually tested at 3 K to verify their performance.

A gold-plated copper enclosure housing the PCB acts as a heat sink for proper thermalization of the chip to the 3 K plate in the fridge, as shown in Fig. 6.15. Indium foils were sandwiched between the die and the enclosure to maximize the contact surface area and minimize thermal resistance. Thanks to its high malleability compared to other metals, Indium can compensate for the mismatch of the thermal expansion between the two mating surfaces (silicon and gold) at cryogenic temperatures.

To monitor the die temperature, on-chip diodes were placed across the chip, as shown in Fig. 6.16. These diodes are calibrated using an external silicon diode temperature sensor (with an accuracy of 0.25 K) mounted close to the enclosure, with the chip powered down. Figure 6.16 shows the junction and plate temperature

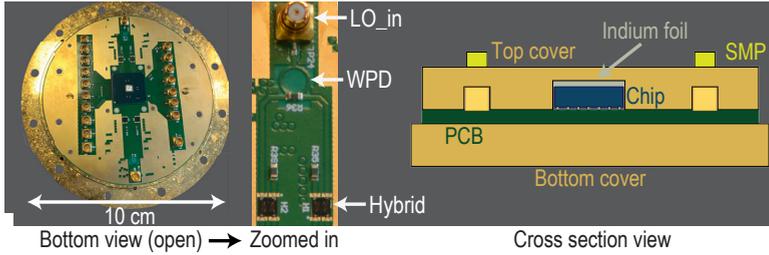


Figure 6.15: From left-to-right: the bottom view of the opened gold-plated copper enclosure, a zoom-in on the PCB showing the WPD and discrete hybrids for the LO input, and a drawing of the cross-section of the full enclosure.

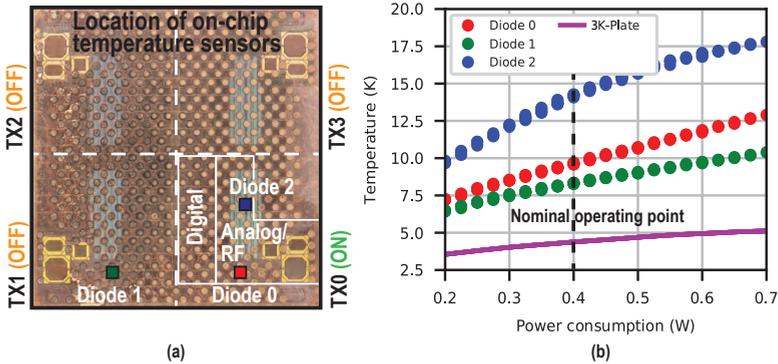


Figure 6.16: (a) Placement of on-chip diodes and (b) measured die temperature versus power consumption.

as a function of the chip power consumption, which is varied by changing the clock frequency and the supply voltage of the digital circuitry. Although the die self-heating increases significantly with power consumption, the plate temperature is only slightly affected. As the dilution unit is connected to a separate plate with an independent pulse tube cooler, the qubit temperature is not affected.

6.5.2. Electrical characterization

While the functionality of all four transmitters has been verified, the performance of one transmitter is reported in the following.

Figure 6.17 shows the power consumption of the various circuit blocks at 1 GHz clock frequency. The digital back-end dominates the power consumption due to the lack of clock gating in a substantial part of the memory and would increase further with clock speed. Hence, to limit the temperature increase of the fridge plate, the chip is operated at a maximum clock frequency of 1 GHz, limiting the available data bandwidth to 1 GHz. The analog power consumption is dominated by the output drivers due to high-linearity requirements at fairly high output power and the support of a 50 Ω load. The total power consumption of 12 mW/qubit would allow the control of > 320 qubits in a state-of-the-art dilution refrigerator,

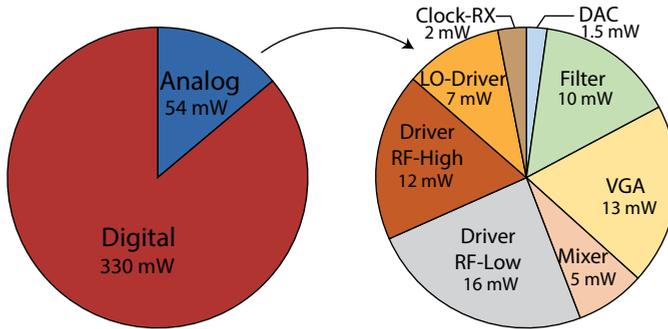


Figure 6.17: Power consumption breakdown, resulting in a total power consumption per qubit of $(330 \text{ mW} + 54 \text{ mW}) / 32 \text{ qubits} = 12 \text{ mW/qubit}$.

over only 10 RF lines, with a single SPI interface wired to room temperature. This is well beyond the number of qubits available in the largest solid-state quantum processor today [13]. Moreover, this work presents a first implementation of the controller, and further power reduction is possible as 1) significant margins were taken during the design to ensure functionality, 2) large output power and frequency ranges were included to support multiple qubit technologies, and 3) the currently dominating digital power consumption could be reduced by e.g. clock gating. With such optimizations, scaling to thousands of qubits is expected to be possible in the near term, while a larger cooling power is expected to extend the scaling in the longer term [31]. Thanks to the integrated digital controller, an external data rate of only $\sim 1 \text{ kb/s}$ over a single trigger line is required, allowing scaling to a large number of controllers sharing a single high-speed connection to room temperature. Moreover, thanks to the use of FDMA in this work, the number of connections to the quantum processor is reduced by $32\times$. However, supporting millions of qubits in the future with the proposed approach would still require a large number of connections to the quantum processor, but this could be eased by co-integrating the controller and the qubits on the same package or die at the same temperature.

Figure 6.19a shows the measured output power versus frequency at 3K for both the output paths. The flatness of the transfer function is deteriorated due to additional ground inductance introduced in the layout between the output matching network and the on-chip solder bumps.

To quantify the attenuation of the sampling replicas and flatness of the baseband transfer function, the measured output at the baseband monitoring node is shown in Fig. 6.19b. An in-band flatness of 1.5 dB is obtained up to 500 MHz as shown in the inset of Fig. 6.19b.

The SFDR obtained for single and two-tone signals at various output frequencies is shown in Fig. 6.18. From the single-tone spectra shown in Fig. 6.18a and Fig. 6.18b, it can be observed that the SFDR is limited by the image-rejection ratio (IRR of 45 dB obtained after calibration. The SFDR measured for various NCO frequencies over the entire data bandwidth is better than 42 dB as shown in Fig. 6.21b. The achieved LO rejection (LO Rejection (LOR)) does not affect the SFDR, since it

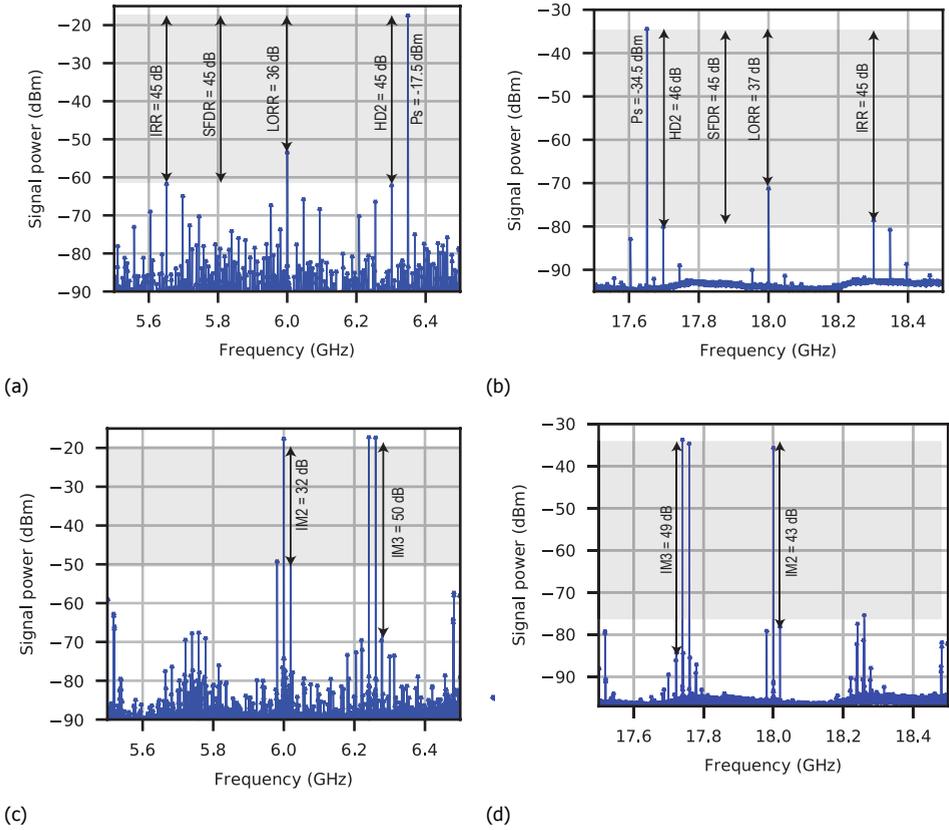


Figure 6.18: Single-tone SFDR at (a) 6.35 GHz and (b) 17.65 GHz. Two-tone output at (c) 6.25 GHz and 6.26 GHz and (d) 17.75 GHz and 17.74 GHz, generated using the 2 DDS banks shown in Fig. 6.3.

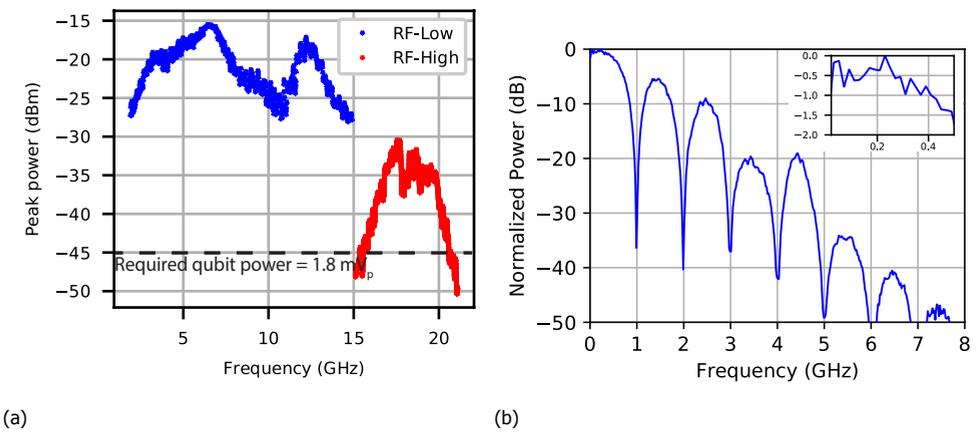


Figure 6.19: (a) Measured RF bandwidth, (b) Transfer function of the baseband output (inset: zoomed in from 0 to 0.5 GHz).

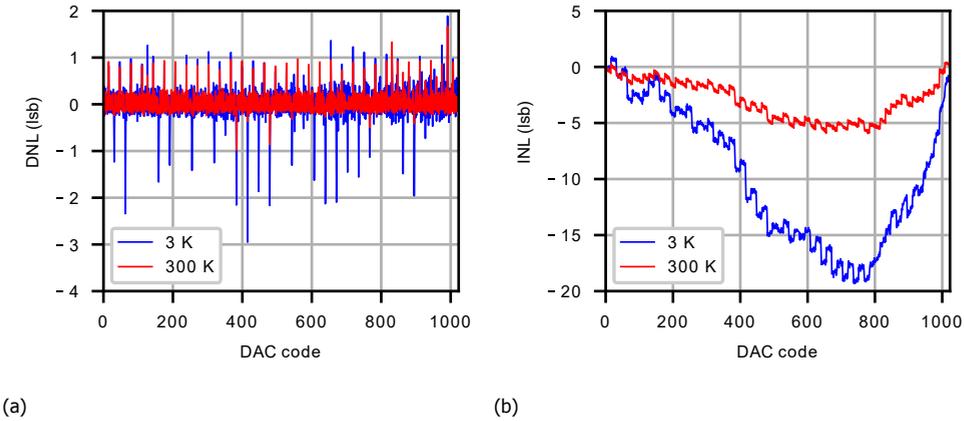


Figure 6.20: DAC linearity at room temperature & 3 K: (a) DNL, (b) INL.

can be avoided by proper choice of the LO frequency.

The SFDR of the two-tone spectrum with a tone spacing of 19 MHz shown in Fig. 6.18c is limited by the Second-order Intermodulation Distortion (IM2) component. Such IM2 can be attributed to the INL of the DAC that shows a quadratic behavior as shown in Fig. 6.20b. This is due to a linear gradient, i.e. systematic mismatch, in the DAC layout that does not use a fully common-centroid layout due to practical layout constraints, but an arrangement only similar to a common-centroid one. This systematic mismatch increases at 3 K. Moreover, random mismatch is degraded at 3 K, as can be seen in the DNL plot in Fig. 6.20a [41]. The large jumps in the DNL plot correspond to the unary element transitions in the segmented DAC. The measured Third-order Intermodulation Distortion (IM3) component with a two-tone spacing of 10 MHz is better than 47 dBc at the highest output power over the entire RF-Low bandwidth as shown in the Fig. 6.21c.

The measured SNR at the maximum output power over a 25 MHz bandwidth is greater than 48 dB as shown in Fig. 6.21b complying with the system requirements presented earlier. Figure 6.21a shows the typically observed noise spectral density around 6 and 18 GHz.

Engineering the pulse shape is critical for addressing multiple qubits over a frequency multiplexed line [120] as the shape of the pulse provides a trade-off between the speed of operation on the addressed qubit versus unwanted energy leaking into the unaddressed qubits. To demonstrate the pulse shaping capabilities of the chip, various pulse envelopes were applied at different offset frequencies as depicted in Fig. 6.22, which shows the time domain (at baseband frequency) and frequency domain response of the chip output.

6.5.3. Comparison with state-of-the-art

Table 6.2 summarizes the performance of the chip. Compared to the state-of-the-art [137], this work incorporates a wideband RF output to support multiple qubit

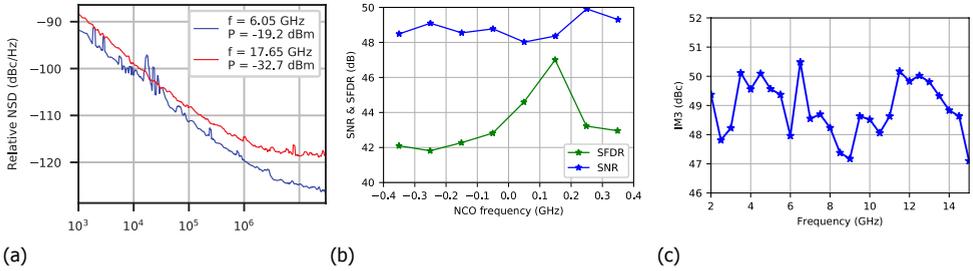


Figure 6.21: (a) Measured noise spectral density around 6 and 18GHz, (b) SNR & Single-tone SFDR versus NCO frequencies at 5GHz, (c) IM3 for a fixed NCO frequency over the entire RF-Low band.

6

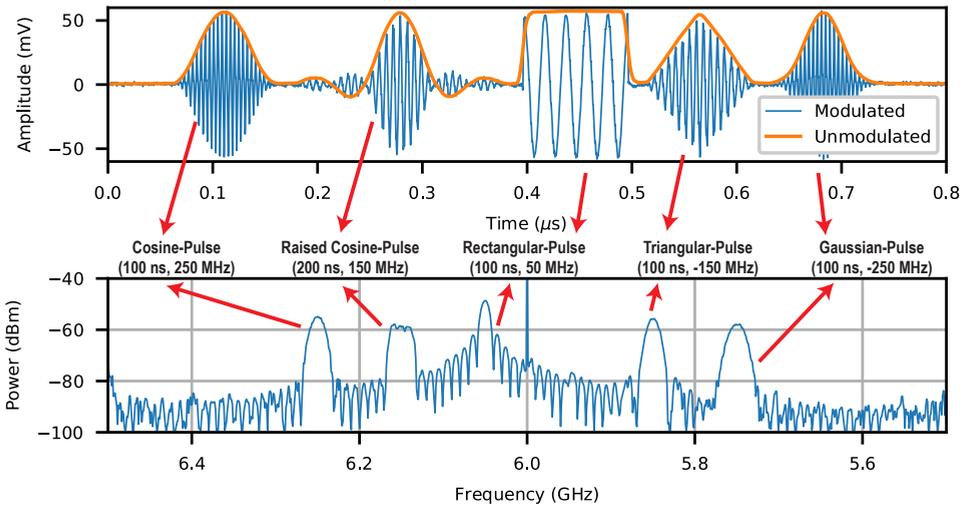


Figure 6.22: Pulse shaping; Top: measured time-domain signal at baseband output, Bottom: upconverted output spectrum.

technologies, frequency multiplexing for scalability with low power consumption and a digitally intensive back-end with an arbitrary-waveform generation memory of more than 40k points and the support of an instruction set for low-latency quantum-algorithm execution.

Table 6.2: Controller comparison table.

Parameter	This work [210]	JSSC'19[137]	RSI'17[26]	RT setup
Operating temperature	3 K	3 K	300 K	300 K
Compatible qubit platforms	Spin qubits & Transmons	Transmons [§]	Transmons [§]	Spin qubits [§]
Qubit frequency	2-20 GHz	4-8 GHz	-	<20 GHz
Channels	32 (128 with 4 TX)	1	4	1
FDMA	Yes, SSB	No	Yes, SSB	No
Data Bandwidth	1 GHz	400 MHz	960 MHz	520 MHz
SFDR [*]	>42 dB	N/A	-	70 dB
IM3 ($\Delta f = 10$ MHz)	>47 dB	N/A	-	-
SNR (Bandwidth (BW)=25 MHz)	48 dB	N/A	-	-
Image & LO leakage calibration	On-chip	Off-chip	Yes	-
Phase correction	Yes	No	No	No
Fidelity (expected) ^{&}	99.99 %	-	-	-
Waveform/ Instructions	Upto 40960 pts AWG	Fixed 22 pts symmetric	-	16M points AWG
Instruction set	Yes	No	Yes	Yes
Power/TX	Analog: 1.7 mW/qubit [@] Digital: 330 mW [#]	Analog <2 mW/qubit	-	850 W
Chip area/TX Technology	4 mm ² Intel 22 nm FinFET CMOS	1.6 mm ² 28 nm bulk CMOS	Discrete	Rack-mount
			-	-

* single tone excluding residual LO leakage

@ including LO & clock driver, only RF-Low active ($f_{LO} = 6$ GHz)

can be reduced with clock gating

§ no information is available regarding compatibility with other qubit technologies

& extrapolated from electrical performance and simulations
reported SFDR, IM3 and SNR values are valid for the entire 2-20 GHz range

6.6. Conclusion

By leveraging their very large scale of integration, cryogenic CMOS circuits can help solve the interconnect bottleneck between the quantum processor and its control electronics, thus enabling to scale up the number of qubits in quantum computers. In this chapter, this has been demonstrated with the design of a scalable cryo-CMOS integrated controller for the multiplexed control of upto 4×32 qubits. The cryogenic microwave signal generator demonstrated here comprises an integrated digital controller that can translate qubit gate operations into the microwave signals necessary for the execution of quantum algorithms. The achieved power efficiency (12 mW/qubit) enabled by a digitally-intensive architecture and the frequency multiplexing allows for operating the chip at 3 K within the cooling capabilities of standard cryogenic refrigerators. The ideas presented in this chapter pave the way towards large scale-quantum computers exploiting control electronics and qubits operating in close proximity at a similar cryogenic temperature. As shown in the measurements, the spectral purity of the generated signals is sufficient for a 99.99 % theoretical fidelity. In the next chapter, the controller will be operated together with the qubits, and qubit benchmarking experiment featuring this controller are described.

7

Benchmarking Horse Ridge with Qubits

In the previous chapter (Chapter 6), a scalable cryogenic electronic interface chip implemented in CMOS operating at 3 K optimized for controlling 128 qubits was presented, with a measured electrical performance indicating the possibility of high-fidelity (99.99 %) qubit control. An important next step is to test the overall performance of the quantum control chip in driving real qubits.

In Section 7.1 the measurement setup used to measure the overall performance is presented. Next, Section 7.2 discusses the experiments performed on a single qubit to verify the cryo-controller's effectiveness in providing coherent qubit control at different qubit frequencies. Sections 7.3 and 7.4 discuss the experiments performed on two qubits without and with the two-qubit interaction, respectively. Finally, conclusions are presented in Section 7.5.

7.1. Setup

As a benchmark of performance, the cryo-controller is used to coherently control a two-qubit quantum processor based on single-electron spin-qubits. In order to benchmark the limits of the controller, the qubits should be kept at ~ 20 mK, where the qubits are most coherent and the demands on the controller are highest. Hence, the quantum processor die is mounted on a PCB (Fig. 7.1a) operated at the base temperature of the dilution refrigerator. Note, that in the future, high-temperature qubits could also be used with this setup.

The quantum processor is made of a double quantum dot electrostatically confined in a 10 nm-thick $^{28}\text{Si}/\text{SiGe}$ quantum well [226]. A SEM image of the quantum processor is shown in Fig. 7.1b. By tuning the voltage on plunger gates LP (blue) and RP (red), two single electrons are locally accumulated underneath each gate,

Parts of this chapter have been published in IEEE Journal of Solid-State Circuits [210] and Nature [225].

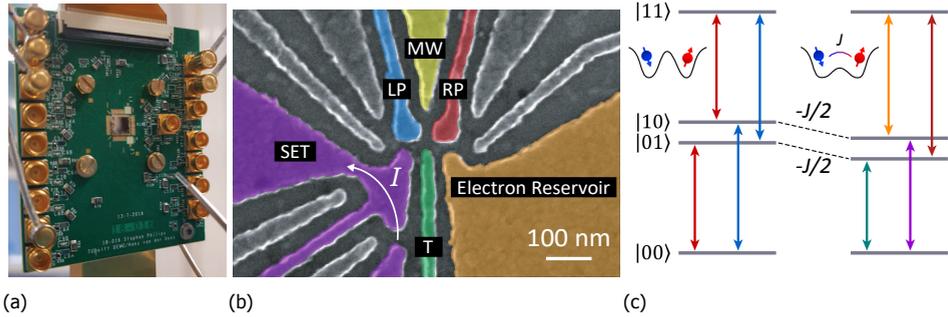


Figure 7.1: (a) PCB hosting the quantum processor chip. (b) false-colored Scanning Electron Microscope (SEM) image of the quantum processor indicating the various gates. (c) Energy level diagram without (left) and with (right) exchange coupling (J).

acting as the 2 single-electron spin qubits. By applying an external magnetic field of 380 mT, combined with the longitudinal magnetic field induced by a micro-magnet on top of the double quantum dot, the qubit states can be encoded into the Zeeman split states of the two electrons, where spin-up is used as $|1\rangle$ and spin-down is used as $|0\rangle$. The resonance frequencies of Qubit 1 (Q_1 , underneath gate LP) and Qubit 2 (Q_2 , underneath gate RP) are 13.62 GHz and 13.51 GHz, respectively, under these conditions ¹. Rotations around the X and Y axes are implemented by sending microwave bursts with the microwave phase controlling the rotation axis. The microwave bursts are applied to gate MW (yellow) resulting in a fast-oscillating electric field which drives electric-dipole spin resonance enabled by the transverse magnetic field gradient from the micro-magnet [227]. Rotations around the Z axis (phase control) are achieved by changing the reference phase in the cryo-controller [156].

The two-qubit interaction is mediated by the exchange coupling (J) between the two spins [77], controlled by gate T (green). Its effect here is to shift the anti-parallel spin states down in energy [182], as shown in Fig. 7.1c. As a result, the resonance frequency of each qubit depends on the state of the other qubit when the coupling is on (low tunnel barrier between the dots), allowing controllable conditional operations on each qubit via narrow-band microwave bursts [228, 229].

The SET next to the quantum dots is voltage biased and the SET signal (I) is converted to a voltage signal through a transimpedance amplifier and digitized by a digitizer card after a 10 kHz analog low-pass filter. The SET signal is sensitive to the charge occupation of the quantum dots, allowing binary single-shot readout of the qubit states via spin-to-charge conversion [230]. The state of qubit Q_2 is directly read out by spin-selective tunneling to the electron reservoir. The state of qubit Q_1 is mapped onto Q_2 by a *CROT* gate (which has been (re)initialized), and is then read out by measuring Q_2 .

A simplified diagram of the rest of the measurement setup is shown in Fig. 6.14b.

¹Different plunger gate voltage and external magnetic field settings are used in the experiments presented in Section 7.2

A programmable mechanical switch at 3 K is used to connect gate MW either to a VSG at room temperature or to the cryo-controller at 3 K. Moreover, the setup contains an FPGA that configures the cryo-controller, programs the various memories inside the cryo-controller (see Section 6.3), and controls the start of the execution of the instruction list (each with up to 2048 instructions). Switching between different instruction lists and synchronization with the rest of the equipment is controlled by two trigger lines from the AWG to the FPGA. One trigger starts the execution of the instruction list that is programmed in the cryo-controller, for performing repeated measurements, while the other trigger loads the next instruction list from the FPGA memory into the cryo-controller's instruction list. After the last list specified in the FPGA memory has been programmed, the first instruction list is loaded again, and a parameter inside the cryo-controller (for instance the NCO frequency, pulse duration, or Z-correction) is reconfigured by the FPGA, as required for e.g. sweeping the tone frequency or pulse duration during pulse calibration.

7.2. Single-Qubit Experiments

In this section, the experiments used to verifying the cryo-controller's effectiveness in providing qubit control in both the low-frequency and high-frequency range are described. In these experiments, only qubit Q_2 (underneath gate RP in Fig. 7.1b) is used.

7.2.1. Rabi oscillation experiment

To demonstrate qubit control, the oscillatory behavior of a two-level quantum system can be produced in a Rabi experiment. The amplitude of the pulse applied to the qubit determines the speed of rotation, i.e the Rabi frequency. By applying pulses with increasing duration, the qubit angle of rotation is increased, producing a typical oscillating pattern when read out. In this experiment, the qubit is first initialized to state $|0\rangle$, then excited by a rectangular microwave pulse with a given duration and finally the quantum state is read out. By varying the pulse duration and averaging the results over multiple runs, a Rabi frequency of 1 MHz and 400 kHz at 13.4 GHz (RF-low output) and 17.5 GHz (RF-high output), respectively, has been measured (Fig. 7.2).

The visibility of the adopted Elzerman readout [68], i.e. the difference between the highest and lowest probability obtained after readout, is affected by noise on the qubit device gates. To improve the readout visibility, a BPF with 2 GHz pass band has been added to the chip output to remove out-of-band spectral content. As shown in Fig. 7.2, this resulted in an improved readout visibility compared to what we previously reported in [209], comparable to what is obtained with the room-temperature control. A measurement of the electron temperature comparing the room-temperature control with the cryo-controller setup indicates that the output noise of the cryo-controller is now low enough to not affect the electron temperature (Fig. 7.3). The discrete fixed-frequency BPF that is currently used, could be replaced by a Surface Acoustic Wave (SAW) filter on the PCB or by an on-chip higher-order reconstruction filter and/or a passive filter at the mixer output, when the frequency

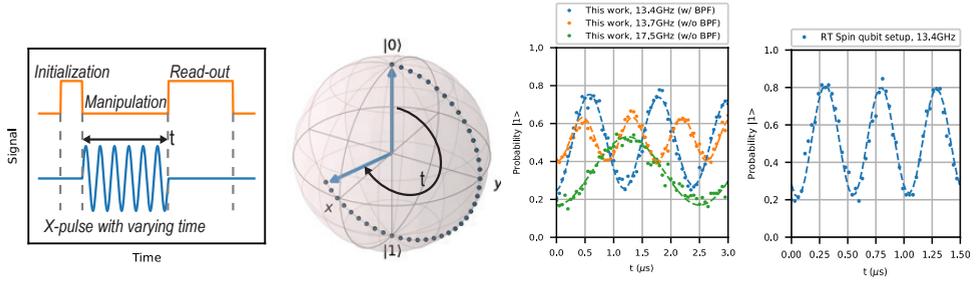


Figure 7.2: Left: explanation of the Rabi oscillation experiment and the signals involved. Right: Rabi oscillations measured at both 13.7 and 17.5GHz obtained using the presented controller, with improvement in readout visibility at 13.4GHz thanks to the BPF, along with the Rabi oscillation obtained using the room-temperature spin qubit setup.

of the qubits is known at design time.

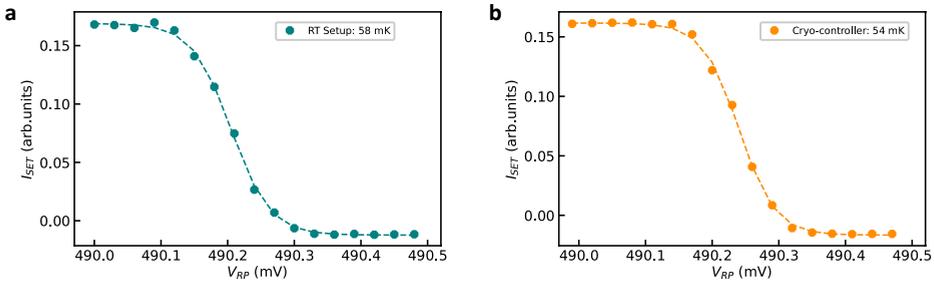


Figure 7.3: The measured electron temperature when the quantum device is connected to (a) the room-temperature VSG and (b) to the cryo-controller (at zero magnetic field). The amplified SET current (I_{SET}) as a function of the plunger gate voltage (V_{RP}) measured at the charge transition between $(1,0)$ and $(1,1)$. The electron temperatures are extracted by fitting the curves with the Fermi-Dirac distribution, with a lever arm of 0.172 meV/V .

7.2.2. Ramsey-style experiment

To demonstrate coherent qubit control over two axes, a Ramsey-style experiment is carried out [15]. Here, the qubit is initialized to state $|0\rangle$ and two rotations around the X axis are then applied ($R_X(\frac{\pi}{2})$) sandwiched by a Z-gate of varying angle from 0° to 360° ($R_Z(\theta)$). This resulted in a cosinusoidal variation in the measured $|1\rangle$ probability (Fig. 7.4), as expected. The X rotation is implemented by a microwave rectangular pulse at 13.7GHz with a duration directly proportional to the rotation angle, while the Z-rotation is implemented by updating the reference phase of the NCO (applying a digital phase offset), which continuously keeps track of this phase evolution. The experimental data closely tracking the theoretical expectation proves coherent qubit control and the capability of correctly executing any type of single-qubit gate.

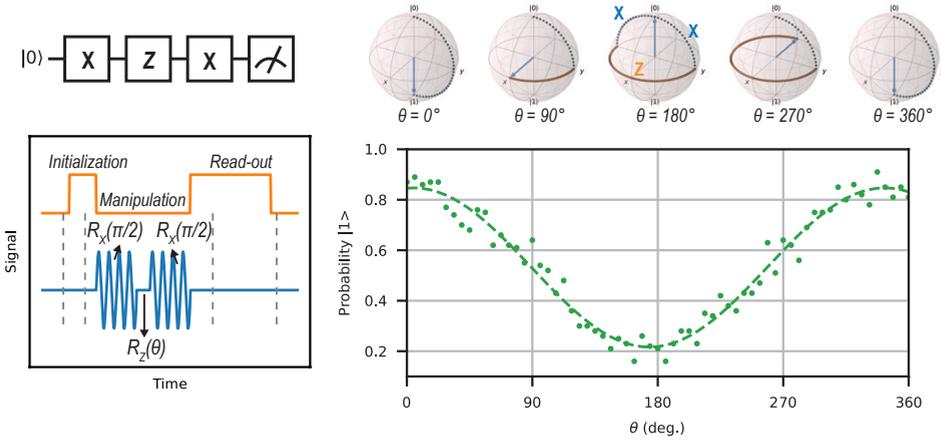


Figure 7.4: Schematic explanation and measurement results of the Ramsey-style experiment.

7.2.3. Pulse shaping experiment

The amplitude and phase modulation capabilities of the controller allow the chip to generate arbitrary waveforms to precisely shape the spectral content of the pulse used to manipulate the qubits, as shown in Fig. 7.5a. Figure 7.5b shows the measured response of Q_2 (which can be considered as the untargeted adjacent qubit) to an on- and off-resonance microwave burst with rectangular versus Gaussian envelope, both calibrated to invert the qubit state when the drive is on-resonance with the qubit. The presence of sidebands in the spectrum of the rectangular pulse are reflected in the qubit response, and confirm the strict requirement of using shaped pulses for preventing cross-talk when controlling multiple qubits spaced closely in frequency sharing a single drive-line (see Section 3.3.2).

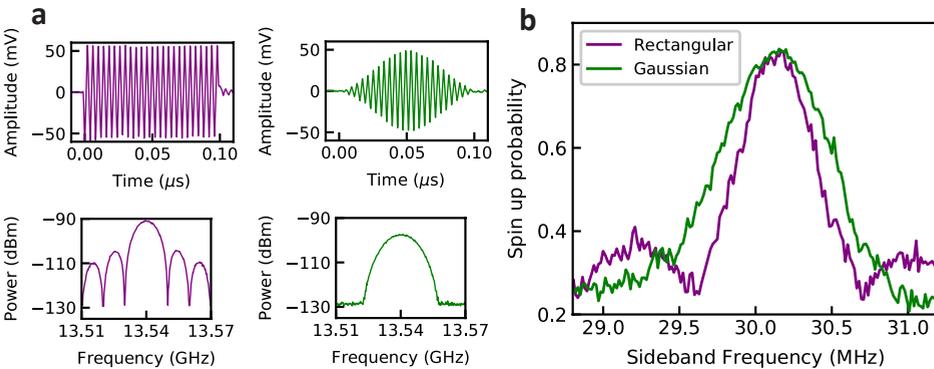


Figure 7.5: (a) Rectangular (purple) and Gaussian (green) shaped bursts before up-conversion and the corresponding spectra after up-conversion. (b) Qubit response for different burst envelopes, obtained when sweeping the NCO frequency around the qubit resonance across a span of ~ 3 MHz with a resolution of 15 kHz.

7.3. Single-Qubit Performance Benchmark

In the following experiments, both qubits are present but are left uncoupled, while performing the various characterizations with the cryo-controller.

7.3.1. Two qubit Rabi oscillation experiment

First, the functionality of the cryo-controller for controlling both qubits while uncoupled is tested. The LO frequency is set to 13.54 GHz. The qubit resonances are found by sweeping one single-sideband tone generated by one NCO (Fig. 7.6a). The frequency of Q_1 is offset from the LO by ~ 24 MHz and Q_2 by ~ -90 MHz. In two separate measurements, the individual Rabi oscillations of the qubits driven by the cryo-controller are obtained, as shown in Figs. 7.6b and 7.6c (T_1 is in the order of 100 ms). The output frequency of two NCOs are set to the frequencies of Q_1 and Q_2 , respectively, but only one NCO is active each time. Next, one NCO from each bank is used to generate a tone on resonance with each of the two qubits and drive simultaneous Rabi oscillations on both qubits, as shown in Figs. 7.6b and 7.6c. Note that the qubits have a different frequency and can hence be addressed independently. Compared to the individual driving experiments, the decay is much faster in the simultaneous Rabi oscillations. This decay arises mainly from the residual coupling between the two qubits.

7.3.2. AllXY and Quantum State Tomography experiment

The pulses for single-qubit rotations are precisely calibrated using the AllXY sequence [231]. In the AllXY experiment, 21 different pairs of single-qubit gates from the set $\{I, X, Y, X^2, Y^2\}$ are applied to a qubit initialized to $|0\rangle$. Here I is the identity operation, X and Y are $\pi/2$ rotations around the X and Y axis respectively, and X^2 and Y^2 are π rotations. The final state Z-projection $\langle\sigma_z\rangle$ takes values from $\{-1, 0, +1\}$ under perfect operations (shown as the gray shaded areas in Fig. 7.7a). Any miscalibration in the amplitude, frequency or phase of the pulse results in deviations from the ideal outcome (hatched bars in Fig. 7.7a). In addition, we reconstruct the trajectory of an X^2 gate by performing quantum-state tomography [232] at incremental burst times of a rectangular microwave signal (Fig. 7.7b). The AllXY and quantum-state tomography results indicate that the single-qubit gate set is well calibrated, offering a good starting point for benchmarking the gate fidelity.

7.3.3. Randomized benchmarking experiment

The gate fidelity is a crucial metric to express the performance of a quantum processor and its classical controller. Here, single-qubit randomized benchmarking [233, 234] is used to compare the performance of the cryo-controller with the conventional room temperature setup, which consists of an arbitrary waveform generator (Tektronix 5014C) and a vector signal generator (Keysight E8267D). In the randomized benchmarking experiment, sequences of increasing numbers of randomly selected Clifford operations are applied to the qubit (Q_2), followed by a final Clifford operation that returns the qubit to its initial state in the ideal case. For each data point in Fig. 7.8, 32 different sequences are randomly sampled and each is re-

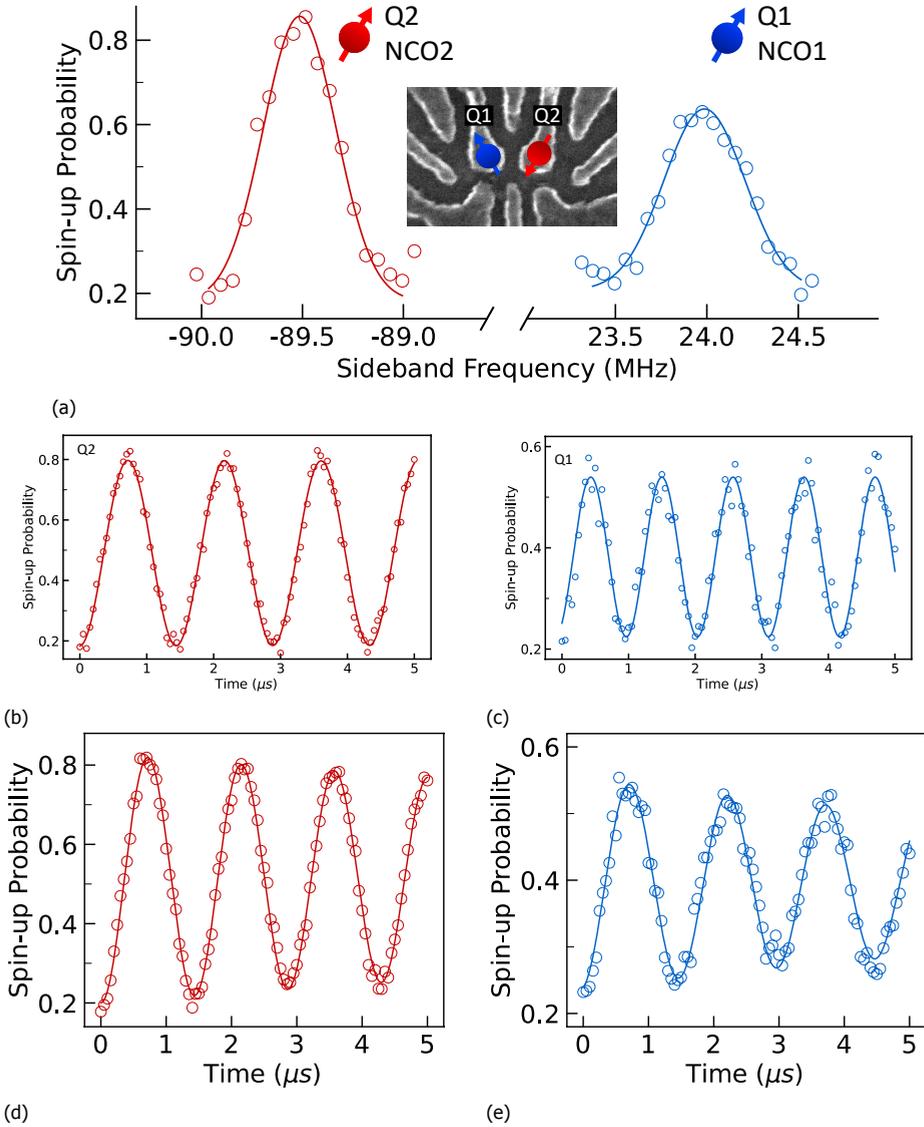


Figure 7.6: (a) Spectra showing the qubit resonances. Inset: SEM image indicating the qubits' positions. Frequency-multiplexed control producing an individual Rabi oscillation on Q_2 (b) and Q_1 (c), and a simultaneous Rabi oscillation of Q_2 (d) and Q_1 (e).

