TU Delft

Delft University of Technology

Cryogenic mixed-signal readout electronics for quantum computers

Kiene, G.

DOI
10.4233/uuid:4d31ae48-6d08-40ad-86f9-0097e641d63a

Publication date
2024

Document Version
Final published version

Citation (APA)
Kiene, G. (2024). *Cryogenic mixed-signal readout electronics for quantum computers*. [Dissertation (TU Delft), Delft University of Technology]. https://doi.org/10.4233/uuid:4d31ae48-6d08-40ad-86f9-0097e641d63a

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.

This work is downloaded from Delft University of Technology.
For technical reasons the number of authors shown on this cover page is limited to a maximum of 10.

Cryogenic mixed-signal readout electronics for quantum computers

Cryogenic mixed-signal readout electronics for quantum computers

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology,
by the authority of the Rector Magnificus, prof. dr. ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates
to be defended publicly on
Monday 25 November 2024 at 10.00 o'clock

by

Gerd KIENE

Master of Science in Physics,
Universität Heidelberg, Germany,
born in Nienburg (Weser), Germany.

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector Magnificus,	Chairperson
Dr. F. Sebastiano,	Delft University of Technology, promotor
Dr. M. Babaie,	Delft University of Technology, promotor
Independent members:	
Prof. dr. H. Blum,	RWTH Aachen, Germany
Dr. ir. K. Bult,	Ethernovia B.V.
Prof. dr. G. Ferrari,	Politecnico di Milano, Italy
Prof. dr. ir. S. Hamdioui,	Delft University of Technology
Prof. dr. ir. L. Vandersypen,	Delft University of Technology

Keywords:	Cryogenic electronics, Cryo-CMOS, SAR ADC, dynamic amplifier, low-frequency noise, device characterization, SET readout
Printed by:	Gildeprint
Cover:	Photo by Gerd Kiene, design by Helena Stern

ISBN 978-94-6384-674-5

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

Contents

1

Introduction

1.1 Why quantum computing

Classical computing power has exploded in the last decades, enabled by tremendous advancements in silicon processing technology. For a large part of this time, the number of transistors that could viably be integrated per chip has doubled every two years [1]. To a certain degree, these improvements continue today [2]. However, further advancements in processing technology are increasingly technologically challenging and expensive due to approaching physical limits. Given the slowing growth, a number of complex computational problems of great practical interest are expected to continue to evade our classical computational abilities. For a subset of these challenging problems, quantum computers are expected to provide significant benefits [3].

One such application where employing quantum computers could lead to a large speedup compared to the use of classical computing hardware is performing quantum simulations [4, 5]. In quantum simulations, a well-controlled quantum system, like for example a quantum computer, is used to explore the properties of another, less accessible, system. Concretely, this could be e.g. applied to simulations of chemical reactions [6].

Benefits of quantum computers are also expected in running certain specialized algorithms known as quantum algorithms. Notable examples of such algorithms include Shor's algorithm for prime factorization [7] and Grover's algorithm for unstructured database search [8]. These algorithms offer an almost exponential and quadratic speedup compared to their classical counterparts, respectively. The faster execution is a result of the use of quantum entanglement and quantum superposition to operate in parallel on an input state - leading to large speed gains if combined with algorithms that harness this parallelism effectively [3]. The operations in quantum computers are performed on quantum bits, commonly known as qubits.

Currently, however, these expected benefits of quantum computers can not be accessed given the lack of two main attributes of present-day machines: limited accuracy in the operations performed on the qubits and limited numbers of qubits. The accuracy of qubit operations, like gates and readout, is generally referred to as fidelity. For example, the best-achieved control fidelities for spin-qubits today are on the order of 99.9% [9], not sufficient to allow for reliable direct execution of algorithms. The reasons for the limited

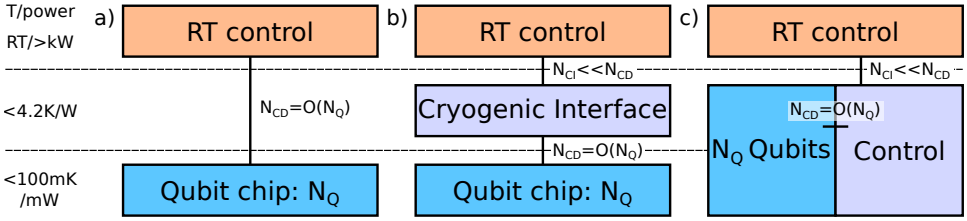

Figure 1.1: a) Schematic of traditional quantum computing setup, b) with a cryogenic electrical interface, c) with co-integration of the cryogenic interface. N_{CD}, N_{CI}, and N_Q are the number of direct control lines, number of lines to the interface, and number of qubits, respectively.

fidelity can be found in noise and interference corrupting the fragile quantum information stored in the qubit. To address the problem of limited fidelity, quantum error correction schemes, such as surface codes [10], have been proposed. In these error correction schemes, however, multiple *physical qubits* encode a single *logical qubit*. As operations can only be performed with the logical qubits, this results in much greater demands on the total number of qubits with certain classes of problems possibly needing millions of physical qubits [11], while the state-of-the-art is currently mere hundreds [12].

1.2 An electrical interface for quantum computers

Due to the fragile nature of the qubits, many of the proposed qubit platforms require operation at deep cryogenic, often below 100 mK, temperatures. Reaching these very low temperatures is possible with dilution refrigerators, such as, for example, [13]. The cooling power currently available at these very low temperatures is typically in the order of mW. This forces the interfacing electronic instrumentation to be operated at room temperature (RT) where cooling power is abundant, see figure 1.1 a). The electronics used in today's experimental setups are often general-purpose instruments, but an increasing number of RT interfaces targeted at the qubits specifically are available [14, 15].

In scaling these systems to higher qubit numbers, however, a wiring and complexity bottleneck exists due to the interconnect [16]: if requiring one or even multiple control lines per qubit, the total number of direct control lines (N_{CD}) scales with the number of qubits (N_Q) as $N_{CD} = O(N_Q)$. This corresponds to millions of control lines for a fully scaled quantum computer. These lines have to span orders of magnitude in temperature, be several meters long to integrate with the dilution refrigerator and carry highly sensitive signals. This is expected to lead to severe wiring problems, affecting system reliability, which is already commonplace in today's much smaller qubit systems.

To address these issues, cryogenic electronic interfaces for quantum computers have been proposed, e.g., in [17, 18]. As the cooling power at deep-cryogenic temperatures is severely limited, most of these interfaces operate between 1 K and 4.2 K. At these elevated temperatures, a cooling power of on the order of a few watt is available in common dilution refrigerators. This leads to the situation in figure 1.1 b), with potentially much reduced wiring from RT, but the same number of wires still interfacing the quantum chip from the intermediate cryogenic interface. So while these intermediate cryogenic interfaces address part of the problem, a difficult part of the bottleneck still exists in the interface

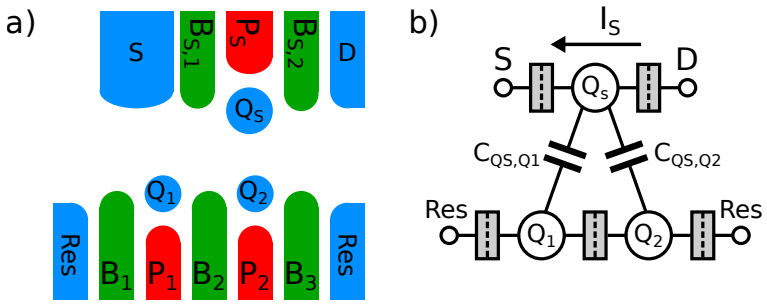

Figure 1.2: a) double quantum dot layout sketch showing the quantum dots (Q_X), the reservoirs (S, D, Res) as well as the controlling plunger and barrier gates (P_X and B_X), b) schematic representation, highlighting the signal current I_S as well as the coupling between the quantum dots.

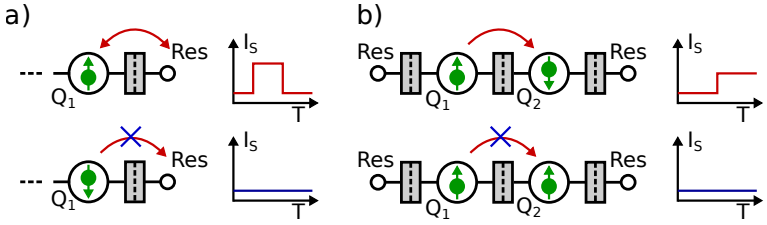

Figure 1.3: Sketch of electron movement for a) energy-selective readout (figure adopted from [27]), b) Pauli-spin-blockade readout. Show on the side for each case is the corresponding sensor response, which is smaller in case b) due to the dipole nature of the signal.

from cryogenic to deep-cryogenic temperature.

One of the most promising suggested solutions to this problem is to operate the qubits at an elevated temperature between 100 mK and 4.2 K and integrate the control electronics alongside it [19], shown in figure 1.1 c). The qubits continue to require the same number of control lines, but these are now much easier to realize: they span no (or only a minor) temperature gradient, and their length may be on the order of cm to µm for co-packaging or integration on the same die, respectively. To achieve this co-integration, semiconductor-based spin-qubits are uniquely positioned. This is due to their demonstrated capability of operating at elevated temperatures [20–22] and their compatibility with industrial CMOS manufacturing [23]. While current state-of-the-art setups for spin qubits are comparatively small arrays [24], numerous proposals have been made on how a scaled system of these could be realized, e.g. [25, 26].

1.3 Readout of spin qubits

In figure 1.2 a), an example layout sketch of a typical double quantum spin qubit system is shown. Two quantum dots ($Q_{1/2}$) are coupled to one sensor dot (Q_S), with the dots confined by the barrier voltages (B_x) and their occupation number controlled by the plungers (P_x). The quantum information is encoded in the spins of electrons captured in $Q_{1/2}$. In figure 1.2 b), a schematic representation of the same configuration is shown. More in-depth

information on quantum dots and spins in quantum dots can be found in [28, 29], and additional information on quantum information processing in [3].

Given the minute signal size of the magnetic moment associated with the electron spin, direct detection of the spins in $Q_{1/2}$ is infeasible. To resolve this problem, indirect schemes have been devised to convert the spin information into the movement of charge for which detection schemes exist. The most common strategies for performing this conversion can be divided into two categories: energy-selective readout and Pauli-spin-blockade readout. In tunnel-frequency selective readout [27], the spin is placed in a strong magnetic field, resulting in an energy splitting. It is then pulsed relative to a reservoir such that one spin state has a high tunnel-frequency to the reservoir and the other a low one. If now waiting for an amount of time that makes the high-rate transition very likely to happen and the low-rate transition very unlikely, the spin can be determined by detecting the electron movement. The resulting vacancy is filled by an electron with low spin from the reservoir tunneling back into the now unoccupied dot. The resulting temporal charge displacement corresponds to a full electron charge, as depicted in figure 1.3 a). In Pauli-spin-blockade [30, 31], again in a magnetic field, the Pauli-exclusion principle is used for determining the spin information. If pulsing electrons from different dots together onto a single dot, this transition is energetically disfavored if the spins are parallel. As a result, only if the spins are in the anti-parallel state, there is a movement of charge, shown in figure 1.3 b). This movement generates a dipole-like signal, generally smaller than the full transition observed in tunnel-frequency-based readout.

In either case, however, a minute charge movement needs recording. To pick up this weak signal, highly sensitive sensors need to be used. Examples of such sensors include quantum point contacts (QPC) as used in [27], and most commonly, single electron transistors (SET) [32]. Both sensors convert changes in the surrounding charge configuration to a change in resistance between a pair of terminals. The sensor dot in the example layout in figure 1.2 is an example of an SET. A different detection scheme exists for energy-selective readout by using gate-based sensors [33, 34]. In these, the change in impedance based on an allowed or dis-allowed transition of the electron is recorded on the plunger gates. Regardless of the scenario, a spin dependent impedance forms the interface to the classical readout electronics. Two main approaches to implement the readout electronics can be identified: RF-readout and DC-readout. The subsequent section describes the electronics necessary for both of these interfaces.

1.4 Readout electronics for spin qubits

In figure 1.4, a proposed cryogenic drive and readout interface for a spin-qubit quantum computer is shown, corresponding to a concrete realization of figure 1.1 b) or, if assuming same-temperature operation, c). The driver section comprises microwave signal generators [35], pulse generators [18], and DC biasing [36]. In the following, an emphasis is placed on the readout interface, specifically on the electronics needed for the two above-mentioned main readout strategies, namely DC- and RF-readout.

In the DC-readout of an SET, the sensor is typically directly connected to the readout amplifier, as depicted in the top part of figure 1.4. A commonly used circuit in this context is the transimpedance amplifier (TIA) [37]. A TIA applies a voltage bias to the variable sensor impedance and converts the resulting current into a voltage that can then be processed to

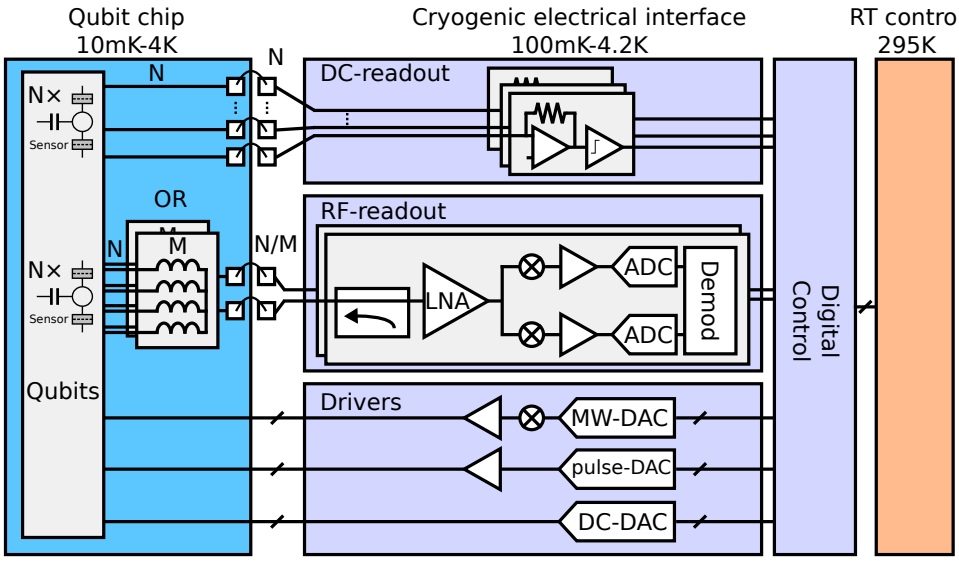

Figure 1.4: Overview over the proposed cryogenic electronic interface for spin qubits, extended from [17]. Both DC- and RF-readout are shown for demonstration. In practical systems, however, only one of them will need to be implemented.

determine the qubit state. The advantages of DC-readout mostly stem from the relative simplicity of the approach: only few components running at speeds not much higher than the bit-rate are required. This results in, under circumstances discussed in chapter 6, very low power consumption. Also, the components carry small layout footprints, making this readout strategy especially relevant for an eventual co-integration. As a disadvantage of DC-readout, one connection between electronics and qubits is needed for each sensor, as shown in figure 1.4 with N qubits requiring N connections. In a two-chip solution, this counter-balances the small layout footprint partially: the space requirements of one chip pad per sensor make the DC-readout interface effectively much larger. It might be possible to partially alleviate this by advanced packaging solutions like flip-chip assembly on a silicon interposer. Also, the performance of DC-readout is very sensitive to parasitic capacitance on the input node, placing tight constraints on the wiring. If the wiring needs to span different temperatures, this results in a much-diminished speed of the readout.