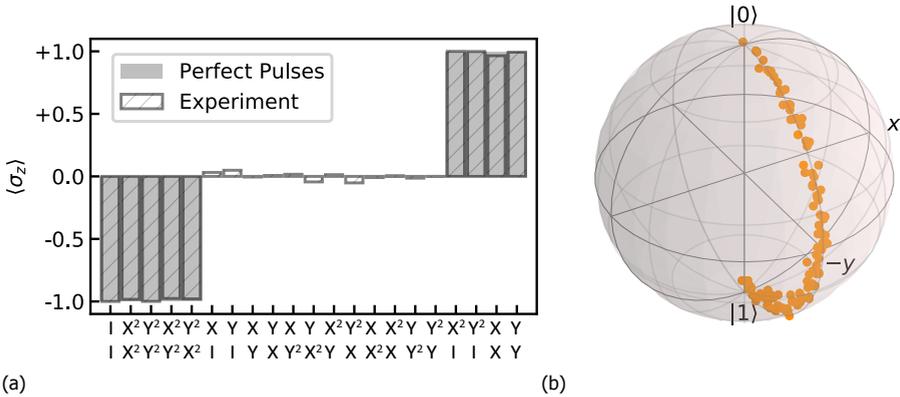


Figure 7.7: (a) $\langle \sigma_z \rangle$ of Q_2 measured after an AllXY sequence. The output power is calibrated to achieve a ~ 1 MHz Rabi frequency (the same applies to the quantum-state tomography and randomized benchmarking experiments). The visibility is normalized by removing the readout error. (b) Trajectory of the state of Q_2 under an X^2 gate reconstructed by quantum-state tomography. Orange data points indicate the qubit state after incrementing microwave burst times.

peated 200 times. Envelopes of all gates to be used are uploaded to the envelope memory, and saved as instructions. The random sequences are constructed by updating the instruction list through the FPGA (Section 7.1). Exactly the same random sequences are used in a randomized benchmarking experiment using the room-temperature setup. We find an average single-qubit gate fidelity of $99.71 \pm 0.03\%$ with the room-temperature setup and $99.69 \pm 0.02\%$ with the cryo-controller. The fidelity in the two cases are equal within the error bars and well above the threshold for fault-tolerance [8], with the infidelity limited by the qubit. These experiments demonstrate the high signal quality from the cryo-controller as well as its capability of generating complex sequences.

7.4. Two-Qubit Experiments

7.4.1. Coupled qubits resonance driving experiment

To further test the programmability of the cryo-controller, we use it to implement two-qubit logic in the quantum processor. Taking advantage of the frequency shift of each qubit conditional on the state of the other qubit (Fig. 7.1c), controlled-rotation (*CROT*) gates are used as the native two-qubit gates. These are achieved by frequency selective addressing [228, 229], thus demanding 2 NCOs per qubit (Fig. 7.9).

7.4.2. Two-qubit quantum algorithm experiment

A π -rotation at the higher or lower frequency implements the canonical controlled-NOT (*CNOT*) gate or the zero-controlled-NOT (*Z-CNOT*) gate, respectively, up to a single-qubit $\pi/2$ Z-rotation on the control qubit. Due to cross-talk, an additional phase correction in the form of a Z-rotation is needed. All Z rotations are imple-

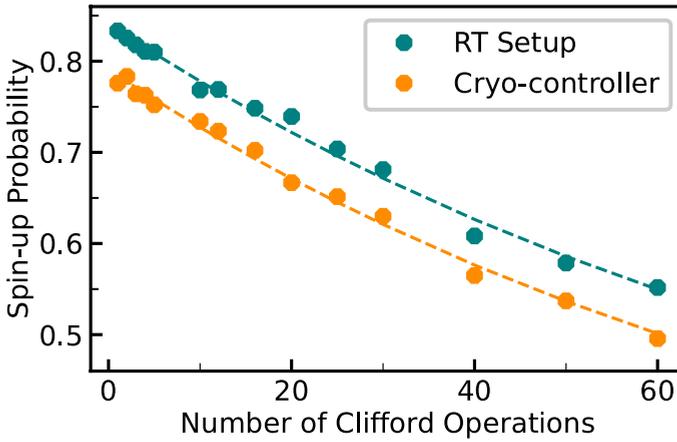


Figure 7.8: Randomized benchmarking of Q_2 performed by the cryo-controller and the room temperature setup. The orange data points are offset by -0.05 to facilitate comparison of the two traces.

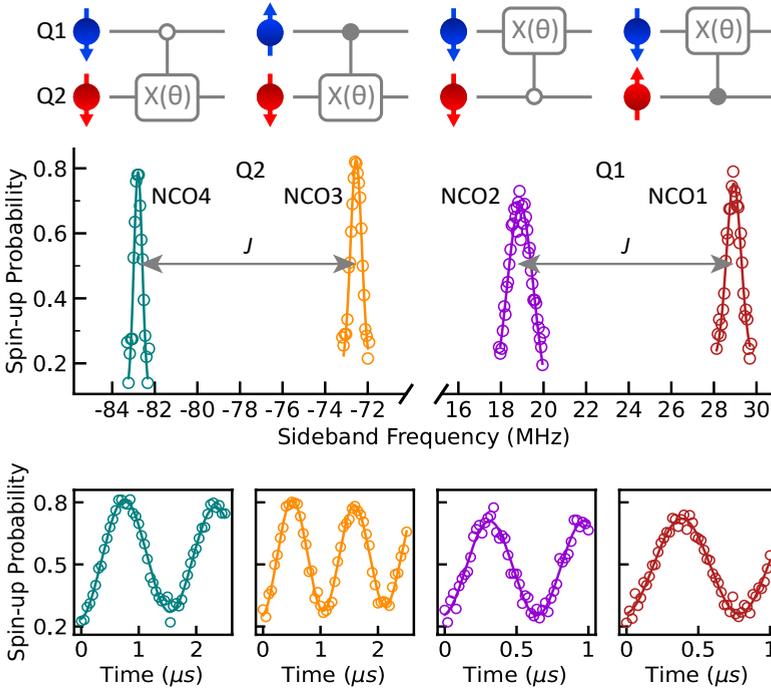


Figure 7.9: Two-qubit logic with the cryo-controller. The middle panel shows the spectra of two qubits obtained using the cryo-controller with the exchange coupling (J) between the qubits turned on. Selective excitation of each of the four resonances can be used for implementing various two-qubit controlled-rotation gates, shown in the upper panel. The lower panels (shared Y-axis labels) show the Rabi oscillations at each frequency.

mented by updating the reference phase of the NCO. Single-qubit gates are implemented by addressing both frequencies of the same qubit sequentially. Making use of four NCOs, the cryo-controller is programmed to run the two-qubit Deutsch–Josza algorithm, which determines whether a function gives constant or balanced outcomes [235]. The two constant (balanced) functions that map one input bit on one output bit are implemented by the $CNOT$ and $Z-CNOT$ (I and X^2) operations (see Fig. 7.10 for the implementation of these gates). Here, Q_1 is chosen to be the output qubit and Q_2 to be the input qubit. Fig. 7.10 shows the pulse sequence and the measurement results, where the constant (balanced) functions lead to a high probability for measuring the data qubit as $|1\rangle$ ($|0\rangle$), as expected. Empirically, we attribute the remaining errors mostly to charge noise in the presence of a finite J . This experiment highlights the ability to program the cryo-controller with arbitrary sequences of operations.

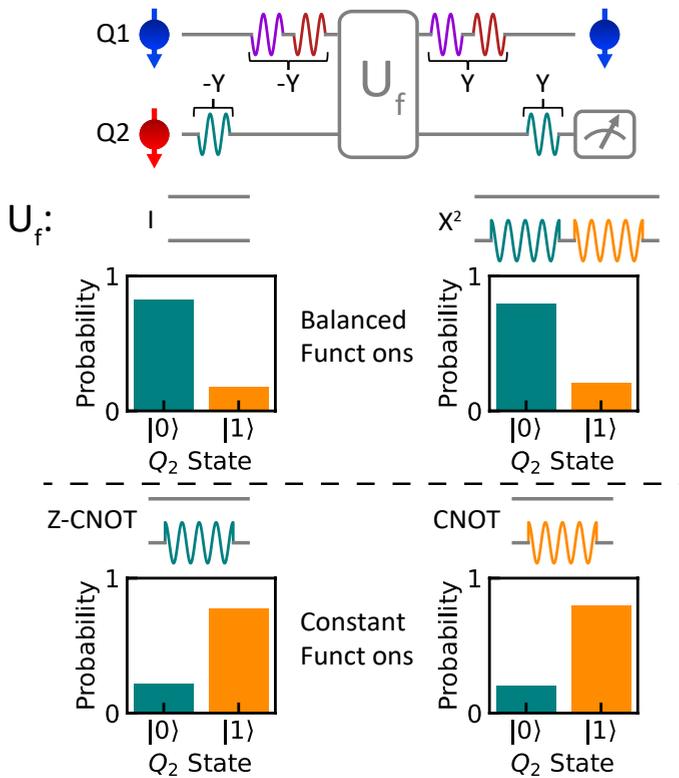


Figure 7.10: Programming a quantum processor with the cryo-controller. (Top) Pulse sequences of the Deutsch–Josza algorithm programmed into the cryo-controller and (Bottom) measured probabilities of the output qubit state (Q_2) after running the algorithm. A constant function is composed of either a $CNOT$ or a $Z-CNOT$ gate, which consists of a $CROT$ gate on Q_2 and a phase correction on Q_1 (not plotted). Only the lower frequency (green branch, $Z-CROT$) is used for the $-Y$ and Y gate on Q_2 because Q_1 (ideally) starts from and ends up in $|0\rangle$. The visibility of Q_2 is normalized by removing the readout error [225].

7.5. Conclusion

In this chapter, the effectiveness of Horse Ridge in the control of a spin qubit quantum processor has been demonstrated through various experiments. Most notably, the randomized benchmarking experiment shows that the same gate fidelity is obtained using the cryo-controller and the room-temperature setup using general-purpose equipment, with the infidelity limited by the qubit, and running a full quantum algorithm using cryo-CMOS control shows the versatility of the presented controller. The cryo-controller allows for even much more complex sequences, containing up to 2048 instructions for each of the four transmitters. Each instruction defines a microwave burst at one of 32 independent frequencies with an amplitude and phase profile that can be arbitrarily shaped. This quantum-classical architecture can thus be directly applied to multi-qubit algorithms and noisy intermediate-scale quantum devices [236]. Moreover, the footprint of just 4 mm², a power consumption of 384 mW, the ability to integrate multiple transmitters on one die, and operation at 3 K, highlight the promise of cryo-controllers to address key challenges in building a large-scale quantum computer.

8

Conclusion

8.1. Research overview

Quantum computers have gained huge interest in the past few years because they can potentially solve problems intractable by classical computers. However, state-of-the-art quantum processors, in any state-of-the-art technology, are limited to less than 100 non-ideal, noisy qubits. Although researchers are busy looking for alternative practical uses of those non-ideal small processors – the so-called Noisy Intermediate-Scale Quantum (NISQ) technology [236] –, the number of qubits is too low to achieve large-scale fault-tolerant quantum computing capable of solving non-trivial problems, such as molecular simulations. Throughout this dissertation, we set out to (partially) answer the difficult multi-disciplinary question: how to scale up quantum computers? More specifically, we focused on how to scale the classical electronic interface required to control and read out the quantum processor operating at cryogenic temperatures. To that end, we presented the design and experimental validation of a scalable cryo-CMOS microwave signal generator for controlling 128 single-electron spin qubits or transmons. To reach this goal, several challenges, as identified in Chapter 1, had to be overcome.

The first challenge has been to identify the optimal controller temperature, technology and scaling strategy. To give clarity towards that goal, the requirements on the control electronics for various solid-state qubit technologies were identified (Chapter 2). Superconducting qubits and single-electron spin qubits, which are among the most promising candidates for large-scale quantum computing, require similar control signals for performing single-qubit operations (modulated microwave bursts) and two-qubit operations (short pulses), as reflected in the similarities in state-of-the-art setups used to control the quantum processors in those technologies. Moreover, while initial setups relied on large, power-hungry general-purpose instruments achieving top-notch (but not necessarily required) performance, a general trend towards tailor-made room-temperature electronics optimized for size, power, and cost, has started. One of the greatest challenges left to overcome is

the wiring complexity imposed by current quantum processors that require about one control signal per qubit, while operating at cryogenic temperatures. The growing number of cables present a heat load to the dilution fridge, incur latencies in the control signals, and reduce the reliability and scalability of the setup. In order to limit wiring complexity, two main approaches can be adopted: multiplexing control and readout signals, or moving the electronic interface closer to the qubits by operating it at cryogenic temperatures. Both of these approaches have been adopted in this work, combining a CMOS controller operating around 4 K, and controlling a quantum processor at the same or lower temperature using multiplexed control signals. For near future solutions, CMOS is most appropriate due to the proven functionality at cryogenic temperatures, its maturity, and its compatibility with VLSI designs. Similarly, the 4 K-stage seems most appropriate for near future solutions due to its proximity to the qubits while offering significant cooling power.

Secondly, there has been a lack of complete and analytically derived specifications for the classical interface, as the electronic interface has been typically built from general-purpose components known to be good enough. When designing a scalable controller operating at cryogenic temperatures, one cannot afford to significantly over-design as that would unnecessarily increase the controller's power consumption, thus limiting its scalability. In Chapter 3, we methodically derived the effect of circuit non-idealities in the classical controller on the qubit fidelity for all possible operations, i.e., single-qubit gates, two-qubit gates, and readout. This includes the effect of signal inaccuracies in the frequency, voltage, and time domain, and covers static and dynamic, systematic and random errors. Only with such a full set of specifications potential bottlenecks can be properly identified and optimized controllers be designed. Most interestingly, a case study targeting a 99.9 % average gate fidelity (sufficient for fault-tolerant quantum computing, while ambitious for most of today's quantum processors) confirms that setups using general-purpose instruments are over-designed for such a fidelity. Mapping the specifications onto already existing room-temperature CMOS circuits shows that the specifications can be met with an expected power consumption in the order of 1 mW/qubit.

As a final challenge, there is scarcity of verification methods and tools for the design flow of large-scale quantum computers spanning several technology domains. Specifically, no tool was available supporting both the simulation of classical electronics and quantum systems. To this end, SPINE (Spin Emulator) has been developed (Chapter 4), encompassing a Hamiltonian solver that can be directly used in the industry standard for CMOS circuit design, Cadence®, for direct simulation of the quantum processor along with the electrical circuits. A quantum/electronic interface co-design methodology is proposed that covers the entire flow from the definition of the target quantum processor till the system-level verification of the electronic interface with that quantum processor. SPINE can effectively be used along all steps in the design flow, as exemplified with various design examples outlined in Chapter 4. Its effectiveness is furthermore demonstrated in the design of the scalable cryo-CMOS microwave signal generator presented in this dissertation.

In Chapter 5, the previously presented results were used to choose the optimal architecture for the microwave signal generator supporting high-fidelity (99.99 %)

multi-qubit (32 in a 2 GHz bandwidth) control. Such a microwave signal generator is an integral part of the electronic interface, as it is used for performing single-qubit operations in a quantum processor based on single-electron spin-qubits or transmons. Frequency-multiplexing is assumed as starting point to deal with the wiring complexity challenge. However, no specific temperature is assumed for the design, as it does not affect the choice of architecture. It is found that the most suitable architecture uses an I/Q-mixer-based SSB analog/RF front-end as it is expected to be most power efficient, while allowing high flexibility in the output frequency range, making this solution suitable for multiple qubit technologies. For the back-end, an NCO-based DDS system is found to be most suitable due to the relatively low hardware cost, both in terms of required memory and the number of digital multipliers. Moreover, with such a design, coherent qubit operations are easily implemented thanks to the NCOs keeping track of the phase of the qubits. Throughout Chapter 5, the signal specifications (obtained following the methods outlined in Chapter 3), are translated into system-level requirements for the selected architecture. For the DDS back-end, a direct link is made between the number of bits in various parts of the system and the resulting gate fidelity, resulting in a design optimized for the target fidelity. For the resulting digital system, the power consumption is estimated to be around 5 mW per qubit in a 22-nm CMOS node, with 80% of the power consumed in the NCOs. From the noise and linearity requirements of the analog/RF front-end, an additional ~ 0.5 mW/qubit is expected, assuming a $50\ \Omega$ output load. This result is comparable to the original power consumption estimate of 1 mW/qubit for a cryo-CMOS controller, with the additional power attributed to the DDS system included here to offer greater flexibility in signal generation during the first development cycles. Power could be saved later when the target qubit technology reaches maturity.

In Chapter 6, the implementation of the proposed controller in 22-nm FinFET operating at 3 K is presented. In addition to a bare DDS-based modulator, the presented digital system includes a digital controller. This extra step is essential to minimize the interface to the room-temperature equipment, while the employed frequency-multiplexing minimizes the interface to the quantum processor. The digital controller can translate qubit gate operations into the microwave signals necessary for the execution of quantum algorithms, following a very basic instruction set for each qubit (with a maximum program length of 2048 instructions). Moreover, the digital controller takes care of the phase corrections required to correct for the AC-Start shift occurring in relatively narrow-band systems that use frequency-multiplexed qubit control. For the analog baseband circuitry, a current-mode design is adopted to achieve the required bandwidth and linearity; it comprises a 10-bit current-steering DAC capable of operating at 2.5 GS/s, an active current-mode gm-C filter implementing a 2nd-order Chebyshev-I filter with 1.8 GHz cut-off frequency, and a current-mode VGA, followed by a double-balanced Gilbert-cell mixer that up-converts the signal to the desired qubit frequency band. Two output bands are used to cover the very wide output frequency range from 2 to 20 GHz. In general, extra tuning knobs, safety margins and limited device stacking are employed to ensure proper operation at cryogenic temperatures. Finally, 4 controllers, each supporting

32 qubits, are implemented on a single die with an area of 16 mm^2 . Cryogenic measurements showed a power consumption of 1.7 mW/qubit for the analog circuitry, in line with previous expectations, but a large digital power consumption of 330 mW at 1 GHz . The large digital power consumption can be explained by the additionally integrated digital controller, which in the current generation lacks clock gating while many of the memory cells are implemented using relatively power-hungry flipflops, leaving plenty of room for future improvements. However, this power consumption limited the testing to clock frequencies up to 1 GHz . The entire 2 to 20 GHz output frequency band was verified to output sufficient power to drive qubits, while showing a typical $\text{SFDR} > 42 \text{ dB}$, $\text{IM3} > 47 \text{ dB}$, and SNR of 48 dB in a 25 MHz bandwidth. These results are theoretically sufficient for a 99.99% average gate fidelity. Overall, this design, supporting a total of 128 qubits, demonstrates that cryogenic CMOS circuits can help solve the interconnect bottleneck between the quantum processor and its control electronics, thus enabling to scale up the number of qubits in quantum computers. The achieved power efficiency (12 mW/qubit) enabled by a digitally-intensive architecture and the frequency multiplexing allows for operating the chip at 3 K within the cooling capabilities of standard cryogenic refrigerators. This paves the way to large scale-quantum computers exploiting control electronics and qubits operating in close proximity at a similar cryogenic temperature.

In Chapter 7, the overall performance of the quantum control chip in driving real single-electron spin-qubits is demonstrated. First, to validate the controller, Rabi and Ramsey experiment were presented showing the controller's ability to control, and more importantly *coherently* control qubits. Next, experiments using off-resonance microwave bursts demonstrated the effectiveness of the controller's pulse shaping capabilities to suppress qubit rotations due to off-resonance bursts, as required in a frequency-multiplexed setup, an essential ingredient for this design to be considered scalable. The demonstration of both the Deutsch-Jozsa quantum algorithm and experiments typically used during the bring-up of a quantum processor, such as the AIXY and Quantum State Tomography, shows that the controller is flexible enough to be used both during multi-qubit quantum algorithm execution, as well as during the bring-up phase thanks to the versatile integrated digital controller. The performed randomized-benchmark experiment shows that the same gate fidelity ($\sim 99.7\%$) is obtained using the cryo-controller and the room-temperature setup using general-purpose equipment, with the infidelity limited by the qubit. While it can not be confirmed with an actual quantum processor that the presented controller supports the targeted 99.99% average gate fidelity, all measurements support the possibility of achieving this fidelity with a quantum processor with higher fidelity. Finally, the small footprint, power consumption, the ability to integrate multiple transmitters on one die, and operation at 3 K , demonstrate the scalability of the presented solution and highlight the promise of cryo-controllers to address the key challenges in building a large-scale quantum computer.

Overall, it can be concluded that a scalable electronic interface relies on two essential ingredients: 1) the integration of a controller at cryogenic temperatures to relieve the interconnect to room-temperature; 2) multiplexing of the control signals to relieve the interconnect to the quantum processor while improving the controller's

power efficiency, thus allowing it to operate within the power budget available at cryogenic temperatures.

8.2. Main contributions

The main contributions presented in this thesis, and described in detail in the previous section, are summarized below:

- The impact of the non-idealities of the electronic interfaces on the qubit fidelity has been analyzed and quantified, thus allowing to draft clear specifications and to perform error budgeting in the electronics given a target fidelity [102];
- The tools and the methodology for co-design and co-simulation of quantum processors and their electrical interfaces have been developed [191, 192] [source: <https://github.com/QE-Lab/SPINE>];
- The architecture for a scalable frequency-multiplexed microwave cryo-CMOS driver has been designed and experimentally demonstrated, showing performance, power consumption, form factor and operating temperature enabling large-scale quantum-processor control [194, 209, 210]¹;
- The capabilities offered by cryogenic control has been benchmarked by programming benchmarking protocols and a quantum algorithm on a two-qubit spin qubit processor using the proposed cryo-CMOS microwave driver [225]².

8.3. Future work

While this work presents the first step towards a scalable electronic interface that can be used in future fault-tolerant quantum-computers, a few suggestions for future work are listed below:

- The design presented in this dissertation only focuses on the single-qubit gates, while this represents only a small portion of the electronic interface, which should also include two-qubit gates, readout and biasing. A fully integrated electronic interface should be developed to support the scaling to larger quantum computers.
- While in this dissertation the focus was on an electronic interface implemented in CMOS, operating around 3 to 4 K, and it seems a good solution for implementing a digital-intensive quantum algorithm execution controller, other parts of the electronic interface could achieve superior performance if implemented in another technology (e.g. LNAs in SiGe technology), or operating at a different temperature. Hybrid solutions, with different parts of the electronic interface operating at different temperatures and possibly in different technologies should be further investigated. Especially with future high-temperature

¹This work was equally contributed by J.P.G. van Dijk and B. Patra.

²This work was equally contributed by J.P.G. van Dijk, B. Patra and X. Xue.

qubits, the optimal division of the electronic interface over the temperature stages might change.

- While this dissertation focused on transmons and single-electron spin-qubits, there are many viable solid-state quantum technologies, some of which require simpler control signals (Chapter 2, e.g. singlet-triplet qubits), which could potentially simplify the electronic interface. To help scale quantum processors based on different quantum technologies, supporting electronic interfaces, tailor made to their requirements, should be developed. To enable their design, again specifications for the controller should first be derived, extending the work presented in Chapter 3 to other qubit technologies, such as singlet-triplet qubits.
- Extensive effort in various fields is required to speed-up the development of tailor-made electronic interfaces. While the SPINE toolset has proven useful, faster and better integrated co-simulators could significantly speed-up the development of electronic interfaces. Moreover, it could lead to quantum processors or qubit operations optimized for the limited capabilities and performance of a power-optimized electronic interface. Similarly, more accurate cryogenic device models are required to develop more efficient controllers and they will help speed-up the design and measurement time by removing the need for many tuning networks. Finally, tighter control on quantum processor technologies will help reduce the signal ranges (power, frequency, duration etc.) that need to be supported.
- Related to the previous point, if tight specifications on the signal ranges are known at design time, possibly different architectures or circuit designs will be preferable. As an example, the previously discarded MRZ DAC analog front-end could be more appropriate if the qubit frequencies are well known in advance. Similarly, for the digital back-end, the choice of an NCO per qubit might not be most efficient due to the high power consumption of a digital accumulator (depending on the technology and operating frequency). Alternatively, a single NCO could be used, and the phase of all other qubits could be derived using an appropriate multiplication factor, effectively changing the need of an always-running NCO into a digital multiplier that is only active when a signal needs to be generated. In general, architectural choices should be re-evaluated every time based on the quantum processor and technology at hand, if it's desirable to have the most power-efficient solution.
- Thermal effects of the cryogenic controller should be properly modeled, to ensure proper cooling of the entire cryogenic electronic interface, and to understand the exact behavior over temperature. In the presented cryogenic controller, significant effort was spent on the PCB design and the design of an appropriate enclosure, but the chip's die temperature was still significantly higher than the temperature on the 3K stage. Local heating could cause variability in the signal and increased noise levels.

- More effort is required towards the co-integration of quantum processors and local electronics (e.g. multiplexing) on either the same die or package to further increase the scalability of the quantum computer. Tighter integration could also remove the need for 50-ohm terminated lines (or other line impedances could be selected), potentially leading to significant power savings and lower self-heating.
- With the presented cryogenic controller there is still a significant size mismatch between the controller die area and the quantum processor size, which should be resolved to allow for the future co-integration. On the one hand, the controller could be made more compact, and on the other hand, research could be done targeting at e.g. increasing the effective pitch between qubits [110].
- Lastly, the current design can be further optimized for power, e.g. by using the appropriate memory cells, implementing clock gating etc. Moreover, area savings are possible, e.g. by integrating the currently off-chip hybrids and filters.

A

Derivations of Qubit Control Specifications

In this appendix, we present detailed derivations for the relations shown in Chapter 3, along with additional closed-form expressions. Appendix A.1 discusses general derivations, Appendix A.2 the derivations for the single-qubit operation, Appendix A.3 for the two-qubit operation, and Appendix A.4 for qubit read-out.

Unless otherwise stated, all angles (θ , ϕ etc.) are in radians.

A.1. General Derivations

The average fidelity of a unitary operation on an n -dimensional complex Hilbert space can be calculated as [170]:

$$F_{\text{av}} = \frac{n + \left| \text{Tr} \left[U_{\text{ideal}}^\dagger \cdot U_{\text{real}} \right] \right|^2}{n \cdot (n + 1)}. \quad (\text{A.1})$$

where U represents the unitary operation performed.

Using the relation between the average gate fidelity (F_{av}) and process fidelity (F) [169]:

$$F_{\text{av}} = \frac{1 + n \cdot F}{n + 1}, \quad (\text{A.2})$$

the process fidelity can be obtained from the unitary operations as:

$$F = \frac{1}{n^2} \cdot \left| \text{Tr} \left[U_{\text{ideal}}^\dagger \cdot U_{\text{real}} \right] \right|^2, \quad (\text{A.3})$$

Parts of this chapter have been published in Physical Review Applied [102].

which is used throughout this work, in line with [145, 147, 148]. In case of random variations of the unitary operation U_{real} , the expected value of the fidelity will be evaluated:

$$\langle F \rangle = \int_{-\infty}^{\infty} F(x)f(x)dx, \quad (\text{A.4})$$

where $F(x)$ is the fidelity for a random parameter x , and $f(x)$ is the probability density function of x .

In case of Gaussian distributed noise with zero mean and standard deviation σ :

$$f(x) = \frac{e^{-\frac{x^2}{2\sigma^2}}}{\sqrt{2 \cdot \pi \cdot \sigma^2}}. \quad (\text{A.5})$$

For this distribution, and certain expressions for $F(x)$, the integral in Eq. (A.4) has a simple solution.

In case $F(x) = 1 - c \cdot x^2$:

$$\langle F \rangle = \int_{-\infty}^{\infty} (1 - c \cdot x^2) \frac{e^{-\frac{x^2}{2\sigma^2}}}{\sqrt{2 \cdot \pi \cdot \sigma^2}} dx \quad (\text{A.6})$$

$$= 1 - c \cdot \sigma^2. \quad (\text{A.7})$$

In case $F(x) = 1 - c \cdot x^4$:

$$\langle F \rangle = \int_{-\infty}^{\infty} (1 - c \cdot x^4) \frac{e^{-\frac{x^2}{2\sigma^2}}}{\sqrt{2 \cdot \pi \cdot \sigma^2}} dx \quad (\text{A.8})$$

$$= 1 - 3 \cdot c \cdot \sigma^4. \quad (\text{A.9})$$

A.1.1. Numerical Simulations

The numerical simulations performed in case of microwave pulses with Gaussian envelopes use finite time steps to approximate the unitary operation:

$$U_{\text{real}} \approx \prod_{n=N}^0 e^{-iH(i\Delta t)\Delta t}, \quad (\text{A.10})$$

where $H(t)$ is the Hamiltonian at time t . The time step Δt is chosen to be constant and sufficiently small in order not to affect the simulation results. A 10 times over-sampling with respect to the signal's carrier frequency was found to give accurate results.

A.1.2. Equivalent Noise Bandwidth

The ENBW of a filter is defined as the bandwidth of the ideal brick-wall filter with the same peak gain of the original filter that results in the same integrated output

noise, under the assumption that only white noise is present at the filter input. For a filter transfer function $G(\omega)$, the ENBW can be computed as:

$$ENBW = \int_0^\infty \left| \frac{G(\omega)}{G_{max}} \right|^2 d\omega \quad (\text{A.11})$$

The ENBW is used to simplify calculations when the noise in the band of interest can be approximated as white noise.

A.2. Derivations for Single-Qubit Operation

Starting from the Hamiltonian that describes a single electron in the lab frame under microwave excitation ($\hbar = 1$)

$$H_{lab} = -\omega_0 \cdot \frac{\sigma_z}{2} + \omega_{ESR}(t) \cdot \frac{\sigma_x}{2}, \quad (\text{A.12})$$

where σ_x , σ_y and σ_z are the Pauli operators. For a microwave signal $\omega_{ESR}(t) = 2 \cdot \omega_R \cdot \cos(\omega_{mw}t + \phi)$, the Hamiltonian can be made time-independent by moving to a reference frame that rotates with a frequency ω_{mw} around the z-axis:

$$H_{ref} = -\omega_{mw} \cdot \frac{\sigma_z}{2} \quad (\text{A.13})$$

$$U_{ref}(t) = e^{-iH_{ref}t}. \quad (\text{A.14})$$

The Hamiltonian in the rotating frame follows as:

$$H_R(t) = U_{ref}(t)^\dagger \cdot (H_{lab} - H_{ref}) \cdot U_{ref}(t) \quad (\text{A.15})$$

$$= \frac{\omega_R}{2} \cdot \begin{bmatrix} \frac{\omega_{mw} - \omega_0}{\omega_R} & e^{i\phi} + e^{-i\phi} e^{-2i\omega_{mw}t} \\ e^{-i\phi} + e^{i\phi} e^{2i\omega_{mw}t} & \frac{\omega_{mw} - \omega_0}{\omega_R} \end{bmatrix}. \quad (\text{A.16})$$

By neglecting the high frequency oscillations ($2 \cdot \omega_{mw} \cdot t$), the time-independent Hamiltonian in the rotating wave approximation is obtained:

$$H = \frac{1}{2} \cdot \begin{bmatrix} \omega_{mw} - \omega_0 & \omega_R \cdot e^{i\phi} \\ \omega_R \cdot e^{-i\phi} & \omega_{mw} - \omega_0 \end{bmatrix} \quad (\text{A.17})$$

$$= (\omega_{mw} - \omega_0) \cdot \frac{\sigma_z}{2} + \omega_R \cdot \left[\cos(\phi) \cdot \frac{\sigma_x}{2} - \sin(\phi) \cdot \frac{\sigma_y}{2} \right]. \quad (\text{A.18})$$

However, the RWA is only valid when the Rabi frequency is sufficiently lower than the Larmor frequency ($\omega_R \ll \omega_0$). Figure A.1 shows the fidelity of a qubit rotation as obtained from a numerical simulation of the full Hamiltonian (Eq. (A.12)). The fidelity has been computed with respect to the ideal rotation that would result from the Hamiltonian in the RWA (Eq. (A.18)). This plot clearly shows that a sufficiently high ratio of Larmor frequency to Rabi frequency is required for this approximation to hold.

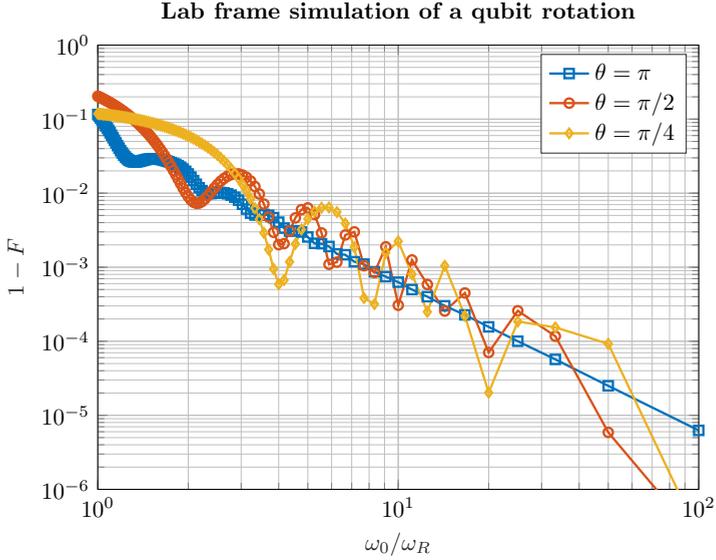


Figure A.1: The numerically simulated fidelity of a qubit rotation around the X-axis as a function of the ω_0/ω_R -ratio, without using the rotating wave approximation, for various rotation angles θ .

Finally, the ideal unitary operation is obtained by evaluating $U = e^{-iH \cdot T}$ for an ideal microwave signal ($\omega_{\text{mw}} = \omega_0$):

$$U_{\text{ideal}} = e^{-i\theta \cdot \left[\cos(\phi) \cdot \frac{\sigma_x}{2} - \sin(\phi) \cdot \frac{\sigma_y}{2} \right]} \quad (\text{A.19})$$

$$= \cos\left(\frac{\theta}{2}\right) \cdot I - i \cdot \sin\left(\frac{\theta}{2}\right) \cdot \left[\cos(\phi) \cdot \sigma_x - \sin(\phi) \cdot \sigma_y \right], \quad (\text{A.20})$$

where I is the identity matrix, and $\theta = \omega_R \cdot T$ the rotation angle and ϕ defines the rotation axis. This unitary operator is used to evaluate the process fidelity (Eq. (A.3) with $n = 2$) of any non-ideal operation U_{real} attempting to implement the same rotation angle around the same rotation axis.

A.2.1. Inaccuracies

Due to inaccuracies in the Hamiltonian or in the timing, the implemented operation $U_{\text{real}} = e^{-iH_{\text{real}} \cdot T_{\text{real}}}$ has a reduced fidelity.

In case of a microwave frequency inaccuracy $\omega_{\text{mw}} = \omega_0 + \Delta\omega_{\text{mw}}$, the fidelity follows (for any rotation angle/axis) as:

$$F = \frac{\left[\sin\left(\frac{\theta}{2}\right) \cdot \sin\left(\frac{\theta}{2} \cdot \sqrt{\alpha^2 + 1}\right) + \sqrt{\alpha^2 + 1} \cdot \cos\left(\frac{\theta}{2}\right) \cdot \cos\left(\frac{\theta}{2} \cdot \sqrt{\alpha^2 + 1}\right) \right]^2}{\alpha^2 + 1}, \quad (\text{A.21})$$

where $\alpha = \Delta\omega_{\text{mw}}/\omega_R$ is the error relative to the Rabi frequency. Taking the Taylor

series expansion of Eq. (A.21) with respect to α leads to:

$$F \approx 1 - \frac{1 - \cos(\theta)}{2} \cdot \alpha^2 + \mathcal{O}(\alpha^4). \quad (\text{A.22})$$

In case of a microwave phase inaccuracy $\phi = \phi_{\text{ideal}} + \Delta\phi$, the fidelity follows (for any rotation angle/axis) as:

$$F = \left[\cos(\alpha) \cdot \sin^2\left(\frac{\theta}{2}\right) + \cos^2\left(\frac{\theta}{2}\right) \right]^2, \quad (\text{A.23})$$

where $\alpha = \Delta\phi$ is the absolute error. Taking the Taylor series expansion of Eq. (A.23) with respect to α again leads to:

$$F \approx 1 - \frac{1 - \cos(\theta)}{2} \cdot \alpha^2 + \mathcal{O}(\alpha^4). \quad (\text{A.24})$$

In case of an inaccuracy in the microwave amplitude $\omega_R = \omega_{R,\text{ideal}} + \Delta\omega_R$, the fidelity follows (for any rotation angle/axis) as:

$$F = \cos^2\left(\frac{\theta}{2} \cdot \alpha\right), \quad (\text{A.25})$$

where $\alpha = \Delta\omega_R/\omega_{R,\text{ideal}}$ is the relative error. Taking the Taylor series expansion of Eq. (A.25) with respect to α leads to:

$$F \approx 1 - \left(\frac{\theta}{2}\right)^2 \cdot \alpha^2 + \mathcal{O}(\alpha^4). \quad (\text{A.26})$$

In case of an inaccuracy in the microwave duration $T = T_{\text{ideal}} + \Delta T$, the fidelity follows (for any rotation angle/axis) again as:

$$F = \cos^2\left(\frac{\theta}{2} \cdot \alpha\right), \quad (\text{A.27})$$

where $\alpha = \Delta T/T_{\text{ideal}}$ is the relative error. Taking the Taylor series expansion of Eq. (A.27) with respect to α will lead to Eq. (A.26).

A Z-rotation can be obtained without applying a signal to the qubit, simply by updating the reference frame, i.e. the phase of the microwave oscillator. The operation is given by:

$$U = \begin{bmatrix} e^{i\phi/2} & 0 \\ 0 & e^{-i\phi/2} \end{bmatrix}, \quad (\text{A.28})$$

where for the ideal operation U_{ideal} , the rotation angle $\phi = \phi_{\text{ideal}}$, and for the implemented operation U_{real} , the rotation angle $\phi = \phi_{\text{ideal}} + \Delta\phi$ has an error $\Delta\phi$. The fidelity follows (for any rotation angle) as:

$$F = \cos^2\left(\frac{\Delta\phi}{2}\right). \quad (\text{A.29})$$

Taking the Taylor series expansion of Eq. (A.29) with respect to $\Delta\phi$ leads to:

$$F \approx 1 - \frac{1}{4} \cdot \Delta\phi^2 + \mathcal{O}(\Delta\phi^4). \quad (\text{A.30})$$

A.2.2. Quasi-static Noise

The fidelity formulas in Eqs. (A.22), (A.24) and (A.26) all follow the relation $F(x) = 1 - c \cdot x^2$ for which the expected fidelity in case of Gaussian distributed inaccuracies (noise) is known (Eq. (A.7)). For all cases but the case of a microwave frequency inaccuracy, the expected fidelity can also be evaluated using the exact fidelity formula (no series expansion), which is more accurate for large amounts of noise.

In case of Gaussian distributed microwave phase noise $\phi = \mathcal{N}(\phi_{\text{ideal}}, \sigma_\phi^2)$, the expected fidelity follows (for any rotation angle/axis) as:

$$F = \frac{1}{2} \cdot \left(1 + e^{-2 \cdot \alpha^2}\right) \sin^4\left(\frac{\theta}{2}\right) + \cos^4\left(\frac{\theta}{2}\right) + \frac{1}{2} \cdot e^{-\frac{\alpha^2}{2}} \cdot \sin^2(\theta), \quad (\text{A.31})$$

where $\alpha = \sigma_\phi$ is the standard deviation of ϕ .

In case of Gaussian distributed microwave amplitude noise $\omega_R = \mathcal{N}(\omega_{R,\text{ideal}}, \sigma_{\omega_R}^2)$, the expected fidelity follows (for any rotation angle/axis) as:

$$F = \frac{1}{2} + \frac{1}{2} \cdot e^{-\frac{1}{2} \cdot \alpha^2 \theta^2}, \quad (\text{A.32})$$

where $\alpha = \sigma_{\omega_R} / \omega_{R,\text{ideal}}$ is the relative standard deviation of ω_R .

In case of Gaussian distributed timing variations $T = \mathcal{N}(T_{\text{ideal}}, \sigma_T^2)$, the expected fidelity follows (for any rotation angle/axis) again as:

$$F = \frac{1}{2} + \frac{1}{2} \cdot e^{-\frac{1}{2} \cdot \alpha^2 \theta^2}, \quad (\text{A.33})$$

where $\alpha = \sigma_T / T_{\text{ideal}}$ is the relative standard deviation of T .

A.2.3. Noise Filtering

In [147, 148] it is shown that when writing the total Hamiltonian $H(t) = H_c(t) + H_0(t)$ as the sum of a noise-free Hamiltonian $H_c(t)$ and a generalized noise Hamiltonian $H_0(t) = \beta_x(t)\sigma_x + \beta_y(t)\sigma_y + \beta_z(t)\sigma_z$, in first order approximation the expected process fidelity (from here on simply denoted with F) can be written as:

$$F = 1 - \frac{1}{2 \cdot \pi} \sum_{i,j,k=x,y,z} \int_{-\infty}^{\infty} S_{ij}(\omega) \cdot \frac{M_{jk}(\omega) \cdot M_{ik}^*(\omega)}{\omega^2} \cdot d\omega, \quad (\text{A.34})$$

with $S_{ij}(\omega)$ the cross-power spectral density between the random variables $\beta_i(t)$ and $\beta_j(t)$. The factors $M_{ij}(\omega)$, i.e. the control matrices in the frequency domain, depend on the control propagator $U_c(t) = e^{-i \cdot H_c(t) \cdot t}$:

$$M_{ij}(\omega) = -\frac{i \cdot \omega}{2} \int_0^T \text{Tr} \left[U_c^\dagger(t) \sigma_i U_c(t) \sigma_j \right] \cdot e^{i \cdot \omega \cdot t} dt. \quad (\text{A.35})$$

For the single-qubit operation, Eq. (A.18) is used to represent the noise-free Hamiltonian $H_c(t)$. The high-frequency noise sources of interest are fluctuations in the microwave frequency $\omega_{\text{mw}}(t) = \omega_{\text{mw,nom}} + \delta\omega_{\text{mw}}(t)$, and fluctuations in

the microwave envelope $\omega_R(t) = \omega_{R,\text{nom}} + \delta\omega_R(t)$. It is safe to assume that the fluctuations $\delta\omega_{\text{mw}}(t)$ and $\delta\omega_R(t)$ (with power spectral densities $S_{\text{mw}}(\omega)$ and $S_R(\omega)$, respectively) are statistically independent due to the different nature of the noise source.

With these assumptions, the generalized noise Hamiltonian follows as:

$$H_0(t) = \delta\omega_R(t) \cdot \left[\cos(\phi) \cdot \frac{\sigma_x}{2} - \sin(\phi) \cdot \frac{\sigma_y}{2} \right], \quad (\text{A.36})$$

in case of amplitude noise, and in case of frequency noise as:

$$H_0(t) = \delta\omega_{\text{mw}}(t) \cdot \frac{\sigma_z}{2}, \quad (\text{A.37})$$

and the non-zero cross-power spectral densities $S_{ij}(\omega)$ are found as:

$$S_{xy}(\omega) = -\frac{1}{4} \cdot S_R(\omega) \cdot \sin(\phi) \cdot \cos(\phi) \quad (\text{A.38})$$

$$S_{yx}(\omega) = -\frac{1}{4} \cdot S_R(\omega) \cdot \sin(\phi) \cdot \cos(\phi) \quad (\text{A.39})$$

$$S_{xx}(\omega) = \frac{1}{4} \cdot S_R(\omega) \cdot \cos(\phi)^2 \quad (\text{A.40})$$

$$S_{yy}(\omega) = \frac{1}{4} \cdot S_R(\omega) \cdot \sin(\phi)^2, \quad (\text{A.41})$$

in case of amplitude noise, and in case of frequency noise as:

$$S_{zz}(\omega) = \frac{1}{4} \cdot S_{\text{mw}}(\omega). \quad (\text{A.42})$$

Eq. (A.34) can now be evaluated to find the expected fidelity in case of amplitude noise (Eq. (A.43)) and microwave frequency noise (Eq. (A.44)).

$$F = 1 - \frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{S_R(\omega)}{\omega_R^2} \cdot |H_R(\omega)|^2 \cdot d\omega \quad (\text{A.43})$$

$$F = 1 - \frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{S_{\text{mw}}(\omega)}{\omega_R^2} \cdot |H_{\text{mw}}(\omega)|^2 \cdot d\omega, \quad (\text{A.44})$$

where $|H_R(\omega)|^2$ and $|H_{\text{mw}}(\omega)|^2$ are the amplitude responses of the filter functions for the respective type of noise:

$$|H_R(\omega)|^2 = \frac{\sin\left(\alpha \cdot \frac{\theta}{2}\right)^2}{\alpha^2} \quad (\text{A.45})$$

$$|H_{\text{mw}}(\omega)|^2 = \frac{[1 - \cos(\theta) \cdot \cos(\alpha \cdot \theta)] \cdot (\alpha^2 + 1) - 2 \cdot \alpha \cdot \sin(\theta) \cdot \sin(\alpha \cdot \theta)}{2 \cdot (\alpha^2 - 1)^2}, \quad (\text{A.46})$$

A

where $\alpha = \frac{\omega}{\omega_R}$, is the frequency normalized to the Rabi frequency.

Note that for the microwave amplitude noise, the frequency axis is defined relative to the Rabi frequency, i.e., the bandwidth relevant for amplitude noise is proportional to the nominal amplitude. However, in case of white noise, Eq. (A.43) simplifies to Eq. (A.25) with $(\Delta\omega_R/\omega_R)^2 = \sigma_{\omega_R}^2/\omega_R^2$ where $\sigma_{\omega_R}^2$ is the noise power. As can be seen, for a certain fidelity the required SNR ($\omega_R^2/\sigma_{\omega_R}^2$) in the qubit's band of sensitivity is fixed.

In case of wideband additive noise ($\delta\omega_{\text{add}}(t)$), the signal can be better modeled as $\omega_{\text{ESR}}(t) = 2 \cdot [\omega_R \cdot \cos(\omega_{\text{mw}} \cdot t) + \delta\omega_{\text{add}}(t)]$. The lab frame Hamiltonian (Eq. (A.12)) follows as:

$$H_{\text{lab}} = -\omega_0 \cdot \frac{\sigma_z}{2} + [\omega_R \cdot \cos(\omega_{\text{mw}} \cdot t) + \delta\omega_{\text{add}}(t)] \cdot \sigma_x. \quad (\text{A.47})$$

Evaluating this Hamiltonian in the rotating frame (Eqs. (A.13) to (A.15) with $\omega_{\text{mw}} = \omega_0$) leads to (after taking the RWA):

$$H = \omega_R \cdot \frac{\sigma_x}{2} + \delta\omega_{\text{add}}(t) \cdot [\cos(t \cdot \omega_0) \cdot \sigma_x + \sin(t \cdot \omega_0) \cdot \sigma_y]. \quad (\text{A.48})$$

The first part of this equation is the noise-free Hamiltonian ($H_c(t) = \omega_R \cdot \frac{\sigma_x}{2}$), while the remainder forms the generalized noise Hamiltonian, with:

$$\beta_x(t) = \delta\omega_{\text{add}}(t) \cdot \cos(t \cdot \omega_0) \quad (\text{A.49})$$

$$\beta_y(t) = \delta\omega_{\text{add}}(t) \cdot \sin(t \cdot \omega_0) \quad (\text{A.50})$$

$$\beta_z(t) = 0. \quad (\text{A.51})$$

Assuming the additive noise $\delta\omega_{\text{add}}(t)$ has a power spectral density $S_{\text{add}}(\omega)$, the non-zero cross-power spectral densities $S_{ij}(\omega)$ are found as the Fourier transform of the cross-correlations $R_{ij}(\tau)$:

$$R_{ij}(\tau) = \int_{-\infty}^{\infty} \beta_i(t) \cdot \beta_j(t + \tau) dt, \quad (\text{A.52})$$

which evaluate to:

$$R_{xy}(\tau) = \frac{1}{2} \cdot \sin(\omega_0 \cdot \tau) \cdot R_{\text{add}}(\tau) \quad (\text{A.53})$$

$$R_{yx}(\tau) = -\frac{1}{2} \cdot \sin(\omega_0 \cdot \tau) \cdot R_{\text{add}}(\tau) \quad (\text{A.54})$$

$$R_{xx}(\tau) = \frac{1}{2} \cdot \cos(\omega_0 \cdot \tau) \cdot R_{\text{add}}(\tau) \quad (\text{A.55})$$

$$R_{yy}(\tau) = \frac{1}{2} \cdot \cos(\omega_0 \cdot \tau) \cdot R_{\text{add}}(\tau), \quad (\text{A.56})$$

where $R_{\text{add}}(\tau)$ is the auto-correlation of $\delta\omega_{\text{add}}(t)$, i.e. the Fourier transform of $S_{\text{add}}(\omega)$. The $\sin(\omega_0 \cdot \tau)$ or $\cos(\omega_0 \cdot \tau)$ modulates the spectrum $S_{\text{add}}(\omega)$, leading

to:

$$S_{xy}(\omega) = \frac{1}{4} \cdot [S_{\text{add}}(\omega + \omega_0) - S_{\text{add}}(\omega - \omega_0)] \quad (\text{A.57})$$

$$S_{yx}(\omega) = \frac{1}{4} \cdot [S_{\text{add}}(\omega - \omega_0) - S_{\text{add}}(\omega + \omega_0)] \quad (\text{A.58})$$

$$S_{xx}(\omega) = \frac{1}{4} \cdot [S_{\text{add}}(\omega + \omega_0) + S_{\text{add}}(\omega - \omega_0)] \quad (\text{A.59})$$

$$S_{yy}(\omega) = \frac{1}{4} \cdot [S_{\text{add}}(\omega + \omega_0) + S_{\text{add}}(\omega - \omega_0)]. \quad (\text{A.60})$$

And the expected fidelity evaluates to (using the symmetry of the power spectral density):

$$F = 1 - \frac{1}{\pi} \int_0^\infty \frac{S_{\text{add}}(\omega - \omega_0)}{\omega_R^2} \cdot |H_{\text{add}}(\omega)|^2 \cdot d\omega, \quad (\text{A.61})$$

with:

$$|H_{\text{add}}(\omega)|^2 = |H_{\text{R}}(\omega)|^2 + |H_{\text{mw}}(\omega)|^2. \quad (\text{A.62})$$

A.2.4. Jitter

Besides the microwave amplitude, also the signal duration T is subject to random variations, i.e., jitter. This period jitter is determined by the single-sideband phase noise $S_\phi(f)$ of the reference clock (period T_{clk}) used to set the duration [172, 177]:

$$\sigma_T = \frac{T_{\text{clk}}}{\pi} \sqrt{\int_{f_{\text{min}}}^\infty S_\phi(f) \cdot \sin^2(\pi \cdot f \cdot T) \cdot df}. \quad (\text{A.63})$$

This integral shows that the phase noise, which generally rolls-off with frequency (Fig. 3.6), is filtered by a high-pass filter ($\sin^2(\pi \cdot f \cdot T)$) with the corner frequency set by the duration T . The resulting variations in the duration, with standard deviation σ_T , lead to an infidelity that can be estimated using Eq. (A.27) with $\Delta T/T = \sigma_T/T$ assuming Gaussian distributed jitter.

A.2.5. Idle Gate

This section discusses other processes that cause the state of the qubit to degrade during an idle period lasting T_{nop} , as indicated in Fig. A.2.

In case no microwave signal is applied to the qubits, residual noise on the drive line can still affect the qubits. Consider the Hamiltonian in the lab frame (Eq. (A.12)) with $\omega_{\text{ESR}}(t) = 2 \cdot \omega_{\text{R}_n}(t)$, where $\omega_{\text{R}_n}(t)$ is the noise signal with spectral density $S_{\text{R}_n}(\omega)$. This Hamiltonian can again be split into a noise-free Hamiltonian ($H_c = -\omega_0 \cdot \sigma_z/2$) and a generalized noise Hamiltonian $H_0(t) = \omega_{\text{R}_n}(t) \cdot \sigma_x$. It follows that $S_{xx}(\omega) = S_{\text{R}_n}(\omega)$ and all other cross-spectral densities are zero.

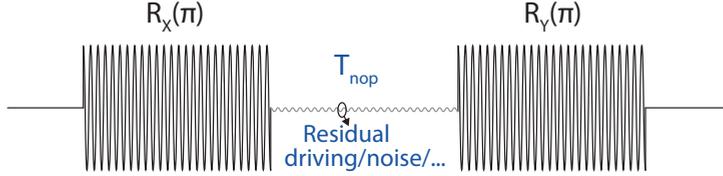


Figure A.2: The state of a qubit is affected during idle times between operations due to, e.g., residual driving on the ESR-line.

Equation (A.34) can now be evaluated, leading to:

$$F = 1 - \frac{1}{\pi} \int_{-\infty}^{\infty} S_{R_n}(\omega) \cdot G(\omega) \cdot d\omega \quad (\text{A.64})$$

$$G(\omega) = \frac{\sin^2 \left[\frac{T(\omega + \omega_0)}{2} \right]}{(\omega + \omega_0)^2} + \frac{\sin^2 \left[\frac{T(\omega - \omega_0)}{2} \right]}{(\omega - \omega_0)^2}. \quad (\text{A.65})$$

Using the symmetry of the spectrum $S_{R_n}(\omega)$ and by assuming $T = \theta/\omega_R$ is the time that would be needed to rotate a qubit by an angle θ when applying the signal amplitude for a Rabi frequency ω_R :

$$F = 1 - \frac{1}{\pi} \int_0^{\infty} \frac{S_{R_n}(\omega)}{\omega_R^2} \cdot |H_n(\omega)|^2 \cdot d\omega, \quad (\text{A.66})$$

where

$$|H_n(\omega)|^2 = 2 \cdot \left(\frac{\omega_R}{\omega - \omega_0} \right)^2 \cdot \sin^2 \left(\frac{\theta}{2} \frac{\omega - \omega_0}{\omega_R} \right), \quad (\text{A.67})$$

which represents the amplitude response of a sinc-shaped band-pass filter centered around ω_0 . For $\omega_0 = 0$ this would be a low-pass filter with equivalent noise bandwidth $ENBW_n = \omega_R \cdot \pi/|\theta|$ and DC-gain $|H_n(0)|^2 = \theta^2/2$. Therefore, a brickwall approximation of the amplitude response of Eq. (A.67), which would be a good approximation in case of white noise, is:

$$|H_n(\omega)|^2 \approx \begin{cases} \theta^2/2 & |\omega - \omega_0| \leq \omega_R \cdot \pi/|\theta| \\ 0 & \text{elsewhere} \end{cases}. \quad (\text{A.68})$$

Frequency inaccuracies also affect the qubits while no operation is performed. In a frame rotating with the oscillator's frequency, the qubit appears to rotate with a frequency $\Delta\omega$ representing the frequency inaccuracy. The fidelity of an identity operation for a duration T_{nop} evaluates to:

$$F = \cos^2 \left(\frac{\Delta\omega \cdot T_{nop}}{2} \right), \quad (\text{A.69})$$

for which the Taylor series expansion follows as:

$$F = 1 - \frac{T_{\text{nop}}^2}{4} \cdot \Delta\omega^2 + \mathcal{O}(\Delta\omega^4). \quad (\text{A.70})$$

A residual spurious tone driving the qubit for a duration T_{nop} while not intended would also reduce its fidelity. Using the rotating frame Hamiltonian of Eq. (A.18) with $\omega_0 = \omega_{\text{mw}}$ and $\omega_{\text{R}} = \omega_{\text{R,spur}}$ results in a fidelity of an identity operation of:

$$F = \cos^2\left(\frac{\omega_{\text{R,spur}} \cdot T_{\text{nop}}}{2}\right), \quad (\text{A.71})$$

for which the Taylor series expansion follows as:

$$F = 1 - \frac{T_{\text{nop}}^2}{4} \cdot \omega_{\text{R,spur}}^2 + \mathcal{O}(\omega_{\text{R,spur}}^4). \quad (\text{A.72})$$

A.2.6. Frequency Multiplexing

Driving a certain qubit at a frequency $\omega_0 = \omega_{\text{mw}}$ can also influence another qubit at a frequency $\omega_{0,\text{other}} = \omega_0 + \omega_{0,\text{space}}$ separated by $\omega_{0,\text{space}}$. To simplify the analysis, again a time-independent Hamiltonian is obtained by moving to a frame rotating with ω_{mw} (Eqs. (A.13) to (A.15)). In this frame, the other qubit appears to rotate around the z-axis with a frequency $\omega_{0,\text{space}}$. The ideal operation, an identity, can be described in this frame as:

$$U_{\text{ideal}} = e^{-i\omega_{0,\text{space}} \cdot \frac{\sigma_z}{2} \cdot T}, \quad (\text{A.73})$$

whereas the real operation follows as (for $\phi = 0$):

$$U_{\text{real}} = e^{-i\left[\omega_{0,\text{space}} \cdot \frac{\sigma_z}{2} + \omega_{\text{R,other}} \cdot \frac{\sigma_x}{2}\right] \cdot T}. \quad (\text{A.74})$$

Recall that the Rabi frequency $\omega_{\text{R,other}}$ is related to the amplitude of the driving magnetic field by $\omega_{\text{R,other}} = A_{\text{other}} \cdot \gamma_e/2$. In general, the required microwave amplitude for a certain Rabi frequency can be different for the two qubits involved. Therefore $\omega_{\text{R,other}}$ in the equation above can be considered as the amplitude driving the other qubit, while ω_{R} is introduced as the amplitude driving a rotation $\theta = \omega_{\text{R}} \cdot T$ in a duration T on the intended qubit. Equation (A.74) can then be rewritten as:

$$U_{\text{real}} = e^{-i\theta \cdot \left[\frac{\omega_{0,\text{space}}}{\omega_{\text{R}}} \cdot \frac{\sigma_z}{2} + \frac{\omega_{\text{R,other}}}{\omega_{\text{R}}} \cdot \frac{\sigma_x}{2}\right]}. \quad (\text{A.75})$$

The fidelity follows as ($\alpha = \frac{\omega_{0,\text{space}}}{\omega_{\text{R}}}$ and $\beta = \frac{\omega_{\text{R,other}}}{\omega_{\text{R}}}$):

$$F = \frac{\left|\sqrt{\alpha^2 + \beta^2} \cdot \cos\left(\frac{\theta}{2}\sqrt{\alpha^2 + \beta^2}\right) \cdot (1 + \cos(\theta\alpha) + i \sin(\theta\alpha)) + \alpha \cdot \sin\left(\frac{\theta}{2}\sqrt{\alpha^2 + \beta^2}\right) \cdot (i - i \cos(\theta\alpha) + \sin(\theta\alpha))\right|^2}{4(\alpha^2 + \beta^2)}. \quad (\text{A.76})$$

Recall from the fidelity formula (Eq. (A.3)) that the fidelity is unity in case $U_{\text{ideal}}^\dagger \cdot U_{\text{real}} = I$ (I is the Identity). To gain more insight into the infidelity of the other qubit,

the decomposition $U_{\text{ideal}}^\dagger U_{\text{real}} = x \cdot \sigma_x + y \cdot \sigma_y + z \cdot \sigma_z + l \cdot I$ is made. Since $|l|^2 = F$ and $|x|^2 + |y|^2 + |z|^2 + |l|^2 = 1$, the infidelity equals $1 - F = |x|^2 + |y|^2 + |z|^2$ and has contributions from X, Y and Z-rotations. Since the Z-rotations can easily be removed by a software update of the reference frame, the residual infidelity contributions ($1 - F_{\text{corr}}$) can be found from the decomposition as:

$$|x|^2 + |y|^2 = \frac{\beta^2}{\alpha^2 + \beta^2} \cdot \sin^2\left(\frac{\theta}{2} \cdot \sqrt{\alpha^2 + \beta^2}\right). \quad (\text{A.77})$$

For $\alpha^2 \gg \beta^2$ (valid for the cases of interest: sufficient qubit spacing and $\omega_{\text{R,other}} \leq \omega_{\text{R}}$):

$$F_{\text{corr}} \approx 1 - \frac{\beta^2}{\alpha^2} \cdot \sin^2\left(\frac{\theta}{2} \cdot \alpha\right) \geq 1 - \frac{\beta^2}{\alpha^2}, \quad (\text{A.78})$$

which shows a close resemblance to the power of the Fourier transform of the rectangular microwave envelope: $4/\alpha^2 \cdot \sin^2\left(\frac{\theta}{2} \cdot \alpha\right)$.

A.3. Derivations for Two-Qubit Operation

In this section, and the next, the Hamiltonian will be extended with the singlet and triplet states, which are defined as:

$$|S\rangle = 1/\sqrt{2} [|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle] \quad (\text{A.79})$$

$$|T_0\rangle = 1/\sqrt{2} [|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle] \quad (\text{A.80})$$

$$|T_-\rangle = |\downarrow\downarrow\rangle \quad (\text{A.81})$$

$$|T_+\rangle = |\uparrow\uparrow\rangle. \quad (\text{A.82})$$

As discussed in the main text, the following Hamiltonian is used for our 2-qubit system ($\hbar = 1$) [66, 181, 182]:

$$H = \begin{bmatrix} -\omega_0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{\delta\omega_0}{2} & 0 & 0 & t_0 & t_0 \\ 0 & 0 & -\frac{\delta\omega_0}{2} & 0 & -t_0 & -t_0 \\ 0 & 0 & 0 & \omega_0 & 0 & 0 \\ 0 & t_0 & -t_0 & 0 & U - \epsilon & 0 \\ 0 & t_0 & -t_0 & 0 & 0 & U + \epsilon \end{bmatrix}. \quad (\text{A.83})$$

A.3.1. Hamiltonian Eigenenergies

The Hamiltonian of Eq. (A.83) has six eigenvalues. However, as the quantum state is encoded in the spin state, the two eigenvalues related to the single-dot singlet states are not analyzed in the following. The remaining four eigenvalues (ω_{λ_i}) of the Hamiltonian (Eq. (A.83)) have been analyzed for various Larmor frequency differences $\delta\omega_0$. However, independent of the choice of $\delta\omega_0$:

$$\omega_{\lambda_1} = -\omega_0 \quad (\text{A.84})$$

$$\omega_{\lambda_4} = \omega_0. \quad (\text{A.85})$$

First, the case of $\delta\omega_0 = 0$ is analyzed. By taking the 2nd-order Taylor series expansion around $t_0 = 0$, i.e., the tunnel coupling small compared to the charging energy, the eigenenergies are found as:

$$\omega_{\lambda_2} = 0 \quad (\text{A.86})$$

$$\omega_{\lambda_3} = -\frac{4 \cdot U \cdot t_0^2}{U^2 - \epsilon^2}. \quad (\text{A.87})$$

Now, the special case of $\delta\omega_0 = \sqrt{2} \cdot t_0$ is analyzed. By again taking the 2nd-order Taylor series expansion around $t_0 = 0$, the eigenenergies are found as:

$$\omega_{\lambda_2} = \frac{\delta\omega_0}{2} - \frac{2 \cdot U \cdot t_0^2}{U^2 - \epsilon^2} \quad (\text{A.88})$$

$$\omega_{\lambda_3} = -\frac{\delta\omega_0}{2} - \frac{2 \cdot U \cdot t_0^2}{U^2 - \epsilon^2}. \quad (\text{A.89})$$

When removing the Larmor precession:

$$\omega'_{\lambda_2} = -\frac{2 \cdot U \cdot t_0^2}{U^2 - \epsilon^2} \quad (\text{A.90})$$

$$\omega'_{\lambda_3} = -\frac{2 \cdot U \cdot t_0^2}{U^2 - \epsilon^2}. \quad (\text{A.91})$$

Comparing Eqs. (A.86) and (A.87) with Eqs. (A.90) and (A.91), it appears that $\omega_{\lambda_2} + \omega_{\lambda_3}$ is independent of the choice of $\delta\omega_0$, but that the ratio $\omega'_{\lambda_2}/\omega'_{\lambda_3}$ changes.

A more in depth analysis shows that the exact, and generally valid, solutions to ω_{λ_2} and ω_{λ_3} are found by solving Eq. (A.92). Under the assumptions $\omega_{\lambda_i} \ll t_0, U$ and $U^2 - \epsilon^2 \gg t_0^2, \delta\omega_0^2$ and $\delta\omega_0^2 \ll 8 \cdot t_0^2$, this equation can be simplified to Eq. (A.93).

$$\omega_{\lambda_i}^4 - 4 \cdot \omega_{\lambda_i}^3 \cdot U + \omega_{\lambda_i}^2 \cdot (-\delta\omega_0^2 - 4 \cdot \epsilon^2 - 16 \cdot t_0^2 + 4 \cdot U^2) + 4 \cdot \omega_{\lambda_i} \cdot U (\delta\omega_0^2 + 8 \cdot t_0^2) + 4 \cdot \delta\omega_0^2 (\epsilon^2 - U^2) = 0 \quad (\text{A.92})$$

$$\omega_{\lambda_i}^2 (-4 \cdot \epsilon^2 + 4 \cdot U^2) + 32 \cdot \omega_{\lambda_i} \cdot U \cdot t_0^2 + 4 \cdot \delta\omega_0^2 (\epsilon^2 - U^2) \approx 0 \quad (\text{A.93})$$

As a result, for the case $\delta\omega_0 < \sqrt{2} \cdot t_0$, the eigenenergies are found as:

$$\omega_{\lambda_2} = -\frac{2 \cdot U \cdot t_0^2}{U^2 - \epsilon^2} + \frac{1}{2} \cdot \sqrt{\delta\omega_0^2 + \left(\frac{4 \cdot U \cdot t_0^2}{U^2 - \epsilon^2}\right)^2} \quad (\text{A.94})$$

$$\omega_{\lambda_3} = -\frac{2 \cdot U \cdot t_0^2}{U^2 - \epsilon^2} - \frac{1}{2} \cdot \sqrt{\delta\omega_0^2 + \left(\frac{4 \cdot U \cdot t_0^2}{U^2 - \epsilon^2}\right)^2}. \quad (\text{A.95})$$

To show the validity of the simplification of Eq. (A.92) to Eq. (A.93), Fig. A.3 compares the eigenvalues of the Hamiltonian (Eq. (A.83)), with the eigenvalues as obtained from Eqs. (A.94) and (A.95).

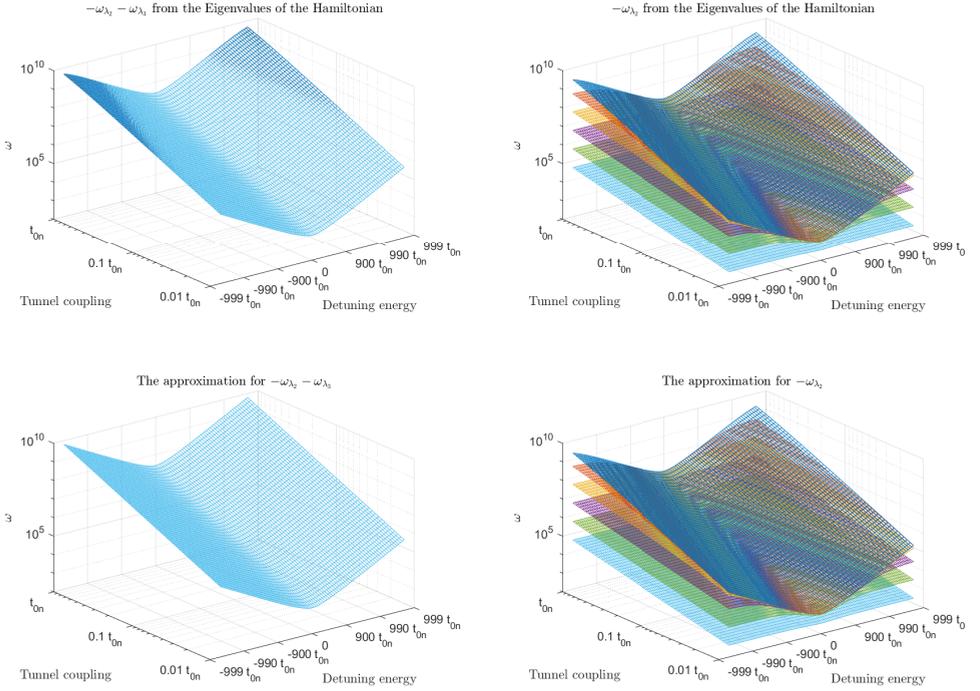


Figure A.3: The 2-qubit operation speeds ω_{λ_2} and ω_{λ_3} versus the interdot tunnel coupling and detuning. A nominal tunnel coupling t_{0n} of 1 GHz is used. The plots on the top show the eigenvalues of the Hamiltonian (Eq. (A.83)), whereas the plots on the bottom show the approximation of Eqs. (A.90) and (A.91). The different colors are used for different values of $\delta\omega_0$ ($\delta\omega_0 = \sqrt{2} \cdot t_0/10^n$, with n ranging from 0 for the blue curves up to 5 for the cyan curves).

Hence, in general:

$$\omega_{\lambda_1} = -\omega_0 \quad (\text{A.96})$$

$$\omega_{\lambda_2} \approx \begin{cases} \frac{-\omega_{\text{op}} + \sqrt{\delta\omega_0^2 + \omega_{\text{op}}^2}}{2} & 0 \leq \delta\omega_0 < \sqrt{2} \cdot t_0 \\ \frac{-\omega_{\text{op}} + \delta\dot{\omega}_0}{2} & \delta\omega_0 = \sqrt{2} \cdot t_0 \end{cases} \quad (\text{A.97})$$

$$\omega_{\lambda_3} \approx \begin{cases} \frac{-\omega_{\text{op}} - \sqrt{\delta\omega_0^2 + \omega_{\text{op}}^2}}{2} & 0 \leq \delta\omega_0 < \sqrt{2} \cdot t_0 \\ \frac{-\omega_{\text{op}} - \delta\dot{\omega}_0}{2} & \delta\omega_0 = \sqrt{2} \cdot t_0 \end{cases} \quad (\text{A.98})$$

$$\omega_{\lambda_4} = \omega_0, \quad (\text{A.99})$$

where

$$\omega_{\text{op}} = 4 \cdot t_0^2 \cdot \frac{U}{U^2 - \epsilon^2}. \quad (\text{A.100})$$

A.3.2. The C-Phase Gate

To perform a C-phase gate, the control parameter must change adiabatically. An adiabatic change of the control implies that if the qubit state was an eigenvector (stationary state or eigenstate) of the Hamiltonian, it remains an eigenstate of the new Hamiltonian after the control change. As a result, the ideal adiabatic operation to the desired operating point can be described as:

$$U_{\text{in}} = V, \quad (\text{A.101})$$

where V contains the eigenvectors of the Hamiltonian in the desired operating point H_{op} , assuming the Hamiltonian is expressed in a basis formed by the eigenvectors in the original operating point.

Moving back from the desired operating point to this original point can then be described as:

$$U_{\text{out}} = V^{-1}. \quad (\text{A.102})$$

Finally, the operation in the desired operating point can be described as:

$$U_{\text{op}} = e^{-i \cdot H_{\text{op}} \cdot T} = V \cdot e^{-i \cdot D \cdot T} \cdot V^{-1}, \quad (\text{A.103})$$

where the eigenvalue decomposition of $H_{\text{op}} = V \cdot D \cdot V^{-1}$ has been used (D is a diagonal matrix containing the eigenvalues of H_{op}).

In total, an adiabatic operation (U) consisting of moving adiabatically to the desired operating point, operating for a while in this point, and moving back adiabatically, can be simplified as:

$$U = U_{\text{out}} \cdot U_{\text{op}} \cdot U_{\text{in}} \quad (\text{A.104})$$

$$= V^{-1} \cdot V \cdot e^{-i \cdot D \cdot T} \cdot V^{-1} \cdot V \quad (\text{A.105})$$

$$= e^{-i \cdot D \cdot T}, \quad (\text{A.106})$$

which is then a diagonal matrix that only depends on the eigenenergies of the Hamiltonian:

$$U_{cz,lab}(t) = \begin{bmatrix} e^{-i \cdot t \cdot \omega_{\lambda_1}} & 0 & 0 & 0 \\ 0 & e^{-i \cdot t \cdot \omega_{\lambda_2}} & 0 & 0 \\ 0 & 0 & e^{-i \cdot t \cdot \omega_{\lambda_3}} & 0 \\ 0 & 0 & 0 & e^{-i \cdot t \cdot \omega_{\lambda_4}} \end{bmatrix}. \quad (\text{A.107})$$

In the rotating frame this becomes the unitary operation describing the C-phase gate (Eq. (3.13)) with rotation angle $\theta_{cz} = -(\phi_{Z,A} + \phi_{Z,B}) = -(\omega_{\lambda_2} + \omega_{\lambda_3}) \cdot t = \omega_{op} \cdot t$.

As long as the adiabatic parts are not too slow, most of the operation occurs at the desired operating point as the exchange interaction is strongest there (Eq. (3.12)). Consequently, only the effects of signal inaccuracy and noise on U_{op} will be considered. Since for the exchange gate the control parameter changes diabatically, again the analysis will be limited to U_{op} only.

In case of an inaccuracy in the control parameters, the operation is still described by the diagonal matrix of Eq. (A.107), and takes the form of Eq. (3.13) in the rotating frame, however with different angles $\phi_{Z,A}$ and $\phi_{Z,B}$. The fidelity of an inaccurate operation follows as (Eq. (3.3)):

$$\begin{aligned} F &= \frac{3}{8} + \cos(\phi_{Z,A,ideal} - \phi_{Z,A,real} - \phi_{Z,B,ideal} + \phi_{Z,B,real}) \\ &+ \frac{2}{8} \cos(\phi_{Z,A,ideal} - \phi_{Z,A,real}) \\ &+ \frac{2}{8} \cos(\phi_{Z,B,ideal} - \phi_{Z,B,real}), \end{aligned} \quad (\text{A.108})$$

where $\phi_{Z,A,ideal}$ and $\phi_{Z,B,ideal}$ are the acquired phases in case of no inaccuracy, and $\phi_{Z,A,real}$ and $\phi_{Z,B,real}$ are the acquired phases in case of an inaccuracy in the control parameter.

Evaluating this formula for inaccuracies in duration ($T_{real} = T + \Delta T$), tunnel coupling ($t_{0,real} = t_0 + \Delta t_0$), and detuning ($\epsilon_{real} = \epsilon + \Delta \epsilon$), for the different operating points, lead to the infidely formulas as summarized in Table 3.3, when taking the 2nd-order Taylor series expansion to the inaccuracy (except for the case $\epsilon = 0$ for which a 4th-order Taylor series expansion is used).

A.3.3. The Exchange Gate

In case the control parameter changes diabatically, as required for the exchange gate, U_{in} and U_{out} approximate the identity matrix, and the resulting 2-qubit operation is described by Eq. (A.103). The relevant eigenenergies (in D) are given in Appendix A.3.1.