RF-readout requires a more complex setup, see the middle of figure 1.4. For a typical example of a reflectometry readout, e.g. [33], the sensor is placed in a matching network as its impedance typically is far from the interconnecting cables. This network is then excited from the readout circuit via a directional coupler, the reflection being received by a low noise amplifier (LNA) [17, 38, 39]. The amplified signal is then either directly digitized or first down-converted and then digitized as in [18]. The digitization is performed by an analog-to-digital converter (ADC). Advantages of RF-readout include very high readout speeds [40, 41] and the potential to use frequency-multiplexing, reducing the number of connections to the qubit chip [42]. The potential for frequency multiplexing makes RF-readout especially relevant in the medium-term, as fewer connections to deep-cryogenic

temperatures ease the issues in the bottleneck described in figure 1.1 b). This is depicted in figure 1.4 by showing an example with M qubits being multiplexed on a single line, now requiring only N/M connections for N qubits. Major disadvantages are the need for numerous high-speed components, some of which, for example, the inductor in the resonator tank, require significant chip area and are difficult to scale. The power efficiency of the entire chain, limited by the necessary high-speed operations at much higher frequencies than the bit-rate, can also be a challenge for RF-readout compared to DC-readout.

CMOS processes are especially promising for the implementation of interface electronics at deep-cryogenic temperatures, which are termed cryo-CMOS in this context. This is because of their superior transistor scaling, allowing for the implementation of power-effective digital functionality, as well as options for eventual co-integration with qubits [23]. CMOS performance is generally sustained at low temperatures, with both threshold voltage and conductivity rising [43, 44]. Throughout this thesis the impact of the cryogenic environment on circuit design is discussed, especially in chapter 2. Other technologies, like for example SiGe bipolar transistors, offer great performance for certain analog functions, for example in LNAs, and could be used alongside cryo-CMOS.

1.5 State-of-the-Art Readout Electronics

Various suggestions for integrated DC-readout circuits have been reported in literature; examples are: current comparators in [45, 46], a transimpedance amplifier in [37], and a charge-sampling integrator in [47]. These designs are either slow compared to the speed on the order of µs readout times achieved by RF readout, or in the case of [47], similarly power-hungry. Alternative solutions using SiGe technology [48] or HBT-based designs [49] have shown advantageous power consumption at high speed when compared to CMOS realizations, but the designs showed little integration. To aid the development of competitive integrated DC-readout systems, a systematic approach to optimize these systems is delivered in chapter 6 in this thesis.

In RF-readout, much work has been reported on LNAs in CMOS [17, 38, 39] and in SiGe [50], but the reports on suitable wide-band cryogenic ADCs have been scarce. The 7.5-bit cryo-CMOS SAR ADC within the 4 K cryo-CMOS SoC in [18] was tested at sampling speeds up to 400 MS/s as part of a spin-qubit readout system but its power consumption and design details have not been reported. In chapters 2 to 4, efficient high-speed ADC designs specifically targeted for the RF-readout interface are presented.

Efficient solutions in the design of all parts in the cryogenic electronic interface, including the readout circuits, rely heavily on a comprehensive understanding of the device characteristics. The absence of good models necessitates additional design considerations like wider tuning ranges - that require performance compromises - and increased risk for design failures. The circuit designs in this thesis have been performed without access to cryogenic simulation models. To improve the understanding of cryogenic device behavior, DC [51], RF [52] and noise [53, 54] characterization and modeling have been performed. In chapter 5, an extensive low-frequency noise characterization of a bulk-CMOS process is contributed to this. This characterization is the first to include significant statistics and systematic area scaling considerations at cryogenic temperatures. Further, in chapter 3, noise measurements on a comparator give additional evidence about the missing noise-scaling of broad-band noise at cryogenic temperatures.

1.6 Objective and Scope

The overarching objective of this thesis is to contribute to the development of a cryogenic electronic readout interface for silicon spin-based quantum computers. With this context in mind, answers to several questions are attempted in this thesis:

1. How can an efficient cryo-CMOS analog-to-digital converter for RF-readout be realized?
2. Is it possible to preserve the efficiency of such a converter when combining it with a driver?
3. How does the low-frequency noise behavior of CMOS devices evolve when cooling from room temperature to cryogenic temperature?
4. What are the limits to which DC-readout of SET charge sensors can be pushed? And how is it possible to approach these limits in practice?

1.7 Organization of the dissertation

To answer these distinct questions, the thesis is organized as follows. Chapters 2 to 4 present SAR ADC designs: starting with the first high-speed CMOS ADC operating at cryogenic temperature in chapter 2, an analysis of the noise of comparators and ADC efficiency improvements in chapter 3 and finally the inclusion of an efficient integrating driver in chapter 4. Further, contributions to low-frequency noise characterization of transistors at cryogenic temperature are found in chapter 5. Then, considerations about DC-readout for spin qubits are presented in chapter 6. Finally, in chapter 7, the thesis concludes with a summary and future directions.

2

A 1-GS/s 6–8-B Cryo-CMOS SAR ADC for Quantum Computing

This chapter presents a two-times interleaved, loop-unrolled SAR analog-to-digital converter (ADC) operational from 300 down to 4.2 K. The 6–8-bit resolution and the sampling speed up to 1 GS/s are targeted at digitizing the multi-channel frequency-multiplexed input in a spin-qubit reflectometry readout for quantum computing. To optimize the circuit for the altered device behavior at cryogenic temperatures, a modified common-mode switching scheme is adopted as well as a flexible calibration. The design is implemented in 40-nm CMOS technology and achieves 36.2-dB signal to noise and distortion ratio (SNDR) for Nyquist input at 4.2 K while maintaining a Walden figure of merit (FOM_W) of 200 pJ/conv-step (for a 10.8-mW power consumption), including the clock receiver, and 15 pJ/conv-step (for a 0.8-mW power consumption) for just the core ADC. With these specifications, the ADC can support the simultaneous readout of 20 qubit channels with a power consumption of 0.5 mW/qubit, thus advancing toward the full integration of the cryogenic readout for future large-scale quantum processors.

This chapter has been published as *Gerd Kiene, Ramon WJ Overwater, Alessandro Catania, Aishwarya Gunaputi Sreenivasulu, Paolo Bruschi, Edoardo Charbon, Masoud Babaie, Fabio Sebastiano, "A 1-GS/s 6–8-b Cryo-CMOS SAR ADC for Quantum Computing." IEEE Journal of Solid-State Circuits (2023)* [55].

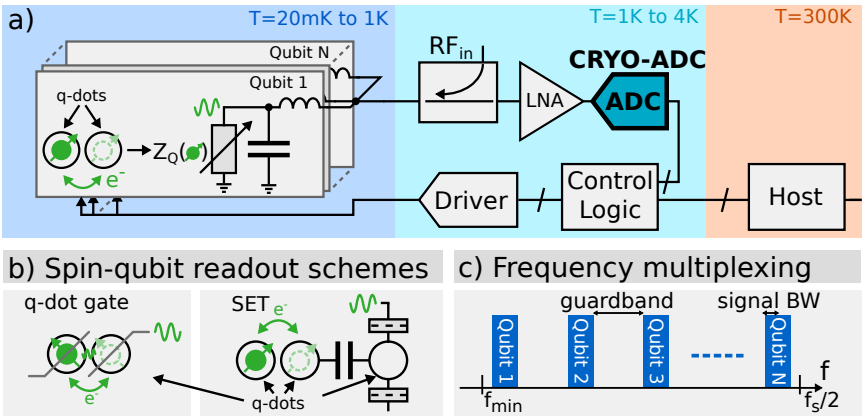

Figure 2.1: The proposed cryo-CMOS ADC within the RF-reflectometry readout for spin qubits.

2.1 Introduction

While RF-readout schemes for spin qubits have been described broadly in chapter 1, here we delve deeper into it for details that are crucial for the ADC designs in this and the following chapters. The task at hand is to read a quantum sensor, for instance the capacitance of an electrode coupled to a quantum dot [33, 34] or the resistance of a single electron transistor (SET) [41], as depicted in figure 2.1 b). For all sensor types, however, a quantum-state-dependent impedance Z_Q is exposed by the sensor. In the RF-readout, this impedance is matched to the 50 Ω cable impedance by an LC network and is probed by an RF excitation. This excitation can be applied via a directional coupler, see figure 2.1 a), or via a transmission measurement as in [34]. The reflected signal is subsequently amplified by a LNA and then digitized by an ADC, either directly or after I/Q mixing. The downsides of such an RF-readout, when compared to the DC-readout, are a higher component count (matching network, directional coupler), and higher power consumption for various sub-blocks, such as the LNA and the ADC. Those disadvantages can be partially mitigated and counterbalanced by frequency multiplexing (figure 2.1 c)): here, not only one but multiple matching networks are placed in parallel, thus allowing for sharing both the RF components and a single cable over multiple qubits, consequently lowering the power consumption per qubit. In the context of RF readout, moving the digitization to cryogenic temperatures is crucial to enable a compact and reliable system: it avoids routing high-frequency sensitive analog signals through the several stages of a cryogenic refrigerator, and potentially allows closing the algorithm execution loop completely inside the cryostat, see figure 2.1.

For the digitization of the input signal, a moderate resolution (≤ 8 bits) is sufficient (see section 2.2), while the sampling speed must be maximized to allow for more frequency-multiplexed qubit channels, thus lowering the readout power per qubit. Specifically, as the system's cooling power is on the order of several W, we require circuits to consume below a few mW per qubit for a system with thousands of qubits.

Among possible candidates for a cryogenic ADC, superconducting ADCs [57, 58] are

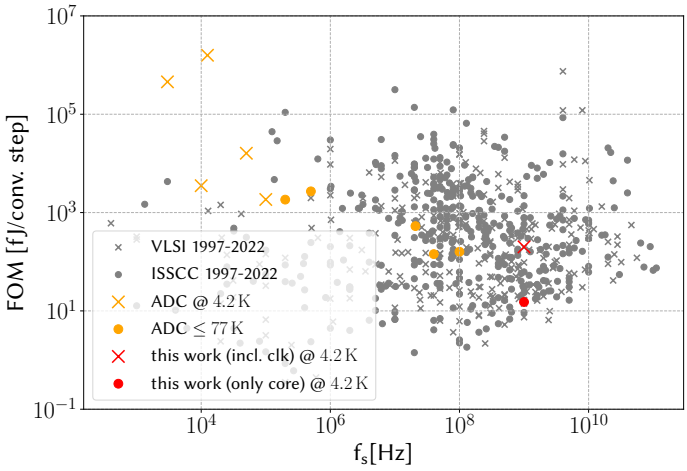

Figure 2.2: Cryo-CMOS ADC survey on basis of [56], showing speed and efficiency of previous RT and cryogenic ADCs.

inherently cryogenic and offer high speed at a low power consumption but are incompatible with the integration in a cryo-CMOS SoC. A cryogenic FPGA-based ADC has demonstrated a sampling rate of 1.2 GS/s and offers great flexibility [59], but its power efficiency is much lower than RT CMOS ADCs, e.g., [60]. Previously reported cryo-CMOS ADCs tested at liquid-helium temperature (4 K) [61–63] showed insufficient sampling speeds and low power efficiency due to the adoption of mature technologies and different application requirements. ADCs tested at liquid nitrogen temperature (77 K) [64–66] are not suitable for below-1-K spin-qubit applications and reach a maximum speed of only 100 MS/s. The 7.5 bit cryo-CMOS SAR ADC within the 4 K cryo-CMOS SoC in [18] was tested at sampling speeds up to 400 MS/s as part of a spin-qubit readout system but its power consumption and design details have not been reported.

To cover the unexplored area in the cryo-CMOS ADC design space shown in figure 2.2, this chapter presents a 4.2 K 1-GS/s ADC with energy efficiency comparable to RT state-of-the-art ADCs and with an expected power consumption of only 0.5 mW per qubit when employed in spin-qubit readout. To achieve this performance, design techniques specifically targeting cryogenic operation and optimization have been employed, including cryogenic-aware comparator optimization and offset-calibration design, ad-hoc CDAC switching scheme, and thick-oxide front-end and clocking, coupled with experimental techniques optimized for cryogenic characterization.

The chapter is organized as follows. In section 2.2, we model the system to derive the specifications for the ADC. The ADC circuit design is described in section 2.3, including the design of peripheral blocks for testing and supply. The ADC testing with emphasis on cryogenic testing is reported in section 2.4 and conclusions are drawn in section 2.5.

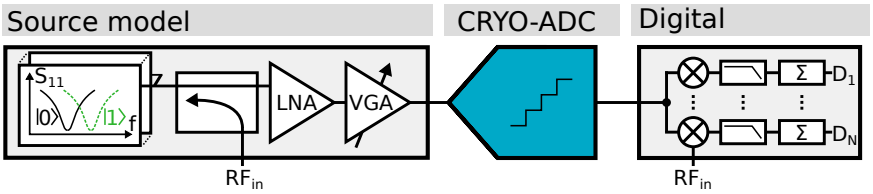

Figure 2.3: System model used for specifications.

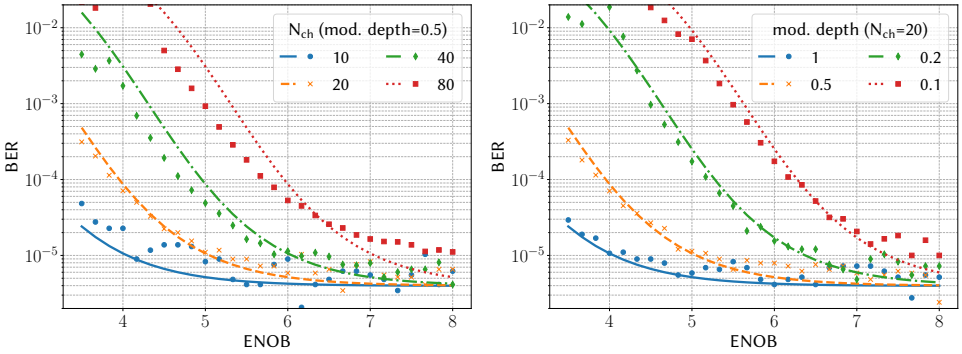

Figure 2.4: BER performance as function of ADC resolution for different channel numbers and different modulation depths. Simulation results and the model in (2.2) are shown with the markers and the continuous lines, respectively. The SNR at the ADC input signal is 13 dB.