For $\delta\omega_0 = 0$ and small t_0 (taking the Taylor series expansion), the 4×4 relevant

entries of the eigenvector matrix can be approximated as:

$$V^{-1} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} & 0 \\ 0 & -\frac{t_0}{U+\epsilon} & \frac{t_0}{U+\epsilon} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (\text{A.109})$$

Equation (A.103) can now be evaluated, leading to:

$$U_{J'}(t) \approx \begin{bmatrix} e^{-i \cdot t \cdot \omega_{\lambda_1}} & 0 & 0 & 0 \\ 0 & \frac{e^{-i \cdot t \cdot \omega_{\lambda_2}} + e^{-i \cdot t \cdot \omega_{\lambda_3}}}{2} & \frac{e^{-i \cdot t \cdot \omega_{\lambda_2}} - e^{-i \cdot t \cdot \omega_{\lambda_3}}}{2} & 0 \\ 0 & \frac{e^{-i \cdot t \cdot \omega_{\lambda_2}} - e^{-i \cdot t \cdot \omega_{\lambda_3}}}{2} & \frac{e^{-i \cdot t \cdot \omega_{\lambda_2}} + e^{-i \cdot t \cdot \omega_{\lambda_3}}}{2} & 0 \\ 0 & 0 & 0 & e^{-i \cdot t \cdot \omega_{\lambda_4}} \end{bmatrix}. \quad (\text{A.110})$$

In the rotating frame this becomes the unitary operation describing the Exchange gate (Eq. (3.14)) with rotation angle $\theta_J = -\omega_{\lambda_3} \cdot t = \omega_{\text{op}} \cdot t$. In case $\theta_J = \pi$, a SWAP operation is obtained.

In case of an inaccuracy in the control parameters, the operation is still described by this unitary matrix, however with a different angle θ_J . The fidelity of an inaccurate operation follows as (Eq. (3.3)):

$$F = \frac{5}{8} + \frac{3}{8} \cdot \cos(\theta_{J,\text{ideal}} - \theta_{J,\text{real}}), \quad (\text{A.111})$$

where $\theta_{J,\text{ideal}}$ is the acquired rotation angle in case of no inaccuracy, and $\theta_{J,\text{real}}$ is the acquired rotation angle in case of an inaccuracy in the control parameter.

Evaluating this formula for inaccuracies in duration ($T_{\text{real}} = T + \Delta T$), tunnel coupling ($t_{0,\text{real}} = t_0 + \Delta t_0$), and detuning ($\epsilon_{\text{real}} = \epsilon + \Delta \epsilon$), lead to the infidely formulas as summarized in Table 3.3, when taking the 2nd-order Taylor series expansion to the inaccuracy (except for the case $\epsilon = 0$ for which a 4th-order Taylor series expansion is used).

A.3.4. Idle Gate

Evaluating the fidelity (Eq. (3.3)) of the 2-qubit operations (Eq. (3.13) and Eq. (3.14)) with respect to an identity operation as ideal operation ($U_{\text{ideal}} = I$), leads to

$$F_I = \frac{3}{8} + \cos(\phi_{Z,A} - \phi_{Z,B}) + \frac{2}{8} \cdot \cos(\phi_{Z,A}) + \frac{2}{8} \cdot \cos(\phi_{Z,B}) \quad (\text{A.112})$$

and

$$F_I = \frac{5}{8} + \frac{3}{8} \cdot \cos(\theta_J), \quad (\text{A.113})$$

for the C-phase gate and exchange gate respectively.

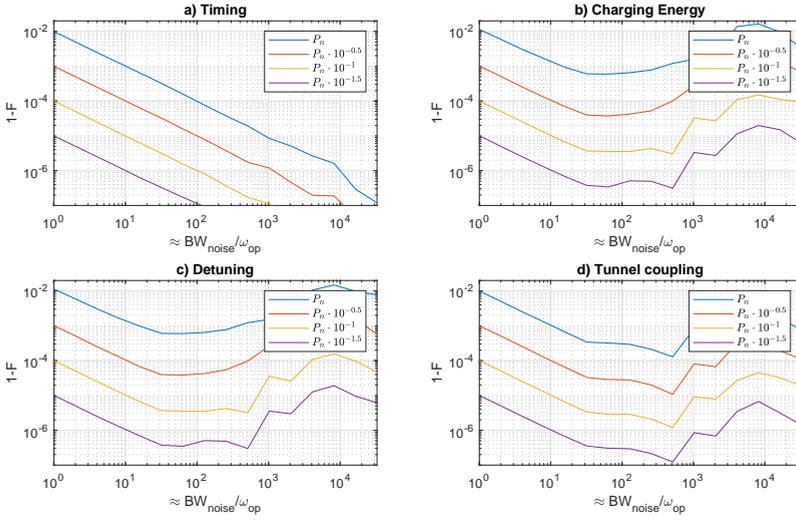


Figure A.4: Numerical simulation of the noise sensitivity for a C-Phase gate at 90% detuning. The total integrated noise is kept constant, while the simulation bandwidth was changed, thereby indirectly changing the noise bandwidth. This simulation was repeated for different amounts of noise power P_n .

Assuming a total acquired phase $\theta_{cZ} = -(\phi_{Z,A} + \phi_{Z,B})$, and using the value of $\phi_{Z,B}$ as summarized in Table 3.3 for different values of $\delta\omega_0$, simplifies Eq. (A.112) to:

$$F_I = \begin{cases} 1 - \frac{3}{16} \cdot \theta_{cZ}^2 & \delta\omega_0 = 0 \\ 1 - \frac{7-4\sqrt{2}}{16} \cdot \theta_{cZ}^2 & \delta\omega_0 = \omega_{op} \\ 1 - \frac{1}{16} \cdot \theta_{cZ}^2 & \delta\omega_0 = \sqrt{2} \cdot t_0 \end{cases}, \quad (\text{A.114})$$

after taking the 2nd order Taylor series expansion to θ_{cZ} .

Taking the 2nd order Taylor series expansion to θ_j in Eq. (A.113) leads to:

$$F_I = 1 - \frac{3}{16} \cdot \theta_j^2. \quad (\text{A.115})$$

A.3.5. The noise sensitivity

In an initial numerical simulation of the noise sensitivity, a C-Phase gate was considered at a detuning of 90% of the charging energy. A two-dimensional sweep was performed, sweeping both the total integrated noise power and the noise bandwidth (by changing the simulation bandwidth). The results are shown in Fig. A.4 when the noise is applied to respectively the timing (Fig. A.4a), the charging energy (Fig. A.4b), the detuning (Fig. A.4c) and the tunnel rate (Fig. A.4d). When the total integrated noise power remains the same and the noise bandwidth is increased, the effect on the qubit is expected to decrease if the qubit is only sensitive to low frequency noise. While this holds for relatively low frequencies, the qubit is again

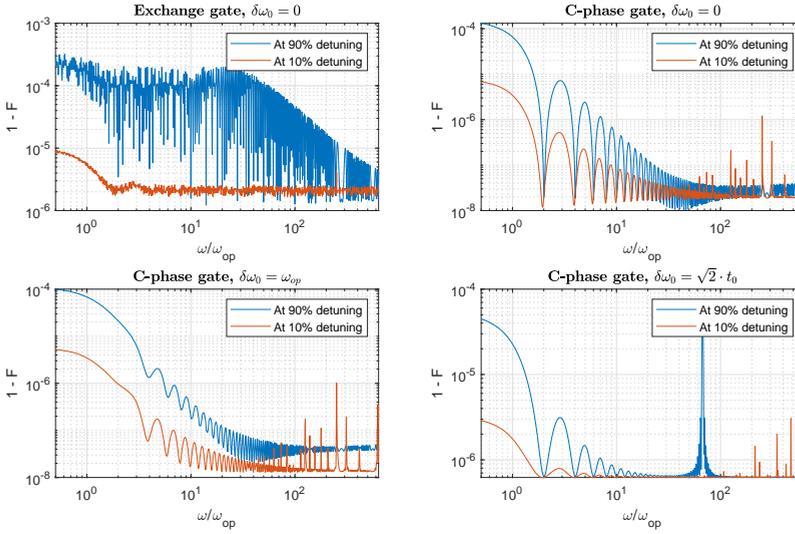


Figure A.5: A numerical simulation of the sensitivity of the qubit to detuning noise of different frequencies for different operating regimes. For this simulation, a sinewave with RMS amplitude of 0.1 % of the charging energy is applied on the detuning control.

affected more when high frequency noise is introduced. An exception is for the timing as only the total duration matters.

In order to get an indication of the noise transfer function for detuning noise, a simulation is performed where a small-amplitude sinewave is superimposed on the detuning to emulate narrow-band noise on the detuning. By sweeping the frequency of this sinewave, the sensitivity to signals/noise at different frequencies can be determined. This simulation is performed for all operating regimes discussed in the main text (Exchange gate and C-Phase gate with 3 different Larmor frequency differences), for both operation at 90 % detuning and 10 % detuning (in this case the amplitude of the sinewave is increased 10 \times as a lower sensitivity is expected). The result of these simulations are shown in Fig. A.5. A similar simulation was performed with a sinewave superimposed on the tunnel rate instead. In this case, the same amplitude of the sinewave is used for operation at 90 % detuning and 10 % detuning. The result of these simulations are shown in Fig. A.6.

From both Figs. A.5 and A.6, it can be seen that generally the qubit is most sensitive to low-frequency noise with a bandwidth of $\sim \omega_{op}$. However, signals/noise at high frequencies can also significantly affect the qubit. This is most clearly seen in the simulations of the C-Phase gate with $\delta\omega_0 = \sqrt{2} \cdot t_0$ in case of 90 % detuning where a peak is seen at $\sim 67 \cdot \omega_{op}$. This frequency, and the ones of the other peaks, correspond to allowed energy transitions (resonances) in the quantum system.

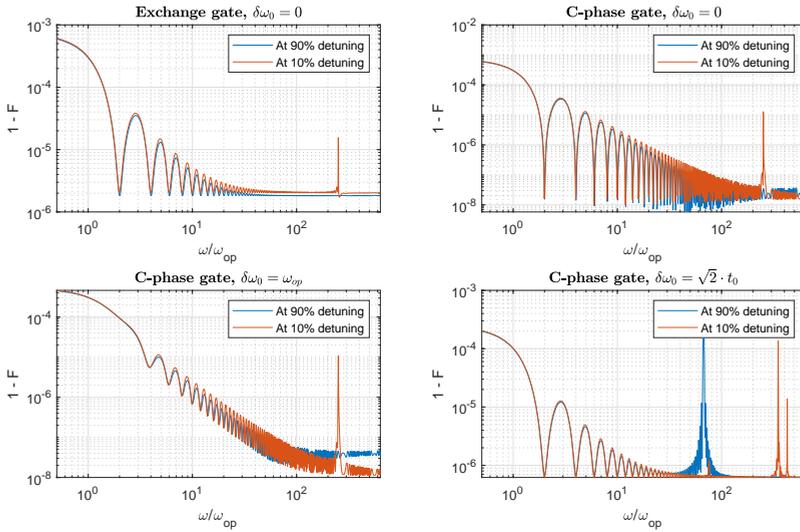


Figure A.6: A numerical simulation of the sensitivity of the qubit to tunnel rate noise of different frequencies for different operating regimes. For this simulation, a sinewave with RMS amplitude of 1% of the tunnel rate is applied on the tunnel rate control.

A.4. Qubit Read-out

For the fidelity, here it is assumed that the post-measurement qubit state is of interest as well. As a result, the fidelity can be formulated as:

$$F = P_{\text{charge}} \cdot [P_{\text{sense}} \cdot P_{\text{detect}} + (1 - P_{\text{sense}}) \cdot (1 - P_{\text{detect}})], \quad (\text{A.116})$$

where also a sensing error ($1 - P_{\text{sense}}$) together with a detection error ($1 - P_{\text{detect}}$) could lead to the correct outcome, assuming these probabilities are uncorrelated. However, when larger fidelities are targeted, and hence smaller errors can be tolerated, a good approximation to Eq. (A.116) is given by:

$$F \approx P_{\text{charge}} \cdot P_{\text{sense}} \cdot P_{\text{detect}}. \quad (\text{A.117})$$

The contribution P_{charge} can be found from the system Hamiltonian. The Hamiltonian of Eq. (A.83) is extended with the lowest-energy triplet states (spaced E_{ST} from the singlet energy level). For the Hamiltonian, only the charge states with one electron in each dot and two electrons in the right dot are considered, i.e. in the basis $\Psi = [|\uparrow, \uparrow\rangle, |\uparrow, \downarrow\rangle, |\downarrow, \uparrow\rangle, |\downarrow, \downarrow\rangle, |0, \uparrow\uparrow\rangle, |0, \uparrow\downarrow\rangle, |0, \downarrow\uparrow\rangle, |0, \downarrow\downarrow\rangle]$:

$$H = \begin{bmatrix} -\omega_0 & 0 & 0 & 0 & \sqrt{2} \cdot t_0 & 0 & 0 & 0 \\ 0 & \frac{\delta\omega_0}{2} & 0 & 0 & 0 & \sqrt{2} \cdot t_0 & 0 & 0 \\ 0 & 0 & -\frac{\delta\omega_0}{2} & 0 & 0 & 0 & \sqrt{2} \cdot t_0 & 0 \\ 0 & 0 & 0 & \omega_0 & 0 & 0 & 0 & \sqrt{2} \cdot t_0 \\ \sqrt{2} \cdot t_0 & 0 & 0 & 0 & U - \epsilon + E_{ST} - \omega_0 & 0 & 0 & 0 \\ 0 & \sqrt{2} \cdot t_0 & 0 & 0 & 0 & U - \epsilon + \frac{E_{ST}}{2} & \frac{E_{ST}}{2} & 0 \\ 0 & 0 & \sqrt{2} \cdot t_0 & 0 & 0 & \frac{E_{ST}}{2} & U - \epsilon + \frac{E_{ST}}{2} & 0 \\ 0 & 0 & 0 & \sqrt{2} \cdot t_0 & 0 & 0 & 0 & U - \epsilon + E_{ST} + \omega_0 \end{bmatrix}. \quad (\text{A.118})$$

To estimate P_{detect} , we assume Gaussian distributed noise and a simple measurement discrimination by comparison with a threshold $I_t = I_s/2$ halfway 0 and I_s (i.e. the signal for the two charge configurations to distinguish). For that case:

$$P_{\text{detect}} = \frac{1}{2} + \frac{1}{2} \cdot \operatorname{erf}\left(\frac{I_t}{\sigma_i \sqrt{2}}\right). \quad (\text{A.119})$$

For an integration time T_{read} , the filter transfer function is given by:

$$H_{\text{read}}(\omega) = e^{-\frac{1}{2}i\omega \cdot T_{\text{read}}} \cdot \frac{2}{\omega} \cdot \sin\left(\frac{1}{2} \cdot \omega \cdot T_{\text{read}}\right). \quad (\text{A.120})$$

This filter function has an equivalent noise bandwidth $ENBW = 1/(2 \cdot T_{\text{read}})$. The standard deviation of the noise σ_i can be approximated by integrating the noise power spectral density $S_i(f)$ (assumed flat) in the ENBW of the filter. The standard deviation of the noise follows as $\sigma_i^2 = S_i/2/T_{\text{read}}$, leading to:

$$P_{\text{detect}} = \frac{1}{2} + \frac{1}{2} \cdot \operatorname{erf}\left(\frac{I_s/2}{\sqrt{S_i/T_{\text{read}}}}\right). \quad (\text{A.121})$$

B

SPINE User Manual

S PINE – SPIN Emulator – is a tool for the co-simulation of classical electrical signals with spin-based quantum processors, and was originally introduced in [191, 192]. The tool was extensively used in deriving the impact of classical control electronics on the fidelity of a single-electron spin qubit [102]. However, the toolset can be directly extended to other qubit technologies.

B.1. Introduction

B.1.1. Simulation Platforms

The simulator is implemented in different platforms, with different features/limitations as discussed in the following. Note that only the Verilog-A implementation can be used with Cadence® for the co-simulation with electronic circuits.

MATLAB

As SPINE was originally written in MATLAB, this implementation contains all features present in the other platforms. Leveraging the power of MATLAB, it has better plotting functions (e.g., 3D plot) and additional solvers (`solver_expm`, `solver_taylor_sparse_approx`).

This implementation was tested on Windows 10 Pro 64-bit (1903) running MATLAB R2018a 64-bit.

C++

The C++ implementation of the simulator can in principle be used on any host OS, without the need for MATLAB to be installed on the system. Additionally, by defining `MKL`, the Intel® Math Kernel Library (with or without multi-threading [requires MPI]) is used to optimize most calculations. Moreover, there is explicit control over the used precision, with single precision floating point arithmetic the default, unless

`DOUBLE_PRECISION` is defined (advised). Finally, when running on Microsoft Windows, plotting functions are available, which can be enabled by defining `PLOT`. All matrices are stored in row-major order.

This implementation was tested on Windows 10 Pro 64-bit (1903) with MKL version 2019.5.281, compiled using Visual Studio 2017 (v141) and Windows SDK Version 10.0.17763.0.

Verilog-A

The Verilog-A implementation is meant for the simulation of the quantum processor in an electrical circuit simulator supporting Verilog-A models. Because of the limitations of the Verilog-A language, only the following modules are available:

- `spine_qubit1`: a port of the simulator for one single-electron spin qubit in the lab frame (`system_1_spin`) solved using the available analytical solution (`solver_analytical_xz`).
- `spine_qubit2`: a port of the simulator for two single-electron spin qubits, each with singlet state (`system_2_spin_2_singlet`), solved using a Taylor series expansion (`solver_taylor`).

This implementation was tested on a CentOS 5.11 server (linux kernel 2.6.18-410.el5) running Cadence® IC6.1.5.500.6 (32-bit) with Spectre version 7.2.0.307.isr10. More details of the Verilog-A implementation can be found in Appendix B.8.

B.2. SPINE

B.2.1. Hamiltonian Simulation

The evolution of the quantum processor's state, captured as a vector $|\psi\rangle$, can be described by the multiplication with a unitary matrix U : $|\psi(t)\rangle = U \cdot |\psi(0)\rangle$. In general, finding this operation U involves solving the time-dependent Schrödinger equation:

$$i\hbar \cdot \frac{\partial |\psi(t)\rangle}{\partial t} = H(t) \cdot |\psi(t)\rangle, \quad (\text{B.1})$$

where $H(t)$ is the Hamiltonian describing the system. In general, an approximate solution can be more easily found by simulation. Such a Hamiltonian simulation relies on the fact that the solution to a time-independent Hamiltonian H_n is trivially:

$$U_n = e^{-i/\hbar \cdot H_n \cdot t}. \quad (\text{B.2})$$

Then, by approximating the time-dependent Hamiltonian as many time-independent Hamiltonians H_n that are valid only for a short duration t , which are all applied subsequently, the overall operation can be found:

$$U_{\text{overall}} = \prod_{n=N}^1 U_n = U_N \cdot \dots \cdot U_2 \cdot U_1. \quad (\text{B.3})$$

This process is called trotterization. Note that the order of the matrix multiplication matters. Moreover, for simplicity $\hbar = 1$ throughout SPINE.

B.2.2. SPINE Simulator Core

The pseudo-code of the core of the SPINE simulator is shown below:

```
function simulate(inHamiltonian, outOperation, solver)
     $U \leftarrow I$ 
    while inHamiltonian( $H$ ,  $dt$ ) do
         $dU \leftarrow \text{solver}(H \cdot dt)$ 
         $U \leftarrow dU \cdot U$ 
        outOperation( $U$ )
    end while
end function
```

This function calculates the overall operation U at any time, given the user-provided Hamiltonian H and time step dt at that time (through the callback function *inHamiltonian*) and passes the resulting U , at that time, back to the user (through the callback function *outOperation*). The user needs to provide a *solver*(arg) that calculates $e^{-i \cdot \text{arg}}$.

If instead of a simulation of the full quantum operation, only a simulation of the quantum state is desired, the faster simulation as shown below is executed:

```
function simulate(inHamiltonian, outOperation, solver, state)
    while inHamiltonian( $H$ ,  $dt$ ) do
         $dU \leftarrow \text{solver}(H \cdot dt)$ 
         $state \leftarrow dU \cdot state$ 
        outOperation( $state$ )
    end while
end function
```

This function calculates the quantum state *state* at any time, given the user-provided Hamiltonian H and time step dt at that time (through the callback function *inHamiltonian*) and passes the resulting *state*, at that time, back to the user (through the callback function *outOperation*). The user needs to provide a *solver*(arg) that calculates $e^{-i \cdot \text{arg}}$.

The advantages of such a generic setup using callbacks are as follows:

- Allows for dynamic time step size.
- There is no requirement for a predetermined number of points or simulation time.
- There is no requirement for a predetermined signal; it can be generated based on previous simulation points if desired.
- There is freedom in what to do with the simulated operation/state at any time instance, e.g. plot or save every N^{th} point.

B.2.3. Functions

simulate

The function prototype of *simulate* in MATLAB and C++, respectively, is:

```
1 function simulate(dim, inHamiltonian, outOperation, solver, varargin)
```

```

1 void simulate(unsigned int dim,
2             bool (*inHamiltonian)(complex * H, realnum * dt),
3             void (*outOperation)(complex * U),
4             void (*solver)(unsigned int dim, complex * H, complex * dU));
5 void simulate(unsigned int dim,
6             bool (*inHamiltonian)(complex * H, realnum * dt),
7             void (*outOperation)(complex * U),
8             void (*solver)(unsigned int dim, complex * H, complex * dU),
9             complex * state);

```

The dimension of the Hamiltonian should be passed in `dim`. Additionally, handles to the functions `inHamiltonian`, `outOperation` and `solver` should be passed. For the C++ implementation, additional optimized functions are available when using real-valued Hamiltonians (`arg` is real). In both the MATLAB and C++ implementation, an additional optional argument `state` can be passed in which case only the quantum state evolution, instead of the full quantum operation, is simulated, assuming initial state `state`.

`inHamiltonian`

The function prototype of `inHamiltonian` in MATLAB and C++, respectively, is:

```

1 function [run, H, dt] = inHamiltonian()
1 bool inHamiltonian(complex * H, realnum * dt);

```

For the MATLAB implementation, `run`, specifies whether the simulation should continue or not, whereas this is returned as boolean variable in the C++ implementation. For the C++ implementation, the memory regions passed as arguments (`H` and `dt`) are filled by the function `inHamiltonian`. For the C++ implementation, an additional optimized function is available when using real-valued Hamiltonians.

Several Hamiltonians for different spin systems are provided in the MATLAB package `spine.systems` and the C++ namespace `spine::systems`. For more information, see Appendix B.3.

`outOperation`

The function prototype of `outOperation` in MATLAB and C++, respectively, is:

```

1 function outOperation(U)
1 void outOperation(complex * U);

```

Several helper functions for plotting the operation/state and calculating the fidelity of the operation are provided in MATLAB as `spine.plot*`(`)` and `spine.fidelity()`(`)`, respectively, and in C++ as `spine::Plot()` and `spine::fidelity()`, respectively. For more information, see Appendix B.4.

`solver`

The function prototype of `solver` in MATLAB and C++, respectively, is:

```

1 function dU = solver(dim, arg)
1 void solver(unsigned int dim, complex * arg, complex * dU);

```

The dimension of the argument `arg` should be passed in `dim`. For the C++ implementation, the memory region passed as argument (`dU`) is filled by the function `solver`. For the C++ implementation, additional optimized functions are available when using real-valued Hamiltonians (`arg` is real).

Several solvers are provided in the MATLAB package `spine.solvers` and the C++ namespace `spine::solvers`. For more information, see Appendix B.5.

B.3. Implemented Spin Systems

For the simulation of single-qubit gates on an isolated quantum dot, a good approximation is obtained when only considering the spin-up and spin-down states of the electron, i.e.: $|\psi\rangle = \alpha_0 \cdot |\uparrow\rangle + \alpha_1 \cdot |\downarrow\rangle$, where $|\uparrow\rangle$ and $|\downarrow\rangle$ can be considered the qubit states $|0\rangle$ and $|1\rangle$ respectively. The state vector only contains 2 complex numbers, and the matrices involved in the simulation are 2×2 .

The Hamiltonian of the spin qubit with the energy levels split by the Zeeman energy $\hbar\omega_0$ under excitation by a signal $x(t)$ in a perpendicular magnetic field is given by (setting $\hbar = 1$):

$$H(t) = -\omega_0 \frac{\sigma_z}{2} + x(t) \frac{\sigma_x}{2} \quad (\text{B.4})$$

where σ_x and σ_z are the X and Z Pauli matrices. In general, the signal $x(t)$ is a sinusoidal signal with frequency ω_0 and varying amplitude.

For the simulation of two qubits, it would seem sufficient to simply take the qubit state as $|\psi\rangle = \alpha_{00} \cdot |\uparrow\uparrow\rangle + \alpha_{01} \cdot |\uparrow\downarrow\rangle + \alpha_{10} \cdot |\downarrow\uparrow\rangle + \alpha_{11} \cdot |\downarrow\downarrow\rangle$ with a Hamiltonian $H = H_A \oplus H_B$ where H_i describes the physics of a single quantum dot. This is indeed sufficient for the simulation of single-qubit operations on multiple isolated qubits. However, for the simulation of two-qubit gates at least one more energy level should be included in the simulation that is responsible for the qubit interactions in the physical system: $|\psi\rangle = \alpha_{00} \cdot |\uparrow\uparrow\rangle + \alpha_{01} \cdot |\uparrow\downarrow\rangle + \alpha_{10} \cdot |\downarrow\uparrow\rangle + \alpha_{11} \cdot |\downarrow\downarrow\rangle + \alpha_{S0} \cdot |S0\rangle$ where $|S0\rangle$ describes the lowest energy state (a Singlet state) where both electrons have moved into one of the quantum dots. The corresponding system Hamiltonian, now also of size 5×5 , is given by [66]:

$$H = \begin{bmatrix} -\frac{\omega_{0,A} + \omega_{0,B}}{2} & 0 & 0 & 0 & 0 \\ 0 & -\frac{\omega_{0,A} - \omega_{0,B}}{2} & 0 & 0 & t_0 \\ 0 & 0 & \frac{\omega_{0,A} - \omega_{0,B}}{2} & 0 & -t_0 \\ 0 & 0 & 0 & \frac{\omega_{0,A} + \omega_{0,B}}{2} & 0 \\ 0 & t_0 & -t_0 & 0 & U - \epsilon \end{bmatrix} \quad (\text{B.5})$$

where for simplicity the driving term $x(t)$ has been left out. The tunnel coupling t_0 and detuning ϵ are generally time-varying signals. This Hamiltonian allows for the simulation of the controlled-Z gate at a detuning ($\epsilon \approx U$).

However, for the simulation of e.g. a controlled-Z gate at no detuning, or a SWAP gate, the state vector has to be expanded with the $|0S\rangle$ state, describing that both electrons can also go into the other quantum dot. For more accurate simulations,

as also required for the simulation of a qubit measurement by pauli-spin blockade, even higher energy levels have to be included.

Due to the fact that it is not sufficient to only simulate two energy levels for a single quantum bit, the simulation complexity grows very rapidly for larger qubit systems. Considering an isolated system (the number of electrons does not change), where only the lowest states with 2 electrons in a single dot are considered (the Singlet states), the length of the state vector (Dimension) grows as:

Qubits	Dimension
2	6
3	20
4	70
5	252
6	924

In MATLAB the following systems are currently available for simulation (in the package `spine.systems`):

- `spine.systems.system_spin`
 - `spine.systems.system_1_spin`
 - `spine.systems.system_1_spin_rwa`
 - `spine.systems.system_2_spin_1_singlet`
 - `spine.systems.system_2_spin_2_singlet`
 - `spine.systems.system_2_spin_1_singlet_triplet`
 - `spine.systems.system_2_spin_2_singlet_triplet`
 - `spine.systems.system_n_spin_n_singlet`
- `spine.systems.system_1_singlet_triplet`
- `spine.systems.system_dispersive_readout`

In C++ the following systems are currently available for simulation (in the namespace `spine::systems`):

- `spine::systems::system_spin`
 - `spine::systems::system_1_spin`
 - `spine::systems::system_1_spin_rwa`
 - `spine::systems::system_2_spin_1_singlet`
 - `spine::systems::system_2_spin_2_singlet`
 - `spine::systems::system_2_spin_1_singlet_triplet`
 - `spine::systems::system_2_spin_2_singlet_triplet`
 - `spine::systems::system_n_spin_n_singlet`
- `spine::systems::system_1_singlet_triplet`
- `spine::systems::system_dispersive_readout`

B.3.1. `system_spin`

Each of the `system*_spin` classes is derived from the `system_spin` base class that contains most of the functionality for simulating any number of *single-electron spin qubits* with any number of energy levels per dot considered (the 0 and 1 computational states, with optional singlet and triplet energy levels).

This base class contains the storage and general getters/setters for the Hamiltonian properties (Larmor frequency, Rabi frequency, charging energy, singlet-triplet energy splitting) and control variables (microwave signal, detuning signal, tunnel control signal) for every dot.

Additionally, it contains helper functions `initialize` and `measure` (`measureST`) to initialize a state vector to the ground state of the system and to measure the state of every qubit in the X,Y,Z basis (or the singlet and triplet occupancy probability of the dot), respectively. Optionally, a time t can be passed to the function `measure` to measure the state in the rotating frame at time t instead of the lab frame. These functions, however, rely on an implementation of the following functions in the derived class:

- `getIndex`: given the desired state of every quantum dot, passed as argument, returns the corresponding index in the state vector/Hamiltonian.
- `getIndexMeasurement`: given a dot and the state of interest, passed as argument, returns all indices in the state vector/Hamiltonian where that dot is in the desired state.
- `getDimension`: returns the dimension of the state vector/Hamiltonian.

On top of the `measure` functions, additional `plot` functions are provided to plot the X,Y,Z measurement probability for every dot in a Bloch sphere (MATLAB) or a simple 2D plot (C++, Windows only), see Appendix B.4.

`system_1_spin`

This class contains the 2×2 real-valued Hamiltonian of 1 single-electron spin-qubit in a single quantum dot, considering only the energy levels of the 0 and 1 states.

`system_1_spin_rwa`

This class contains the 2×2 complex Hamiltonian of 1 single-electron spin-qubit in a single quantum dot, in the rotating frame with rotating wave approximation, considering only the energy levels of the 0 and 1 states.

`system_2_spin_1_singlet`

This class contains the 5×5 real-valued Hamiltonian of 2 single-electron spin-qubits in a double quantum dot, considering only the energy levels of the 0 and 1 states. Additionally, the singlet state energy level of one of the dots is included to allow for basic 2-qubit operations. For simplicity, a single microwave drive line and a common Rabi frequency is assumed for the two dots.

`system_2_spin_2_singlet`

This class contains the 6×6 real-valued Hamiltonian of 2 single-electron spin-qubits in a double quantum dot, considering the energy levels of the 0 and 1 states, and the singlet state for both quantum dots. For simplicity, a single microwave drive line and a common Rabi frequency and charging energy are assumed for the two dots. Moreover, only the relative detuning of the two dots is considered.

`system_2_spin_1_singlet_triplet`

This class contains the 8×8 real-valued Hamiltonian of 2 single-electron spin-qubits in a double quantum dot, considering only the energy levels of the 0 and 1 states. Additionally, the energy levels of the singlet and 3 triplet states of one of the dots is included to allow for basic 2-qubit operations and simulation of Pauli-spin blockade readout. For simplicity, a single microwave drive line and a common Rabi frequency is assumed for the two dots.

`system_2_spin_2_singlet_triplet`

This class contains the 12×12 real-valued Hamiltonian of 2 single-electron spin-qubits in a double quantum dot, considering the energy levels of the 0 and 1 states, singlet state and 3 triplet states for both quantum dots. For simplicity, a single microwave drive line and a common Rabi frequency, charging energy and singlet-triplet energy splitting are assumed for the two dots. Moreover, only the relative detuning of the two dots is considered.

`system_n_spin_n_singlet`

This class contains the real-valued Hamiltonian of N single-electron spin-qubits in N quantum dots ($N \geq 2$). For each quantum dot the energy levels of the 0 and 1 states, and the singlet state is simulated.

B.3.2. `system_1_singlet_triplet`

This class contains the 2×2 real-valued Hamiltonian of 1 singlet-triplet qubit in a double quantum dot, considering only the energy levels of the 0 and 1 states.

Similar as for the `system_spin` class, this class contains the storage and general getters/setters for the Hamiltonian properties (magnetic field gradient) and control variables (exchange interaction).

Additionally, it contains helper functions `initialize` and `measure` to initialize a state vector to the ground state of the system and to measure the state of the qubit in the X,Y,Z basis, respectively. On top of the `measure` function, additional `plot` functions are provided to plot the X,Y,Z measurement probability for the dot in a Bloch sphere (MATLAB) or a simple 2D plot (C++, Windows only).

B.3.3. `system_dispersive_readout`

Finally, the master equation governing dispersive readout has been rewritten in the form of a complex 2×2 non-Hermitian 'Hamiltonian' (not a true Hamiltonian as the resulting operation is not unitary). The implemented master equation is given by

[237]:

$$\frac{dP_1(t)}{dt} + \Gamma_0 P_1(t) = \Gamma_+(t) \quad (\text{B.6})$$

where P_1 is the probability of the electron being in the dot and Γ_+ is the tunnel rate:

$$\Gamma_+(t) = \frac{\Gamma_0}{1 + e^{\Delta E(t)/k_B/T}} \quad (\text{B.7})$$

where k_B is the Boltzmann constant, T the temperature, $\Delta E(t)$ the time-dependent energy difference and Γ_0 is the constant tunnel rate away from the degeneracy.

Similar as for the `system_spin` class, this class contains the storage and general getters/setters for the Hamiltonian properties (electron temperature, tunnel rate and lever arm) and control variables (gate voltage).

Additionally, it contains helper functions `initialize` and `measure` to set the initial probability and to measure the probability respectively. On top of the `measure` function, additional `plot` functions are provided to plot the measurement probability of the electron being in the dot in a simple 2D plot (MATLAB, Windows only for C++).

B.4. Helper Functions

For calculating the fidelity of a 2×2 unitary operation, the following functions are available in MATLAB and C++ respectively:

```
1 function F = fidelity(dim, U, varargin)
1 realnum fidelity(unsigned int dim, complex * U, complex * Uideal);
2 realnum fidelity(unsigned int dim, complex * U, realnum theta, realnum phi)
;
```

where the dimension of the unitary (i.e. 2) needs to be passed in the argument `dim`, and the unitary operation in `U`. Next, either a single argument must be passed containing the ideal unitary operation `Uideal`, or two arguments follow: the rotation angle `theta` and rotation axis `phi` of the ideal rotation.

B.4.1. Plotting in MATLAB

For plotting in MATLAB, the following function draws a simple 3D Bloch sphere:

```
1 function plotBlochSphere()
```

This function is for instance used in the base class `system_spin` (see Appendix B.3.1) in the `plot` function, which takes the following arguments:

```
1 function plot(obj, state_or_U, t, varargin)
```

where `state_or_U` is either the quantum state or the unitary operation (in which case an initial ground state initialization is assumed to determine the state to visualize). The argument `t` is used to determine the measurement probability in the lab frame, which is shown by default. By passing a first optional argument `plot_lab_style`, the measurement probability in the lab frame is also shown as either an arrow (`plot_lab_style = 1`) or a full trace (`plot_lab_style = 2`). By passing a second

optional argument `plot_st`, additional plots of the singlet-triplet state occupancy of the dot are generated. When `plot_st = 1`, only the singlet state occupancy is plotted; when `plot_st = 2`, the expected number of electrons in the dot is plotted.

Finally, the following function plots the eigenenergy diagram of a double dot spin system (i.e. the energy levels of the various states versus detuning):

```
1 function plotEigenEnergies(system, varargin)
```

where the system is passed in the argument `system`. By default the detuning is swept from -1.25 to 1.25 times the charging energy with 1001 points. By passing a second optional argument `points` this number can be changed.

B.4.2. Plotting in C++

Note: these plotting functions are only available under Windows and require `PLOT` to be defined!

The class `Plot` is used for plotting and has the following constructors:

```
1 Plot(LPCWSTR name, unsigned int points = 100);
2 Plot(LPCWSTR name, unsigned int points, double * xdata, double * ydata,
   BYTE r, BYTE g, BYTE b, double xmin, double xmax, double ymin, double
   ymax);
3 Plot(LPCWSTR name, unsigned int points, double * ydata, BYTE r, BYTE g,
   BYTE b, double xmin, double xmax, double ymin, double ymax);
```

where `name` is the title of the plotting window, `points` is the number of points (expected) to be plotted. Optionally, a curve can directly be plotted by passing the `ydata` and optionally the `xdata` along with the desired plotting color (`r`, `g`, `b`) and plot window limits (`xmin`, `xmax`, `ymin`, `ymax`).

An additional curve can be added through the method:

```
1 void add(BYTE r = 0, BYTE g = 0, BYTE b = 0, double xmin = 0, double xmax =
   0, double ymin = 0, double ymax = 1);
```

which again takes the desired plotting color and plot window limits. Next, a handle to the curve can be obtained through the method:

```
1 PlotSeries * get(unsigned int n);
```

where `n` is the curve number (starting from 0 for the first curve that was added). For each of these curves, points can later be added, and the window limits can be changed, through the following methods of `PlotSeries`:

```
1 void add(double x, double y);
2 void xlim(double xmin, double xmax);
3 void ylim(double ymin, double ymax);
```

When using the latter 3 functions, a redraw must be forced by calling the `redraw()` method of the `Plot` class.

These plotting functions are for instance used in the base class `system_spin` (see Appendix B.3.1) in the `plot*` methods:

```
1 void plotSetup(unsigned int points = 100, double xmin = 0, double xmax = 0)
   ;
2 void plotAddU(complex * U, realnum t = 0, bool plot_lab = false);
3 void plotAdd(complex * state, realnum t = 0, bool plot_lab = false);
4 void plot();
```

From these functions, `plotSetup` can optionally be called before starting a simulation to setup up the number of points that will be plotted and the x-scale (alternatively, the x-scale will be adjusted with every point added), which can significantly speed-up the plotting. The functions `plotAddU` and `plotAdd` are used to add a plot point for the simulated unitary operation (again assuming an initial ground state initialization) or quantum state, equivalent to the MATLAB function `plot(obj, state_or_U, t, varargin)` (see Appendix B.4.1). The argument `t` is used to determine the measurement probability in the rotating frame, which is shown by default. By passing setting the optional argument `plot_lab` to `true`, the measurement probability in the lab frame is also shown as a full trace. Unlike for the MATLAB implementation, an additional call to the `plot` method is required to update the actual graphics with the added points when desired.

Note: After all simulations/calculations are finished, the endless loop `Plot::run()` must be called, which handles the Windows GUI message loop that takes care of the GUI actions, such as plot window resizing/closing.

B.5. Implemented Solvers

In MATLAB the following solvers are currently available (in the package `spine.solvers`):

- `spine.solvers.solver_analytical_xz`
- `spine.solvers.solver_expm`
- `spine.solvers.solver_diagonalization`
- `spine.solvers.solver_taylor`
- `spine.solvers.solver_taylor_sparse_approx`

In C++ the following solvers are currently available (in the namespace `spine::solvers`):

- `spine::solvers::solver_analytical_xz`
- `spine::solvers::solver_diagonalization`
- `spine::solvers::solver_taylor`

B.5.1. `solver_analytical_xz`

This solver uses the analytical solution available in the case of a real-valued 2×2 Hamiltonian of the following form:

$$H = a \cdot \sigma_x + b \cdot \sigma_z \quad (\text{B.8})$$

where σ_x and σ_z are the Pauli-X and Pauli-Z matrices, respectively, and a and b are real numbers.