2.2 System modeling and specifications

The use-case to derive the ADC specification is a receiver for the RF-readout scheme, as shown in figure 2.3. Each of the quantum sensors is placed in a matching network designed to match its nominal impedance, typically 0.1 MΩ to 1 MΩ for an SET, to the 50 Ω cable impedance. N_{ch} of these networks, each tuned to a different frequency, are placed in parallel and excited at their resonance via a directional coupler. For the 'zero' state of the quantum-state-dependent impedance Z_Q, we assume a matching condition; for the 'one' state, we assume signal reflection, resulting in a change of amplitude and phase of the signal at the LNA output. The modulation depth μ_{mod} of this amplitude-modulated signal can be significantly affected by the spread of the component values in the matching network (L, C) and the accuracy of the SET tuning. As those parameters vary between different systems and the modulation depth impacts the receiver specifications, we will derive the specifications taking into account a practical range of values for the modulation depth.

As the magnitude of the excitation needs to be low not to disturb the quantum sample, e.g., -110 dBm at the matching network input in [34], the input-referred noise power of the LNA is typically dominating the system noise. Therefore, we assume white noise over the signal bandwidth in the following. After the LNA, a VGA adjusts the amplitude to occupy the full-scale of the following ADC.

As the model concentrates on the ADC, we abstract the signal chain preceding it as an

ideal signal source producing N_{ch} frequency-multiplexed channels equally spaced over the ADC bandwidth. Although both amplitude and phase information are typically used to minimize the demodulation errors [39], the modulation parameters greatly vary for the different applications. Thus, without loss of generality, we assume a purely amplitude-modulated pulse data superimposed to a white-noise background dominated by the LNA. The demodulation after the ADC is performed by a digital coherent receiver.

The main performance specification is the resulting bit error rate (*BER*), which, for binary amplitude modulation, is given by

$$BER = Q(\sqrt{SNR}) \tag{2.1}$$

where Q is the 'Q-function'. The SNR for each of the channels is then given by the ratio of RMS signal power and the sum of the quantization-noise within the channel bandwidth f_{ch} and the noise from the signal source, computed at the output of the digital demodulator following the ADC, can be computed as:

$$SNR = \frac{\frac{\mu_{mod}^2}{N_{ch}^2}}{\frac{\mu_{mod}^2}{N_{ch}^2 SNR_{source}} + \frac{2f_{ch}}{6f_s 2^{2ENOB}}} \tag{2.2}$$

where *ENOB* is the effective number of bits of the ADC, we approximate the quantization noise to be white, and SNR_{source} is the input-signal-limited SNR appearing in case of an ideal ADC. As the readout error rate directly translates into a lower bound for the qubit readout *infidelity*, a *BER* below 10^{-5} is expected not to limit the system's performance in typical error correction schemes [10]. Therefore, we assume an input SNR of 13 dB for the ADC and compare the influence of the ADC ENOB on the attainable BER. The ADC quantization noise should be designed to be negligible in the total receiver noise budget, as a higher resolution is expected to be cheaper in terms of power consumption than a lower LNA's noise floor.

For a fixed energy per conversion in the ADC, the power consumption per qubit is independent of the ADC sample rate, since the sample rate scales proportionally to the number of qubit channels. However, a higher sample rate fitting more qubit channels is preferred, as it would reduce the number of required receivers and optimize the compactness of the whole quantum computer. A 1-GS/s ADC sampling rate is then chosen as this is expected to allow for a highly power-efficient ADC in the target technology (40-nm CMOS) [60, 67]. We assume a per-channel data rate of 1 MHz, a modulation depth of 0.5, a channel count $N_{ch} = 20$ and an input SNR of 13 dB unless otherwise noted. The channels are assumed to be equally distributed from 100 MHz to 500 MHz, with the minimum frequency being motivated by the size-requirements of the matching network. figure 2.4 shows the simulation results for the above model for a range of channel counts and modulation depth settings. Diminishing returns are reached for up to 40 channels for an ENOB above 6 bit, while 80 channels require up to 8 bit. When decreasing the modulation depth, the requirements change significantly: a modulation depth of 0.1 requires a resolution of at least 8 bit for 20 channels.

Based on these results, the target specifications are 6-8 bit ENOB at 1 GS/s. This would allow 20 channels with a modulation depth above 0.2 without limiting the achievable BER. If budgeting 0.5 mW per qubit, this requires a FoM_W better than 156 fJ/conv-step.

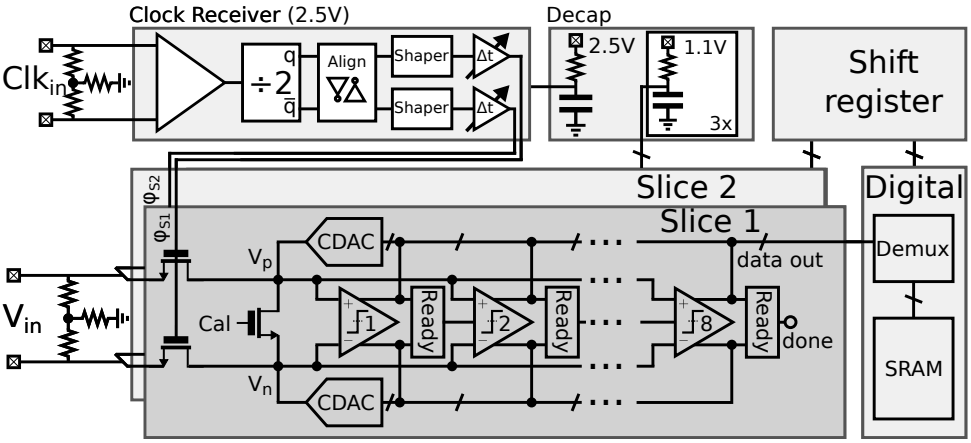

Figure 2.5: Chip-level block-diagram.

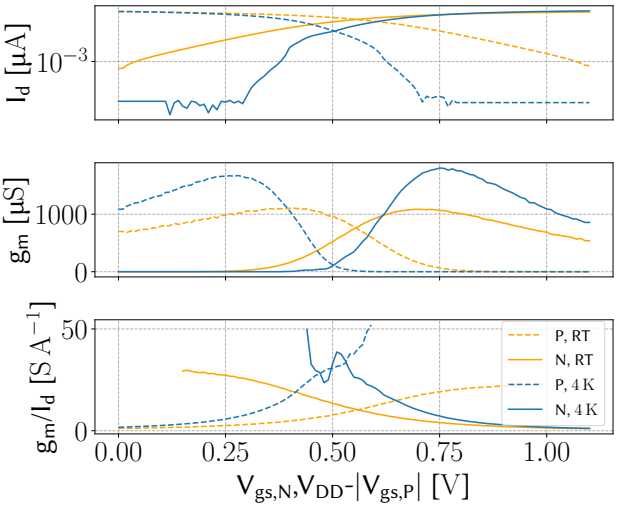

Figure 2.6: DC-characteristics at V_{ds} = 50 mV, NMOS W/L = 6x1.2 µm/40 nm; PMOS W/L = 6x2.4 µm/40 nm. The x-axis shows $V_{gs,N}$ and $V_{dd} - |V_{gs,P}|$ for all the curves related to the NMOS and the PMOS, respectively.

2.3 ADC architecture and component design

We have chosen the SAR architecture for the ADC, as it combines two essential features for our application: first, it offers good power efficiency, even for sampling speeds above 100 MS/s [60]; second, it is robust against the expected variations in device behavior at cryogenic temperature thanks to its predominantly digital and capacitor-based operation. An overview of the proposed SAR ADC is shown in figure 2.5. To minimize the time for non-critical comparator decisions, the ADC core is asynchronous [68], and employs

the loop-unrolled technique [67]. In such an architecture, each decision is performed by a different comparator that directly drives the capacitive DAC (CDAC) control. This minimizes the logic delay in the critical loop as the ready logic and DAC settling now run in parallel, allowing for higher sampling speed. The final comparator used in 8-bit mode is designed for lower thermal noise, as it carries the largest probability of noise-critical decisions [69]. Two identical cores are time-interleaved for a further increase in speed. All comparators are used for 8 bit conversions, while the LSB comparators are disabled for lower-resolution settings, effectively ending the conversion earlier.

2.3.1 Cyrogenic circuit design

No cryogenic simulation models were available for the target process at the design time. To achieve a robust design, we had to account for various changes in device behavior without using an exact model.

The most consequential changes when operating at cryogenic temperature occur in the transistor's DC characteristics. Both threshold voltage and mobility increase significantly at cryogenic temperatures [44]. The threshold voltage increases by about 120 mV in the adopted technology, while g_m approximately doubles for the same drain current, see the measurement of minimum length NMOS and PMOS devices in figure 2.6. We also observe a significant $>$ 2x increase in g_m/I_d when biasing in deep sub-threshold, but in a much lower bias range, due to the much increased sub-threshold slope. To approximate this behavior in circuit simulation, we simulated all circuits at $-55\,^{\circ}$C using the model provided by the foundry. For additional functional verification we included voltage sources mimicking the threshold voltage increase in series with crucial transistors, e.g., in the sampling front-end. Additionally, degradation of transistor matching is reported in [70]. While the increase in ΔV_{th} is small, the current gain mismatch $\Delta\beta$ increases by 20% and requires a slightly widened range in the comparator offset calibration (see section III.B). A moderate increase of about 10% is expected for the speed of digital circuits [71, 72]. This relatively minor change is due to the counterbalancing effects of the increases in both threshold and mobility. Although the absolute temperature decreases by 70×, the broadband noise of the active devices is expected to reduce by only a much smaller factor, as the white noise is believed to be caused by a mix of temperature-independent shot noise and thermal noise [73]. In this design, we conservatively assumed no improvement in noise, as this did not affect the design apart from a slightly enlarged comparator for the final decision.

Apart from the effects on the transistors, the change in temperature also influences the used passive components. Most relevant for this design is the effect on metal-oxide-metal capacitors, that have been characterized in [74]. While the absolute value of the capacitance is almost unaffected, the quality factor increases by 5× thanks to the improvement in metal conductivity and higher substrate resistance, which must be taken into account for the decoupling capacitors. For implementation of all resistors on this work we used unsilicided n-type polysilicon due to its stability at cryogenic temperatures. In [75] a change in resistance of $\approx$ 5% was measured between cryogenic and room temperature.

2.3.2 Comparator

The Strong-Arm comparator [76], shown in figure 2.7, offers a good balance between power efficiency and speed for moderate-resolution SAR designs [60]. In figure 2.8, the comparator

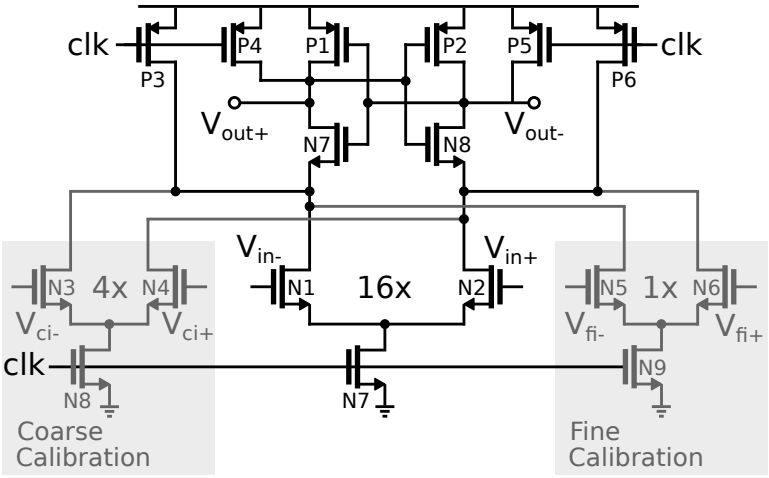

Figure 2.7: Strong-Arm comparator with calibration pairs.

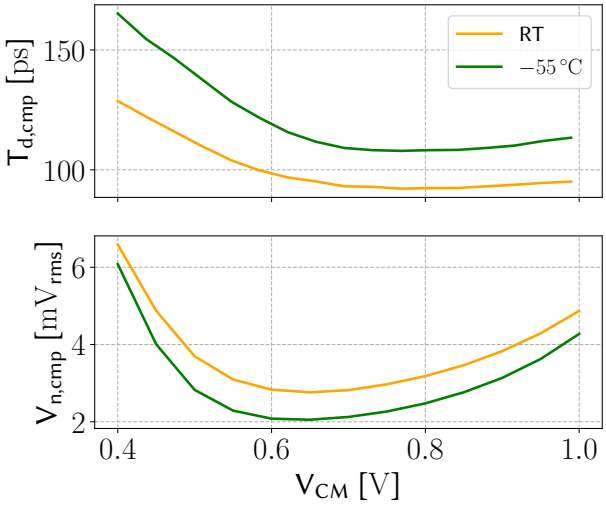

Figure 2.8: Comparator delay ($T_{d,cmp}$) and input referred noise ($V_{n,cmp}$) as a function of the comparator input common-mode V_{cm}.

simulated speed and noise performance is shown as function of the input common mode V_{CM} at RT and −55 °C. The common mode at the calibration pairs (N3-N6, see section 2.3.3) was kept constant to 800 mV, while the differential input was 1 mV. With increasing V_{CM}, the input pairs overdrive voltage V_{OD} rises, initially resulting in increased speed due to faster activation of the latch (N7, N8, P1, P2 in figure 2.7). This improvement levels off with higher V_{CM}, due to reduction in input pair gain outweighing the improvement in latch activation speed, creating an optimum range. At −55 °C, the comparator is overall

slower, with a stronger onset of slow-down at decreased V_{CM} due to increased V_{th}. For the noise, we observe a similar optimum at moderate V_{CM}. At low V_{CM}, the constant calibration pair V_{CM} dominates the noise, while at high V_{CM} the amplification of the input pair drops and the latch noise becomes important. At cryogenic temperatures, we expect a further shift of these curves towards higher V_{CM}, due to further increase in V_{th}, but also increased speed due to the higher mobility (see figure 2.6). Due to the uncertainty in the device noise behavior at cryogenic temperatures, we designed the comparator to meet the noise requirements even at room temperature, thus resulting in a moderate over-design at cryogenic temperature.

After sampling, the ADC common mode is predominantly affected by the DAC switching scheme (see section 2.3.4) and the comparator kickback after the decision. Every time one of the 8 comparators in the loop-unrolled ADC decides, the clock feed-through via its NMOS input transistors (N1, N2 in figure 2.7) reduces the common-mode voltage. The decreased V_{CM} for later decisions will cause an associated slow-down of the decisions. As discussed above, this effect is significantly more pronounced at cryogenic temperatures with its increased threshold voltage. Such a slow-down could be mitigated by increasing the input common-mode voltage. This, however, comes at the cost of reduced available swing at the ADC input due to headroom limitations in the ADC driver. To allow for a around mid-rail common-mode input voltage and still alleviate the slow-down due to the increased threshold, we adopted a variable common-mode switching scheme detailed in the following section.

2.3.3 Comparator calibration

Due to the loop-unrolled nature of the core ADC, individual comparator offsets cause distortion [67], thus requiring calibration of the comparator offsets. As the circuit will be operating in a highly temperature-controlled environment, we adopt a single foreground calibration. The offset calibration is performed via additional coarse (N3,4, sized at 1/4 of N1,2) and fine calibration pairs (N5,6, sized 1/16 of N1,2) in parallel with the main comparator pair, as shown in figure 2.7. Using two separate pairs relaxes the requirements of the calibration DAC used to generate the controlling voltages $V_{ci+,-}$ and $V_{fi+,-}$. A resistive-ladder calibration DAC is shared among all comparators and among both fine and coarse calibration pairs, see figure 2.9. The voltages for each calibration pair are tapped from the ladder by a separate set of switches. Sign reversal is possible via an additional switch that swaps the positive and negative voltages. The nominal DAC resolution is 3b and 4b for the coarse and fine pair, respectively. Missing codes between coarse and fine ranges are prevented by creating a 1b overlap. To avoid crosstalk between the calibration pairs of different comparators, decoupling capacitors are added after the selection switch. The resistive DAC ladder consumes up to 30 µA of static current, which is negligible in the total ADC power budget. To account for the increased mismatch at cryogenic temperatures [70], the calibration range needs to be widened by 20% compared to RT requirements, primarily to account for the increase in β mismatch. To set the calibration range, 3b resistive DACs defining the reference voltages V_{cm+} and V_{cm-}, labeled $R_{DAC,-}$ and $R_{DAC,+}$ in figure 2.9, are added.

As mentioned above, the comparator common mode varies during the conversion. The comparator offset is a function of the input common mode as this changes the relative weight

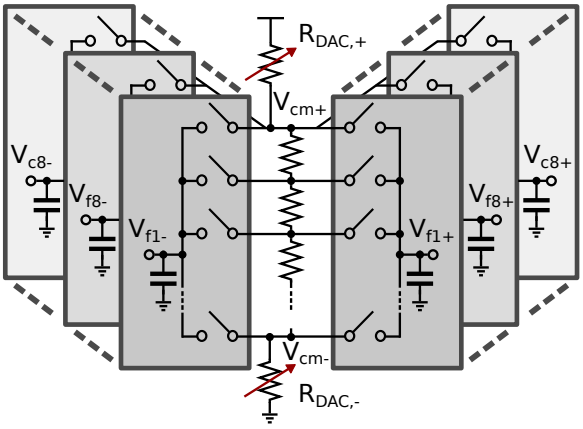

Figure 2.9: Comparator calibration DAC.

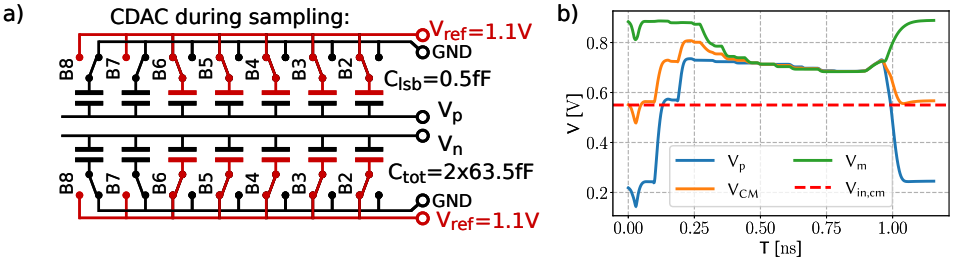

Figure 2.10: a) Main C-DAC in reset, b) simulated common-mode transient during the conversion.

of the input-pair and latch mismatch. To calibrate the comparators at the correct common-mode voltage, a switch shorting the CDAC top-plate (see figure 2.5, labeled *Cal*) is activated and then normal conversions are performed. The differential input voltage generated by the changing DAC codes during the conversion is now decaying to zero via the shorting switch, while the common-mode voltage varies as described in section 2.3.4. To ensure sufficient settling of the differential input voltage in between comparator decisions, the shorting switch must have an on-resistance below 500 Ω to settle the input error to less than 0.25 LSB according to simulations. For reliably achieving this low on-resistance at cryogenic temperatures, the shorting switch is implemented using thick-oxide (2.5 V) devices, as standard-transistor resistance would be heavily affected by the increased threshold voltage at mid-rail voltages. The comparator output is therefore determined by the sign of the comparator offset voltage. As the individual decisions will be noisy, many conversions can be averaged, and, based on this sign information, a linear search algorithm can find the optimum calibration setting for all comparators in parallel.

2.3.4 CDAC

The CDAC is implemented using a custom-designed capacitor array. The unit capacitors are 500 aF in size and arranged in a common-centroid layout.

As discussed above, we want to implement a variable common mode in the DAC switching to counteract the comparator slow-down for low common mode at cryogenic temperatures. Variable common modes have been employed before in RT designs [77–79]. To increase V_{CM} during conversion, the MSB and MSB-1 bottom plates are reset to ground during tracking while the rest of the bottom plates are reset to V_{ref}, as shown in figure 2.10 a). Upon the decision of the MSB comparator after the top-plate sampling, one side of the DAC switches the MSB capacitor accordingly to V_{ref}. Thereby, the differential input amplitude is decreased in accordance with the binary search algorithm, but also V_{CM} increases (figure 2.10 b). This increase in V_{CM} allows for a higher conversion speed in the MSB-1 decision. The same procedure is repeated for the MSB-1 capacitor. Afterwards, all the remaining decisions are switch-down as the speed advantage deriving from the increased V_{CM} becomes marginal and the NMOS used for switching down requires a lower driver fan-out thanks to the smaller necessary transistor. It needs to be noted that this variable V_{CM} is possible only due to the loop-unrolled architecture of the ADC core, as otherwise the variable comparator offset would cause distortion.

Due to the uncertainty in the device models, there was a risk of incomplete DAC settling in the case of higher speed-up in the ready logic at cryogenic temperatures than in the settling. To avoid an over-design to mitigate this risk, an optional additional delay has been implemented in the ready logic, but was kept deactivated during measurement. It has been used in calibration mode to allow for additional settling time of the differential signal across the input-shorting switch.

2.3.5 Sampling front-end

Both the input and clock signals are terminated differentially on chip with 100 Ω. At the center tap of the input termination, a 1.5 kΩ resistor is added to absorb the common-mode kickback from the track-and-hold switches in the ADC front-end.

The sampling front-end is implemented using thick-oxide NMOS transistors with an additional half-size charge injection cancellation pair. The choice of thick-oxide transistors for sampling rather than the usual pass-gate, which would be sufficient at RT for the speed and linearity required here [60], is motivated by the increased threshold voltage at cryogenic temperature. This increase leads to an estimated 100 mV dead zone around mid-rail for pass-gates at cryogenic temperatures, which would require exponentially larger transistors for sufficient settling. For design robustness and simplicity, we choose the thick-oxide sampling switch over alternative solutions, such as bootstrapping.

2.3.6 Clock receiver

The sampling switches are controlled by a thick-oxide-based clock receiver. The thick-oxide nature of the front-end switches motivates implementing the clock receiver using thick-oxide devices, as a supply domain crossing would necessitate significant power to ensure low-enough jitter, as well as additional alignment calibration.

After the on-chip clock termination, pseudo-differential self-biased inverters followed by a differential-to single-ended converter are used to generate a CMOS-level clock. The

single-ended full-rate clock signal is subsequently divided by two and the resulting 180°-phase-shifted clock phases are aligned by cross-coupled inverters, see figure 2.5. The sampling pulses are then shaped to optimize sampling and conversion time. By default, the ADC sampling pulse duration is shaped to be 650 ps for a 1 GHz input clock. To allow for more time for the conversion for the low-resolution settings, the sampling time can be shortened by up to 215 ps in steps of 80 ps by delaying the rising edge of the pulse while leaving the falling (jitter sensitive) edge unaffected.

Due to interleaving, artifacts as analysed in [80] can occur. The two-times interleaving employed here causes DC and $F_s/2$ spurs due to offset, which are not detrimental for our application, as well as a $F_s/2$-F_{in} spur due to gain and timing mismatch. The timing mismatch can be calibrated in the clock receiver by a variable capactive load per slice in the fan-out towards the sampling switch. The timing calibration is simulated to have a 1.5 ps step and a maximum range of 45 ps. At cryogenic temperature, both the calibration range and the step are expected to slightly decrease due to the increased carrier mobility [44]. By design, the gain interleaving spur is below the noise floor, so no calibration is included for it.

2.3.7 Digital back-end

Due to the 3 m-long cables in the available cryogenic measurement setup, real-time streaming of the ADC output is unfeasible. To evaluate the performance of the proposed ADC at cryogenic temperatures, the data is stored in an on-chip memory and read out slowly in a second step. As storage medium, we chose an on-chip 290-kb SRAM generated using a compiler supplied by the foundry. The SRAMs specified speed at RT was not high enough to store the samples at full speed, and an additional margin was necessary for eventual speed degradation of the SRAM at cryogenic temperatures as these are outside the SRAM specified operational range. For these reasons, we choose to operate the SRAM at $F_S/8$, with an 8× wider parallel data interface. For the digital implementation flow, an additional safety margin of 50 ps was used in the hold and setup constraints to ensure correct operation at cryogenic temperatures, similar to the precautions taken in [35].

2.3.8 Supply decoupling

The ADC's analog supplies as well as the digital supply are decoupled by on-chip capacitors. As the chip is wire-bonded, these capacitors are prone to oscillation with the bondwire inductance. As this effect is more pronounced at cryogenic temperatures due to the higher quality factor of the metal MOM capacitors, all supplies are degenerated with a 10 Ω resistor in series with the bondwire. Also this low-value resistor is realized with unsilicided n-type polysilicon for its temperature stability. To realize the low value and comply with the current-density requirements, the degeneration resistors are implemented with a width of 200 µm. Furthermore, numerous bondwires are used in parallel for the ground connection to reduce its inductance. The ground connection is not degenerated.

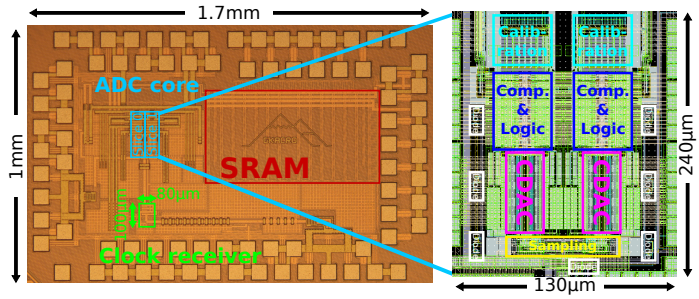

Figure 2.11: Chip micrograph and layout details.

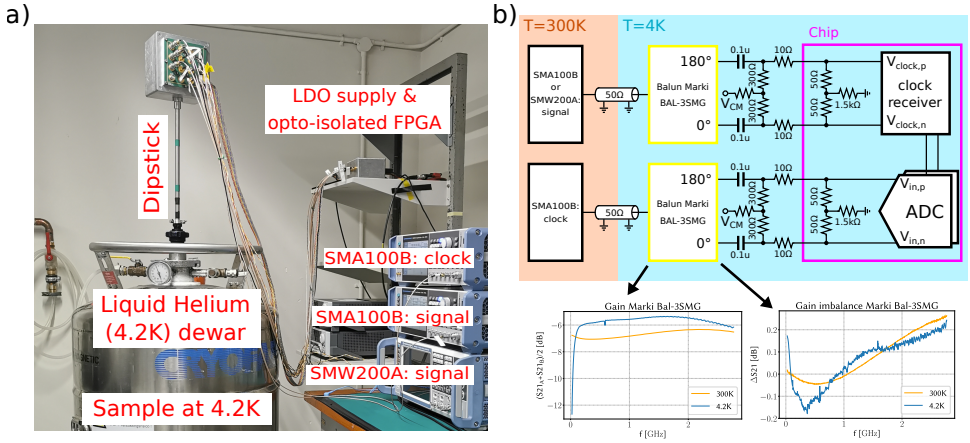

Figure 2.12: a) Cryogenic measurement setup, b) measurement board details and balun performance.

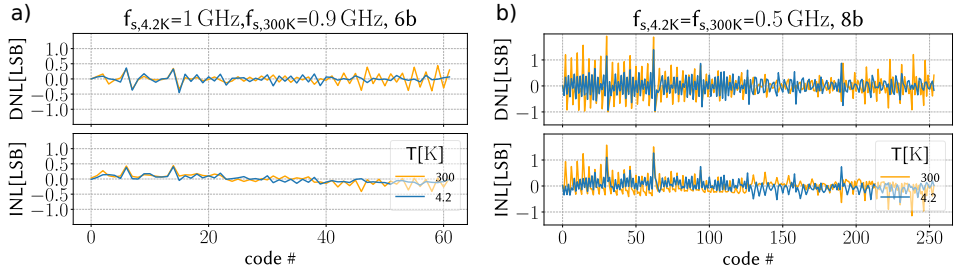

Figure 2.13: DNL and INL.

2.4 Measurement results

The proposed ADC is manufactured in a 40 nm LP process, see micrograph and layout in figure 2.11. The core ADC and the clock receiver occupy an area of 130 µm×240 µm and 100 µm×80 µm, respectively. The chip is wire-bonded and tested in a dip-stick setup, shown in figure 2.12. In the dip-stick, the components to be tested are inserted directly into liquid helium to provide a 4.2 K environmental temperature. The offset-calibration algorithm is implemented off-chip as part of the measurement routine. The 1.1 V and 2.5 V nominal supply voltages are used for all measurements. The data underlying the plots shown here can be found under [81].

The ADC has a differential interface for both input signal and clock. When generating differential signals at RT and delivering them through the dipstick to cryogenic temperature, a differential phase relationship between the signals can not be accurately ensured. To solve this issue, we tested RT baluns and found the Marki BAL-3SMG works well at cryogenic temperatures, see measurements in figure 2.12. The only limitation is a much increased loss at cryogenic temperatures for signals below 50 MHz. This is acceptable as the target signal and clock frequencies for our application lie above this limit.

INL and DNL measurements are shown in figure 2.13. Here we observe < 0.5 LSB DNL / INL for the 6 bit case in a) and a significant degradation to < 1.5 LSB DNL / INL for the 8 bit case b). This issue can be traced to the comparator schematic. The tail node at the drain of N7 in figure 2.7 is not reset with the other nodes. This leads to the input pair acting as a source follower, effectively implementing a 'max hold' on the tail node. Depending on the history of the DAC voltages, this tail node gives a different starting condition to the comparator operation, resulting in offset and producing the linearity limitations. To alleviate this, the tail of the comparator can be reset using a pull-up PMOS in a future redesign. Due to this issue, the optimal calibration settings could not be found in a straightforward binary search but needed manual fine-tuning. As predicted, the calibration range needed to be enlarged (by 25%) at 4.2 K with respect to RT via $R_{DAC,+}$ and $R_{DAC,-}$.

In figure 2.14, we show a summary of dynamic measurements performed on the core. In a), the output spectrum at F_s=1 GHz and T=4.2 K in 6-bit mode is shown, resulting in SNDR=35.2 dB. No spur is visible at the mirror of the input frequency, suggesting a sufficiently accurate timing calibration. As SFDR = 46.9dB at 1 GS/s, the residual skew between the two ADC slices after calibration is estimated to be at least below 3 ps. Therefore, the ADC shows sufficient performance for our target application of frequency-multiplexing 20 qubit channels with a modulation depth of 0.5. For the 8 bit mode in b), at F_s=0.5 GHz and T=4.2 K, the SNDR is limited to 42.7dB due to the linearity limitation found in the INL/DNL measurement. Still, the higher SNR allows the coverage of use-cases requiring higher ENOB at a lower speed.