B.5.2. `solver_expn`

This solver uses the `expm` function available in MATLAB to evaluate the matrix exponential.

B.5.3. `solver_diagonalization`

This solver calculates the matrix exponential using the eigenvalue decomposition. The Hamiltonian is assumed to be Hermitian (symmetric in case of a real-valued Hamiltonian).

B.5.4. `solver_taylor`

This solver approximates the matrix exponential using the Taylor series expansion upto an order that is set by the MATLAB global variable `solver_taylor_accuracy` or the C++ local variable `accuracy`. For increased accuracy, it uses the scaling-and-squaring method with scaling factor 2^N where N is set by the MATLAB global variable `solver_taylor_scaling` or the C++ local variable `scaling`. Default values are provided.

B.5.5. `solver_taylor_sparse_approx`

Same as the `solver_taylor`, but uses sparse matrices (available in MATLAB only). To ensure efficient use of sparse matrices, during the squaring step, all elements with magnitude below a certain level (set by the MATLAB global variable `solver_taylor_tolerance`) are discarded. While this reduces the simulation accuracy, it allows for significantly faster simulation of larger spin qubit systems.

B.6. Examples

Currently, equivalent examples are provided for the MATLAB and C++ implementation demonstrating each of the currently implemented spin systems:

- `example_1_spin`
- `example_1_spin_rwa`
- `example_2_spin_1_singlet`
- `example_2_spin_2_singlet`
- `example_2_spin_1_singlet_triplet`
- `example_2_spin_2_singlet_triplet`
- `example_n_spin_n_singlet`
- `example_1_singlet_triplet`
- `example_dispersive_readout`

B.6.1. `example_1_spin`

This example demonstrates a π -rotation on a single-electron spin qubit using a rectangular envelope. Every $N_{\text{plot}}^{\text{th}}$ unitary operation is plotted in the Bloch sphere, in both the lab frame and the rotating frame. Finally, the fidelity of the operation is calculated from the ideal rotation angle/axis and printed.

B.6.2. `example_1_spin_rwa`

This example demonstrates a $\pi/2$ -rotation along the X-axis followed by a $\pi/2$ -rotation along the Y-axis on a single-electron spin qubit in the rotating frame using a rectangular envelope. Every $N_{\text{plot}}^{\text{th}}$ unitary operation is plotted in the Bloch sphere, in the rotating frame (with additional arrow). Finally, the fidelity of the operation is calculated from the ideal unitary operation and printed.

B.6.3. `example_2_spin_1_singlet`

In this example, a controlled-NOT (CNOT), based on a 2-qubit CZ-gate at detuning, is demonstrated, where for the first CNOT the control qubit is in the ground state, and for the second CNOT the control qubit is in the excited state. Gaussian envelopes are used for the microwave control signal, which is frequency multiplexed over the 2 qubits. Every $N_{\text{plot}}^{\text{th}}$ unitary operation is plotted in the Bloch sphere, in both the lab frame (shown as an arrow) and the rotating frame.

B.6.4. `example_2_spin_2_singlet`

In this example, a SWAP gate is demonstrated at detuning for a double dot. Additionally, a simulation of the quantum state, instead of the unitary operation, is used, as the initial state is chosen differently for both qubits. Every $N_{\text{plot}}^{\text{th}}$ state is plotted in the Bloch sphere, in the rotating frame.

B.6.5. `example_2_spin_1_singlet_triplet` and `example_2_spin_2_singlet_triplet`

These examples demonstrate Pauli-spin blockade readout of a double quantum dot. Initially both qubits are in the ground state, and no charge is transferred when adiabatically detuning the quantum dots, whereas later charge is transferred as one qubit is rotated into the excited state. Every $N_{\text{plot}}^{\text{th}}$ unitary operation is plotted in the Bloch sphere, in the rotating frame.

B.6.6. `example_n_spin_n_singlet`

In this example 4 qubits are simulated and every $N_{\text{plot}}^{\text{th}}$ state is plotted, with the following instructions executed in sequence:

1. Rotate qubit 1 from 0 to 1 using a microwave signal
2. SWAP between qubit 1 and qubit 2 using the tunnel coupling
3. Rotate qubit 3 and qubit 4 in plane using a microwave signal
4. CZ-gate between qubit 1/qubit 3 and qubit 2/qubit 4

5. Rotate qubit 3 and qubit 4 out plane using a microwave signal, out of phase
6. Pauli-spin blockade measurement of qubit 3 with respect to qubit 4
7. Rotate qubit 1 using a microwave signal

B.6.7. [example_1_singlet_triplet](#)

In this example, a π -rotation along the Y-axis is demonstrated on a singlet-triplet qubit using a composite pulse comprising 2 rectangular pulses with different amplitude and duration. Additionally, this example demonstrates the use of a variable timestep, as both rectangular pulses are simulated with a different timestep. Every $N_{\text{plot}}^{\text{th}}$ unitary operation is plotted in the Bloch sphere, in the lab frame. Finally, the fidelity is calculated and printed given the ideal rotation angle and axis.

B.6.8. [example_dispersive_readout](#)

In this example, dispersive gate readout is demonstrated. A sinusoidal gate voltage is applied and the resulting dot occupancy is simulated. Finally, from this occupancy, the expected gate current is calculated and plotted. The resulting waveforms depend on the drive frequency relative to the tunnel coupling (capacitive vs. resistive regimes).

B.7. Templates

Below are the templates for a typical SPINE program written in MATLAB and C++. In these templates, <SYSTEM>, <SET>, <VALUE>, <POINTS> and <SOLVER> should be adjusted to the needs.

B.7.1. MATLAB

```

1 % Global variables
2 global nsim;
3 global Nsim;
4 global system;
5
6 % Create the system
7 system = spine.systems.<SYSTEM>();
8 system.<SET>(<VALUE>);
9
10 % Simulate the system
11 nsim = 0;
12 Nsim = <POINTS>;
13 spine.simulate(system.getDimension(), @inHamiltonian, @outOperation, @spine
    .solvers.<SOLVER>);
14
15 function [run, H, timestep] = inHamiltonian()
16 global nsim;
17 global Nsim;
18 global system;
19
20 if (nsim < Nsim)
21

```

```

22 % Set the signal at this time instance
23 system.<SET>(<VALUE>);
24
25 % Update the Hamiltonian accordingly
26 H = system.updateHamiltonian();
27
28 % Provide the current timestep, and continue the simulation
29 timestep = <Timestep>;
30 run = true;
31
32 else
33 run = false;
34 H = [];
35 timestep = [];
36 end
37 end
38
39 function outOperation(U)
40 global nsim;
41 global Nsim;
42
43 % Plot, print, store
44 ...
45
46 % Continue the simulation
47 nsim = nsim + 1;
48 if (nsim == Nsim)
49
50 % Last point
51 ...
52
53 end
54
55 end

```

B.7.2. C++

```

1 // Includes
2 #include "spine/simulate.h"
3 #include "spine/math.h"
4 #include "spine/systems/<SYSTEM>.h"
5 #include "spine/solvers/<SOLVER>.h"
6
7 // Namespaces
8 using namespace spine::math;
9
10 // Function prototypes
11 bool inHamiltonian(realnum * H, realnum * dt);
12 void outOperation(complex * U);
13
14 // Private global variables/constants
15 static spine::systems::<SYSTEM> spin_system = spine::systems::<SYSTEM>();
16 static const realnum dt = (realnum) 10e-12;
17 static int Nsim, nsim;
18
19 int main(void)
20 {

```

```

21 // Create the system
22 spin_system.<SET>(<VALUE>);
23
24 // Simulate the system
25 nsim = 0;
26 Nsim = <POINTS>;
27 spine::simulate(spin_system.getDimension(), inHamiltonian, outOperation,
    spine::solvers::<SOLVER>);
28 }
29
30 bool inHamiltonian(realnum * H, realnum * dt)
31 {
32 if (nsim < Nsim)
33 {
34 // Set the signal at this time instance
35 spin_system.<SET>(<VALUE>);
36
37 // Update the Hamiltonian accordingly
38 spin_system.updateHamiltonian(H);
39
40 // Provide the current timestep, and continue the simulation
41 *dt = ::dt;
42 return true;
43 }
44 return false;
45 }
46
47 void outOperation(complex * U)
48 {
49 // Plot, print, store
50 ...
51
52 // Continue the simulation
53 nsim++;
54 if (nsim == Nsim)
55 {
56 // Last point
57 ...
58 }
59 }

```

B.8. Verilog-A

As advanced electrical circuit simulators use quasi-static time-domain solvers, they provide a favorable environment for the inclusion of a time-discrete Hamiltonian simulation. Using Cadence® as a framework, the quantum physical system is included in the electrical simulation as a module that takes as input the control signals for the quantum system and outputs the quantum operation (Figs. B.1 and B.2).

In the Verilog-A implementation, only modules emulating either one single-electron spin qubit (Fig. B.1) or a system of two coupled single-electron spin qubits (Fig. B.2) are currently available:

- `spine_qubit1`: a port of the simulator for one single-electron spin qubit in the lab frame (`system_1_spin`) solved using the available analytical solution (`solver_analytical_xz`).

- `spine_qubit2`: a port of the simulator for two single-electron spin qubits, each with singlet state (`system_2_spin_2_singlet`), solved using a Taylor series expansion (`solver_taylor`).

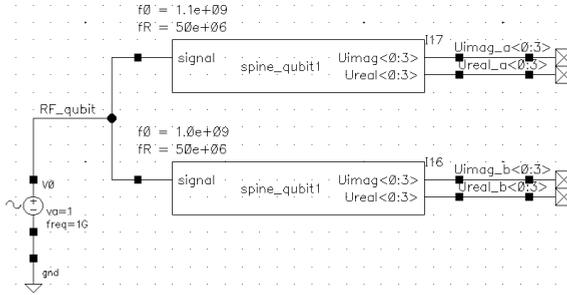


Figure B.1: Two modules, each emulating one single-electron spin qubit, have been instantiated as a verilog-A module in the electrical circuit simulator; the two qubits are uncoupled and cannot be entangled.

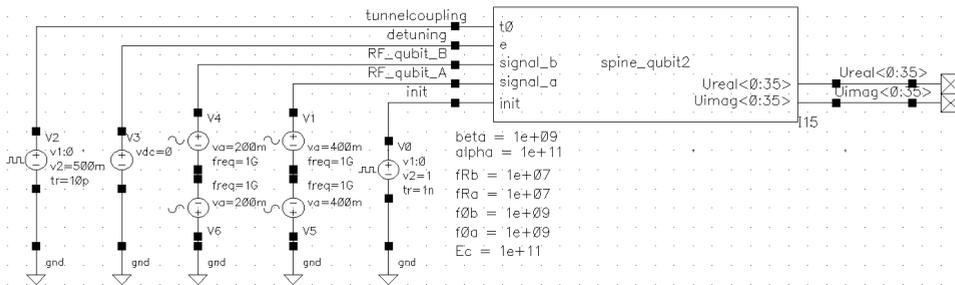


Figure B.2: A system of two coupled single-electron spin qubits is included as a verilog-A module in the electrical circuit simulator.

The inputs to the Verilog-A blocks are the electrical signals applied to the quantum processor (Figs. B.1 and B.2; `signal` is the RF-signal, `e` for detuning the quantum dots and `t0` to control the tunnel coupling) and `init` to reset the operation to the identity matrix. The resulting complex operation U , which can be used to calculate the operation fidelity, is available at the output with separated real (`Ureal<>`) and imaginary (`Uimag<>`) parts in row-major order. Parameters of the physical system are set as a module parameter when instantiating the module in the circuit schematic (Figs. B.1 and B.2; E_C is the charging energy of the quantum dot, f_0 the spin precession frequency, f_R the rotation frequency at 1-V RF-signal¹, β the tunnel coupling at 1 V and α the detuning energy at 1 V).

While the time step control of a transient simulation is managed by the circuit

¹Note that typically much smaller signal amplitudes are used, and hence f_R should be set to a high value to get the desired Rabi frequency at this smaller signal amplitude.

simulator, which relaxes the time step when tolerable, a maximum time step is set by the Verilog-A module to ensure accurate simulation of the quantum physics.

C

Horse Ridge measurement setup

In this appendix, the details of the measurement setup used to test Horse Ridge (Chapters 6 and 7) are disclosed.

C.1. Horse Ridge Motherboard

Horse Ridge is flipchip packaged in a 19×19 mm BGA package with 324 balls, most of which are for grounding. On the package, 2 nF local decoupling capacitors are placed for the various supplies. A motherboard hosting the packaged chip has been designed to operate at cryogenic temperatures implementing additional supply decoupling (330 μ F, 100 μ F and 22 μ F tantalum capacitors, and 1 μ F and 47 nF NP0-type ceramic capacitors) and cryogenic LDOs (Fig. C.1). The LDO design is based on the design presented in [238], with added frequency compensation to achieve stability at cryogenic temperatures. When integrated in the dilution refrigerator, copper lines ($\sim 3\Omega$) are used to supply the input current, and more resistive phosphor bronze lines are used to supply the reference and to measure back the LDO output voltage (Fig. C.1).

Horse Ridge has various on-chip temperature sensing diodes available, 8 of which have the cathode directly pad-accessible (anode grounded). As long resistive lines connect the motherboard to the rest of the setup (Appendix C.2), a force/sense connection is implemented on the PCB, and only half of the diodes are connected due to the limited wiring available in the setup (Fig. C.1).

Two on-chip Multiplexer (MUX)es allow the measurement of various on-chip voltages/currents, and the application of an external reference current bypassing the on-chip reference generator. As the MUX output can be brought into high-impedance mode, the outputs from the different transmitters can be shared. For

the external reference current, local 100 nF NP0-type ceramic decoupling capacitors are placed, and individual connections are maintained.

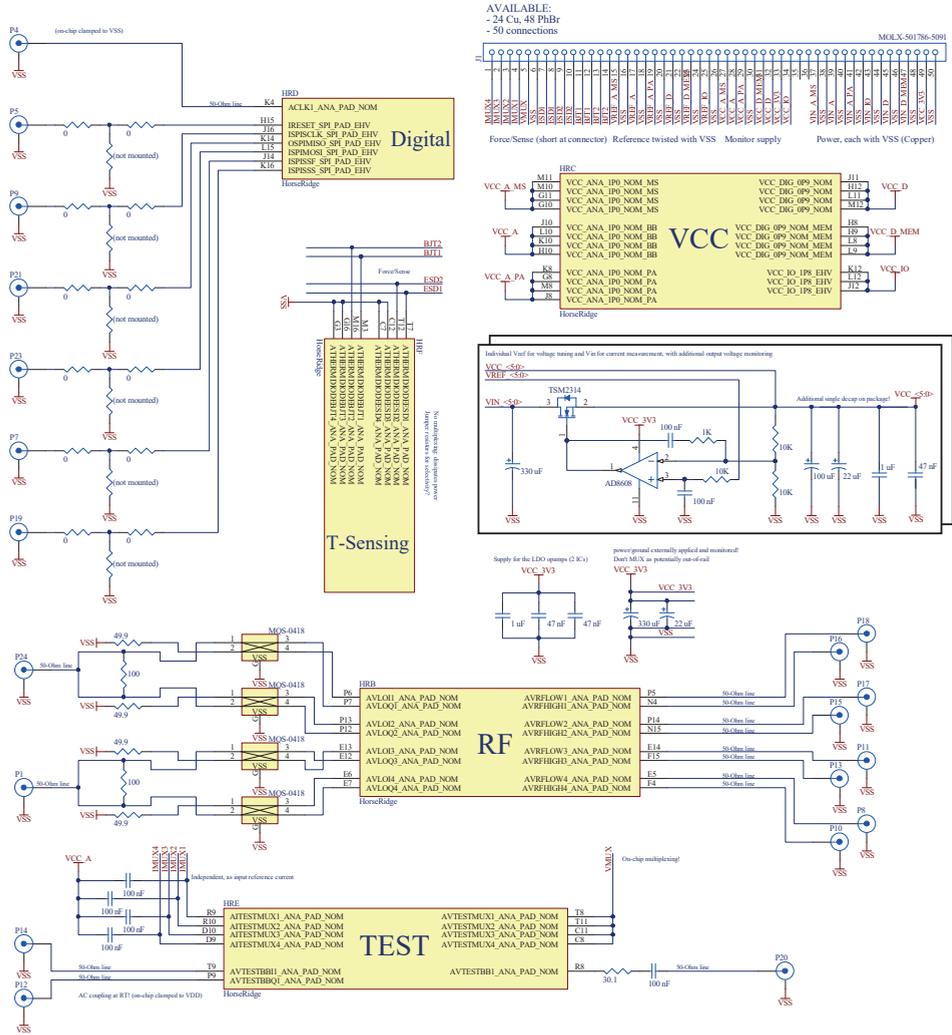


Figure C.1: The schematic of the Horse Ridge motherboard.

Transmission lines are used to feed a signal into the baseband input (bypassing the on-chip DAC) and to monitor the baseband output signal (after the on-chip VGA). The baseband input has on-chip 50-Ω termination, but the baseband output requires off-chip series termination (Fig. C.1). Similarly, 50-Ω transmission lines are used to interface to the external clock (2.5 GHz) and high-speed SPI interface (upto 200 MHz).

All RF outputs are routed to individual SMP connectors. However, the LO signals

are shared per pair of transmitters, through a Wilkinson power divider. Commercial quadrature hybrid couplers (Marki Microwave MQS-0418), tested at cryogenic temperatures, are used to obtain the required quadrature LO signals.

An overview of the DC FFC-type connector, and the intended allocation of dilution refrigerator wires, is given in Fig. C.1. All high-frequency lines (baseband input/output, digital interface, clock, LO inputs, and RF outputs) are accessible through SMP connectors and are to be connected to regular or superconducting coax lines in the dilution refrigerator.

This board is used for room-temperature testing, dipstick testing and testing in the dilution refrigerator (see Appendix C.2). The connectors, FFC and SMP, are chosen for compatibility with standard dilution refrigerator setups. The PCB size is restricted to a 10 cm diameter circle in the fridge. The 'side flaps' of the PCB can be removed to obtain a 6 cm wide board for use in the dipstick, as restricted by the width of the Helium barrel neck.

The PCB design is shown in Fig. C.2. Six layers are used, dedicated (from top to bottom) to high-frequency RF signals, ground, analog supplies, ground, digital supplies, and finally the LDO circuitry and digital interface routing at the bottom of the PCB. Standard FR-4 material is used, except for the top layer which is using RT/duroid 6002 microwave substrate, which was proven to work well at cryogenic temperatures [239].

To maximize thermal contact to the enclosure (see Chapter 6), as much copper grounding as possible was exposed by removing most of the solder-mask at the top and bottom of the PCB (Fig. C.2). Furthermore, heat transfer was improved by placing as many ground vias as possible, and placing most of the decoupling capacitors further away from the chip.

C.2. Generic Measurement Setup

The Horse Ridge motherboard can be used in 3 different measurement environments: room-temperature testing, dipstick testing and testing in the dilution refrigerator, through the use of different adapter boards and cabling (Fig. C.3). The dilution refrigerator setup uses the Matrix Module as in [240].

As the Horse Ridge motherboard requires several supplies and provides multiple measurement points, a breakout board containing LDOs and multiplexers has been designed to reduce the number of bench top supplies and SMUs. The breakout board is controlled by the FPGA, the Opal Kelly XEM7320, which also handles the communication with Horse Ridge over its digital interface. The schematic of the breakout board is shown in Fig. C.4. All connections to the breakout board can be made in 3 ways, as for the 3 different setups: through an FFC flatcable connector (room temperature setup), Samtec TFM connectors (dipstick setup), or MCX connectors (dilution refrigerator setup). For simplicity, the latter is not shown in the schematic.

The LDOs on this board provide the input voltages for the LDOs on the Horse Ridge motherboard, and their 3.3V supply voltage, along with a supply for the isolator, I/O extender and multiplexers used for multiplexing the SMUs. The LDO output voltage is tunable through a trim potentiometer. Mechanical switches at the

C

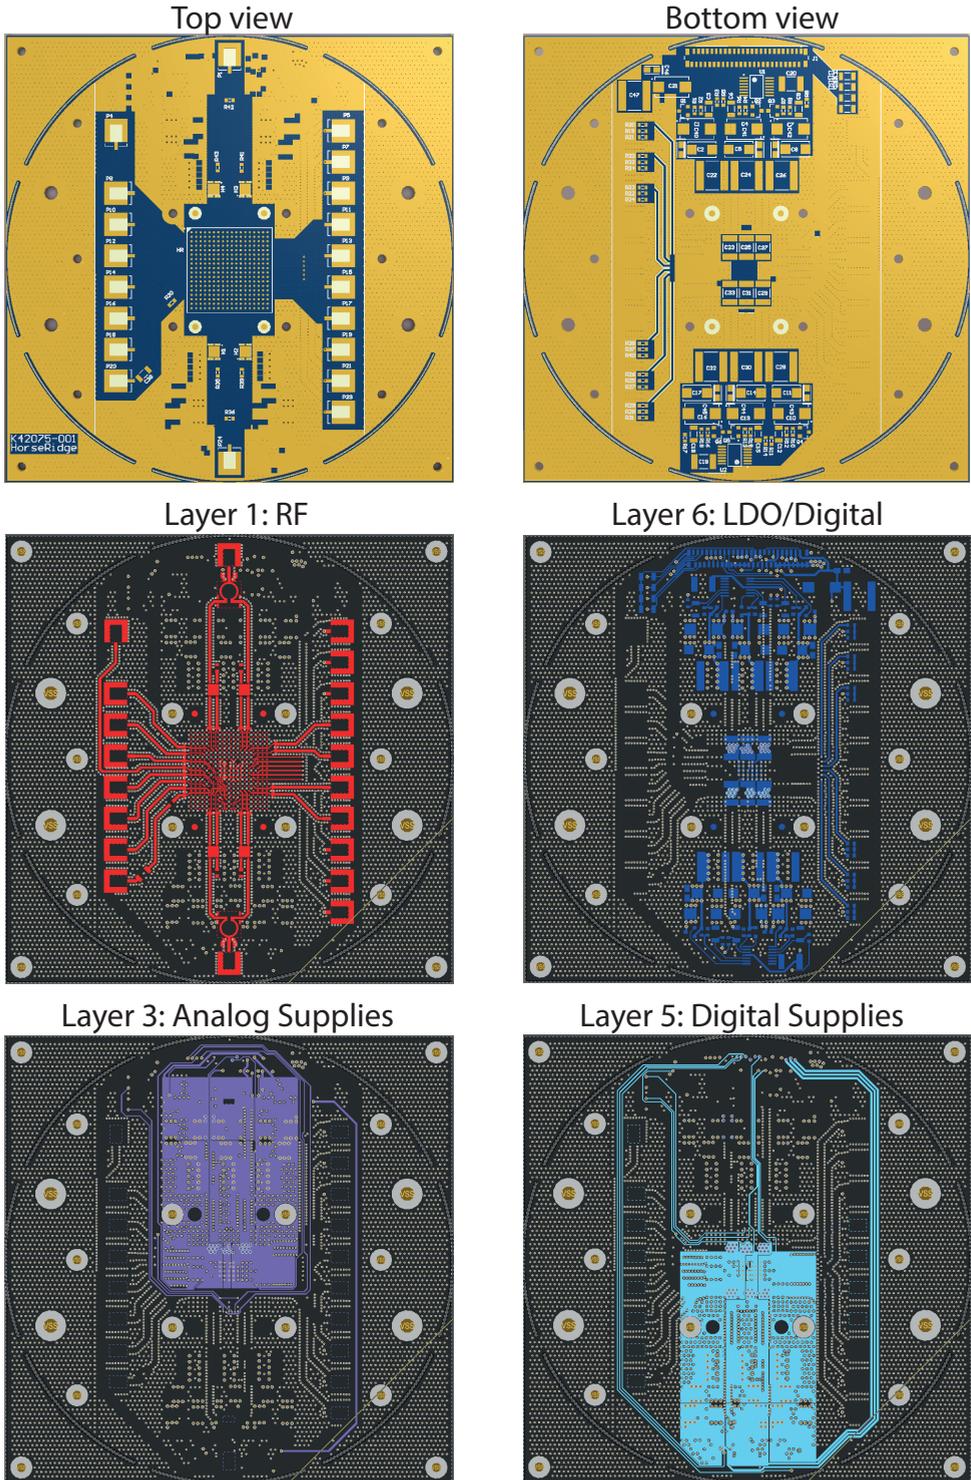


Figure C.2: The PCB design of the Horse Ridge motherboard showing the relevant layers (ground fill hidden).

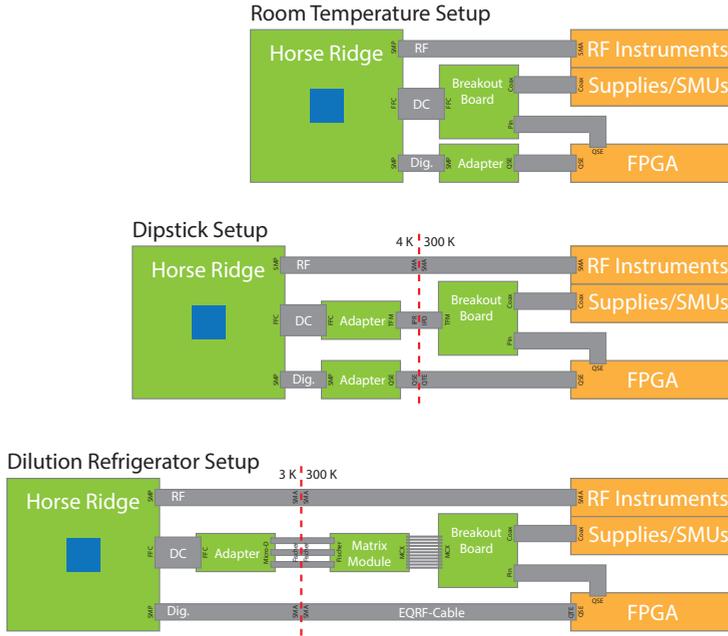


Figure C.3: The Horse Ridge motherboard in the 3 different measurement environments.

LDO output select either this voltage, or an externally applied voltage (through SMU CH A), for the purpose of measuring the supply current.

Besides that, SMU CH A can be used to measure the output voltages, and local VSS, of the LDOs on the Horse Ridge motherboard. SMU CH B is multiplexed between the different pad accessible temperature sensing diodes, and finally SMU CH C can be used to measure the on-chip voltages/currents or supply an external reference current. The reference voltages for the LDOs on the Horse Ridge motherboard are generated by a DAC, that generates voltages with respect to the local VSS monitored from the Horse Ridge motherboard, as to compensate for IR-drops. The DAC can be programmed from the FPGA to perform voltage sweeps.

Importantly, the unprogrammed DAC generates by default an 1.25V output voltage which can result in a supply voltage that is too high for the Horse Ridge chip. Hence, the DAC needs to be programmed before an input voltage is supplied to the Horse Ridge motherboard's LDOs.

A picture of the breakout board is shown in Fig. C.5. The switches, and some of the coax connectors are placed on the bottom side of the PCB.

C.3. Dipstick Design

A dipstick has been designed for the measurement of Horse Ridge that contains 40 DC lines, 10 high-speed lines for the digital interface, and 8 RF lines. High-density connectors have been used on the top box (Aluminium Hammond 1590V) for the

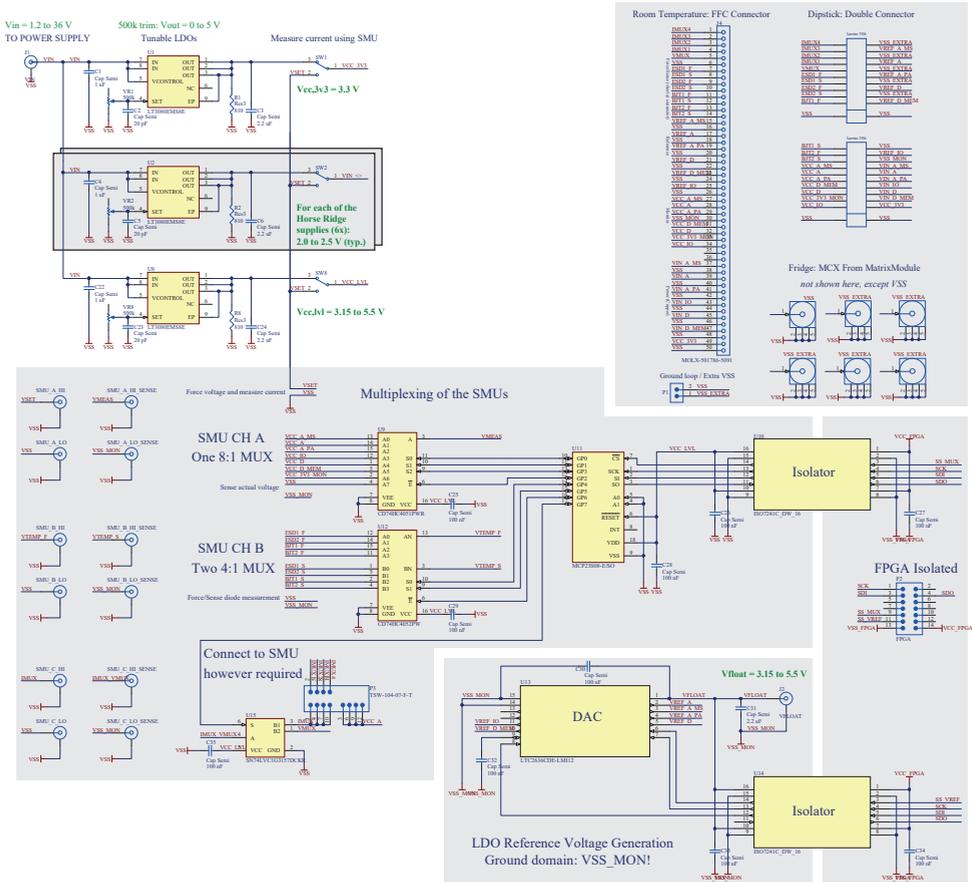


Figure C.4: The schematic of the breakout board.

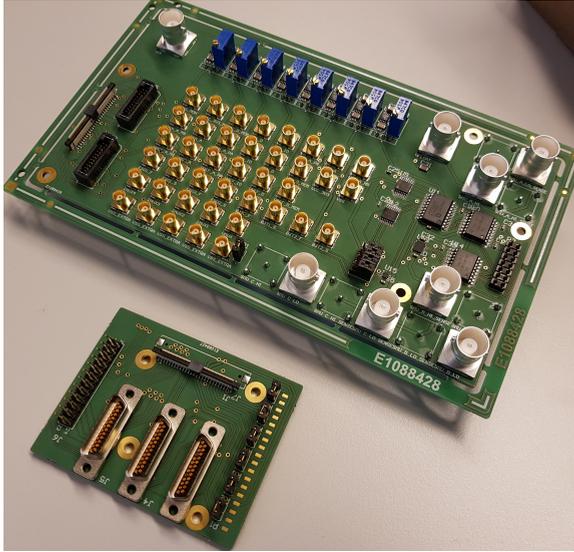


Figure C.5: Picture of the breakout board and adapter board used in the dilution refrigerator setup (Appendix C.4).

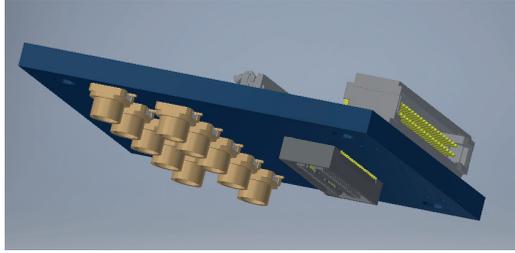
DC lines (Samtec IPR connectors) and high-speed lines (Samtec SQE connector) for quick connect/disconnect. The appropriate adapter boards for these interfaces are efficiently stacked in the dipstick as illustrated in Fig. C.6. The RF lines, on the other hand, use RG178 coax with panel-mount SMA connectors on the top box. Additionally, there are 4 airtight BNC connectors that connect to a Lakeshore temperature sensing diode that can be mounted on the Horse Ridge motherboard.

The mechanical mounting is designed for a Helium barrel with a 7 cm wide neck, and requires the Horse Ridge motherboard PCB to be altered accordingly (see Appendix C.1). Inside the stainless steel 304L pipe, aluminum semi-circles have been placed regularly with different orientation against thermo-acoustic effects in the pipe, and a tap has been placed at the top of the box. Figure C.6 shows the design of the dipstick, and the mounting of the various adapter boards/cables. Aluminum frames are used to hold the SMP cables in place. Additionally, a 7 cm diameter cylinder has been drawn as reference for the Helium barrel neck.

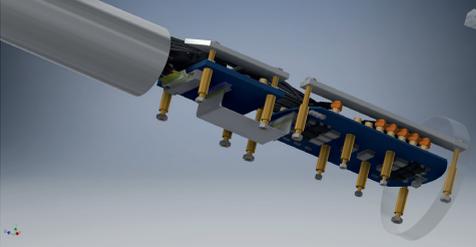
C.4. Fridge Integration

As discussed in Chapter 6, a gold-plated annealed copper enclosure has been designed for the Horse Ridge motherboard that mounts into a fixed holder inside the dilution refrigerator. The fixed holder connects the SMPs of the Horse Ridge motherboard, i.e. the RF lines and high-speed digital interface, to the appropriate ports in the fridge. The various RF outputs are connected through a Radiall SP6T relay that allows for the selection of a certain output, as indicated in Fig. C.7. Moreover, a Radiall DPDT relay controls whether the RF output is applied to the qubit, or can

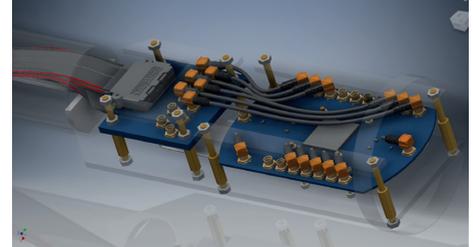
The stacked adapter boards



Bottom side (frame removed)



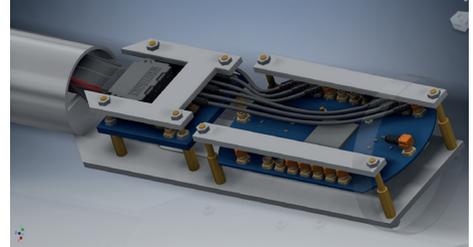
Top side (frame removed)



Bottom side (with frame)



Top side (with frame)



Top box



Sample side (all cabling)



Top box (see through)



Sample side (see through)

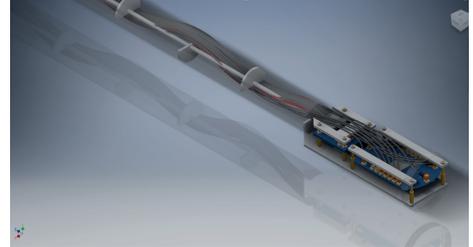


Figure C.6: The design of the dipstick.

be monitored at room temperature (in which case the qubit can be controlled from room temperature as well). Bandpass filters (Marki FB1310 and FB1840) are placed at the TX2 outputs to improve the qubit readout visibility when using the standard Elzerman readout. These filters were tested at cryogenic temperatures and show minimal change in transfer function over temperature.

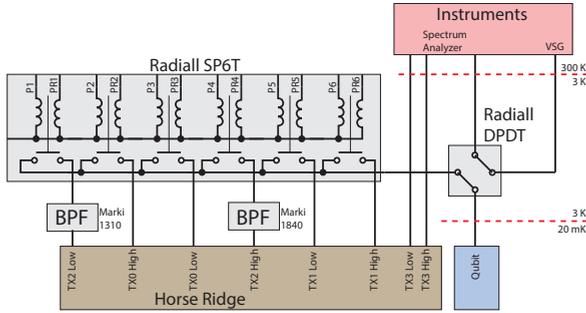


Figure C.7: The routing of the Horse Ridge RF outputs through the fridge.

A PCB was designed for the control of the Radiall relays (Fig. C.8), that in turn uses G6K-2P-Y relays to implement the H-bridge, as only low-power discrete NMOS transistors were previously cryogenically validated. The transistors are controlled by a discrete shift register to reduce the number of wires.

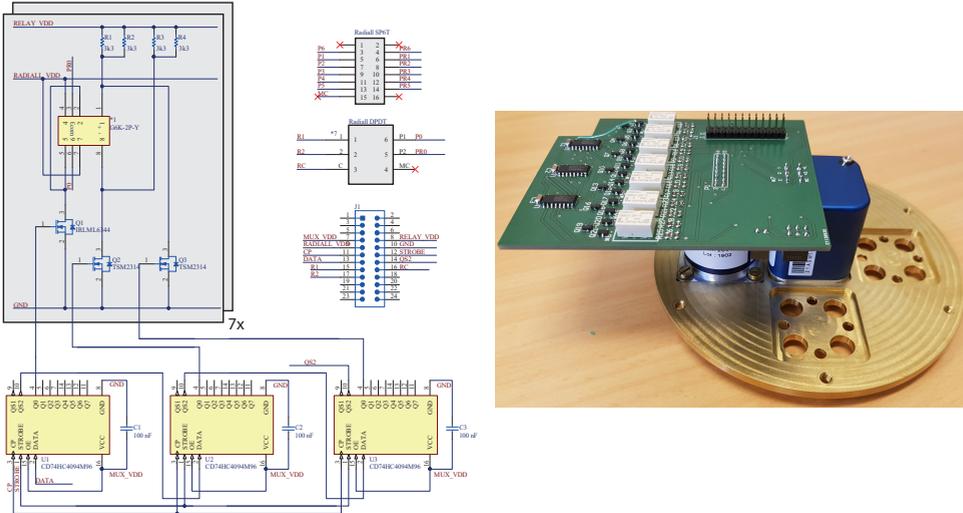


Figure C.8: Left: schematic of the relay control PCB; right: photo of the PCB mounted with the relays.

The control and DC lines from the relay control PCB and the Horse Ridge motherboard go to the appropriate adapter board (see Fig. C.5) mounted on the same plate, that adapts to the Micro-D connectors as used in the dilution refrigerator. Care has been taken in twisting the appropriate lines (as indicated in Fig. C.1), and



for the purpose of synchronization with the rest of the qubit measurement setup. The application of the *execute* trigger starts the execution of a certain command or list of commands. Additionally, the FPGA provides a pulse modulation equivalent of the output stage of Horse Ridge for the duration of the trigger pulse, by enable/disable the Horse Ridge power amplifier (PA). The application of the *sweep* trigger loads the next list of commands into Horse Ridge as has been set up in the FPGA SRAM (programmable over USB). After the last list, the first list is loaded again, and a parameter inside Horse Ridge (for instance the NCO frequency, pulse duration, or Z-correction) is swept.