In c), we sweep the sampling frequency keeping the input frequency close to Nyquist for different bit settings. For the 6-bit and 7-bit resolution settings, the sampling pulse was shortened in the clock receiver configuration to allow for more time spent in the conversion. Cryogenic and RT performance closely track each other up to moderate frequencies. At high frequency, we observe a speed advantage when operating at cryogenic temperatures. At cryogenic temperature, sampling rates as high as 1.1 GHz are possible but at a slightly reduced SNDR of 35 dB. Also the flexibility of the ADC is shown: at reduced sampling

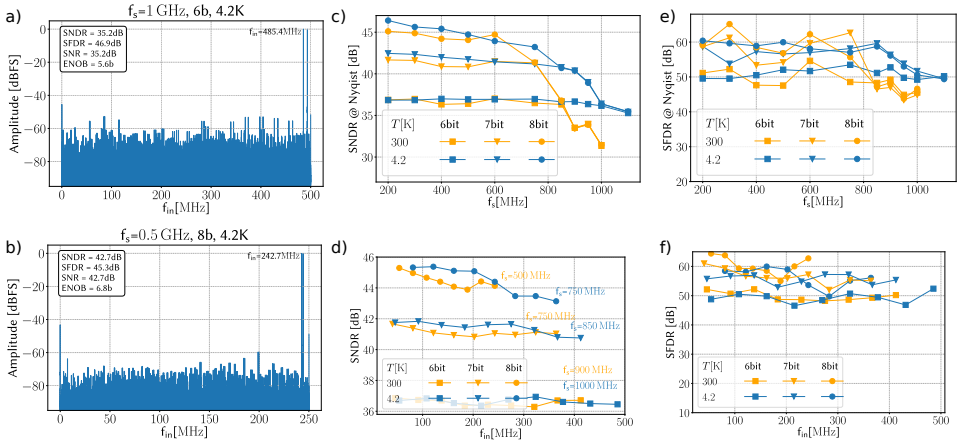

Figure 2.14: Spectra (a, b), SNDR (c, d) and SFDR (e,f) over sampling frequency and input frequency.

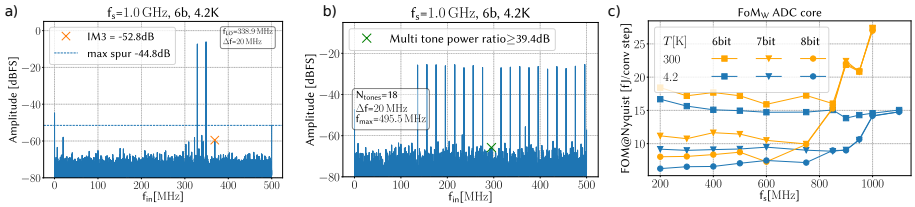

Figure 2.15: a) Two-tone measurement, b) multi-tone measurement, c) FOM_W overview.

speed and higher resolution settings, an SNDR around 45 dB can be achieved. The RT results closely follow the expectations derived from simulations. The speed advantage at cryogenic temperatures can be explained by the increased performance of the logic and the comparators. For the comparators, the threshold voltage increase does not hinder the comparator speed thanks to the CDAC switching scheme while the mobility increases significantly. This allows for a faster comparator decision. To investigate further, we sweep the input frequency for a series of fixed sampling frequencies in d). The ADC's performance is largely constant with respect to the input frequency, indicating no drops in performance for a particular input frequency.

We repeated the same analysis for the SFDR in e) and f), ignoring the DC and F_s spur contributions. At moderate speed, little difference between RT and 4.2 K is observed, while the same improvement as for the SNDR is observed towards high sampling speeds. We can conclude that the front-end linearity is not majorly impacted by cooling the circuit due to the choice of a thick-oxide front-end switch.

In figure 2.15 a), we perform a two-tone measurement showing a -51.4dB third-order intermodulation spur. The dominating interleaving spurs at Nyquist and DC due to the offset and offset mismatch between the slices is not important in our application. The variation in the amplitude between the two tones is attributed to the imperfect calibration of the long measurement cables.

We also tested the multi-tone performance of the ADC in figure 2.15 b) where we removed one peak out of a comb of frequencies generated by a VSG. We observe that no significant tone is generated by intermodulation at the frequency of the missing tone. The other spurious tones are similar in level to the previous test.

Finally, in figure 2.15 c), we plot the FOM_W of the ADC core excluding the clock receiver. For sampling frequencies below 900 MHz, the ADC shows better than 10 fJ/conv-step. performance both at cryogenic and RT. At cryogenic temperature, a FOM_W of 15 fJ/conv-step is achieved at 1GS/s. This is on-par with the performance of RT designs at similar sampling speeds, as shown in table 2.1. When including the thick-oxide clock receiver, this increases to 200 fJ/conv-step.

In figure 2.16, a breakdown of the ADC's power consumption at a sampling speed of 1GS/s is given. Most of the power is dissipated in the thick-oxide clock receiver, which was not the central focus of the design. Replacing it with a bootstrapped switch would reduce this power drastically. In the ADC core, most of the power is dissipated in the logic and the CDAC reference. The comparators contribute only 12% of the total power dissipation, as the noise requirements for an 8 bit ADC allowed using small devices for their implementation. In summary, this power consumption means that, for the case of 20 qubit channels, 0.5 mW is consumed per qubit.

One important consideration in the design of cryogenic chips, in general, is the on-chip temperature, which might largely deviate from the environmental temperature of 4.2 K, e.g., as shown in [35, 82]. As the measurements are performed in a dip-stick, the chip is fully submerged in liquid helium. Although this results in an expected good thermal coupling, to verify the absence of any local hot spots on the chip, an array of temperature-sensing diodes is placed on the chip, distributed over the ADC core, clock receiver, digital and termination resistors, as shown in figure 2.16 a). These P+/N-well diodes are multiplexed via thick-oxide selection switches on the anode side and read via a sense-and-force connection by a source measure unit (SMU) using a current of 1 µA. For calibration, the diode voltage was recorded at a range of environmental temperatures, measured with an external temperature sensor, with the ADC being inactive. Due to the loss of sensitivity in the diodes at lower temperatures in the calibration curve shown in figure 2.16 b), temperatures below 6 K are affected by a non-negligible error. In figure 2.16 c), we show the heat map when operating the ADC with a clock of 1 GHz and Nyquist-rate input signal. Very little global self-heating was observable. This is in line with the results in [82], where for a very small heater dissipating 6.3 mW, a self-heating below 0.5 K was observed at a distance of just 20 µm, which is the minimum distance of the on-chip sensing diodes from the power-dissipating circuit elements. Also, the power consumption in this chip is spread over the area of the clock receiver and ADC, rather than being concentrated in a single point, further decreasing the effect. We can conclude that the global chip temperature was very close to 4.2 K.

In table 2.1 we compare the presented ADC with prior works showing operation at cryogenic and RT. Compared to other cryogenic designs, we have made significant advances in efficiency and speed, while maintaining similar performance to RT state-of-the-art designs in similar technologies.

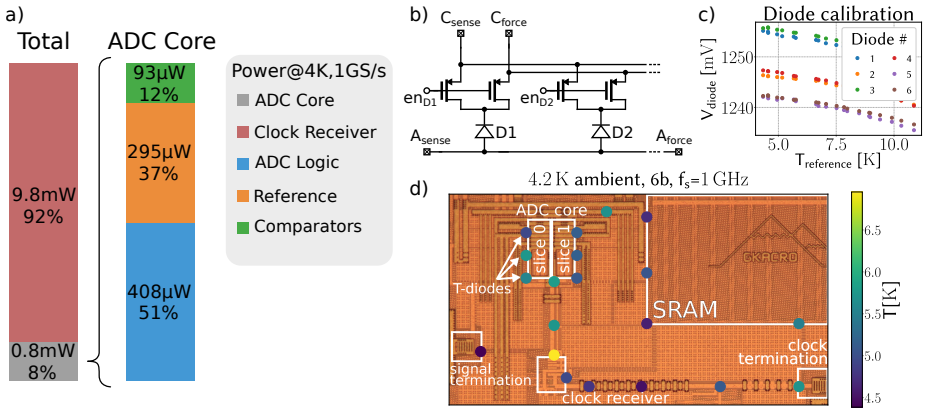

Figure 2.16: a) Measured power breakdown, b) Schematic of the on-chip temperature measurement structure, c) diode calibration with powered-down chip, d) measurement result.

	This work		Okcan, RSI 2010	Creten, JSSC 2009	Jiang, JSSC 2012	Kull, ISSCC 2013	Kull, ISSCC 2017
Temperature [K]	300	4.2	4.4	4.2	300	300	300
Architecture	TI SAR		SAR	Flash	SAR	SAR	PP-SAR
Max sampling rate [MS/s]	900	1000	0.05	0.0125	1250	1300	950
Resolution [bit]	6-8		12	8	6	8	10
Technology [nm]	40		350	700	40	32	14
Supplies [V]	1.1(core), 2.5(clock)		3.3	5.5	1.1	1	0.7
Input range [Vpp]	0.7[2]	0.7[1]	-	-	0.5	0.5	0.5
SNDR@Nyquist [dB]	33.4[2]	36.2[1]	48.7	-	27.7	39.3	50
SFDR [dB]	48.4[2]	48.5[1]	57.9	-	39	49.6	58
Power [mW]	0.7[2,3] (10.3[2,4])	0.8[1,3] (10.6[1,4])	0.297	5.1	5.28	3.1	2.26
FoM$_W$[fJ/c.step]	20[2,3] (260[2,4])	15[1,3] (200[1,4])	16100	1.6 ·10^6 ‡	148	28	8.9
Core area [mm^2]	0.045		2.7[5]	40[5]	0.014	0.0015	0.0016

[1]6b, 1GS/s [2]6b, 0.9GS/s [3]ADC core [4]Core + clock receiver + sampling switch [5]full chip ‡estimated from INL

Table 2.1: Comparison table

2.5 Conclusion

In this chapter, we demonstrated a loop-unrolled ADC operational at 4.2 K. Using a DAC switching scheme that increases the comparator input common mode for most decisions, we reached higher speed at cryogenic temperatures. By providing an ENOB of 5.7 b at 1 GS/s, the ADC is well suited for an RF reflectometry readout for spin qubits, allowing for 20 frequency-multiplexed qubit channels while maintaining a power efficiency of 0.5 mW/qubit. The ADC power efficiency can be further improved by including a more efficient clock receiver, which is possible when replacing the thick-oxide front-end sampling switches with bootstrapped switches. To the best of the authors' knowledge, this design is the fastest sampling cryo-CMOS ADC reported to date. With the demonstrated

performance, the proposed ADC contributes to the progress towards the fully-integrated cryo-CMOS readout required in future large-scale quantum computers.

3

Cryogenic Comparator Characterization and Modeling for a Cryo-CMOS SAR ADC

This chapter reports the experimental characterization and modelling of a stand-alone StrongARM comparator at both room temperature (RT) and cryogenic temperature (4.2 K). The observed 6-dB improvement in the comparator input noise at 4.2 K is attributed to the reduction of the thermal noise and to the suppressed shot noise in the MOS transistors becoming dominant at cryogenic temperature. The proposed model is employed in the design of a loop-unrolled 2× time-interleaved 1-GSa/s 7b SAR ADC for spin-qubit readout. As predicted by the comparator model, the ADC is noise-limited at RT to a SNDR of 38.2 dB at Nyquist input, while this improves to 41.1 dB at 4.2 K, now limited by distortion, thus resulting in the state-of-the-art FoM_W for cryo-CMOS ADC of 20.9 fJ/conv-step.

This chapter has been published as *Gerd Kiene, Aishwarya Gunaputi Sreenivasulu, Ramon WJ Overwater, Masoud Babaie, Fabio Sebastiano, "Cryogenic Comparator Characterization and Modeling for a Cryo-CMOS 7b 1-GSa/s SAR ADC." European Solid State Circuits Conference (ESSCIRC) (2022)* [83].

3.1 Introduction

Noise performance of medium-resolution SAR ADCs, which are well-suited for the RF-readout chain as argued previously in chapter 2, is typically dominated by the comparator noise. This is because the capacitor size dictated by matching considerations is typically already significantly larger than required for the respective kT/C noise level. Such a comparator noise would be expected to drop by ≈ 70× when reducing the temperature from RT to 4.2 K, if purely thermal in nature. However, the characterization of various circuits at cryogenic temperatures suggests that this noise scaling can not be achieved in practice. For example, the noise figure of the LNA in [39] showed only an improvement from 2.5 dB to 0.6 dB, much less than predicted by assuming pure thermal noise. The absence of a comprehensive noise model for cryo-CMOS devices and the lack of direct noise measurements of dynamic latching circuits at cryogenic temperatures hinder the design optimization of cryo-CMOS ADCs. To fill this gap, we report, for the first time, the measurement and the analysis of the noise and offset of a latching comparator at cryogenic temperatures. Based on those findings and improving on the design presented in chapter 2, we optimized the design of a prototype 7b 1-GSa/s SAR ADC, thus verifying in a practical application the validity of the comparator characterization and its modelling . The designed ADC improves significantly in power efficiency compared to the design in chapter 2.

In the following, section 3.2 describes the proposed StrongARM comparator implementation, results and analysis. section 3.3 covers the circuit design of the prototype ADC chip, concluding with its experimental validation.

3.2 Dynamic comparator

The StrongARM comparator is a common choice in medium-resolution ADCs for its high speed, power efficiency and compactness. The proposed implementation (figure 3.1) comprises two additional coarse and fine calibration pairs to allow for an efficient foreground calibration [84]. In the following, a model for this circuit's noise and offset – relevant comparator parameters for ADC design – are derived and compared with experimental results. A specific focus is put on the behavior over input common-mode (CM), as this changes significantly due to both threshold-voltage increase at cryogenic temperature and CM variations on the capacitive DAC during a typical SAR conversion.

3.2.1 Measurements

The stand-alone comparator in figure 3.1 has been fabricated in a 40-nm CMOS process (figure 3.2b). All transistors are implemented as minimum length devices. On-chip input capacitors of 200 fF have been used to minimize the kT/C noise present at the input. A dip-stick setup employing liquid He (figure 3.2c) is used for cryogenic characterization, adopting the fixture shown in figure 3.2a. All DC input lines are filtered to reject interference coupling to the long cables in the dip-stick and are driven by 18b battery-operated RT DAC modules. We sweep the differential input voltage for a given CM and record the comparator decision statistics, from which we extract input noise and offset (figure 3.3a). On all sweeps, the CM voltage of the input and calibration pairs is the same.

When increasing the CM voltage at RT for four samples (figure 3.3b), we observe a monotonic increase of noise that is captured by the room-temperature simulation. When

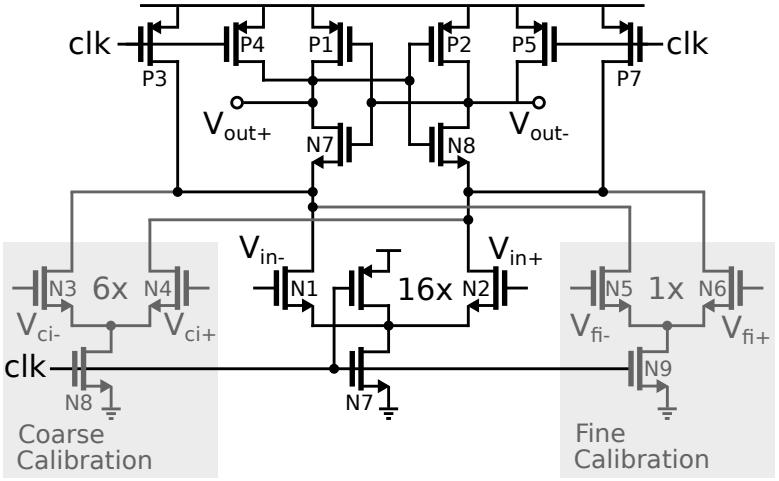

Figure 3.1: Comparator schematic; N3-4 (N5-6) is the coarse (fine) offset-calibration pair with input $V_{ci+,-}$ ($V_{fi+,-}$) provided by the ADC calibration circuit.

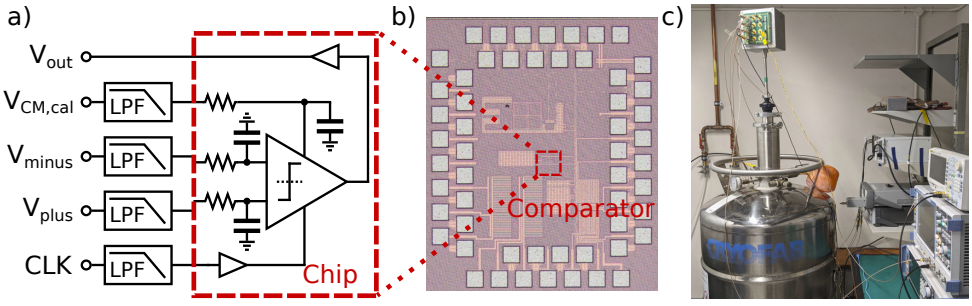

Figure 3.2: Test-chip for comparator characterization: a) schematic, b) chip micrograph, c) measurement setup.

cooling down to 4.2 K, the transistor threshold voltage increases by about 80 mV, resulting in a shift of the curve towards a higher common-mode voltage. Compared to RT, the input noise voltage is about 2× lower and also significantly lower from simulations at −40 °C, the edge of the validity of the foundry supplied model.

The offset voltage (figure 3.3c) also shows a strong dependence on input CM, but no significant change is observed from RT to 4.2 K.

3.2.2 Model derivation

Previous analytical derivations for the noise of regenerative comparators assuming thermal-noise sources suggest that the input-referred noise power scales with absolute temperature [85, 86]:

$$\sigma_{in}^2 \propto T \tag{3.1}$$

However, the experimental characterization does not show the predicted large difference in noise power between RT and 4.2 K. To address the mismatch between the experimental

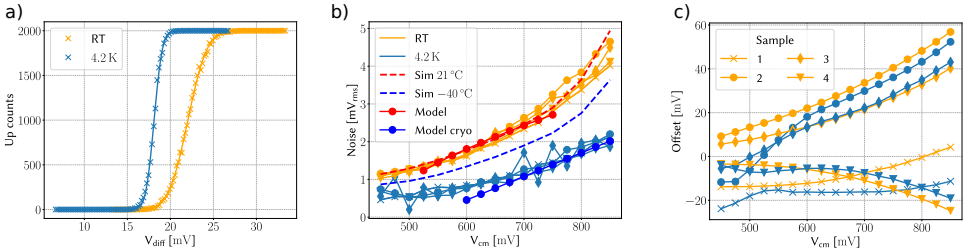

Figure 3.3: Comparator characterization: a) example measurement; b) noise voltage over input CM (V_{cm}) for four measured samples, simulations using the foundry device model, and the model in equation (3.5); c) measured offset over V_{cm} for four samples. Symbols for each sample are the same in b) and c).

results and the commonly accepted theory, we derive the input-referred noise by assuming the transistor noise to be a combination of thermal noise and suppressed shot noise [73].

First, we calculate the gain of the input pair during its amplification phase. This phase lasts from the positive clock edge till the input and calibration pair transistors, N1-N6, have discharged the parasitic capacitance C_p at their drains enough for the latch NMOS transistors (N7 and N8 in figure 3.1) to turn on, initiating exponential latching. The amplification time T_{amp} can be computed considering operation around common-mode as:

$$T_{amp} = \frac{2V_{th,N}C_p}{g_{m,all}V_{od}} \tag{3.2}$$

where $V_{th,N}$ is the NMOS device threshold voltage, $g_{m,all}$ is the combined transconductance of the input and calibration pairs, $V_{od} = V_{gs} - V_{th,N}$ is the input-pair overdrive voltage, and we assume strong-inversion operation for the input transistors. The gain of the input pair can then be written as:

$$A = \frac{g_m}{C_p}T_{amp} = \frac{2W_{in}V_{th,N}}{W_{all}V_{od}} \tag{3.3}$$

where W_{in} and $W_{all} = W_{in} + W_{cal}$ are the width of the input pair and the combined width of input and calibration pairs, respectively. We neglect the MOS output impedance. The power spectral density of the current noise in each half of the comparator is dominated by thermal and suppressed shot noise and it is given by:

$$S_{id} = 2qFI_{CM,all} + 4k_bT\gamma g_{m,all} \tag{3.4}$$

with q is the elementary charge, F the Fano factor, $I_{CM,all}$ the combined common-mode current, k_b the Boltzmann constant and T the absolute temperature. The input-referred noise in the amplification phase is then given by

$$\sigma_{in}^2 = \frac{W_{all}^2V_{od}^2}{2C_pW_{in}^2V_{th,N}}\left(qF + \frac{4kT\gamma}{V_{od}}\right) \tag{3.5}$$

This expression clearly show the increase of the input-referred noise with the input common-mode $V_{cm} = V_{od} + V_{th,N}$ that we observe in the measurement.

According to our model, while thermal noise dominates at RT by contributing 84% of the total noise for $V_{cm} = 650\,\text{mV}$, it becomes negligible at cryogenic temperatures. At

4.2 K, the noise behavior is mainly determined by the shift in threshold voltage and the temperature dependence of the Fano factor. The threshold voltage (extracted from DC characterization of individual transistors) increases from 450 mV at RT to 530 mV at 4.2 K, decreasing the overdrive voltage and therefore also lowering the noise. In practice, the $V_{th,N}$ change results in a shift of the noise curve towards higher V_{cm}.

When assuming a mix of thermal noise with $\gamma = 1$ and shot noise with $F = 0.1$ at RT, we achieve in good fitting between the foundry device model and equation (3.4) in transistor noise simulation. As shown in figure 3.3b, this also results in good match between the proposed model and the experimental data. The region of validity of our model is bounded by the subthreshold region towards lower V_{cm} and by the latch noise becoming relevant for higher V_{cm} as the input-pair gain decreases, as predicted in equation (3.3).

Using the model we can achieve a good fit for the dependency of the noise on V_{cm} even at 4.2 K. However, here the Fano factor must be increased to $F = 0.25$ (only at 4.2 K) to achieve an acceptable match with experimental data. Such an increase in the suppression factor could be attributed to the device behaving closer to an ideal ballistic device at lower temperatures, e.g., due to a decrease in carrier scattering in the channel. Although the shot-noise in nanometer CMOS devices have been extensively studied at RT [87], no comprehensive data on the behavior of F at cryogenic temperatures is available, thus pointing to the need for further cryogenic device characterization and physical modelling. As an outcome, the proposed model successfully captures the circuit noise behavior within the CM region, which is mainly relevant for high-speed ADC design, shown in section 3.3.

The comparator offset due to the input and calibration pairs can be calculated:

$$\Delta V_{in,os} = \left(\Delta V_{th,in} + \frac{\Delta\beta}{\beta} \frac{V_{od,in}}{2} \right) \sqrt{1 + \frac{W_{cal}}{W_{in}}} \tag{3.6}$$

where β is the device current gain, $\Delta V_{th,in}$, $\Delta\beta$ are the standard deviation of mismatch in threshold voltage and current gain of the input pair respectively. As reported in [70], the mismatch in threshold voltage shows a negligible variation at cryogenic temperature while the mismatch in β is expected to increase by 25%, thus explaining the little observed change in figure 3.3c.

3.3 Cryo-CMOS ADC

3.3.1 Circuit design

We verify the StrongARM comparator model in a prototype ADC tested at both RT and 4.2 K. The ADC architecture and most of the circuits are based on the design presented in chapter 2. Here, we only briefly highlight the similarity and differences with respect to that design and refer the reader to chapter 2 for more details. Similar to [84], the circuit is a 2× time-interleaved loop-unrolled (LU) 7b SAR ADC. The 6b capacitive-DAC (CDAC) employs top-plate sampling and is implemented using custom 0.5 fF unit cells. Since the ADC is intended for integration with an input on-chip driver in the target qubit-readout system, non-switched capacitors connected to ground are included in the CDAC to reduce the input range of the ADC to 600 mV_{pp}, thus easing the output swing specifications of the input driver. To further facilitate the input driver, the ADC also implements a reset to V_{cm} before sampling, followed by a DAC switching similar to [84]. This provides higher V_{cm} for

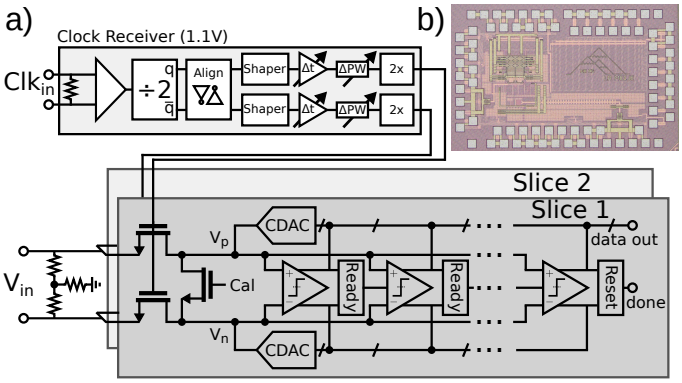

Figure 3.4: ADC protoype: a) block diagram; b) micrograph.

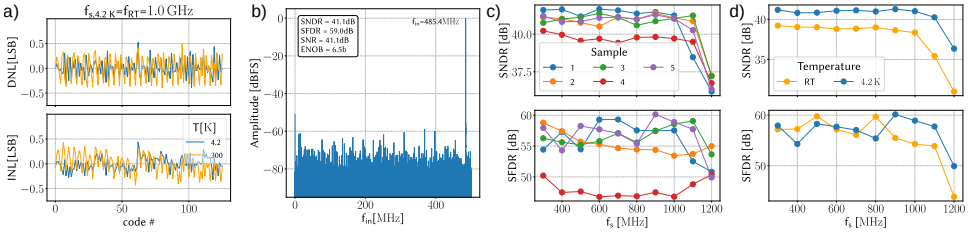

Figure 3.5: ADC characterization: a) INL and DNL for worst sample (sample 4) ; b) output spectrum at 1 GS/s and 4.2 K (sample 5); c) SNDR and SFDR at 4.2 K for all measured samples. SFDR does not include the TI-induced spur at $f_s/2$ d) SNDR and SFDR at RT and 4.2 K for a typical sample (sample 5).

the comparators after the MSB decision to compensate for the increased V_{th} at cryogenic temperature.

To improve power efficiency compared to [84], the entire sampling clock chain has been implemented using thin-oxide devices. At the end of the clock chain, thin-oxide NMOS switches with 2× boosted control voltage are used to sample the input signal. A pulse shaper in the clock chain tunes the duty-cycle of the sampling clock to provide more than 50% of the clock period for conversion. The 2× time-interleaved ADC suffers from interleaving artifacts caused by timing mismatch. To calibrate these, a pair of digitally controlled delay lines are designed to adjust the delay of the sampling clock provided to each slice.

The StrongARM comparator measured and analyzed in section 3.2 is used in the LU-SAR loop. The comparator design is optimized for speed and the ADC is targeted to be comparator-noise limited at RT. Since the comparator offset causes distortion for LU-SAR, the input voltages of the fine and coarse differential pair in figure 3.1 are generated by a resistive DAC to calibrate the offset, similar to [84] but with an optimized scheme with 5b (coarse) and 6b (fine) resolution. During foreground calibration, a switch with boosted static gate voltage is activated to short V_p and V_n (figure 3.4a) to calibrate every comparator at the correct V_{CM} level. Resetting the comparator tail node (drain node of N7 in figure 3.1) eliminates any signal-dependant offset, which would otherwise limit the

	This work		[4]		Kull, ISSCC 2013	Kull, ISSCC 2017
Temperature [K]	300	4.2	300	4.2	300	300
Architecture	TI SAR		TI SAR		SAR	PP-SAR
Max f_s[MS/s]	1000		900	1000	1300	950
Resolution [bit]	7		6-8		8	10
Technology [nm]	40		40		32	14
Supplies [V]	1.1		1.1(core), 2.5(clock)		1	0.7
Input range [Vpp]	0.6		0.7	0.7	0.5	0.5
SNDR@Nyquist [dB]	38.2[2]	41.1[2]	33.4	36.2	39.3	50
SFDR [dB]	>50[2]	>50[2]	48.4	48.5	49.6	58
Power [mW]	1.94[1]		10.3[1]	10.6[1]	3.1	2.26
FoM_w[fJ/c.step]	29.2[1, 2]	20.9[1, 2]	260[1]	200[1]	28	8.9
Core area [mm^2]	0.042[1]		0.045[1]		0.0015	0.0016

[1]Full ADC, including clock receiver [2]typical sample

Table 3.1: Comparison table

achievable calibration accuracy, as was the case in [84].

3.3.2 Measurement

The chip is implemented in a 40-nm process, see the micrograph in figure 3.4b. Five sample dies have been bonded to the test PCBs and characterized in the same dip-stick environment as the comparator.

The comparator offset calibration based on binary search and the timing skew calibration are implemented off-chip. To optimize the linearity, manual calibration fine-tuning is done besides binary search for samples 1 and 3.

In figure 3.5a, the static performance of the worst-case sample (sample 4) is reported. Unlike [84], no signal-dependent residual offset appears and the maximum INL/DNL is below 0.55 LSB for both RT and 4.2 K. The Nyquist-rate input measurement in figure 3.5b shows an SNDR of 41.1dB. Combined with the 1.94 mW power drawn from the nominal 1.1-V supply, this results in a typical Walden FOM_W of 20.9 fJ/conv·step (including the clock receiver), a 10× improvement compared to [84], see section 3.3.2. All tested samples reach 1 GSa/s before dropping in SNDR value, see figure 3.5c. The dynamic performance of sample 4 is limited by the accuracy of the offset calibration. The higher sampling rates at 4.2 K are attributed to the increased speed of the comparator and the digital circuitry in the SAR loop.