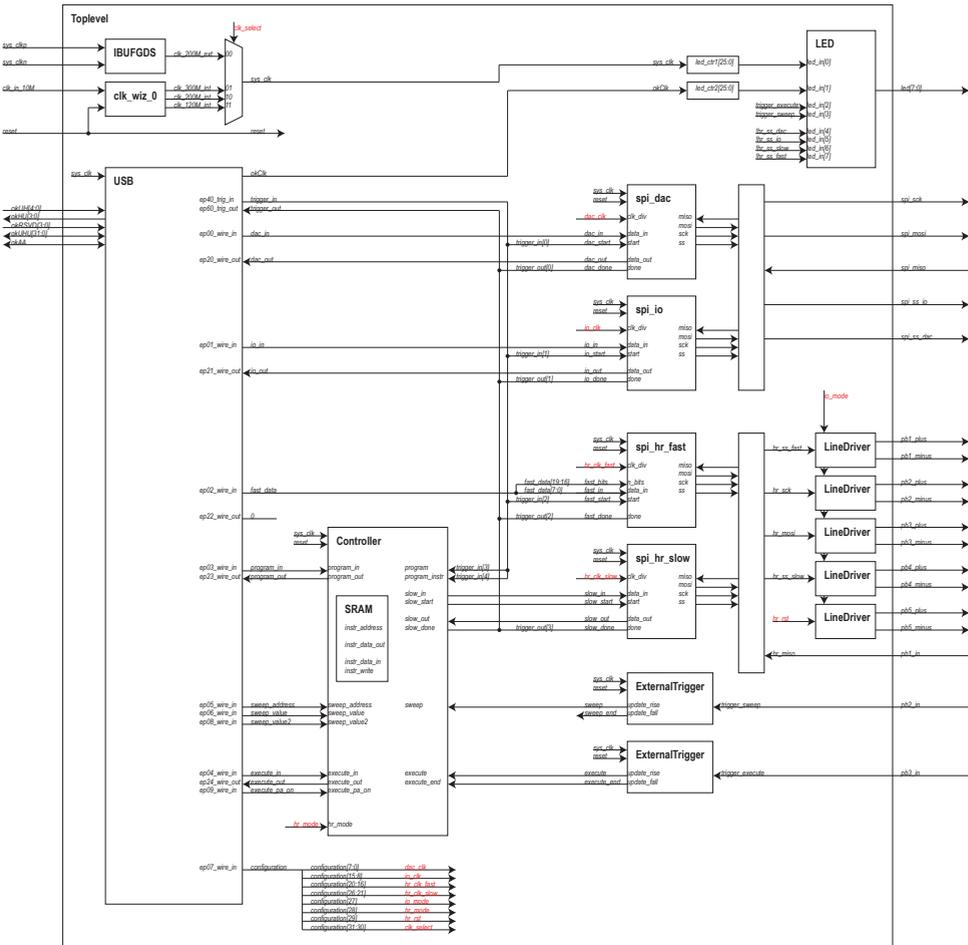


Figure C.11: Architectural diagram of the internals of the FPGA.

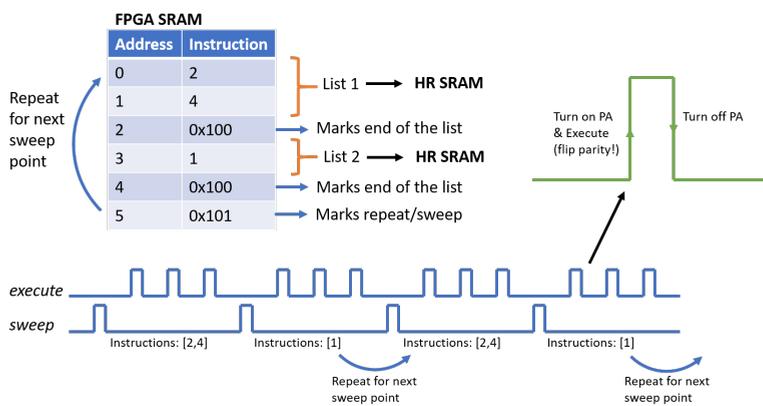


Figure C.12: The functionality of the controller inside the FPGA, and its response to externally applied triggers.

Bibliography

- [1] R. P. Feynman. 'Simulating physics with computers'. In: *International Journal of Theoretical Physics* 21.6-7 (June 1982), pp. 467–488. doi: [10.1007/bf02650179](https://doi.org/10.1007/bf02650179).
- [2] K. M. Svore and M. Troyer. 'The Quantum Future of Computation'. In: *Computer* 49.9 (Sept. 2016), pp. 21–30. doi: [10.1109/mc.2016.293](https://doi.org/10.1109/mc.2016.293).
- [3] E. Gibney. 'Physics: Quantum computer quest'. In: *Nature* 516.7529 (Dec. 2014), pp. 24–26. doi: [10.1038/516024a](https://doi.org/10.1038/516024a).
- [4] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer. 'Elucidating reaction mechanisms on quantum computers'. In: *Proceedings of the National Academy of Sciences* 114.29 (July 2017), pp. 7555–7560. doi: [10.1073/pnas.1619152114](https://doi.org/10.1073/pnas.1619152114).
- [5] L. K. Grover. 'A fast quantum mechanical algorithm for database search'. In: *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing - STOC '96*. ACM. ACM Press, 1996, pp. 212–219. doi: [10.1145/237814.237866](https://doi.org/10.1145/237814.237866).
- [6] P. Shor. 'Algorithms for quantum computation: discrete logarithms and factoring'. In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. Ieee. IEEE Comput. Soc. Press, 1994, pp. 124–134. doi: [10.1109/sfcs.1994.365700](https://doi.org/10.1109/sfcs.1994.365700).
- [7] F. G. Brandao and K. M. Svore. 'Quantum Speed-Ups for Solving Semidefinite Programs'. In: *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, Oct. 2017. doi: [10.1109/focs.2017.45](https://doi.org/10.1109/focs.2017.45).
- [8] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland. 'Surface codes: Towards practical large-scale quantum computation'. In: *Physical Review A* 86.3 (3 Sept. 2012), p. 032324. doi: [10.1103/physreva.86.032324](https://doi.org/10.1103/physreva.86.032324).
- [9] X. Fu, R. N. Schouten, C. G. Almudever, L. DiCarlo, K. Bertels, M. A. Rol, C. C. Bultink, J. van Someren, N. Khammassi, I. Ashraf, R. F. L. Vermeulen, J. C. de Sterke, and W. J. Vlothuizen. 'An experimental microarchitecture for a superconducting quantum processor'. In: *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture - MICRO-50 '17*. ACM. ACM Press, 2017, pp. 813–825. doi: [10.1145/3123939.3123952](https://doi.org/10.1145/3123939.3123952).
- [10] B. Patra, R. M. Incandela, J. P. G. van Dijk, H. A. R. Homulle, L. Song, M. Shahmohammadi, R. B. Staszewski, A. Vladimirescu, M. Babaie, F. Sebastiano, and E. Charbon. 'Cryo-CMOS Circuits and Systems for Quantum Computing Applications'. In: *IEEE Journal of Solid-State Circuits* 53.1 (Jan. 2018), pp. 309–321. issn: 0018-9200. doi: [10.1109/jssc.2017.2737549](https://doi.org/10.1109/jssc.2017.2737549).

- [11] R. Versluis and C. Hagen. 'Quantum computers scale up: Constructing a universal quantum computer with a large number of qubits will be hard but not impossible'. In: *IEEE Spectrum* 57.4 (Apr. 2020), pp. 24–29. doi: [10.1109/mspec.2020.9055969](https://doi.org/10.1109/mspec.2020.9055969).
- [12] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer. 'Gate-count estimates for performing quantum chemistry on small quantum computers'. In: *Physical Review A* 90.2 (2 Aug. 2014), p. 022305. doi: [10.1103/physreva.90.022305](https://doi.org/10.1103/physreva.90.022305).
- [13] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis. 'Quantum supremacy using a programmable superconducting processor'. In: *Nature* 574.7779 (Oct. 2019), pp. 505–510. doi: [10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5).
- [14] IBM Research. *IBM Builds Its Most Powerful Universal Quantum Computing Processors*. May 2017. url: <http://www-03.ibm.com/press/us/en/pressrelease/52403.wss> (visited on 07/06/2017).
- [15] T. F. Watson, S. G. J. Philips, E. Kawakami, D. R. Ward, P. Scarlino, M. Veldhorst, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen. 'A programmable two-qubit quantum processor in silicon'. In: *Nature* 555.7698 (Feb. 2018), pp. 633–637. doi: [10.1038/nature25766](https://doi.org/10.1038/nature25766).
- [16] Zurich Instruments. *HDAWG Arbitrary Waveform Generator*. 2018. url: <https://www.zhinst.com/products/hdawg#specification> (visited on 11/28/2018).
- [17] *Quantum Researchers Toolkit + Labber*. S5060A Solution Brief. Keysight Technologies. Aug. 2017.
- [18] R. D. V. Meter III. 'Architecture of a quantum multicomputer optimized for Shor's factoring algorithm'. In: *arXiv preprint quant-ph/0607065* (2006).
- [19] J. Hornibrook, J. Colless, I. C. Lamb, S. Pauka, H. Lu, A. Gossard, J. Watson, G. Gardner, S. Fallahi, M. Manfra, and D. Reilly. 'Cryogenic Control Architecture for Large-Scale Quantum Computing'. In: *Physical Review Applied* 3.2 (2 Feb. 2015), p. 024010. doi: [10.1103/physrevapplied.3.024010](https://doi.org/10.1103/physrevapplied.3.024010).

- [20] I. D. C. Lamb, J. I. Colless, J. M. Hornibrook, S. J. Pauka, S. J. Waddy, M. K. Frechtling, and D. J. Reilly. 'An FPGA-based instrumentation platform for use at deep cryogenic temperatures'. In: *Review of Scientific Instruments* 87.1 (Jan. 2016), p. 014701. doi: [10.1063/1.4939094](https://doi.org/10.1063/1.4939094).
- [21] H. Homulle, S. Visser, B. Patra, G. Ferrari, E. Prati, C. G. Almudéver, K. Bertels, F. Sebastiano, and E. Charbon. 'CryoCMOS hardware technology a classical infrastructure for a scalable quantum computer'. In: *Proceedings of the ACM International Conference on Computing Frontiers - CF '16*. CF '16. Como, Italy: ACM Press, 2016, pp. 282–287. isbn: 978-1-4503-4128-8. doi: [10.1145/2903150.2906828](https://doi.org/10.1145/2903150.2906828).
- [22] E. Charbon, F. Sebastiano, A. Vladimirescu, H. Homulle, S. Visser, L. Song, and R. M. Incandela. 'Cryo-CMOS for quantum computing'. In: *2016 IEEE International Electron Devices Meeting (IEDM)*. IEEE, Dec. 2016, pp. 13.5.1–13.5.4. doi: [10.1109/iedm.2016.7838410](https://doi.org/10.1109/iedm.2016.7838410).
- [23] E. Charbon, F. Sebastiano, M. Babaie, A. Vladimirescu, M. Shahmohammadi, R. B. Staszewski, H. A. Homulle, B. Patra, J. P. G. van Dijk, R. M. Incandela, L. Song, and B. Valizadehpasha. '15.5 Cryo-CMOS circuits and systems for scalable quantum computing'. In: *2017 IEEE International Solid-State Circuits Conference (ISSCC)*. IEEE, Feb. 2017, pp. 264–265. doi: [10.1109/isscc.2017.7870362](https://doi.org/10.1109/isscc.2017.7870362).
- [24] F. Sebastiano, H. Homulle, B. Patra, R. Incandela, J. P. G. van Dijk, L. Song, M. Babaie, A. Vladimirescu, and E. Charbon. 'Cryo-CMOS Electronic Control for Scalable Quantum Computing'. In: *Proceedings of the 54th Annual Design Automation Conference 2017 on - DAC '17*. DAC '17. Austin, TX, USA: ACM Press, 2017. isbn: 978-1-4503-4927-7. doi: [10.1145/3061639.3072948](https://doi.org/10.1145/3061639.3072948).
- [25] X. Qin, Z. Shi, Y. Xie, L. Wang, X. Rong, W. Jia, W. Zhang, and J. Du. 'An integrated device with high performance multi-function generators and time-to-digital convertors'. In: *Review of Scientific Instruments* 88.1 (Jan. 2017), p. 014702. doi: [10.1063/1.4973725](https://doi.org/10.1063/1.4973725).
- [26] C. A. Ryan, B. R. Johnson, D. Ristè, B. Donovan, and T. A. Ohki. 'Hardware for dynamic quantum computing'. In: *Review of Scientific Instruments* 88.10 (Oct. 2017), p. 104703. doi: [10.1063/1.5006525](https://doi.org/10.1063/1.5006525).
- [27] J. P. G. van Dijk, E. Charbon, and F. Sebastiano. 'The electronic interface for quantum processors'. In: *Microprocessors and Microsystems* 66 (Apr. 2019), pp. 90–101. doi: [10.1016/j.micpro.2019.02.004](https://doi.org/10.1016/j.micpro.2019.02.004).
- [28] D. Reilly. 'Challenges in Scaling-up the Control Interface of a Quantum Computer'. In: *arXiv preprint arXiv:1912.05114* (2019).
- [29] G. Batey, A. J. Matthews, and M. Patton. 'A new ultra-low-temperature cryogen-free experimental platform'. In: *Journal of Physics: Conference Series* 568.3 (Dec. 2014), p. 032014. doi: [10.1088/1742-6596/568/3/032014](https://doi.org/10.1088/1742-6596/568/3/032014).

- [30] L. M. K. Vandersypen, H. Bluhm, J. S. Clarke, A. S. Dzurak, R. Ishihara, A. Morello, D. J. Reilly, L. R. Schreiber, and M. Veldhorst. 'Interfacing spin qubits in quantum dots and donors—hot, dense, and coherent'. In: *npj Quantum Information* 3.1 (Sept. 2017), p. 34. issn: 2056-6387. doi: [10.1038/s41534-017-0038-y](https://doi.org/10.1038/s41534-017-0038-y).
- [31] L. Tavian. *Latest Developments in Cryogenics at CERN*. Tech. rep. CERN-AT-2005-011-ACR. Geneva: CERN, July 2005. url: <http://cds.cern.ch/record/851586>.
- [32] L. Petit, H. G. J. Eenink, M. Russ, W. I. L. Lawrie, N. W. Hendrickx, S. G. J. Philips, J. S. Clarke, L. M. K. Vandersypen, and M. Veldhorst. 'Universal quantum logic in hot silicon qubits'. In: *Nature* 580.7803 (Apr. 2020), pp. 355–359. doi: [10.1038/s41586-020-2170-7](https://doi.org/10.1038/s41586-020-2170-7).
- [33] L. Petit, J. Boter, H. Eenink, G. Droulers, M. Tagliaferri, R. Li, D. Franke, K. Singh, J. Clarke, R. Schouten, V. Dobrovitski, L. Vandersypen, and M. Veldhorst. 'Spin Lifetime and Charge Noise in Hot Silicon Quantum Dot Qubits'. In: *Physical Review Letters* 121.7 (Aug. 2018), p. 076801. doi: [10.1103/physrevlett.121.076801](https://doi.org/10.1103/physrevlett.121.076801).
- [34] L. Song, H. Homulle, E. Charbon, and F. Sebastiano. 'Characterization of bipolar transistors for cryogenic temperature sensors in standard CMOS'. In: *2016 IEEE SENSORS*. IEEE, Oct. 2016. doi: [10.1109/icsens.2016.7808759](https://doi.org/10.1109/icsens.2016.7808759).
- [35] J. Schlee, H. Rodilla, N. Wadefalk, P. Nilsson, and J. Grahn. 'Cryogenic noise performance of InGaAs/InAlAs HEMTs grown on InP and GaAs substrate'. In: *Solid-State Electronics* 91 (Jan. 2014), pp. 74–77. doi: [10.1016/j.sse.2013.10.004](https://doi.org/10.1016/j.sse.2013.10.004).
- [36] J. D. Cressler and A. H. Mantooth. *Extreme environment electronics*. CRC Press, 2013. isbn: 9781138074224.
- [37] R. Kirschman. 'Survey of Low-Temperature Electronics'. In: *Workshop on Low-Temp. Electronics WOLTE11*. 2014.
- [38] S. R. Ekanayake, T. Lehmann, A. S. Dzurak, R. G. Clark, and A. Brawley. 'Characterization of SOS-CMOS FETs at Low Temperatures for the Design of Integrated Circuits for Quantum Bit Control and Readout'. In: *IEEE Transactions on Electron Devices* 57.2 (Feb. 2010), pp. 539–547. doi: [10.1109/ted.2009.2037381](https://doi.org/10.1109/ted.2009.2037381).
- [39] R. M. Incandela, L. Song, H. Homulle, E. Charbon, A. Vladimirescu, and F. Sebastiano. 'Characterization and Compact Modeling of Nanometer CMOS Transistors at Deep-Cryogenic Temperatures'. In: *IEEE Journal of the Electron Devices Society* 6 (2018), pp. 996–1006. doi: [10.1109/jeds.2018.2821763](https://doi.org/10.1109/jeds.2018.2821763).
- [40] H. Homulle, S. Visser, B. Patra, G. Ferrari, E. Prati, F. Sebastiano, and E. Charbon. 'A reconfigurable cryogenic platform for the classical control of quantum processors'. In: *Review of Scientific Instruments* 88.4 (Apr. 2017), p. 045103. doi: [10.1063/1.4979611](https://doi.org/10.1063/1.4979611).

- [41] P. A. 't Hart, J. P. G. van Dijk, M. Babaie, E. Charbon, A. Vladimircscu, and F. Sebastiano. 'Characterization and Model Validation of Mismatch in Nanometer CMOS at Cryogenic Temperatures'. In: *2018 48th European Solid-State Device Research Conference (ESSDERC)*. IEEE, Sept. 2018, pp. 246–249. doi: [10.1109/essderc.2018.8486859](https://doi.org/10.1109/essderc.2018.8486859).
- [42] A. Coskun and J. Bardin. 'Cryogenic small-signal and noise performance of 32nm SOI CMOS'. In: *2014 IEEE MTT-S International Microwave Symposium (IMS2014)*. IEEE, June 2014. doi: [10.1109/mwsym.2014.6848614](https://doi.org/10.1109/mwsym.2014.6848614).
- [43] J. Yoneda, K. Takeda, T. Otsuka, T. Nakajima, M. R. Delbecq, G. Allison, T. Honda, T. Kodera, S. Oda, Y. Hoshi, N. Usami, K. M. Itoh, and S. Tarucha. 'A quantum-dot spin qubit with coherence limited by charge noise and fidelity higher than 99.9%'. In: *Nature Nanotechnology* 13.2 (Dec. 2017), pp. 102–106. doi: [10.1038/s41565-017-0014-x](https://doi.org/10.1038/s41565-017-0014-x).
- [44] J. Johansson, P. Nation, and F. Nori. 'QuTiP: An open-source Python framework for the dynamics of open quantum systems'. In: *Computer Physics Communications* 183.8 (Aug. 2012), pp. 1760–1772. doi: [10.1016/j.cpc.2012.02.021](https://doi.org/10.1016/j.cpc.2012.02.021).
- [45] D. P. DiVincenzo, D. Bacon, J. Kempe, G. Burkard, and K. B. Whaley. 'Universal quantum computation with the exchange interaction'. In: *Nature* 408.6810 (Nov. 2000), pp. 339–342. doi: [10.1038/35042541](https://doi.org/10.1038/35042541).
- [46] N. K. Langford, R. Sagastizabal, M. Kounalakis, C. Dickel, A. Bruno, F. Luthi, D. J. Thoen, A. Endo, and L. DiCarlo. 'Experimentally simulating the dynamics of quantum light and matter at deep-strong coupling'. In: *Nature Communications* 8.1 (Nov. 2017), p. 1715. doi: [10.1038/s41467-017-01061-x](https://doi.org/10.1038/s41467-017-01061-x).
- [47] C. Dickel. *How to make artificial atoms out of electrical circuits – Part II: Circuit quantum electrodynamics and the transmon*. Aug. 2017. url: <http://blog.gatech.nl/index.php/2017/08/13/how-to-%20make-artificial-atoms-out-of-electrical-circuits-part-ii-%20circuit-quantum-electrodynamics-and-the-transmon/> (visited on 06/26/2018).
- [48] D. Kim, Z. Shi, C. B. Simmons, D. R. Ward, J. R. Prance, T. S. Koh, J. K. Gamble, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, and M. A. Eriksson. 'Quantum control and process tomography of a semiconductor quantum dot hybrid qubit'. In: *Nature* 511.7507 (July 2014), pp. 70–74. doi: [10.1038/nature13407](https://doi.org/10.1038/nature13407).
- [49] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. 'Charge-insensitive qubit design derived from the Cooper pair box'. In: *Physical Review A* 76.4 (Oct. 2007), p. 042319. doi: [10.1103/physreva.76.042319](https://doi.org/10.1103/physreva.76.042319).

- [50] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. 'Demonstration of two-qubit algorithms with a superconducting quantum processor'. In: *Nature* 460.7252 (June 2009), pp. 240–244. doi: [10.1038/nature08121](https://doi.org/10.1038/nature08121).
- [51] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis. 'Superconducting quantum circuits at the surface code threshold for fault tolerance'. In: *Nature* 508.7497 (Apr. 2014), pp. 500–503. doi: [10.1038/nature13171](https://doi.org/10.1038/nature13171).
- [52] J. M. Chow, J. M. Gambetta, E. Magesan, D. W. Abraham, A. W. Cross, B. R. Johnson, N. A. Masluk, C. A. Ryan, J. A. Smolin, S. J. Srinivasan, and M. Steffen. 'Implementing a strand of a scalable fault-tolerant quantum computing fabric'. In: *Nature Communications* 5.1 (June 2014), p. 4015. doi: [10.1038/ncomms5015](https://doi.org/10.1038/ncomms5015).
- [53] E. Jeffrey, D. Sank, J. Mutus, T. White, J. Kelly, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Megrant, P. O'Malley, C. Neill, P. Roushan, A. Vainsencher, J. Wenner, A. Cleland, and J. M. Martinis. 'Fast Accurate State Measurement with Superconducting Qubits'. In: *Physical Review Letters* 112.19 (May 2014), p. 190504. doi: [10.1103/physrevlett.112.190504](https://doi.org/10.1103/physrevlett.112.190504).
- [54] D. Ristè, S. Poletto, M.-Z. Huang, A. Bruno, V. Vesterinen, O.-P. Saira, and L. DiCarlo. 'Detecting bit-flip errors in a logical qubit using stabilizer measurements'. In: *Nature Communications* 6.1 (Apr. 2015), p. 6983. doi: [10.1038/ncomms7983](https://doi.org/10.1038/ncomms7983).
- [55] M. Reagor, C. B. Osborn, N. Tezak, A. Staley, G. Prawiroatmodjo, M. Scheer, N. Alidoust, E. A. Sete, N. Didier, M. P. da Silva, E. Acala, J. Angeles, A. Bestwick, M. Block, B. Bloom, A. Bradley, C. Bui, S. Caldwell, L. Capelluto, R. Chilcott, J. Cordova, G. Crossman, M. Curtis, S. Deshpande, T. E. Bouayadi, D. Girshovich, S. Hong, A. Hudson, P. Karalekas, K. Kuang, M. Lenihan, R. Manenti, T. Manning, J. Marshall, Y. Mohan, W. O'Brien, J. Otterbach, A. Papageorge, J.-P. Paquette, M. Pelstring, A. Polloreno, V. Rawat, C. A. Ryan, R. Renzas, N. Rubin, D. Russel, M. Rust, D. Scarabelli, M. Selvanayagam, R. Sinclair, R. Smith, M. Suska, T.-W. To, M. Vahidpour, N. Vodrahalli, T. Whyland, K. Yadav, W. Zeng, and C. T. Rigetti. 'Demonstration of universal parametric entangling gates on a multi-qubit lattice'. In: *Science Advances* 4.2 (Feb. 2018), eao3603. doi: [10.1126/sciadv.aao3603](https://doi.org/10.1126/sciadv.aao3603).
- [56] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm. 'Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits'. In: *Physical Review Letters* 103.11 (Sept. 2009), p. 110501. doi: [10.1103/physrevlett.103.110501](https://doi.org/10.1103/physrevlett.103.110501).

- [57] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta. 'Efficient Z gates for quantum computing'. In: *Physical Review A* 96.2 (Aug. 2017), p. 022330. doi: [10.1103/physreva.96.022330](https://doi.org/10.1103/physreva.96.022330).
- [58] J. M. Chow, A. D. Córcoles, J. M. Gambetta, C. Rigetti, B. R. Johnson, J. A. Smolin, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen. 'Simple All-Microwave Entangling Gate for Fixed-Frequency Superconducting Qubits'. In: *Physical Review Letters* 107.8 (Aug. 2011), p. 080502. doi: [10.1103/physrevlett.107.080502](https://doi.org/10.1103/physrevlett.107.080502).
- [59] J. M. Chow, J. M. Gambetta, A. W. Cross, S. T. Merkel, C. Rigetti, and M. Steffen. 'Microwave-activated conditional-phase gate for superconducting qubits'. In: *New Journal of Physics* 15.11 (Nov. 2013), p. 115012. doi: [10.1088/1367-2630/15/11/115012](https://doi.org/10.1088/1367-2630/15/11/115012).
- [60] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver. 'A quantum engineer's guide to superconducting qubits'. In: *Applied Physics Reviews* 6.2 (June 2019), p. 021318. doi: [10.1063/1.5089550](https://doi.org/10.1063/1.5089550).
- [61] M. Veldhorst, J. C. C. Hwang, C. H. Yang, A. W. Leenstra, B. de Ronde, J. P. Dehollain, J. T. Muhonen, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak. 'An addressable quantum dot qubit with fault-tolerant control-fidelity'. In: *Nature Nanotechnology* 9.12 (Oct. 2014), pp. 981–985. doi: [10.1038/nnano.2014.216](https://doi.org/10.1038/nnano.2014.216).
- [62] R. Maurand, X. Jehl, D. Kotekar-Patil, A. Corna, H. Bohuslavskyi, R. Laviéville, L. Hutin, S. Barraud, M. Vinet, M. Sanquer, and S. D. Franceschi. 'A CMOS silicon spin qubit'. In: *Nature Communications* 7.1 (Nov. 2016), p. 13575. doi: [10.1038/ncomms13575](https://doi.org/10.1038/ncomms13575).
- [63] A. Sammak, D. Sabbagh, N. W. Hendrickx, M. Lodari, B. P. Wuetz, A. Tosato, L. Yeoh, M. Bollani, M. Virgilio, M. A. Schubert, P. Zaumseil, G. Capellini, M. Veldhorst, and G. Scappucci. 'Shallow and Undoped Germanium Quantum Wells: A Playground for Spin and Hybrid Quantum Technology'. In: *Advanced Functional Materials* 29.14 (Jan. 2019), p. 1807613. doi: [10.1002/adfm.201807613](https://doi.org/10.1002/adfm.201807613).
- [64] K. C. Nowack, M. Shafiei, M. Laforest, G. E. D. K. Prawiroatmodjo, L. R. Schreiber, C. Reichl, W. Wegscheider, and L. M. K. Vandersypen. 'Single-Shot Correlations and Two-Qubit Gate of Solid-State Spins'. In: *Science* 333.6047 (Aug. 2011), pp. 1269–1272. doi: [10.1126/science.1209524](https://doi.org/10.1126/science.1209524).
- [65] E. Kawakami, P. Scarlino, D. R. Ward, F. R. Braakman, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen. 'Electrical control of a long-lived spin qubit in a Si/SiGe quantum dot'. In: *Nature Nanotechnology* 9.9 (Aug. 2014), pp. 666–670. doi: [10.1038/nnano.2014.153](https://doi.org/10.1038/nnano.2014.153).
- [66] M. Veldhorst, C. H. Yang, J. C. C. Hwang, W. Huang, J. P. Dehollain, J. T. Muhonen, S. Simmons, A. Laucht, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak. 'A two-qubit logic gate in silicon'. In: *Nature* 526.7573 (Oct. 2015). Letter, pp. 410–414. issn: 0028-0836. doi: [10.1038/nature15263](https://doi.org/10.1038/nature15263).

- [67] D. Loss and D. P. DiVincenzo. 'Quantum computation with quantum dots'. In: *Physical Review A* 57.1 (Jan. 1998), pp. 120–126. doi: [10.1103/physreva.57.120](https://doi.org/10.1103/physreva.57.120).
- [68] J. M. Elzerman, R. Hanson, L. H. W. van Beveren, B. Witkamp, L. M. K. Vandersypen, and L. P. Kouwenhoven. 'Single-shot read-out of an individual electron spin in a quantum dot'. In: *Nature* 430.6998 (July 2004), pp. 431–435. doi: [10.1038/nature02693](https://doi.org/10.1038/nature02693).
- [69] K. Ono. 'Current Rectification by Pauli Exclusion in a Weakly Coupled Double Quantum Dot System'. In: *Science* 297.5585 (July 2002), pp. 1313–1317. doi: [10.1126/science.1070958](https://doi.org/10.1126/science.1070958).
- [70] J. I. Colless, A. C. Mahoney, J. M. Hornibrook, A. C. Doherty, H. Lu, A. C. Gossard, and D. J. Reilly. 'Dispersive Readout of a Few-Electron Double Quantum Dot with Fast rf Gate Sensors'. In: *Physical Review Letters* 110.4 (Jan. 2013), p. 046805. doi: [10.1103/physrevlett.110.046805](https://doi.org/10.1103/physrevlett.110.046805).
- [71] J. J. Pla, K. Y. Tan, J. P. Dehollain, W. H. Lim, J. J. L. Morton, D. N. Jamieson, A. S. Dzurak, and A. Morello. 'A single-atom electron spin qubit in silicon'. In: *Nature* 489.7417 (Sept. 2012), pp. 541–545. doi: [10.1038/nature11449](https://doi.org/10.1038/nature11449).
- [72] J. J. Pla, K. Y. Tan, J. P. Dehollain, W. H. Lim, J. J. L. Morton, F. A. Zwanenburg, D. N. Jamieson, A. S. Dzurak, and A. Morello. 'High-fidelity readout and control of a nuclear spin qubit in silicon'. In: *Nature* 496.7445 (Apr. 2013), pp. 334–338. doi: [10.1038/nature12011](https://doi.org/10.1038/nature12011).
- [73] R. Kalra, A. Laucht, C. D. Hill, and A. Morello. 'Robust Two-Qubit Gates for Donors in Silicon Controlled by Hyperfine Interactions'. In: *Physical Review X* 4.2 (June 2014), p. 021044. doi: [10.1103/physrevx.4.021044](https://doi.org/10.1103/physrevx.4.021044).
- [74] J. T. Muhonen, J. P. Dehollain, A. Laucht, F. E. Hudson, R. Kalra, T. Sekiguchi, K. M. Itoh, D. N. Jamieson, J. C. McCallum, A. S. Dzurak, and A. Morello. 'Storing quantum information for 30 seconds in a nanoelectronic device'. In: *Nature Nanotechnology* 9.12 (Oct. 2014), pp. 986–991. doi: [10.1038/nnano.2014.211](https://doi.org/10.1038/nnano.2014.211).
- [75] F. A. Zwanenburg, A. S. Dzurak, A. Morello, M. Y. Simmons, L. C. L. Holtenberg, G. Klimeck, S. Rogge, S. N. Coppersmith, and M. A. Eriksson. 'Silicon quantum electronics'. In: *Reviews of Modern Physics* 85.3 (July 2013), pp. 961–1019. doi: [10.1103/revmodphys.85.961](https://doi.org/10.1103/revmodphys.85.961).
- [76] J. Levy. 'Universal Quantum Computation with Spin-1/2 Pairs and Heisenberg Exchange'. In: *Physical Review Letters* 89.14 (Sept. 2002), p. 147902. doi: [10.1103/physrevlett.89.147902](https://doi.org/10.1103/physrevlett.89.147902).
- [77] J. R. Petta, A. C. Johnson, J. M. Taylor, E. A. Laird, A. Yacoby, M. D. Lukin, C. M. Marcus, M. P. Hanson, and A. C. Gossard. 'Coherent Manipulation of Coupled Electron Spins in Semiconductor Quantum Dots'. In: *Science* 309.5744 (Sept. 2005), pp. 2180–2184. issn: 0036-8075. doi: [10.1126/science.1116955](https://doi.org/10.1126/science.1116955).

- [78] S. Foletti, H. Bluhm, D. Mahalu, V. Umansky, and A. Yacoby. 'Universal quantum control of two-electron spin quantum bits using dynamic nuclear polarization'. In: *Nature Physics* 5.12 (Oct. 2009), pp. 903–908. doi: [10.1038/nphys1424](https://doi.org/10.1038/nphys1424).
- [79] B. M. Maune, M. G. Borselli, B. Huang, T. D. Ladd, P. W. Deelman, K. S. Holabird, A. A. Kiselev, I. Alvarado-Rodriguez, R. S. Ross, A. E. Schmitz, M. Sokolich, C. A. Watson, M. F. Gyure, and A. T. Hunter. 'Coherent singlet-triplet oscillations in a silicon-based double quantum dot'. In: *Nature* 481.7381 (Jan. 2012), pp. 344–347. doi: [10.1038/nature10707](https://doi.org/10.1038/nature10707).
- [80] M. D. Shulman, O. E. Dial, S. P. Harvey, H. Bluhm, V. Umansky, and A. Yacoby. 'Demonstration of Entanglement of Electrostatically Coupled Singlet-Triplet Qubits'. In: *Science* 336.6078 (Apr. 2012), pp. 202–205. doi: [10.1126/science.1217692](https://doi.org/10.1126/science.1217692).
- [81] X. Wu, D. R. Ward, J. R. Prance, D. Kim, J. K. Gamble, R. T. Mohr, Z. Shi, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, and M. A. Eriksson. 'Two-axis control of a singlet-triplet qubit with an integrated micromagnet'. In: *Proceedings of the National Academy of Sciences* 111.33 (Aug. 2014), pp. 11938–11942. doi: [10.1073/pnas.1412230111](https://doi.org/10.1073/pnas.1412230111).
- [82] J. M. Nichol, L. A. Orona, S. P. Harvey, S. Fallahi, G. C. Gardner, M. J. Manfra, and A. Yacoby. 'High-fidelity entangling gate for double-quantum-dot spin qubits'. In: *npj Quantum Information* 3.1 (Jan. 2017), p. 3. doi: [10.1038/s41534-016-0003-1](https://doi.org/10.1038/s41534-016-0003-1).
- [83] L. Gaudreau, G. Granger, A. Kam, G. C. Aers, S. A. Studenikin, P. Zawadzki, M. Pioro-Ladrière, Z. R. Wasilewski, and A. S. Sachrajda. 'Coherent control of three-spin states in a triple quantum dot'. In: *Nature Physics* 8.1 (Nov. 2011), pp. 54–58. doi: [10.1038/nphys2149](https://doi.org/10.1038/nphys2149).
- [84] J. Medford, J. Beil, J. M. Taylor, S. D. Bartlett, A. C. Doherty, E. I. Rashba, D. P. DiVincenzo, H. Lu, A. C. Gossard, and C. M. Marcus. 'Self-consistent measurement and state tomography of an exchange-only spin qubit'. In: *Nature Nanotechnology* 8.9 (Sept. 2013), pp. 654–659. doi: [10.1038/nnano.2013.168](https://doi.org/10.1038/nnano.2013.168).
- [85] K. Eng, T. D. Ladd, A. Smith, M. G. Borselli, A. A. Kiselev, B. H. Fong, K. S. Holabird, T. M. Hazard, B. Huang, P. W. Deelman, I. Milosavljevic, A. E. Schmitz, R. S. Ross, M. F. Gyure, and A. T. Hunter. 'Isotopically enhanced triple-quantum-dot qubit'. In: *Science Advances* 1.4 (May 2015), e1500214. doi: [10.1126/sciadv.1500214](https://doi.org/10.1126/sciadv.1500214).
- [86] Y.-P. Shim and C. Tahan. 'Charge-noise-insensitive gate operations for always-on, exchange-only qubits'. In: *Physical Review B* 93.12 (Mar. 2016), p. 121410. doi: [10.1103/physrevb.93.121410](https://doi.org/10.1103/physrevb.93.121410).
- [87] F. Setiawan, H.-Y. Hui, J. P. Kestner, X. Wang, and S. D. Sarma. 'Robust two-qubit gates for exchange-coupled qubits'. In: *Physical Review B* 89.8 (Feb. 2014), p. 085314. doi: [10.1103/physrevb.89.085314](https://doi.org/10.1103/physrevb.89.085314).

- [88] Z. Shi, C. B. Simmons, J. R. Prance, J. K. Gamble, T. S. Koh, Y.-P. Shim, X. Hu, D. E. Savage, M. G. Lagally, M. A. Eriksson, M. Friesen, and S. N. Coppersmith. 'Fast Hybrid Silicon Double-Quantum-Dot Qubit'. In: *Physical Review Letters* 108.14 (Apr. 2012), p. 140503. doi: [10.1103/physrevlett.108.140503](https://doi.org/10.1103/physrevlett.108.140503).
- [89] T. S. Koh, J. K. Gamble, M. Friesen, M. A. Eriksson, and S. N. Coppersmith. 'Pulse-Gated Quantum-Dot Hybrid Qubit'. In: *Physical Review Letters* 109.25 (Dec. 2012), p. 250503. doi: [10.1103/physrevlett.109.250503](https://doi.org/10.1103/physrevlett.109.250503).
- [90] Z. Shi, C. B. Simmons, D. R. Ward, J. R. Prance, X. Wu, T. S. Koh, J. K. Gamble, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, and M. A. Eriksson. 'Fast coherent manipulation of three-electron states in a double quantum dot'. In: *Nature Communications* 5.1 (Jan. 2014), p. 3020. doi: [10.1038/ncomms4020](https://doi.org/10.1038/ncomms4020).
- [91] E. Ferraro, M. D. Michielis, G. Mazzeo, M. Fanciulli, and E. Prati. 'Effective Hamiltonian for the hybrid double quantum dot qubit'. In: *Quantum Information Processing* 13.5 (Jan. 2014), pp. 1155–1173. doi: [10.1007/s11128-013-0718-2](https://doi.org/10.1007/s11128-013-0718-2).
- [92] D. Kim, D. R. Ward, C. B. Simmons, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, and M. A. Eriksson. 'High-fidelity resonant gating of a silicon-based quantum dot hybrid qubit'. In: *npj Quantum Information* 1.1 (Oct. 2015), p. 15004. doi: [10.1038/npjqi.2015.4](https://doi.org/10.1038/npjqi.2015.4).
- [93] F. Mallet, F. R. Ong, A. Palacios-Laloy, F. Nguyen, P. Bertet, D. Vion, and D. Esteve. 'Single-shot qubit readout in circuit quantum electrodynamics'. In: *Nature Physics* 5.11 (Sept. 2009), pp. 791–795. doi: [10.1038/nphys1400](https://doi.org/10.1038/nphys1400).
- [94] *4-8 GHz Cryogenic Low Noise Amplifier*. LNF-LNC4-8A. Low Noise Factory. Jan. 2013.
- [95] L. M. K. Vandersypen, J. M. Elzerman, R. N. Schouten, L. H. W. van Beveren, R. Hanson, and L. P. Kouwenhoven. 'Real-time detection of single-electron tunneling using a quantum point contact'. In: *Applied Physics Letters* 85.19 (2004), p. 4394. doi: [10.1063/1.1815041](https://doi.org/10.1063/1.1815041).
- [96] I. T. Vink, T. Nooitgedagt, R. N. Schouten, L. M. K. Vandersypen, and W. Wegscheider. 'Cryogenic amplifier for fast real-time detection of single-electron tunneling'. In: *Applied Physics Letters* 91.12 (Sept. 2007), p. 123512. doi: [10.1063/1.2783265](https://doi.org/10.1063/1.2783265).
- [97] D. J. Reilly, C. M. Marcus, M. P. Hanson, and A. C. Gossard. 'Fast single-charge sensing with a rf quantum point contact'. In: *Applied Physics Letters* 91.16 (Oct. 2007), p. 162101. doi: [10.1063/1.2794995](https://doi.org/10.1063/1.2794995).
- [98] J. M. Hornibrook, J. I. Colless, A. C. Mahoney, X. G. Croot, S. Blanvillain, H. Lu, A. C. Gossard, and D. J. Reilly. 'Frequency multiplexing for readout of spin qubits'. In: *Applied Physics Letters* 104.10 (Mar. 2014), p. 103108. doi: [10.1063/1.4868107](https://doi.org/10.1063/1.4868107).