Comparing the dynamic performance at RT and 4.2 K in figure 3.5d, reveals an improvement of 2-dB SNDR at 4.2 K. At 4.2 K harmonic distortion is not limiting the performance as indicated by the measured SFDR, and the SNDR is in line with the quantization-noise-limited performance of an ADC with 0.5-LSB INL. Therefore, the 6-dB improvement in comparator noise measurements at 4.2 K (figure 3.3b) makes the comparator noise negligible at cryogenic temperature.

3.4 Conclusion

We have reported the first direct noise and offset measurements and modelling on a latching comparator at cryogenic temperatures. The gathered evidence points towards a relatively small improvement in white noise at cryogenic temperature, due to the white noise originating from a mix of thermal and shot noise. The proposed model has been applied to the design of a high-speed cryo-CMOS SAR ADC for spin-qubit readout, resulting in state-of-the-art performance compared to previously reported cryo-CMOS ADCs. In particular, the presented prototype showed a significant improvement in power efficiency with respect to the cryo-CMOS ADC presented in chapter 2.

4

A FIA Sampling Driver enabled by Cryogenic-Aware Body-Biasing

This chapter presents a floating inverter amplifier (FIA) that performs high-linearity amplification and sampling while driving a 2× time-interleaved (TI) SAR ADC, operating from room temperature (RT) down to 4.2 K. The power-efficient FIA samples the continuous-time input signal by windowed integration, thus avoiding the traditional sample-and-hold. Cascode switching, a floating supply and accurate pulse-width timing calibration enable high-speed operation and interleaving. In addition, by exploiting the behavior of CMOS devices at cryogenic temperatures, forward-body-biasing (FBB) is pushed well beyond what is possible at RT to ensure performance down to 4.2 K, and its impact on the performance of cryogenic circuits is analyzed. The resulting ADC, implemented in 40-nm bulk CMOS and including the FIA driver, achieves SNDR=38.7 dB (38.2 dB), SFDR>50 dB (>50 dB), and FOM_W=25.4 fJ/conv-step (31.3 fJ/conv-step) with Nyquist-rate input at 1.0 GS/s (0.9 GS/s) at 4.2 K (RT), respectively.

This chapter has been published as *Gerd Kiene, Ramon Overwater, Masoud Babaie and Fabio Sebastiano, A Cryo-CMOS SAR ADC with FIA Sampling Driver enabled by Cryogenic-Aware Back-Biasing.*

4.1 Introduction

Cryogenic ADCs, as proposed in [18, 88] as well as in chapters 2 and 3, are required for the RF readout of spin qubits. The converters are generally high-speed to provide the wide bandwidth necessary for facilitating frequency-multiplexing and require medium resolution, as discussed in chapter 2. The power dissipation of the circuitry in the RF-readout frontend is strictly constrained by the limited cooling power available in deep-cryogenic environments. This requires a significant focus on power efficiency in the ADC design, leading to cryogenic converters with efficiencies similar to those at room temperature. Nevertheless, these cryogenic converters also need an effective driving circuit, which in prior works was either neglected or implemented with traditional power-hungry settling drivers, e.g. in [18]. This is a substantial shortcoming as these settling drivers can require a power budget even larger than the ADC itself [89]. In this chapter, we will address the challenge of implementing a power-efficient driver for cryo-CMOS ADCs. While this is already a relevant research topic for the field of data converters, it is crucial for the target application of this thesis, as the ADCs presented in chapters 2 and 3 must be integrated within the whole receiving chain presented in 1, and their driver must be integrated in the same cryo-CMOS SoC between the RF frontend and the ADC itself.

As an alternative to settling amplifiers, open-loop dynamic amplifiers have been proposed for their high efficiency combined with high linearity [90]. These dynamic amplifiers have been used as sample-and-hold [91], drivers for ADCs [92–94], and as interstage amplifiers in pipeline ADCs. For the latter, common-mode control has been eased by adopting floating supplies, forming floating inverter amplifiers (FIA) [95, 96]. However, employing dynamic amplifiers at cryogenic temperatures is a daunting task due to the lack of reliable device models and the significant cryogenic increase in threshold voltage V_{th} (0.1/0.18 V for NMOS/PMOS) [97], which prevents biasing power-efficient inverter-based amplifiers in the high-linearity region. Although independently AC-coupling the PMOS and the NMOS could alleviate this, it would limit the usable ADC bandwidth near DC. The increased V_{th} complicates even the adoption of standard techniques, such as pass gates for switching mid-rail voltages [72]. Thus, clock boosting, bootstrapping or high-voltage supply domains [83, 98] are necessary, deteriorating the power efficiency and increasing the design complexity.

To address those issues, we propose the use of cryogenic-aware forward body-bias (FBB). FBB has been used in FDSOI technology to mitigate the cryogenic increase in threshold voltage by applying a large back-gate biasing voltage (up to −5.8 V for PMOS) [71]. Although the control range for the body voltage in bulk technologies is severely limited by the forward conduction of the bulk-source diode, the modeling and the characterization in [99] suggest that a level of control comparable to FDSOI can also be achieved in bulk CMOS, given the lowered forward bias diode leakage at cryogenic temperature [82]. In this work, we employ, for the first time, cryogenic-aware FBB in bulk CMOS to control the V_{th} of individual transistors in a wide range of cryogenic analog circuits, thus enabling the first dynamic ADC driver at cryogenic temperatures. The presented driver and ADC combination achieves high linearity with more than 50 dB SFDR and also a competitive FOM_W=31.3/25.4 fJ/conv-step with Nyquist-rate input at 0.9/1.0 GS/s at RT/4.2 K. These advances are enabled, in addition to the cryogenic-aware FBB, by the use of cascode switching, the adoption of a floating supply, and the use of accurate pulse-width timing

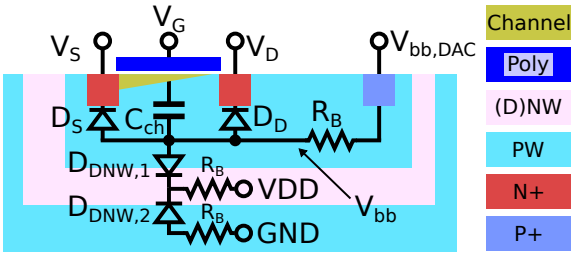

Figure 4.1: Sketch of NMOS in DNW with resistances and diodes considered here.

calibration.

The chapter is organized as follows: after a description of the impact of body-biasing in analog design at cryogenic temperatures (Section II), we describe the amplifier design (Section III), its experimental validation (Section IV), and draw the conclusions in Section V.

4.2 Forward-Body-Biasing (FBB) in cryogenic analog circuit design

FBB primarily affects the transistor's V_{th}. This effect can approximately be described by the body factor ζ

$$\zeta(V_{bb}) = \frac{\partial V_{th}}{\partial V_{bb}}, \tag{4.1}$$

where V_{bb} is the voltage applied via the body contact, as in figure 4.1. In the 40 nm bulk process adopted here, ζ varies between ≈ 0.1 to ≈ 0.35 at 4.2 K when the body-bias is swept from 0 to 1.1 V with an average of ≈ 0.25 V/V [99], which is higher than in common FDSOI technologies with, for example, 0.085 V/V in [71].

While the body contact has been used at RT both as a tuning knob for mitigating mismatch [100], or as additional input [101], the usable range for FBB is much wider at cryogenic temperature thanks to the reduced forward-bias leakage of the bulk-source diode. For 40-nm CMOS, a 5 µm×0.2 µm P+/N-well diode conducts ≈1 nA when forward biased with the full nominal supply voltage (1.1 V) at 4.2 K [82], more than 5 orders of magnitude less than at RT. For more sensitive applications, the diode leakage can be decreased by applying a lower FBB, since the leakage decreases by ≈ 10× for a 100 mV decrease in V_{BB}, as estimated from figure 4.19. With a full FBB $V_{BB} - V_S = V_{dd} = 1.1$ V, the threshold voltage can be shifted by >200 mV in the adopted technology. Combined with the available threshold flavors, this offers a wide range of viable threshold values.

In the following subsections, we analyze two examples of circuits enabled by cryogenic-aware FBB and their limitations. For the analysis, we use data measured at RT and 4.2 K from a characterization chip, as no accurate model for simulation was available for cryogenic behavior of the adopted process at design time. Both circuits will be used in the driver design described in section 4.3.

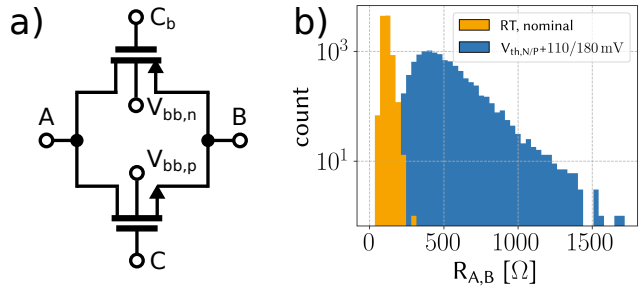

Figure 4.2: a) Pass-gate, b) RT MC simulation of pass-gate resistance at V_{CM} = 550 mV, 10^4 samples, cryogenic behavior only modeled by V_{th} increase with an equivalent series voltage.

4

4.2.1 Pass-gate for fast switching of mid-rail voltages

A pass gate (figure 4.2 a) can be easily designed to switch mid-rail voltages at room temperature, as shown by the limited spread in the monte carlo (MC) simulation of its mid-rail (550 mV) on-resistance (figure 4.2 b). Considering the V_{th} increase of 110/180 mV in NMOS/PMOS measured at cryogenic temperatures in triode for 40 nm devices, and, for simplicity, no further changes in device behavior, the standard deviation of the mid-rail resistance spread increases dramatically by >4×. Assuming no change in the spread of model parameters from RT to cryogenic temperatures may even underestimate the variation, as variability, such as device mismatch [70], typically degrades at cryogenic temperatures. To recover RT performance with traditional methods, the pass gate either needs to be significantly enlarged to contain the spread, or be replaced by a boosted or bootstrapped switch.

Alternatively, applying FBB can bring the V_{th} back to its RT value, or even below, thus reducing the on-resistance as shown in [99]. At mid-rail, the switch will also benefit from the generally increased mobility at cryogenic temperatures [44], allowing for smaller sizing than possible at RT. Although the increase in subthreshold leakage associated with a lower threshold may be a concern. This effect is contained by the about 3× steeper subthreshold slope at cryogenic temperature as reported in [44]. This allows to reduce the transistor threshold voltage even below RT values without deteriorating leakage performance.

4.2.2 DC-coupled linear inverter amplifier

The inverter amplifier, see figure 4.3 a), is a core building block of many efficient amplifier architectures, thanks to its power efficiency obtained by current reuse and the beneficial scaling with technology. At RT, this amplifier is also moderately linear when biased at mid-rail and used in a differential configuration. This is illustrated in figure 4.4, where we show the inverter transconductance ($g_m = g_{m,N} + g_{m,P}$, with $g_{m,N/P}$ the transconductane of the individual transistors) derived from measured I_d (V_d=550 mV) of individual devices. A sizeable linear region can be observed in the differential transconductance $g_{m,diff} = (g_{m,N,1}(V_{in}) + g_{m,P,1}(V_{in})) - (g_{m,N,2}(-V_{in}) + g_{m,P,2}(-V_{in}))$ of an inverter-based pseudo-differential pair figure 4.4b) at the mid-rail point. This breaks down at cryogenic temperatures, where, due to the increased threshold voltage, a significant dip in the g_m is observed, corresponding to a limited linearity. To avoid this dip and recover the linear

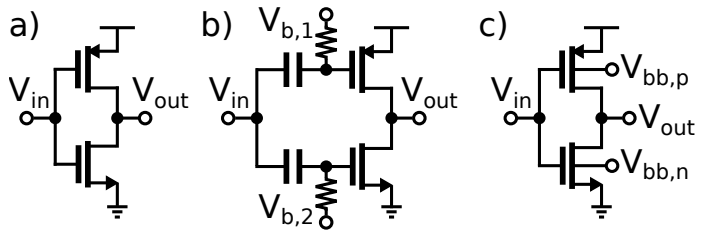

Figure 4.3: Inverter-amplifier a) DC-coupled, b) AC-coupled, c) DC-coupled with body-bias.

behavior, the 4.2 K characteristic needs to be shifted by 100/140 mV for NMOS/PMOS, see figure 4.4. We can now observe similar linearity if comparing the transconductance of the differential pair in figure 4.4b). The mobility increase at lower temperatures does not severely compromise the linearity in the nominal case, as shown in figure 4.4. Although the linearity could degrade over different process corners, RT corner simulations showed the linearity to be robust against process spread. Since no data about the process spread at cryogenic temperature has been reported to the best of the author's knowledge, we assume that the linearity at cryogenic temperature would be comparable to the one predicted by corner simulations at room temperature, as it happens for the particular case shown in Fig. 4.

For implementing this shift, we could use a bias-T as shown in figure 4.3 and applying bias voltages $V_{b,1/2}$, but the amplifier bandwidth would be reduced around DC by the bias-T high-pass characteristic and the signal would suffer attenuation due to the parasitics of the passive network. Alternatively, FBB can shift the transfer characteristics by shifting V_{th} without significantly altering the transistor characteristics. This allows recovering the linearity without introducing any additional components into the signal path and/or limiting the input bandwidth.

4.2.3 Limitations of cryogenic-aware FBB

Applying FBB via the bulk contact may be potentially limited by the high substrate resistance at cryogenic temperature, as indicated by typical N-well resistances up to a few GΩ/□ at 4.2 K [44, 75]. If such large bulk resistance (R_B in figure 4.1) would be effectively present, the applied bias V_{bb} would only set the DC operating point, around which capacitively coupled excitations could alter the bulk potential, causing unexpected effects. For instance, the capacitive coupling via the drain-bulk diode (D_D) could lower the output resistance due to modulation of the bulk potential. If floating the bulk terminal, the size of this effect is about 8% in RT simulation. The influence of the gate in this context is largely reduced due to shielding by the channel. Luckily, the field-dependent ionization might significantly reduce the effective resistance, as soon as potential differences in the order of mV build up over the bulk resistance [102], which is in-line with the steep drop in substrate resistance with increasing bulk current shown in [44]. To mitigate the effects of the unknown substrate resistance, substrate contacts can be placed near the active devices to ensure field-dependent ionization in case of potential differences. We have chosen a contact distance in the order of 1 µm in this design, maximizing the field strength while still allowing for a dense layout.