- [99] E. A. Laird, J. M. Taylor, D. P. DiVincenzo, C. M. Marcus, M. P. Hanson, and A. C. Gossard. 'Coherent spin manipulation in an exchange-only qubit'. In: *Physical Review B* 82.7 (Aug. 2010), p. 075403. doi: [10.1103/physrevb.82.075403](https://doi.org/10.1103/physrevb.82.075403).
- [100] *Arbitrary Waveform Generators AWG5000 Series Data Sheet*. 76W2226014. Tektronix. Mar. 2017.
- [101] *E8267D PSG Vector Signal Generator*. 5989-0697EN. Keysight Technologies. Feb. 2016.
- [102] J. P. G. van Dijk, E. Kawakami, R. Schouten, M. Veldhorst, L. Vandersypen, M. Babaie, E. Charbon, and F. Sebastiano. 'Impact of Classical Control Electronics on Qubit Fidelity'. In: *Physical Review Applied* 12.4 (Oct. 2019), p. 044054. doi: [10.1103/physrevapplied.12.044054](https://doi.org/10.1103/physrevapplied.12.044054).
- [103] J. E. Levy, M. S. Carroll, A. Ganti, C. A. Phillips, A. J. Landahl, T. M. Gurrieri, R. D. Carr, H. L. Stalford, and E. Nielsen. 'Implications of electronics constraints for solid-state quantum error correction and quantum circuit failure probability'. In: *New Journal of Physics* 13.8 (Aug. 2011), p. 083021. doi: [10.1088/1367-2630/13/8/083021](https://doi.org/10.1088/1367-2630/13/8/083021).
- [104] Y. Salathé, P. Kurpiers, T. Karg, C. Lang, C. K. Andersen, A. Akin, S. Krinner, C. Eichler, and A. Wallraff. 'Low-Latency Digital Signal Processing for Feedback and Feedforward in Quantum Computing and Communication'. In: *Physical Review Applied* 9.3 (Mar. 2018), p. 034011. doi: [10.1103/physrevapplied.9.034011](https://doi.org/10.1103/physrevapplied.9.034011).
- [105] J. Lin, F.-T. Liang, Y. Xu, L.-H. Sun, C. Guo, S.-K. Liao, and C.-Z. Peng. 'High Performance and Scalable AWG for Superconducting Quantum Computing'. In: *arXiv preprint arXiv:1806.03660* (2018).
- [106] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf. 'Extending the lifetime of a quantum bit with error correction in superconducting circuits'. In: *Nature* 536.7617 (July 2016), pp. 441–445. doi: [10.1038/nature18949](https://doi.org/10.1038/nature18949).
- [107] S. Asaad, C. Dickel, N. K. Langford, S. Poletto, A. Bruno, M. A. Rol, D. Deurloo, and L. DiCarlo. 'Independent, extensible control of same-frequency superconducting qubits by selective broadcasting'. In: *npj Quantum Information* 2.1 (Aug. 2016), p. 16029. doi: [10.1038/npjqi.2016.29](https://doi.org/10.1038/npjqi.2016.29).
- [108] M. Veldhorst, H. G. J. Eenink, C. H. Yang, and A. S. Dzurak. 'Silicon CMOS architecture for a spin-based quantum computer'. In: *Nature Communications* 8.1 (Dec. 2017), p. 1766. issn: 2041-1723. doi: [10.1038/s41467-017-01905-6](https://doi.org/10.1038/s41467-017-01905-6).
- [109] R. Li, L. Petit, D. P. Franke, J. P. Dehollain, J. Helsen, M. Steudtner, N. K. Thomas, Z. R. Yoscovits, K. J. Singh, S. Wehner, L. M. K. Vandersypen, J. S. Clarke, and M. Veldhorst. 'A crossbar network for silicon quantum dot qubits'. In: *Science Advances* 4.7 (July 2018), eaar3960. doi: [10.1126/sciadv.aar3960](https://doi.org/10.1126/sciadv.aar3960).

- [110] J. M. Boter, J. P. Dehollain, J. P. G. van Dijk, T. Hensgens, R. Versluis, J. S. Clarke, M. Veldhorst, F. Sebastiano, and L. M. K. Vandersypen. 'A sparse spin qubit array with integrated control electronics'. In: *2019 IEEE International Electron Devices Meeting (IEDM)*. IEEE, Dec. 2019. doi: [10.1109/iedm19573.2019.8993570](https://doi.org/10.1109/iedm19573.2019.8993570).
- [111] R. K. Puddy, L. W. Smith, H. Al-Taie, C. H. Chong, I. Farrer, J. P. Griffiths, D. A. Ritchie, M. J. Kelly, M. Pepper, and C. G. Smith. 'Multiplexed charge-locking device for large arrays of quantum devices'. In: *Applied Physics Letters* 107.14 (Oct. 2015), p. 143501. doi: [10.1063/1.4932012](https://doi.org/10.1063/1.4932012).
- [112] S. Schaal, S. Barraud, J. Morton, and M. Gonzalez-Zalba. 'Conditional Dispersive Readout of a CMOS Single-Electron Memory Cell'. In: *Physical Review Applied* 9.5 (May 2018), p. 054016. doi: [10.1103/physrevapplied.9.054016](https://doi.org/10.1103/physrevapplied.9.054016).
- [113] S. Pauka, K. Das, R. Kalra, A. Moini, Y. Yang, M. Trainer, A. Bousquet, C. Cantaloube, N. Dick, G. Gardner, *et al.* 'A cryogenic interface for controlling many qubits'. In: *arXiv preprint arXiv:1912.01299* (2019).
- [114] B. P. Wuetz, P. Bavdaz, L. Yeoh, R. Schouten, H. van der Does, M. Tiggelman, D. Sabbagh, A. Sammak, C. Almudever, F. Sebastiano, *et al.* 'Multiplexed quantum transport using commercial off-the-shelf CMOS at subkelvin temperatures'. In: *arXiv preprint arXiv:1907.11816* (2019).
- [115] S.-Y. Wu, C. Lin, M. Chiang, J. Liaw, J. Cheng, S. Yang, C. Tsai, P. Chen, T. Miyashita, C. Chang, V. Chang, K. Pan, J. Chen, Y. Mor, K. Lai, C. Liang, H. Chen, S. Chang, C. Lin, C. Hsieh, R. Tsui, C. Yao, C. Chen, R. Chen, C. Lee, H. Lin, C. Chang, K. Chen, M. Tsai, K. Chen, Y. Ku, and S. M. Jang. 'A 7nm CMOS platform technology featuring 4th generation FinFET transistors with a 0.027 μ m² high density 6-T SRAM cell for mobile SoC applications'. In: *2016 IEEE International Electron Devices Meeting (IEDM)*. IEEE. IEEE, Dec. 2016, pp. 2–6. doi: [10.1109/iedm.2016.7838333](https://doi.org/10.1109/iedm.2016.7838333).
- [116] R. Versluis, S. Poletto, N. Khammassi, B. Tarasinski, N. Haider, D. Michalak, A. Bruno, K. Bertels, and L. DiCarlo. 'Scalable Quantum Circuit and Control for a Superconducting Surface Code'. In: *Physical Review Applied* 8.3 (Sept. 2017). doi: [10.1103/physrevapplied.8.034021](https://doi.org/10.1103/physrevapplied.8.034021).
- [117] M. Jerger, S. Poletto, P. Macha, U. Hübner, E. Il'ichev, and A. V. Ustinov. 'Frequency division multiplexing readout and simultaneous manipulation of an array of flux qubits'. In: *Applied Physics Letters* 101.4 (July 2012), p. 042604. doi: [10.1063/1.4739454](https://doi.org/10.1063/1.4739454).
- [118] Y. Tokura, W. G. van der Wiel, T. Obata, and S. Tarucha. 'Coherent Single Electron Spin Control in a Slanting Zeeman Field'. In: *Physical Review Letters* 96.4 (Jan. 2006), p. 047202. doi: [10.1103/physrevlett.96.047202](https://doi.org/10.1103/physrevlett.96.047202).

- [119] E. Kawakami, T. Jullien, P. Scarlino, D. R. Ward, D. E. Savage, M. G. Lagally, V. V. Dobrovitski, M. Friesen, S. N. Coppersmith, M. A. Eriksson, and L. M. K. Vandersypen. 'Gate fidelity and coherence of an electron spin in an Si/SiGe quantum dot with micromagnet'. In: *Proceedings of the National Academy of Sciences* 113.42 (Oct. 2016), pp. 11738–11743. doi: [10.1073/pnas.1603251113](https://doi.org/10.1073/pnas.1603251113).
- [120] M. Steffen, L. M. Vandersypen, and I. L. Chuang. 'Simultaneous Soft Pulses Applied at Nearby Frequencies'. In: *Journal of Magnetic Resonance* 146.2 (Oct. 2000), pp. 369–374. issn: 1090-7807. doi: [10.1006/jmre.2000.2178](https://doi.org/10.1006/jmre.2000.2178).
- [121] D. J. Reilly. 'Engineering the quantum-classical interface of solid-state qubits'. In: *npj Quantum Information* 1.1 (Oct. 2015), p. 15011. doi: [10.1038/npjqi.2015.11](https://doi.org/10.1038/npjqi.2015.11).
- [122] S. R. Ekanayake, T. Lehmann, A. S. Dzurak, and R. G. Clark. 'Quantum bit controller and observer circuits in SOS-CMOS technology for gigahertz low-temperature operation'. In: *2007 7th IEEE Conference on Nanotechnology (IEEE NANO)*. IEEE. IEEE, Aug. 2007, pp. 1283–1287. doi: [10.1109/nano.2007.4601417](https://doi.org/10.1109/nano.2007.4601417).
- [123] C. Degenhardt, L. Geck, A. Kruth, P. Vliex, and S. van Waasen. 'CMOS Based Scalable Cryogenic Control Electronics for Qubits'. In: *2017 IEEE International Conference on Rebooting Computing (ICRC)*. IEEE. IEEE, Nov. 2017, pp. 1–4. doi: [10.1109/icrc.2017.8123682](https://doi.org/10.1109/icrc.2017.8123682).
- [124] BlueFors Cryogenics. *XLD Series*. url: <http://www.bluefors.com/newsite/index.php/xld-seriest> (visited on 03/17/2017).
- [125] M. Vahidpour, W. O'Brien, J. T. Whyland, J. Angeles, J. Marshall, D. Scarabelli, G. Crossman, K. Yadav, Y. Mohan, C. Bui, *et al.* 'Superconducting through-silicon vias for quantum integrated circuits'. In: *arXiv preprint arXiv:1708.02226* (2017).
- [126] T. Tighe, G. Akerling, and A. Smith. 'Cryogenic packaging for multi-GHz electronics'. In: *IEEE Transactions on Applied Superconductivity* 9.2 (June 1999), pp. 3173–3176. doi: [10.1109/77.783703](https://doi.org/10.1109/77.783703).
- [127] H. van Weers, G. Kunkel, M. Lindeman, and M. Leeman. 'Niobium flex cable for low temperature high density interconnects'. In: *Cryogenics* 55-56 (May 2013), pp. 1–4. doi: [10.1016/j.cryogenics.2012.10.006](https://doi.org/10.1016/j.cryogenics.2012.10.006).
- [128] N. K. Katam, O. A. Mukhanov, and M. Pedram. 'Superconducting Magnetic Field Programmable Gate Array'. In: *IEEE Transactions on Applied Superconductivity* 28.2 (Mar. 2018), pp. 1–12. doi: [10.1109/tasc.2018.2797262](https://doi.org/10.1109/tasc.2018.2797262).
- [129] R. M. Incandela, L. Song, H. Homulle, F. Sebastiano, E. Charbon, and A. Vladimirescu. 'Nanometer CMOS characterization and compact modeling at deep-cryogenic temperatures'. In: *2017 47th European Solid-State Device Research Conference (ESSDERC)*. IEEE. IEEE, Sept. 2017, pp. 58–61. doi: [10.1109/essderc.2017.8066591](https://doi.org/10.1109/essderc.2017.8066591).

- [130] H. Homulle, L. Song, E. Charbon, and F. Sebastiano. 'The Cryogenic Temperature Behavior of Bipolar, MOS, and DT MOS Transistors in Standard CMOS'. In: *IEEE Journal of the Electron Devices Society* 6.1 (2018), pp. 263–270. doi: [10.1109/jeds.2018.2798281](https://doi.org/10.1109/jeds.2018.2798281).
- [131] N. C. Dao, A. E. Kass, C. T. Jin, and P. H. W. Leong. 'Impact of Series Resistance on Bulk CMOS Current Matching Over the 5–300K Temperature Range'. In: *IEEE Electron Device Letters* 38.7 (July 2017), pp. 847–850. doi: [10.1109/led.2017.2709545](https://doi.org/10.1109/led.2017.2709545).
- [132] A. Beckers, F. Jazaeri, and C.ENZ. 'Characterization and Modeling of 28-nm Bulk CMOS Technology Down to 4.2 K'. In: *IEEE Journal of the Electron Devices Society* 6 (2018), pp. 1007–1018. doi: [10.1109/jeds.2018.2817458](https://doi.org/10.1109/jeds.2018.2817458).
- [133] S. S. Tannu, D. M. Carmean, and M. K. Qureshi. 'Cryogenic-DRAM based memory system for scalable quantum computers'. In: *Proceedings of the International Symposium on Memory Systems - MEMSYS '17*. ACM. ACM Press, 2017, pp. 189–195. doi: [10.1145/3132402.3132436](https://doi.org/10.1145/3132402.3132436).
- [134] S. R. Ekanayake, T. Lehmann, A. S. Dzurak, and R. G. Clark. 'Qubit Control-Pulse Generator Circuits for Operation at Cryogenic Temperatures'. In: *2008 8th IEEE Conference on Nanotechnology*. IEEE. IEEE, Aug. 2008, pp. 472–475. doi: [10.1109/nano.2008.143](https://doi.org/10.1109/nano.2008.143).
- [135] M. T. Rahman and T. Lehmann. 'A Self-Calibrated Cryogenic Current Cell for 4.2 K Current Steering D/A Converters'. In: *IEEE Transactions on Circuits and Systems II: Express Briefs* 64.10 (Oct. 2017), pp. 1152–1156. doi: [10.1109/tcsii.2016.2631489](https://doi.org/10.1109/tcsii.2016.2631489).
- [136] J. Gong, Y. Chen, F. Sebastiano, E. Charbon, and M. Babaie. '19.3 A 200dB FoM 4-to-5GHz Cryogenic Oscillator with an Automatic Common-Mode Resonance Calibration for Quantum Computing Applications'. In: *2020 IEEE International Solid-State Circuits Conference - (ISSCC)*. IEEE, Feb. 2020. doi: [10.1109/isscc19947.2020.9062913](https://doi.org/10.1109/isscc19947.2020.9062913).
- [137] J. C. Bardin, T. White, M. Giustina, K. J. Satzinger, K. Arya, P. Roushan, B. Chiaro, J. Kelly, Z. Chen, B. Burkett, Y. Chen, E. Jeffrey, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, R. Graff, P. Klimov, J. Mutus, M. J. McEwen, M. Neeley, C. J. Neill, E. Lucero, C. Quintana, A. Vainsencher, H. Neven, J. Martinis, T. Huang, S. Das, D. T. Sank, O. Naaman, A. E. Megrant, and R. Barends. 'Design and Characterization of a 28-nm Bulk-CMOS Cryogenic Quantum Controller Dissipating Less Than 2 mW at 3 K'. In: *IEEE Journal of Solid-State Circuits* 54.11 (Nov. 2019), pp. 3043–3060. doi: [10.1109/jssc.2019.2937234](https://doi.org/10.1109/jssc.2019.2937234).
- [138] K. Das and T. Lehmann. 'Low power fast cryogenic CMOS circuit for digital readout of single electron transistor'. In: *2011 IEEE 54th International Midwest Symposium on Circuits and Systems (MWSCAS)*. IEEE. IEEE, Aug. 2011, pp. 1–4. doi: [10.1109/mwscas.2011.6026405](https://doi.org/10.1109/mwscas.2011.6026405).

- [139] H. Homulle. *Cryogenic electronics for the read-out of quantum processors*. 2019. doi: [10.4233/UUID:E833F394-C8B1-46E2-86B8-DA0C71559538](https://doi.org/10.4233/UUID:E833F394-C8B1-46E2-86B8-DA0C71559538).
- [140] M. Mehrpoo, F. Sebastiano, E. Charbon, and M. Babaie. 'A Cryogenic CMOS Parametric Amplifier'. In: *IEEE Solid-State Circuits Letters* 3 (2020), pp. 5–8. doi: [10.1109/lssc.2019.2950186](https://doi.org/10.1109/lssc.2019.2950186).
- [141] A. Ruffino, Y. Peng, F. Sebastiano, M. Babaie, and E. Charbon. 'A Wideband Low-Power Cryogenic CMOS Circulator for Quantum Applications'. In: *IEEE Journal of Solid-State Circuits* (2020), pp. 1–1. doi: [10.1109/jssc.2020.2978020](https://doi.org/10.1109/jssc.2020.2978020).
- [142] H. Homulle, F. Sebastiano, and E. Charbon. 'Deep-Cryogenic Voltage References in 40-nm CMOS'. In: *IEEE Solid-State Circuits Letters* 1.5 (May 2018), pp. 110–113. doi: [10.1109/lssc.2018.2875821](https://doi.org/10.1109/lssc.2018.2875821).
- [143] J. van Staveren, C. G. Almudever, G. Scappucci, M. Veldhorst, M. Babaie, E. Charbon, and F. Sebastiano. 'Voltage References for the Ultra-Wide Temperature Range from 4.2K to 300K in 40-nm CMOS'. In: *ESSCIRC 2019 - IEEE 45th European Solid State Circuits Conference (ESSCIRC)*. IEEE, Sept. 2019. doi: [10.1109/esscirc.2019.8902861](https://doi.org/10.1109/esscirc.2019.8902861).
- [144] E. Leonard, M. A. Beck, J. Nelson, B. Christensen, T. Thorbeck, C. Howington, A. Opremca, I. Pechenezhskiy, K. Dodge, N. Dupuis, M. Hutchings, J. Ku, F. Schlenker, J. Suttle, C. Wilen, S. Zhu, M. Vavilov, B. Plourde, and R. McDermott. 'Digital Coherent Control of a Superconducting Qubit'. In: *Physical Review Applied* 11.1 (Jan. 2019), p. 014009. doi: [10.1103/physrevapplied.11.014009](https://doi.org/10.1103/physrevapplied.11.014009).
- [145] H. Ball, W. D. Oliver, and M. J. Biercuk. 'The role of master clock stability in quantum information processing'. In: *npj Quantum Information* 2.1 (Nov. 2016), p. 16033. doi: [10.1038/npjqi.2016.33](https://doi.org/10.1038/npjqi.2016.33).
- [146] J. Raftery, A. Vrajitoarea, G. Zhang, Z. Leng, S. Srinivasan, and A. Houck. 'Direct digital synthesis of microwave waveforms for quantum computing'. In: *arXiv preprint arXiv:1703.00942* (2017).
- [147] T. Green, H. Uys, and M. J. Biercuk. 'High-Order Noise Filtering in Non-trivial Quantum Logic Gates'. In: *Physical Review Letters* 109.2 (July 2012), p. 020501. doi: [10.1103/physrevlett.109.020501](https://doi.org/10.1103/physrevlett.109.020501).
- [148] T. J. Green, J. Sastrawan, H. Uys, and M. J. Biercuk. 'Arbitrary quantum control of qubits in the presence of universal noise'. In: *New Journal of Physics* 15.9 (Sept. 2013), p. 095004. doi: [10.1088/1367-2630/15/9/095004](https://doi.org/10.1088/1367-2630/15/9/095004).
- [149] P. Rebentrost, I. Serban, T. Schulte-Herbrüggen, and F. K. Wilhelm. 'Optimal Control of a Qubit Coupled to a Non-Markovian Environment'. In: *Physical Review Letters* 102.9 (Mar. 2009), p. 090401. doi: [10.1103/physrevlett.102.090401](https://doi.org/10.1103/physrevlett.102.090401).

- [150] S. M. Anton, C. Müller, J. S. Birenbaum, S. R. O'Kelley, A. D. Fefferman, D. S. Golubev, G. C. Hilton, H.-M. Cho, K. D. Irwin, F. C. Wellstood, G. Schön, A. Shnirman, and J. Clarke. 'Pure dephasing in flux qubits due to flux noise with spectral density scaling as $1/f\alpha$ '. In: *Physical Review B* 85.22 (June 2012), p. 224505. doi: [10.1103/physrevb.85.224505](https://doi.org/10.1103/physrevb.85.224505).
- [151] Y. V. Nazarov. *Quantum noise in mesoscopic physics*. Vol. 97. Springer Science & Business Media, 2012. isbn: 978-94-010-0089-5.
- [152] M. Reed, B. Maune, R. Andrews, M. Borselli, K. Eng, M. Jura, A. Kiselev, T. Ladd, S. Merkel, I. Milosavljevic, E. Pritchett, M. Rakher, R. Ross, A. Schmitz, A. Smith, J. Wright, M. Gyure, and A. Hunter. 'Reduced Sensitivity to Charge Noise in Semiconductor Spin Qubits via Symmetric Operation'. In: *Physical Review Letters* 116.11 (11 Mar. 2016), p. 110402. doi: [10.1103/physrevlett.116.110402](https://doi.org/10.1103/physrevlett.116.110402).
- [153] F. Martins, F. K. Malinowski, P. D. Nissen, E. Barnes, S. Fallahi, G. C. Gardner, M. J. Manfra, C. M. Marcus, and F. Kuemmeth. 'Noise Suppression Using Symmetric Exchange Gates in Spin Qubits'. In: *Physical Review Letters* 116.11 (11 Mar. 2016), p. 116801. doi: [10.1103/physrevlett.116.116801](https://doi.org/10.1103/physrevlett.116.116801).
- [154] N. Didier, E. A. Sete, J. Combes, and M. P. da Silva. 'ac Flux Sweet Spots in Parametrically Modulated Superconducting Qubits'. In: *Physical Review Applied* 12.5 (Nov. 2019), p. 054015. doi: [10.1103/physrevapplied.12.054015](https://doi.org/10.1103/physrevapplied.12.054015).
- [155] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, R. Cleve, and I. L. Chuang. 'Experimental Realization of an Order-Finding Algorithm with an NMR Quantum Computer'. In: *Physical Review Letters* 85.25 (25 Dec. 2000), pp. 5452–5455. doi: [10.1103/physrevlett.85.5452](https://doi.org/10.1103/physrevlett.85.5452).
- [156] L. M. K. Vandersypen and I. L. Chuang. 'NMR techniques for quantum control and computation'. In: *Reviews of Modern Physics* 76.4 (4 Jan. 2005), pp. 1037–1069. doi: [10.1103/revmodphys.76.1037](https://doi.org/10.1103/revmodphys.76.1037).
- [157] J. I. Cirac and P. Zoller. 'Quantum Computations with Cold Trapped Ions'. In: *Physical Review Letters* 74.20 (May 1995), pp. 4091–4094. doi: [10.1103/physrevlett.74.4091](https://doi.org/10.1103/physrevlett.74.4091).
- [158] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt. '14-Qubit Entanglement: Creation and Coherence'. In: *Physical Review Letters* 106.13 (Mar. 2011), p. 130506. doi: [10.1103/physrevlett.106.130506](https://doi.org/10.1103/physrevlett.106.130506).
- [159] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai. 'Coherent control of macroscopic quantum states in a single-Cooper-pair box'. In: *Nature* 398.6730 (Apr. 1999), pp. 786–788. doi: [10.1038/19718](https://doi.org/10.1038/19718).
- [160] Y. Makhlin, G. Schön, and A. Shnirman. 'Quantum-state engineering with Josephson-junction devices'. In: *Reviews of Modern Physics* 73.2 (May 2001), pp. 357–400. doi: [10.1103/revmodphys.73.357](https://doi.org/10.1103/revmodphys.73.357).

- [161] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis. 'State preservation by repetitive error detection in a superconducting quantum circuit'. In: *Nature* 519.7541 (Mar. 2015), pp. 66–69. doi: [10.1038/nature14270](https://doi.org/10.1038/nature14270).
- [162] L. Childress, M. V. G. Dutt, J. M. Taylor, A. S. Zibrov, F. Jelezko, J. Wrachtrup, P. R. Hemmer, and M. D. Lukin. 'Coherent Dynamics of Coupled Electron and Nuclear Spin Qubits in Diamond'. In: *Science* 314.5797 (Oct. 2006), pp. 281–285. doi: [10.1126/science.1131871](https://doi.org/10.1126/science.1131871).
- [163] R. Hanson, L. P. Kouwenhoven, J. R. Petta, S. Tarucha, and L. M. K. Vandersypen. 'Spins in few-electron quantum dots'. In: *Reviews of Modern Physics* 79.4 (Oct. 2007), pp. 1217–1265. doi: [10.1103/revmodphys.79.1217](https://doi.org/10.1103/revmodphys.79.1217).
- [164] S. J. Angus, A. J. Ferguson, A. S. Dzurak, and R. G. Clark. 'Gate-Defined Quantum Dots in Intrinsic Silicon'. In: *Nano Letters* 7.7 (July 2007), pp. 2051–2055. doi: [10.1021/nl070949k](https://doi.org/10.1021/nl070949k).
- [165] D. Zajac, T. Hazard, X. Mi, E. Nielsen, and J. Petta. 'Scalable Gate Architecture for a One-Dimensional Array of Semiconductor Spin Qubits'. In: *Physical Review Applied* 6.5 (5 Nov. 2016), p. 054013. doi: [10.1103/physrevapplied.6.054013](https://doi.org/10.1103/physrevapplied.6.054013).
- [166] J. P. Dehollain, J. J. Pla, E. Siew, K. Y. Tan, A. S. Dzurak, and A. Morello. 'Nanoscale broadband transmission lines for spin qubit control'. In: *Nanotechnology* 24.1 (Dec. 2012), p. 015202. doi: [10.1088/0957-4484/24/1/015202](https://doi.org/10.1088/0957-4484/24/1/015202).
- [167] A. West, B. Hensen, A. Jouan, T. Tanttu, C.-H. Yang, A. Rossi, M. F. Gonzalez-Zalba, F. Hudson, A. Morello, D. J. Reilly, and A. S. Dzurak. 'Gate-based single-shot readout of spins in silicon'. In: *Nature Nanotechnology* 14.5 (Mar. 2019), pp. 437–441. issn: 1748-3395. doi: [10.1038/s41565-019-0400-7](https://doi.org/10.1038/s41565-019-0400-7).
- [168] C. Barthel, M. Kjaergaard, J. Medford, M. Stopa, C. M. Marcus, M. P. Hanson, and A. C. Gossard. 'Fast sensing of double-dot charge arrangement and spin state with a radio-frequency sensor quantum dot'. In: *Physical Review B* 81.16 (16 Apr. 2010), p. 161308. doi: [10.1103/physrevb.81.161308](https://doi.org/10.1103/physrevb.81.161308).
- [169] M. A. Nielsen. 'A simple formula for the average gate fidelity of a quantum dynamical operation'. In: *Physics Letters A* 303.4 (Oct. 2002), pp. 249–252. doi: [10.1016/s0375-9601\(02\)01272-0](https://doi.org/10.1016/s0375-9601(02)01272-0).
- [170] L. H. Pedersen, N. M. Møller, and K. Mølmer. 'Fidelity of quantum operations'. In: *Physics Letters A* 367.1-2 (July 2007), pp. 47–51. issn: 0375-9601. doi: [10.1016/j.physleta.2007.02.069](https://doi.org/10.1016/j.physleta.2007.02.069).
- [171] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser. 'Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms'. In: *Journal of Magnetic Resonance* 172.2 (Feb. 2005), pp. 296–305. doi: [10.1016/j.jmr.2004.11.004](https://doi.org/10.1016/j.jmr.2004.11.004).

- [172] T. C. Carusone, D. Johns, and K. Martin. *Analog Integrated Circuit Design*. WILEY, Dec. 1, 2011. 794 pp. isbn: 0470770104. url: https://www.ebook.de/de/product/14948016/tony_chan_carusone_david_johns_kenneth_martin_analog_integrated_circuit_design.html.
- [173] G. S. Uhrig. 'Keeping a Quantum Bit Alive by Optimized π -Pulse Sequences'. In: *Physical Review Letters* 98.10 (Mar. 2007), p. 100504. doi: [10.1103/physrevlett.98.100504](https://doi.org/10.1103/physrevlett.98.100504).
- [174] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger. 'Optimized dynamical decoupling in a model quantum memory'. In: *Nature* 458.7241 (Apr. 2009), pp. 996–1000. doi: [10.1038/nature07951](https://doi.org/10.1038/nature07951).
- [175] *Solutions for Reducing Phase Noise at RF and Microwave Frequencies*. 5990-7529EN. Keysight Technologies. Aug. 2014.
- [176] J. Phillips and K. Kundert. 'Noise in mixers, oscillators, samplers, and logic an introduction to cyclostationary noise'. In: *Proceedings of the IEEE 2000 Custom Integrated Circuits Conference (Cat. No.00CH37044)*. IEEE. IEEE, 2000, pp. 431–438. doi: [10.1109/cicc.2000.852702](https://doi.org/10.1109/cicc.2000.852702).
- [177] B. Drakhlis. 'Calculate oscillator jitter by using phase-noise analysis'. In: *MICROWAVES & RF* 40.1 (2001), pp. 82–+.
- [178] M. McCoy and L. Mueller. 'Selective Decoupling'. In: *Journal of Magnetic Resonance, Series A* 101.2 (Feb. 1993), pp. 122–130. doi: [10.1006/jmra.1993.1021](https://doi.org/10.1006/jmra.1993.1021).
- [179] E. Kupce and R. Freeman. 'Close Encounters between Soft Pulses'. In: *Journal of Magnetic Resonance, Series A* 112.2 (Feb. 1995), pp. 261–264. issn: 1064-1858. doi: [10.1006/jmra.1995.1043](https://doi.org/10.1006/jmra.1995.1043).
- [180] J. Pauly, P. L. Roux, D. Nishimura, and A. Macovski. 'Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm (NMR imaging)'. In: *IEEE Transactions on Medical Imaging* 10.1 (Mar. 1991), pp. 53–65. issn: 0278-0062. doi: [10.1109/42.75611](https://doi.org/10.1109/42.75611).
- [181] W. A. Coish and D. Loss. 'Exchange-controlled single-electron-spin rotations in quantum dots'. In: *Physical Review B* 75.16 (16 Apr. 2007), p. 161302. doi: [10.1103/physrevb.75.161302](https://doi.org/10.1103/physrevb.75.161302).
- [182] 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 (12 Mar. 2011), p. 121403. doi: [10.1103/physrevb.83.121403](https://doi.org/10.1103/physrevb.83.121403).
- [183] E. A. Laird, J. R. Petta, A. C. Johnson, C. M. Marcus, A. Yacoby, M. P. Hanson, and A. C. Gossard. 'Effect of Exchange Interaction on Spin Dephasing in a Double Quantum Dot'. In: *Physical Review Letters* 97.5 (5 July 2006), p. 056801. doi: [10.1103/physrevlett.97.056801](https://doi.org/10.1103/physrevlett.97.056801).
- [184] G. Turin. 'An introduction to matched filters'. In: *IEEE Transactions on Information Theory* 6.3 (June 1960), pp. 311–329. doi: [10.1109/tit.1960.1057571](https://doi.org/10.1109/tit.1960.1057571).

- [185] *Keysight M9330A Arbitrary Waveform Generator 15-bit, 1.25 GS/s*. 5990-6426EN. Keysight Technologies. July 2014.
- [186] *MODELS WX1281/2C 1.25GS/s Single/Dual Channel Arbitrary Waveform Generators*. WX1282C. Tabor Electronics Inc. Jan. 2015.
- [187] *R&S®SMW200A Vector Signal Generator Specifications*. Version 06.01, January 2018. Rohde & Schwarz. Jan. 2018.
- [188] W.-T. Lin and T.-H. Kuo. 'A Compact Dynamic-Performance-Improved Current-Steering DAC With Random Rotation-Based Binary-Weighted Selection'. In: *IEEE Journal of Solid-State Circuits* 47.2 (Feb. 2012), pp. 444–453. issn: 0018-9200. doi: [10.1109/jssc.2011.2168651](https://doi.org/10.1109/jssc.2011.2168651).
- [189] W.-T. Lin, H.-Y. Huang, and T.-H. Kuo. 'A 12-bit 40 nm DAC Achieving SFDR > 70 dB at 1.6 GS/s and IMD $\leq -61\text{dB}$ at 2.8 GS/s With DEM-DRZ Technique'. In: *IEEE Journal of Solid-State Circuits* 49.3 (Mar. 2014), pp. 708–717. issn: 0018-9200. doi: [10.1109/jssc.2014.2301769](https://doi.org/10.1109/jssc.2014.2301769).
- [190] K. Raczkowski, N. Markulic, B. Hershberg, and J. Craninckx. 'A 9.2-12.7 GHz Wideband Fractional-N Subsampling PLL in 28 nm CMOS With 280 fs RMS Jitter'. In: *IEEE Journal of Solid-State Circuits* 50.5 (May 2015), pp. 1203–1213. issn: 0018-9200. doi: [10.1109/jssc.2015.2403373](https://doi.org/10.1109/jssc.2015.2403373).
- [191] J. P. G. van Dijk, A. Vladimirescu, M. Babaie, E. Charbon, and F. Sebastiano. 'A co-design methodology for scalable quantum processors and their classical electronic interface'. In: *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE. IEEE, Mar. 2018, pp. 573–576. doi: [10.23919/date.2018.8342072](https://doi.org/10.23919/date.2018.8342072).
- [192] J. P. G. van Dijk, A. Vladimirescu, M. Babaie, E. Charbon, and F. Sebastiano. 'SPINE (SPIN Emulator) - A Quantum-Electronics Interface Simulator'. In: *2019 IEEE 8th International Workshop on Advances in Sensors and Interfaces (IWASI)*. IEEE. IEEE, June 2019, pp. 23–28. doi: [10.1109/iwasi.2019.8791334](https://doi.org/10.1109/iwasi.2019.8791334).
- [193] N. Khammassi, I. Ashraf, X. Fu, C. Almudever, and K. Bertels. 'QX: A high-performance quantum computer simulation platform'. In: *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2017*. IEEE, Mar. 2017, pp. 464–469. doi: [10.23919/date.2017.7927034](https://doi.org/10.23919/date.2017.7927034).
- [194] J. P. G. van Dijk, B. Patra, S. Pellerano, E. Charbon, F. Sebastiano, and M. Babaie. 'Designing a DDS-Based SoC for High-Fidelity Multi-Qubit Control'. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 67.12 (Dec. 2020), pp. 5380–5393. doi: [10.1109/tcsi.2020.3019413](https://doi.org/10.1109/tcsi.2020.3019413).
- [195] Y. M. Greshishchev, D. Pollex, S.-C. Wang, M. Besson, P. Flemeke, S. Szilagyi, J. Aguirre, C. Falt, N. Ben-Hamida, R. Gibbins, and P. Schvan. 'A 56GS/S 6b DAC in 65nm CMOS with 256Mbytes memory'. In: *2011 IEEE International Solid-State Circuits Conference*. IEEE, Feb. 2011, pp. 194–196. doi: [10.1109/isscc.2011.5746279](https://doi.org/10.1109/isscc.2011.5746279).

- [196] L. Duncan, B. Dupaix, J. J. McCue, B. Mathieu, M. LaRue, V. J. Patel, M. Teshome, M.-J. Choe, and W. Khalil. 'A 10-bit DC-20-GHz Multiple-Return-to-Zero DAC With >48-dB SFDR'. In: *IEEE Journal of Solid-State Circuits* 52.12 (Dec. 2017), pp. 3262–3275. issn: 0018-9200. doi: [10.1109/jssc.2017.2749441](https://doi.org/10.1109/jssc.2017.2749441).
- [197] J. C. Bardin, E. Jeffrey, E. Lucero, T. Huang, O. Naaman, R. Barends, T. White, M. Giustina, D. Sank, P. Roushan, K. Arya, B. Chiaro, J. Kelly, J. Chen, B. Burkett, Y. Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, R. Graff, P. Klimov, J. Mutus, M. McEwen, A. Megrant, M. Neeley, C. Neill, C. Quintana, A. Vainsencher, H. Neven, and J. Martinis. '29.1 A 28nm Bulk-CMOS 4-to-8GHz i2mW Cryogenic Pulse Modulator for Scalable Quantum Computing'. In: *2019 IEEE International Solid-State Circuits Conference - (ISSCC)*. IEEE, Feb. 2019, pp. 456–458. doi: [10.1109/isscc.2019.8662480](https://doi.org/10.1109/isscc.2019.8662480).
- [198] E. Bechthum, G. I. Radulov, J. Briaire, G. J. G. M. Geelen, and A. H. M. van Roermund. 'A Wideband RF Mixing-DAC Achieving IMD < -82 dBc Up to 1.9 GHz'. In: *IEEE Journal of Solid-State Circuits* 51.6 (June 2016), pp. 1374–1384. issn: 0018-9200. doi: [10.1109/jssc.2016.2543703](https://doi.org/10.1109/jssc.2016.2543703).
- [199] R. Cushing *et al.* 'A Technical Tutorial on Digital Signal Synthesis'. In: *Analog Devices* (1999).
- [200] Z. Chen, J. Kelly, C. Quintana, R. Barends, B. Campbell, Y. Chen, B. Chiaro, A. Dunsworth, A. Fowler, E. Lucero, E. Jeffrey, A. Megrant, J. Mutus, M. Neeley, C. Neill, P. O'Malley, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. White, A. Korotkov, and J. M. Martinis. 'Measuring and Suppressing Quantum State Leakage in a Superconducting Qubit'. In: *Physical Review Letters* 116.2 (2 Jan. 2016), p. 020501. doi: [10.1103/physrevlett.116.020501](https://doi.org/10.1103/physrevlett.116.020501).
- [201] E. Vittoz. 'Low-power design: ways to approach the limits'. In: *Proceedings of IEEE International Solid-State Circuits Conference - ISSCC '94*. IEEE, Feb. 1994, pp. 14–18. doi: [10.1109/isscc.1994.344744](https://doi.org/10.1109/isscc.1994.344744).
- [202] M. J. Pelgrom. *Analog-to-Digital Conversion*. Springer New York, 2013. doi: [10.1007/978-1-4614-1371-4](https://doi.org/10.1007/978-1-4614-1371-4).
- [203] B. Razavi. *RF Microelectronics (2Nd Edition) (Prentice Hall Communications Engineering and Emerging Technologies Series)*. 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall Press, 2011. isbn: 0137134738.
- [204] M. H. Roshan, S. Zaliasl, K. Joo, K. Souri, R. Palwai, W. Chen, S. Pamarti, J. C. Doll, N. Miller, C. Arft, S. Tabatabaei, C. Sechen, A. Partridge, and V. Menon. '11.1 Dual-MEMS-resonator temperature-to-digital converter with 40 K resolution and FOM of 0.12pJK²'. In: *2016 IEEE International Solid-State Circuits Conference (ISSCC)*. IEEE, Jan. 2016. doi: [10.1109/isscc.2016.7417976](https://doi.org/10.1109/isscc.2016.7417976).

- [205] T. Yoo, H. C. Yeoh, Y.-H. Jung, S.-J. Cho, Y. S. Kim, S.-M. Kang, and K.-H. Baek. 'A 2 GHz 130 mW Direct-Digital Frequency Synthesizer With a Non-linear DAC in 55 nm CMOS'. In: *IEEE Journal of Solid-State Circuits* 49.12 (Dec. 2014), pp. 2976–2989. doi: [10.1109/jssc.2014.2359674](https://doi.org/10.1109/jssc.2014.2359674).
- [206] A. M. Alonso, M. Miyahara, and A. Matsuzawa. 'A 12.8-ns-Latency DDFS MMIC With Frequency, Phase, and Amplitude Modulations in 65-nm CMOS'. In: *IEEE Journal of Solid-State Circuits* 53.10 (Oct. 2018), pp. 2840–2849. doi: [10.1109/jssc.2018.2859393](https://doi.org/10.1109/jssc.2018.2859393).
- [207] G. G. Shahidi. 'Chip Power Scaling in Recent CMOS Technology Nodes'. In: *IEEE Access* 7 (2019), pp. 851–856. doi: [10.1109/access.2018.2885895](https://doi.org/10.1109/access.2018.2885895).
- [208] C. Hardy, W. Edelstein, and D. Vatis. 'Efficient adiabatic fast passage for NMR population inversion in the presence of radiofrequency field inhomogeneity and frequency offsets'. In: *Journal of Magnetic Resonance (1969)* 66.3 (Feb. 1986), pp. 470–482. issn: 0022-2364. doi: [10.1016/0022-2364\(86\)90190-3](https://doi.org/10.1016/0022-2364(86)90190-3).
- [209] B. Patra, J. P. G. van Dijk, S. Subramanian, A. Corna, X. Xue, C. Jeon, F. Sheikh, E. Juarez-Hernandez, B. P. Esparza, H. Rampurawala, B. Carlton, N. Samkharadze, S. Ravikumar, C. Nieva, S. Kim, H.-J. Lee, A. Sammak, G. Scappucci, M. Veldhorst, L. M. K. Vandersypen, M. Babaie, F. Sebastiano, E. Charbon, and S. Pellerano. '19.1 A Scalable Cryo-CMOS 2-to-20GHz Digitally Intensive Controller for 4×32 Frequency Multiplexed Spin Qubits/Transmons in 22nm FinFET Technology for Quantum Computers'. In: *2020 IEEE International Solid-State Circuits Conference - (ISSCC)*. IEEE, Feb. 2020. doi: [10.1109/isscc19947.2020.9063109](https://doi.org/10.1109/isscc19947.2020.9063109).
- [210] J. P. G. V. Dijk, B. Patra, S. Subramanian, X. Xue, N. Samkharadze, A. Corna, C. Jeon, F. Sheikh, E. Juarez-Hernandez, B. P. Esparza, H. Rampurawala, B. R. Carlton, S. Ravikumar, C. Nieva, S. Kim, H.-J. Lee, A. Sammak, G. Scappucci, M. Veldhorst, L. M. K. Vandersypen, E. Charbon, S. Pellerano, M. Babaie, and F. Sebastiano. 'A Scalable Cryo-CMOS Controller for the Wideband Frequency-Multiplexed Control of Spin Qubits and Transmons'. In: *IEEE Journal of Solid-State Circuits* 55.11 (Nov. 2020), pp. 2930–2946. doi: [10.1109/jssc.2020.3024678](https://doi.org/10.1109/jssc.2020.3024678).
- [211] C. H. Yang, R. C. C. Leon, J. C. C. Hwang, A. Saraiva, T. Tanttu, W. Huang, J. C. Lemyre, K. W. Chan, K. Y. Tan, F. E. Hudson, K. M. Itoh, A. Morello, M. Pioro-Ladrière, A. Laucht, and A. S. Dzurak. 'Operation of a silicon quantum processor unit cell above one kelvin'. In: *Nature* 580.7803 (Apr. 2020), pp. 350–354. doi: [10.1038/s41586-020-2171-6](https://doi.org/10.1038/s41586-020-2171-6).
- [212] L. C. Camenzind, S. Geyer, A. Fuhrer, R. J. Warburton, D. M. Zumbühl, and A. V. Kuhlmann. 'A spin qubit in a fin field-effect transistor'. In: (Mar. 12, 2021). arXiv: <http://arxiv.org/abs/2103.07369v1> [cond-mat.mes-hall].

- [213] A. Grill, E. Bury, J. Michl, S. Tyaginov, D. Linten, T. Grasser, B. Parvais, B. Kaczer, M. Waltl, and I. Radu. 'Reliability and Variability of Advanced CMOS Devices at Cryogenic Temperatures'. In: *2020 IEEE International Reliability Physics Symposium (IRPS)*. IEEE, Apr. 2020, pp. 1–6. doi: [10.1109/irps45951.2020.9128316](https://doi.org/10.1109/irps45951.2020.9128316).
- [214] P. A. T. Hart, M. Babaie, E. Charbon, A. Vladimirescu, and F. Sebastiano. 'Characterization and Modeling of Mismatch in Cryo-CMOS'. In: *IEEE Journal of the Electron Devices Society* 8 (2020), pp. 263–273. doi: [10.1109/jeds.2020.2976546](https://doi.org/10.1109/jeds.2020.2976546).
- [215] B. Patra, M. Mehrpoo, A. Ruffino, F. Sebastiano, E. Charbon, and M. Babaie. 'Characterization and Analysis of On-Chip Microwave Passive Components at Cryogenic Temperatures'. In: *IEEE Journal of the Electron Devices Society* 8 (2020), pp. 448–456. doi: [10.1109/jeds.2020.2986722](https://doi.org/10.1109/jeds.2020.2986722).
- [216] Y. Wang, C. Hull, G. Murata, and S. Ravid. 'A linear-in-dB analog baseband circuit for low power 60GHz receiver in standard 65nm CMOS'. In: *2013 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*. IEEE, June 2013. doi: [10.1109/rfic.2013.6569567](https://doi.org/10.1109/rfic.2013.6569567).
- [217] A. van den Bosch, M. Borremans, M. Steyaert, and W. Sansen. 'A 10-bit 1-GSample/s Nyquist current-steering CMOS D/A converter'. In: *IEEE Journal of Solid-State Circuits* 36.3 (Mar. 2001), pp. 315–324. doi: [10.1109/4.910469](https://doi.org/10.1109/4.910469).
- [218] C.-H. Lin, F. van der Goes, J. Westra, J. Mulder, Y. Lin, E. Arslan, E. Ayranci, X. Liu, and K. Bult. 'A 12b 2.9GS/s DAC with IM3 ≪−60dBc beyond 1GHz in 65nm CMOS'. In: *2009 IEEE International Solid-State Circuits Conference - Digest of Technical Papers*. IEEE, Feb. 2009. doi: [10.1109/isscc.2009.4977314](https://doi.org/10.1109/isscc.2009.4977314).
- [219] B. Razavi. *Design of Analog CMOS Integrated Circuits*. MCGRAW HILL BOOK CO, Aug. 1, 2000. 704 pp. isbn: 0072380322. url: https://www.ebook.de/de/product/2731239/behzad_razavi_design_of_analog_cmos_integrated_circuits.html.
- [220] S.-G. Lee and J.-K. Choi. 'Current-reuse bleeding mixer'. In: *Electronics Letters* 36.8 (2000), p. 696. doi: [10.1049/e1:20000556](https://doi.org/10.1049/e1:20000556).
- [221] M. Vigilante and P. Reynaert. 'On the Design of Wideband Transformer-Based Fourth Order Matching Networks for E-Band Receivers in 28-nm CMOS'. In: *IEEE Journal of Solid-State Circuits* 52.8 (Aug. 2017), pp. 2071–2082. doi: [10.1109/jssc.2017.2690864](https://doi.org/10.1109/jssc.2017.2690864).
- [222] M. Babaie, F.-W. Kuo, H.-N. R. Chen, L.-C. Cho, C.-P. Jou, F.-L. Hsueh, M. Shahmohammadi, and R. B. Staszewski. 'A Fully Integrated Bluetooth Low-Energy Transmitter in 28 nm CMOS With 36% System Efficiency at 3 dBm'. In: *IEEE Journal of Solid-State Circuits* 51.7 (July 2016), pp. 1547–1565. doi: [10.1109/jssc.2016.2551738](https://doi.org/10.1109/jssc.2016.2551738).

- [223] S. C. Blaakmeer, E. A. M. Klumperink, D. M. W. Leenaerts, and B. Nauta. 'Wideband Balun-LNA With Simultaneous Output Balancing, Noise-Canceling and Distortion-Canceling'. In: *IEEE Journal of Solid-State Circuits* 43.6 (June 2008), pp. 1341–1350. doi: [10.1109/jssc.2008.922736](https://doi.org/10.1109/jssc.2008.922736).
- [224] B. Sell, B. Bigwood, S. Cha, Z. Chen, P. Dhage, P. Fan, M. Giraud-Carrier, A. Kar, E. Karl, C.-J. Ku, R. Kumar, T. Lajoie, H.-J. Lee, G. Liu, S. Liu, Y. Ma, S. Mudanai, L. Nguyen, L. Paulson, K. Phoa, K. Pierce, A. Roy, R. Russell, J. Sandford, J. Stoeger, N. Stojanovic, A. Sultana, J. Waldemer, J. Wan, W. Xu, D. Young, J. Zhang, Y. Zhang, and P. Bai. '22FFL: A high performance and ultra low power FinFET technology for mobile and RF applications'. In: *2017 IEEE International Electron Devices Meeting (IEDM)*. IEEE, Dec. 2017. doi: [10.1109/iedm.2017.8268475](https://doi.org/10.1109/iedm.2017.8268475).
- [225] X. Xue, B. Patra, J. P. G. van Dijk, N. Samkharadze, S. Subramanian, A. Corna, B. P. Wuetz, C. Jeon, F. Sheikh, E. Juarez-Hernandez, B. P. Esparza, H. Rampurawala, B. Carlton, S. Ravikumar, C. Nieva, S. Kim, H.-J. Lee, A. Sammak, G. Scappucci, M. Veldhorst, F. Sebastiano, M. Babaie, S. Pellerano, E. Charbon, and L. M. K. Vandersypen. 'CMOS-based cryogenic control of silicon quantum circuits'. In: *Nature* 593.7858 (May 2021), pp. 205–210. doi: [10.1038/s41586-021-03469-4](https://doi.org/10.1038/s41586-021-03469-4).
- [226] L. M. K. Vandersypen and M. A. Eriksson. 'Quantum computing with semiconductor spins'. In: *Physics Today* 72.8 (Aug. 2019), pp. 38–45. doi: [10.1063/pt.3.4270](https://doi.org/10.1063/pt.3.4270).
- [227] M. Pioro-Ladrière, T. Obata, Y. Tokura, Y.-S. Shin, T. Kubo, K. Yoshida, T. Taniyama, and S. Tarucha. 'Electrically driven single-electron spin resonance in a slanting Zeeman field'. In: *Nature Physics* 4.10 (Aug. 2008), pp. 776–779. doi: [10.1038/nphys1053](https://doi.org/10.1038/nphys1053).
- [228] D. M. Zajac, A. J. Sigillito, M. Russ, F. Borjans, J. M. Taylor, G. Burkard, and J. R. Petta. 'Resonantly driven CNOT gate for electron spins'. In: *Science* 359.6374 (Dec. 2017), pp. 439–442. doi: [10.1126/science.aao5965](https://doi.org/10.1126/science.aao5965).
- [229] W. Huang, C. H. Yang, K. W. Chan, T. Tanttu, B. Hensen, R. C. C. Leon, M. A. Fogarty, J. C. C. Hwang, F. E. Hudson, K. M. Itoh, A. Morello, A. Laucht, and A. S. Dzurak. 'Fidelity benchmarks for two-qubit gates in silicon'. In: *Nature* 569.7757 (May 2019), pp. 532–536. doi: [10.1038/s41586-019-1197-0](https://doi.org/10.1038/s41586-019-1197-0).
- [230] X. Xue, B. D'Anjou, T. F. Watson, D. R. Ward, D. E. Savage, M. G. Lagally, M. Friesen, S. N. Coppersmith, M. A. Eriksson, W. A. Coish, and L. M. K. Vandersypen. 'Repetitive Quantum Nondemolition Measurement and Soft Decoding of a Silicon Spin Qubit'. In: *Physical Review X* 10.2 (Apr. 2020). doi: [10.1103/physrevx.10.021006](https://doi.org/10.1103/physrevx.10.021006).
- [231] M. Reed. 'Entanglement and Quantum Error Correction with Superconducting Qubits'. In: (Nov. 26, 2013). arXiv: <http://arxiv.org/abs/1311.6759v1> [quant-ph].

- [232] J. B. Altepeter, E. R. Jeffrey, and P. G. Kwiat. 'Photonic state tomography'. In: *Advances in Atomic, Molecular, and Optical Physics* 52 (2005), pp. 105–159.
- [233] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. 'Randomized benchmarking of quantum gates'. In: *Physical Review A* 77.1 (Jan. 2008). doi: [10.1103/physreva.77.012307](https://doi.org/10.1103/physreva.77.012307).
- [234] E. Magesan, J. M. Gambetta, and J. Emerson. 'Characterizing quantum gates via randomized benchmarking'. In: *Physical Review A* 85.4 (Apr. 2012). doi: [10.1103/physreva.85.042311](https://doi.org/10.1103/physreva.85.042311).
- [235] 'Rapid solution of problems by quantum computation'. In: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 439.1907 (Dec. 1992), pp. 553–558. doi: [10.1098/rspa.1992.0167](https://doi.org/10.1098/rspa.1992.0167).
- [236] J. Preskill. 'Quantum Computing in the NISQ era and beyond'. In: *Quantum* 2 (Aug. 2018), p. 79. doi: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79).
- [237] M. Gonzalez-Zalba, S. Barraud, A. Ferguson, and A. Betz. 'Probing the limits of gate-based charge sensing'. In: *Nature communications* 6 (2015), p. 6084.
- [238] H. Homulle and E. Charbon. 'Cryogenic low-dropout voltage regulators for stable low-temperature electronics'. In: *Cryogenics* 95 (Oct. 2018), pp. 11–17. doi: [10.1016/j.cryogenics.2018.08.006](https://doi.org/10.1016/j.cryogenics.2018.08.006).
- [239] I. Malo-Gomez, J. Gallego-Puyol, C. Diez-Gonzalez, I. Lopez-Fernandez, and C. Briso-Rodriguez. 'Cryogenic Hybrid Coupler for Ultra-Low-Noise Radio Astronomy Balanced Amplifiers'. In: *IEEE Transactions on Microwave Theory and Techniques* 57.12 (Dec. 2009), pp. 3239–3245. doi: [10.1109/tmmt.2009.2033874](https://doi.org/10.1109/tmmt.2009.2033874).
- [240] R. Schouten. *QT designed instrumentation - matrixrack*. url: <http://qtwork.tudelft.nl/~schouten/matrix/index-matrix.htm> (visited on 03/11/2020).
- [241] J. van Dijk, P. t Hart, G. Kiene, R. Overwater, P. Padalia, J. van Staveren, M. Babaie, A. Vladimirescu, E. Charbon, and F. Sebastiano. 'Cryo-CMOS for Analog/Mixed-Signal Circuits and Systems'. In: *2020 IEEE Custom Integrated Circuits Conference (CICC)*. IEEE, Mar. 2020. doi: [10.1109/cicc48029.2020.9075882](https://doi.org/10.1109/cicc48029.2020.9075882).
- [242] F. Sebastiano, J. van Dijk, P. t Hart, B. Patra, J. van Staveren, X. Xue, C. Almudever, G. Scappucci, M. Veldhorst, L. Vandersypen, A. Vladimirescu, S. Pellerano, M. Babaie, and E. Charbon. 'Cryo-CMOS Interfaces for Large-Scale Quantum Computers'. In: *2020 IEEE International Electron Devices Meeting (IEDM)*. IEEE, Dec. 2020. doi: [10.1109/iedm13553.2020.9372075](https://doi.org/10.1109/iedm13553.2020.9372075).

- [243] M. Mehrpoo, M. Babaie, B. Patra, J. Gong, J. P. G. van Dijk, H. Homulle, G. Kiene, A. Vladimirescu, F. Sebastiano, and E. Charbon. 'Benefits and Challenges of Designing Cryogenic CMOS RF Circuits for Quantum Computers'. In: *2019 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE. IEEE, May 2019, pp. 1–5. doi: [10.1109/iscas.2019.8702452](https://doi.org/10.1109/iscas.2019.8702452).
- [244] F. Sebastiano, H. A. Homulle, J. P. G. van Dijk, R. M. Incandela, B. Patra, M. Mehrpoo, M. Babaie, A. Vladimirescu, and E. Charbon. 'Cryogenic CMOS interfaces for quantum devices'. In: *2017 7th IEEE International Workshop on Advances in Sensors and Interfaces (IWASI)*. IEEE. IEEE, June 2017, pp. 59–62. doi: [10.1109/iwasi.2017.7974215](https://doi.org/10.1109/iwasi.2017.7974215).

Acknowledgements

Fabio Sebastiano, I'm very grateful for all the support and motivations you have given me the last few years. I also want to thank you for having given me the opportunity to work on quantum computing. I have learned a lot from you, both on a technical level, and on other skills such as writing and presenting. I also want to thank you for making the CoolGroup a great and fun place to work in, and for organizing the outings, drinks and BBQs we had.

Edoardo Charbon, I'd also like to thank you for having given me this opportunity, and supporting my PhD. I remember that, while you were very busy, you even managed to spend time on the airplane reviewing one of my papers. Of course I also want to thank you for having given me this opportunity and for the various outings (I will always dearly remember the pizza party with Korean pizzas and fussball), for the many positive feedbacks, and mostly for ensuring that my presentations were always crisp!

Masoud Babaie, although you were not my supervisor, at times it felt like you were, as you were closely involved in most of the work I have done. I have learned a lot from you, and enjoyed working with you. Andrei Vladimirescu, I enjoyed working with you and your enthusiasm for SPINE, as well as your knowledge of good wines and restaurants anywhere in the world.

Bishnu Patra, it goes without saying that I'm very grateful for the long collaboration that we had during our PhDs, working on Horse Ridge, all the way from system design to measurement with qubits. But obviously I'm grateful for much more than that. While driving to Amsterdam together, playing soccer table, ping pong, supporting each other through the many hard times etc. I feel that we became close friends. You are a very smart RF designer, always calm, and uttermost handy in getting things done.

Andrea Ruffino, I had a lot of fun with you, both in Delft and whenever we met on one of our trips. You have great views on life, amazing opinions and I like that you speak your mind. I will also never forget the amazing car seating arrangements we designed for you. Pascal 't Hart, you have brought a lot of fun to the CoolGroup, and created a nice little museum of electronics to which we all gladly contribute. Gerd Kiene, you were a great office mate, ADC designer and cook. I also really enjoyed the many times we went for some exercise. While I was never into running at all, you guys managed to turn me into a runner! Jiang Gong, you are a funny, and very smart guy, capable of great designs, keep that up! Lotte Geck, I had a great time collaborating with you, and in general it was very nice to have you in the group and join our runs. Lizzy Hatfield, I enjoyed our technical conversations, but even more so the non-technical ones, and your love for the US. You single-handedly made the phrase "you can't spell cryogenic without cry" both the best and worst catchphrase in the group! All other group members,

former group members, and guest group members, Milad, Bagas, Job, Luc, Ramon, Pinakin, Allesandro, Rikardur (and I hope I didn't forget anyone else), you all made the group fun to work in, and the group meetings more interesting to attend.

Rosario Incandela, Bahador Valizadeh Pasha, Lin Song, and Harald Homulle, you were all part of the original CoolGroup, and made me feel really welcome in the group. I often fondly think back to the time I was sharing the office with Bishnu, Bahador and Rosario on the 14th floor, and we played futsal on an almost daily basis. Bahador, I think you still hold the record for fastest cannonball shots with futsal! Rosario, you made our trips to China, India, and the many other places we visited a whole lot more enjoyable and fun! Harald, you always truly inspired me with your amazing videos and 3D animations.

Augusto Carimatto, Ashish Sachdeva, and Preethi Padmanabhan, you are all very dear friends to me, and have all helped me through some hard times. Augusto, I don't even know where to start. You have supported me countless times, and we helped each other push to finish our PhDs. You have made me laugh countless times at your crazy jokes, but we also had many deep conversations. You motivated me to start running, and I have enjoyed our many runs and dinners together. You even got me interested in chess! Preethi, I enjoyed all the time we got to spend together before you moved to Switzerland. We have had some great times together, and I really enjoyed the many coffees and good Indian meals we shared. Ashish, you have made me laugh a lot, with all your jokes and amazing actions at the soccer table. It has always been fun and relaxing to talk with you.

I would also like to thank all of the other members of Edoardo's imaging group that stayed in Delft at some point or another during my PhD; Ting, Chao, Esteban (you have tested all of our patience with soccer table, but you are the man to get good steaks for a BBQ!), Michel, Siddharta, Andrada, Arin, Francesco, Pengfei, Scott, Augusto Ximenes, and MJ. We had great pizza nights, several trips and drinks together. While you were not working on quantum computing, it truly felt like we were one group!

Moving across the campus to our friends in physics; Andrea Corna, you were a lot of fun to work with and have taught us a lot about cryogenics, and helped us a lot in getting the fridge up and running for Horse Ridge. Xiao Xue and Nodar Samkharadze, it was great working with you on Horse Ridge. I have learned a lot of practical stuff from you, and have gotten great respect for all physicists that manage to build a quantum processor or operate it; without the two of you we would not have gotten this far! Jelmer Boter and JP Dehollain, I enjoyed our many conversations and brainstorming sessions on future quantum processors with integrated electronics. It was nice to learn from you about the physics, and teach you in return about electronics. Alice Xu, it was fun sharing a lab and fridge with you. You always had a smile on your face that cheered us up. Erika Kawakami, I have learned so much from you about quantum physics and qubit control in my first few months that have helped me throughout my entire PhD, thank you.

Lieven Vandersypen, you were very supportive throughout my PhD, and I'm very grateful for that. You have patiently tried to teach us about qubit control for weeks on, provided great feedback on our papers, and enabled us to test our chip

with qubits. Similarly, Menno Veldhorst, you have taught me a lot about qubit control, and especially with the perspective of scaling. You have repeatedly positively challenged our ideas and kept our minds sharp. Not to be forgotten, Richard Versluis and Raymond Schouten, we have had many great conversations regarding implementing electronics for quantum computing. I'd also like to thank the many technicians that have helped us, especially the Raymond's in providing whatever we needed, and Olaf for helping us out whenever we had problems with our dilution fridge.

I'd also like to extend my gratitude to the many people from Intel that have helped turn Horse Ridge into a success, especially Stefano Pellerano and Sushil Subramanian. It was really great and easy to work with you, but perhaps even better to share some drinks with you to celebrate our successes. I also want to thank the people working in the Amsterdam office, who made our time in Amsterdam a lot more fun, comparing Bishnu and me to Leonard and Sheldon from the big bang theory, and making us the perfect independent judges on their bake-offs. I'd also like to thank Jim Clarke and Jim Held, for the interesting discussions we had, and supporting our project.

I'm also thankful for the many engaging discussions I've had with Koen Bertels, Carmina Almudever, Xiang Fu, Lingling Lao, and others in the Quantum Computer Architectures Group. Moreover, I'd like to thank Dennis Kaandorp. It was fun discussing my master thesis work with you, and seeing you improve upon it!

In general, Pascal, Job, Luc, Ramon, and the many people from Qutech, you have kept my knowledge of the dutch language alive, something that I almost forgot working in such a large internationally diverse group! It was great to practice my mother tongue with you once in a while.

Finally, I'd like to thank my family and friends that have supported me for so many years and have helped me through this long journey. A very special thanks to Renita Danarianti, who has made my life better in so many ways since the day I met her, and who has supported and motivated me countless times: kamu terbaik!

Curriculum Vitæ

Jeroen Petrus Gerardus VAN DIJK

02-08-1992 Born in Olst, The Netherlands.

Education

- 2004–2010 Atheneum, Cum Laude
Christelijke Scholengemeenschap De Lage Waard,
Papendrecht, The Netherlands
- 2010–2013 B.Sc. Electrical Engineering, Cum Laude
Delft University of Technology, Delft, The Netherlands
- 2013–2016 M.Sc. Electrical Engineering, Microelectronics, Cum Laude
Delft University of Technology, Delft, The Netherlands
Thesis: An On-Chip Noise Thermometer
Promotor: Prof. dr. K.A.A. Makinwa
Supervisor: Dr. F. Sebastiano

Awards

- 2018 DATE Best IP Award
2021 ISSCC 2020 Evening Session Award
2021 ISSCC 2020 Jan Van Vessel Award for Outstanding European Paper

List of Publications

Journal papers

- X. Xue, B. Patra, J. P. G. van Dijk, N. Samkharadze, S. Subramanian, A. Corna, B. P. Wuetz, C. Jeon, F. Sheikh, E. Juarez-Hernandez, B. P. Esparza, H. Rampurawala, B. Carlton, S. Ravikumar, C. Nieva, S. Kim, H.-J. Lee, A. Sammak, G. Scappucci, M. Veldhorst, F. Sebastiano, M. Babaie, S. Pellerano, E. Charbon, and L. M. K. Vandersypen. 'CMOS-based cryogenic control of silicon quantum circuits'. In: *Nature* 593.7858 (May 2021), pp. 205–210. doi: [10.1038/s41586-021-03469-4](https://doi.org/10.1038/s41586-021-03469-4) (equal contribution)
- J. P. G. V. Dijk, B. Patra, S. Subramanian, X. Xue, N. Samkharadze, A. Corna, C. Jeon, F. Sheikh, E. Juarez-Hernandez, B. P. Esparza, H. Rampurawala, B. R. Carlton, S. Ravikumar, C. Nieva, S. Kim, H.-J. Lee, A. Sammak, G. Scappucci, M. Veldhorst, L. M. K. Vandersypen, E. Charbon, S. Pellerano, M. Babaie, and F. Sebastiano. 'A Scalable Cryo-CMOS Controller for the Wideband Frequency-Multiplexed Control of Spin Qubits and Transmons'. In: *IEEE Journal of Solid-State Circuits* 55.11 (Nov. 2020), pp. 2930–2946. doi: [10.1109/jssc.2020.3024678](https://doi.org/10.1109/jssc.2020.3024678) (equal contribution)
- J. P. G. van Dijk, B. Patra, S. Pellerano, E. Charbon, F. Sebastiano, and M. Babaie. 'Designing a DDS-Based SoC for High-Fidelity Multi-Qubit Control'. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 67.12 (Dec. 2020), pp. 5380–5393. doi: [10.1109/tcsi.2020.3019413](https://doi.org/10.1109/tcsi.2020.3019413) (equal contribution)
- J. P. G. van Dijk, E. Kawakami, R. Schouten, M. Veldhorst, L. Vandersypen, M. Babaie, E. Charbon, and F. Sebastiano. 'Impact of Classical Control Electronics on Qubit Fidelity'. In: *Physical Review Applied* 12.4 (Oct. 2019), p. 044054. doi: [10.1103/physrevapplied.12.044054](https://doi.org/10.1103/physrevapplied.12.044054)
- J. P. G. van Dijk, E. Charbon, and F. Sebastiano. 'The electronic interface for quantum processors'. In: *Microprocessors and Microsystems* 66 (Apr. 2019), pp. 90–101. doi: [10.1016/j.micpro.2019.02.004](https://doi.org/10.1016/j.micpro.2019.02.004)
- L. Geck, J. P. G. van Dijk, F. Sebastiano, S. van Waasen, and S. Heinen. 'Influence of control circuit noise on quantum computation accuracy with GaAs qubits'. (in preparation)
- J. M. Boter, J. P. Dehollain, J. P. G. van Dijk, Y. Xu, T. Hensgens, R. Versluis, H. W. L. Naus, J. S. Clarke, M. Veldhorst, F. Sebastiano, and L. M. K. Vandersypen. 'The *spider-web* array — A sparse spin qubit array'. (in preparation)

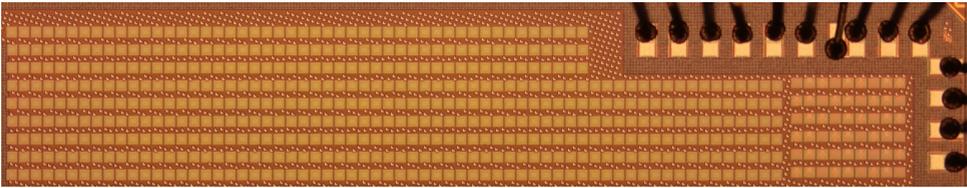
- B. Patra, R. M. Incandela, J. P. G. van Dijk, H. A. R. Homulle, L. Song, M. Shahmohammadi, R. B. Staszewski, A. Vladimirescu, M. Babaie, F. Sebastiano, and E. Charbon. 'Cryo-CMOS Circuits and Systems for Quantum Computing Applications'. In: *IEEE Journal of Solid-State Circuits* 53.1 (Jan. 2018), pp. 309–321. issn: 0018-9200. doi: [10.1109/jssc.2017.2737549](https://doi.org/10.1109/jssc.2017.2737549)

Conference papers

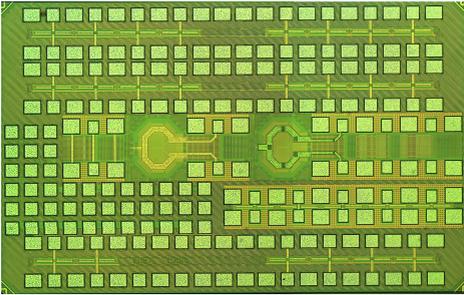
- B. Patra, J. P. G. van Dijk, S. Subramanian, A. Corna, X. Xue, C. Jeon, F. Sheikh, E. Juarez-Hernandez, B. P. Esparza, H. Rampurawala, B. Carlton, N. Samkharadze, S. Ravikumar, C. Nieva, S. Kim, H.-J. Lee, A. Sammak, G. Scappucci, M. Veldhorst, L. M. K. Vandersypen, M. Babaie, F. Sebastiano, E. Charbon, and S. Pellerano. '19.1 A Scalable Cryo-CMOS 2-to-20GHz Digitally Intensive Controller for 4×32 Frequency Multiplexed Spin Qubits/Transmons in 22nm FinFET Technology for Quantum Computers'. In: *2020 IEEE International Solid-State Circuits Conference - (ISSCC)*. IEEE, Feb. 2020. doi: [10.1109/isscc19947.2020.9063109](https://doi.org/10.1109/isscc19947.2020.9063109) (equal contribution)
- J. P. G. van Dijk, A. Vladimirescu, M. Babaie, E. Charbon, and F. Sebastiano. 'SPINE (SPIN Emulator) - A Quantum-Electronics Interface Simulator'. In: *2019 IEEE 8th International Workshop on Advances in Sensors and Interfaces (IWASI)*. IEEE. IEEE, June 2019, pp. 23–28. doi: [10.1109/iwasi.2019.8791334](https://doi.org/10.1109/iwasi.2019.8791334)
- J. P. G. van Dijk, A. Vladimirescu, M. Babaie, E. Charbon, and F. Sebastiano. 'A co-design methodology for scalable quantum processors and their classical electronic interface'. In: *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE. IEEE, Mar. 2018, pp. 573–576. doi: [10.23919/date.2018.8342072](https://doi.org/10.23919/date.2018.8342072)
- J. van Dijk, P. t Hart, G. Kiene, R. Overwater, P. Padalia, J. van Staveren, M. Babaie, A. Vladimirescu, E. Charbon, and F. Sebastiano. 'Cryo-CMOS for Analog/Mixed-Signal Circuits and Systems'. In: *2020 IEEE Custom Integrated Circuits Conference (CICC)*. IEEE, Mar. 2020. doi: [10.1109/cicc48029.2020.9075882](https://doi.org/10.1109/cicc48029.2020.9075882)
- F. Sebastiano, J. van Dijk, P. t Hart, B. Patra, J. van Staveren, X. Xue, C. Almudever, G. Scappucci, M. Veldhorst, L. Vandersypen, A. Vladimirescu, S. Pellerano, M. Babaie, and E. Charbon. 'Cryo-CMOS Interfaces for Large-Scale Quantum Computers'. In: *2020 IEEE International Electron Devices Meeting (IEDM)*. IEEE, Dec. 2020. doi: [10.1109/iedm13553.2020.9372075](https://doi.org/10.1109/iedm13553.2020.9372075)
- J. M. Boter, J. P. Dehollain, J. P. G. van Dijk, T. Hensgens, R. Versluis, J. S. Clarke, M. Veldhorst, F. Sebastiano, and L. M. K. Vandersypen. 'A sparse spin qubit array with integrated control electronics'. In: *2019 IEEE International Electron Devices Meeting (IEDM)*. IEEE, Dec. 2019. doi: [10.1109/iedm19573.2019.8993570](https://doi.org/10.1109/iedm19573.2019.8993570)

- M. Mehrpoo, M. Babaie, B. Patra, J. Gong, J. P. G. van Dijk, H. Homulle, G. Kiene, A. Vladimirescu, F. Sebastiano, and E. Charbon. 'Benefits and Challenges of Designing Cryogenic CMOS RF Circuits for Quantum Computers'. In: *2019 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE. IEEE, May 2019, pp. 1–5. doi: [10.1109/iscas.2019.8702452](https://doi.org/10.1109/iscas.2019.8702452)
- P. A. ' Hart, J. P. G. van Dijk, M. Babaie, E. Charbon, A. Vladimircscu, and F. Sebastiano. 'Characterization and Model Validation of Mismatch in Nanometer CMOS at Cryogenic Temperatures'. In: *2018 48th European Solid-State Device Research Conference (ESSDERC)*. IEEE. IEEE, Sept. 2018, pp. 246–249. doi: [10.1109/essderc.2018.8486859](https://doi.org/10.1109/essderc.2018.8486859)
- E. Charbon, F. Sebastiano, M. Babaie, A. Vladimirescu, M. Shahmohammadi, R. B. Staszewski, H. A. Homulle, B. Patra, J. P. G. van Dijk, R. M. Incandela, L. Song, and B. Valizadehpasha. '15.5 Cryo-CMOS circuits and systems for scalable quantum computing'. In: *2017 IEEE International Solid-State Circuits Conference (ISSCC)*. IEEE. IEEE, Feb. 2017, pp. 264–265. doi: [10.1109/isscc.2017.7870362](https://doi.org/10.1109/isscc.2017.7870362)
- F. Sebastiano, H. Homulle, B. Patra, R. Incandela, J. P. G. van Dijk, L. Song, M. Babaie, A. Vladimirescu, and E. Charbon. 'Cryo-CMOS Electronic Control for Scalable Quantum Computing'. In: *Proceedings of the 54th Annual Design Automation Conference 2017 on - DAC '17*. DAC '17. Austin, TX, USA: ACM Press, 2017. isbn: 978-1-4503-4927-7. doi: [10.1145/3061639.3072948](https://doi.org/10.1145/3061639.3072948)
- F. Sebastiano, H. A. Homulle, J. P. G. van Dijk, R. M. Incandela, B. Patra, M. Mehrpoo, M. Babaie, A. Vladimirescu, and E. Charbon. 'Cryogenic CMOS interfaces for quantum devices'. In: *2017 7th IEEE International Workshop on Advances in Sensors and Interfaces (IWASI)*. IEEE. IEEE, June 2017, pp. 59–62. doi: [10.1109/iwasi.2017.7974215](https://doi.org/10.1109/iwasi.2017.7974215)

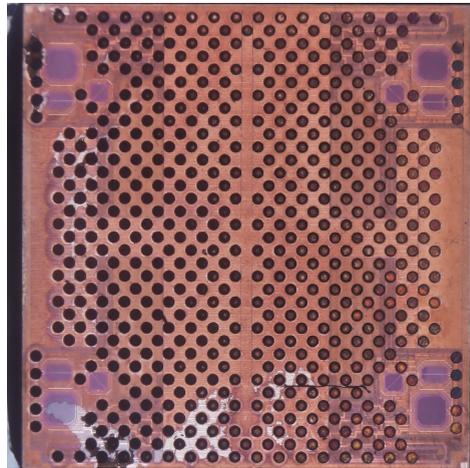
Chip Gallery



(a) Test structures for mismatch characterization at cryogenic temperatures, implemented in ST 40nm technology [ESSDERC2018].



(b) Test structures for RF-characterization at cryogenic temperatures, implemented in TSMC 40nm technology (with B. Patra and H. Homulle).



(c) Horse Ridge, implemented in Intel 22nm FinFET (22FFL) technology [ISSCC2020] (with B. Patra and Intel Corporation).