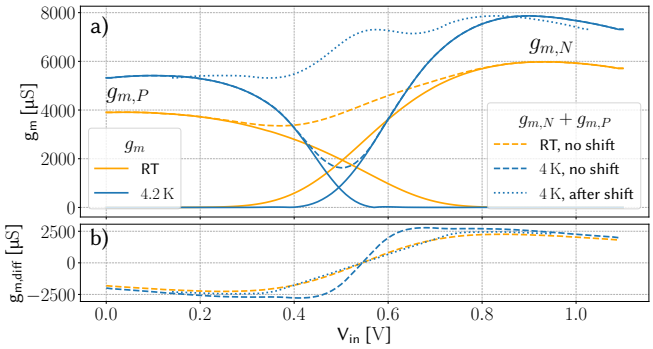

Figure 4.4: a) Transconductance (g_m) of an inverter amplifier and of its individual PMOS and NMOS the x-axis corresponds to the inverter input voltage as shown in figure 4.3. The g_m is derived from measured I_d of individual NMOS/PMOS with L=100 nm, W=1.2/2.4 µm and 6 fingers. b) Transconductance of an inverter based differential pair.

The application of body-bias is restricted by the available process. It is applicable to planar bulk technologies with a triple-well option, as well as to FDSOI technologies. Effective FBB is precluded in FinFET technologies, as these generally have a very low body factor and are therefore ill-suited for adopting body-bias [103].

If circuits employing FBB must operate both at RT and cryogenic temperatures, measures must be taken to ensure correct operation, especially when using high FBB values. For instance, to avoid excessive diode leakage at RT, the body potential must be switched depending on the operating temperature, or DACs adjusting the body-bias are required. This is not an issue for the target application in quantum-computer interfaces, which always operate at cryogenic temperatures.

If using a triple-well layout for minimizing leakage paths when employing FBB, additional area might be necessary due to the design rules of such processes, see, e.g., the layout in figure 4.13 c). Especially the distance of a deep-N-well (DNW) to an N-well (NW) of different potential typically carries a significant distance requirement. The additional area may also cause increased parasitic capacitance due to necessary routing between now spaced transistors, which may be critical for parasitic-sensitive scenarios like the input of a latching comparator. To avoid this space constraint, the PMOS can be placed in the DNW surrounding the NMOS P-well (PW). While reducing the required extra area to a minimum, this leads to some additional leakage via the P-well/N-well diode if the PMOS transistors inside the DNW are also using FBB. Additionally, this would also imply using the same body-bias for all PMOS transistors sharing the DNW.

4.3 Architecture and circuit design

The acquisition front-end in figure 4.5 comprises the ADC core with its two time-interleaved slices A, B driven by the FIA. All body biases used in the amplifier are static and generated by the on-chip DAC. The FIA and ADC are clocked by the timing generator synthesizing all necessary timing signals from a single full-rate clock signal. The front-end operates in three phases on each slice in alternation, see figure 4.6: First, during T_R, the slice is reset

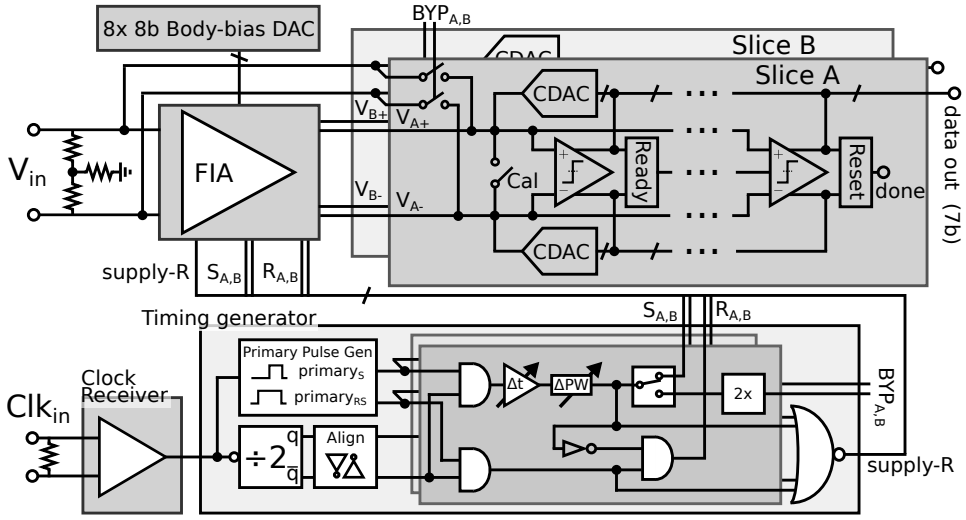

Figure 4.5: Acquisition front-end block diagram.

and its input settled to V_{CM}. Second, during T_S, the differential input signal V_{in} is amplified via windowed integration on the top-plate of the ADC sampling capacitor, $V_{A/B,+/-}$. Finally, during T_{conv}, the amplified signal is converted by the slice to the output word. The slices are 7-bit SAR ADCs that are loop-unrolled for speed and equipped with foreground calibration for the comparator offset. The slices' design is identical to [83] except for changes in the timing circuitry necessary to integrate the amplifier and a slightly increased capacitive DAC (CDAC) to retain the input voltage range of 600 $mV_{pp,d}$ after adding the amplifier parasitics. In an optional bypass-mode included to verify the ADC stand-alone performance, the FIA is disabled and the input is directly sampled on the DAC top-plates by clock-boosted sampling transistors W=1.5 µm to minimize feed-through), similar to [83].

Our target front-end specification required >50dB SFDR, >38dB SNDR when operating at a conversion rate of ≥1GS/s, see chapter 2. As the ADC slices described in chapter 3 meet these specifications, the following sections focus on the driver design. Although the target application requires only cryogenic operation, the chip was designed also for RT operation to allow RT characterization, thus easing the chip testing, and also to showcase the state-of-the-art RT performance of the proposed architecture, which can be employed also in other non-cryogenic applications.

4.3.1 Core FIA

The core differential amplifier, see figure 4.7 for the schematic and table 4.1 for the device sizes, uses the same set of amplifying inverters (M_1-M_4) for driving both ADC slices. Instantiating a separate amplifier for each slice would not result in a direct power penalty due to the fully dynamic operation but would require an extended amount of inter-slice calibration. The inverters are designed to deliver an output current signal for windowed integration, rather than settling to a voltage for the associated benefits in power efficiency

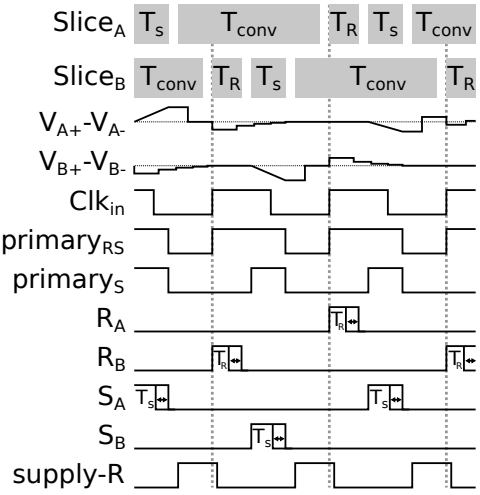

Figure 4.6: Timing diagram.

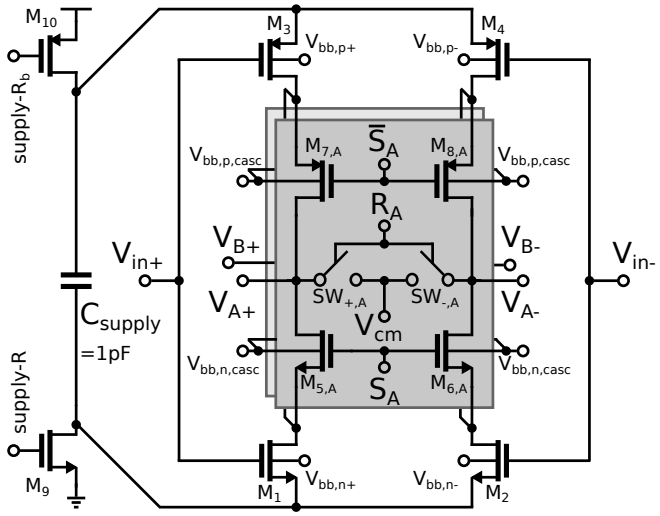

Figure 4.7: Proposed floating inverter amplifier (FIA)

[90, 94]. Therefore M_1-M_4 are chosen with a length of 100 nm to increase the intrinsic gain of the amplifying transistors, approximating an integrating behavior.

Interleaving of the shared inverters is implemented by a separate set of cascodes ($M_{5/6,A/B}$-$M_{7/8,A/B}$) and pass-gate reset switches ($SW_{+/-,A/B}$) for each of the two slices (A, B) [93]. First, during T_R, see figure 4.6, $SW_{+/-,A/B}$, controlled by $R_{A/B}$, reset the output of the amplifier to V_{CM}. In case of a metastability event causing the previous ADC slice conversion time (T_{conv}) to extend up to T_R, the *data out* bits are latched in their incomplete state and the CDAC undergoes a forced reset to avoid propagating the error to the following

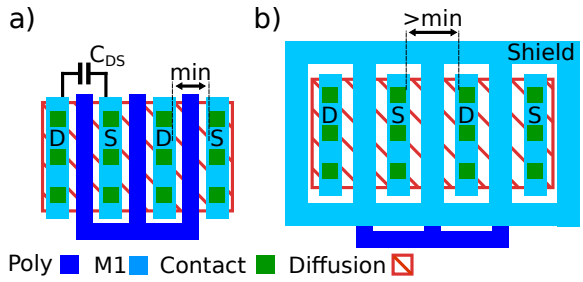

Figure 4.8: Cascode layout a) standard, b) proposed layout with improved isolation

Table 4.1: FIA core sizing

Device	W_{finger} [µm]	L [nm]	$N_{fingers}$
$M_{1/2}$	1.2	100	12
$M_{3/4}$	2.4	100	12
$M_{5/6,A/B}$	1.2	120	12
$M_{7/8,A/B}$	2.4	120	12
$M_{SW,A/B,P}$	3	40	8
$M_{SW,A/B,N}$	3	40	8
M_9	1.2	40	12
M_{10}	2.4	40	12

conversion. Then, during T_S, the cascodes connecting to the target slice are turned on using the $S_{A/B}$ signal and the input signal is integrated on the cap-DAC top-plate. This windowed-integration operation during T_S dictates the circuit transfer function, which can be approximated as [94]:

$$|H(f)| = \frac{g_m T_S}{C_{DAC}} sinc(\pi T_S f) \tag{4.2}$$

where g_m is the differential-inverter transconductance, and C_{DAC} is the load capacitance. In addition to the limited intrinsic gain of the devices, deviations from this *sinc* shape are caused by the stray capacitance at the drain of the input transistors [94]. Both the cascodes as well as the reset switches contribute charge to the output node due to charge injection and clock feed-through. This charge signal is predominantly input signal independent common-mode, with a minor differential contribution creating a slight increase in offset. The duration of $R_{A/B}$ and $S_{A/B}$ can be configured in the timing generator, see section 4.3.2. $S_{A/B}$ is shorter than 400 ps, leading to an output attenuation below 7% for a 0.5 GHz input compared to the DC gain, which is acceptable in the scope of our application. At the end of T_S, the slice conversion T_{conv} and supply reset *supply-R* are triggered. During *supply-R*, the amplifier's floating-supply capacitor C_{supply} is reset via M_9/M_{10} to ground/V_{dd}, respectively. The process continues at the next clock edge with a reset on the other slice.

The choice of a floating supply allows for the stable definition of the output common mode without using a power-hungry full-rate common-mode feedback circuit [95, 96]. Since C_{supply} is disconnected from the ground/V_{dd} supply during T_S, it acts as a floating

battery-like supply. As the current is now sourced from this floating supply, the amplifier has (ideally) no common-mode drive capability, and can therefore not alter the output common mode that was reset to V_{CM} during T_R. Both V_{CM} and V_{in} are nominally set to 550 mV, with the amplifier gain showing only minor variations withing a ±25 mV common-mode range in RT simulations. In practice, the amplifier is not fully floating due to the parasitic capacitance of C_{supply} and the core transistors towards the AC ground. The amplifiers common-mode specifications are especially important for the loop-unrolled ADC driven here, as the architecture has poor common-mode rejection caused by the common-mode dependence of the comparator offset [83]. Also the floating supply reduces the common mode gain to 0.5 in RT simulation for a small power overhead, while it would equal the differential gain without any common-mode control. With a full-scale differential output signal, the amplifier produces a 4 mV common-mode signal in extracted RT simulations, resulting in negligible comparator offset variation. This common mode signal is caused by second-order distortion in the signal inputs, that is canceled in the differential signal domain. The C_{supply} is designed to be large (1.3 pF), compared to the load cap (113 fF), largely avoiding the degenerative effect of the floating supply to enable a larger gain and sustained bandwidth during amplification. We did not target the narrow high-linearity condition outlined in [95] in favor of robustness, as the achieved linearity is sufficient for the application. The amplifier shows robust linearity performance over corners and temperature within the validity of the RT device models. The amplifier shows robust linearity performance over corners and temperature within the validity of the RT device models. This robustness, in combination with the analysis in section 4.2.2 showing how a linearity comparable to RT can be reached at cryogenic temperatures by means of a threshold shift, was used to extrapolate the cryogenic linearity behavior after application of FBB. A detailed analysis of the linearity of capacitively degenerated inverter amplifiers can be found in [95].

The cascode-sampling scheme used here replaces an otherwise needed sampling switch at the output, while also providing a small boosting of the inverter output impedance. The limitation in boosting is caused by the cascodes' operation close to triode due to the full-swing $S_{A/B}$ control signals. As the cascodes are not shared between slices A, B, mismatch in them causes differences in impedance boosting. This in turn adds a small gain error that can calibrated by the timing generator, see below. A downside of implementing interleaving with the cascodes is the introduced inter-slide feed-through via C_{DS} during $S_{A/B}$ onto the top-plate of slice B/A. This feed-through happens during the sensitive conversion phase $T_{conv,B/A}$. To address this, different strategies can be employed: to cancel the feed-through, an additional pair of cross-coupled always-off transistors could be employed as done for the switches in [104] but at the cost of significant additional capacitive load and layout complexity. In [105], the coupling capacitance was minimized by spacing the source and drain contacts apart, thus minimizing the coupling capacitance. Here, we pursue a third approach for isolation, by increasing the diffusion-contact-to-gate distance of the cascode transistors to allow for metal shielding above the gate, see figure 4.8. RT simulations shows negligible feed-through due to the residual coupling through C_{DS}. In addition to implementing the interleaving, also turning the amplifier off during *supply-R* is ensured by the cascodes being open outside $S_{A/B}$. This removes the need for additional switches at the source of the input transistors M_1-M_4 used in [95, 96], which can cause additional source

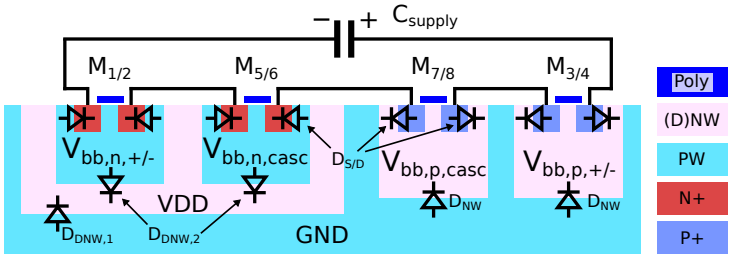

Figure 4.9: Well layout in the amplifier

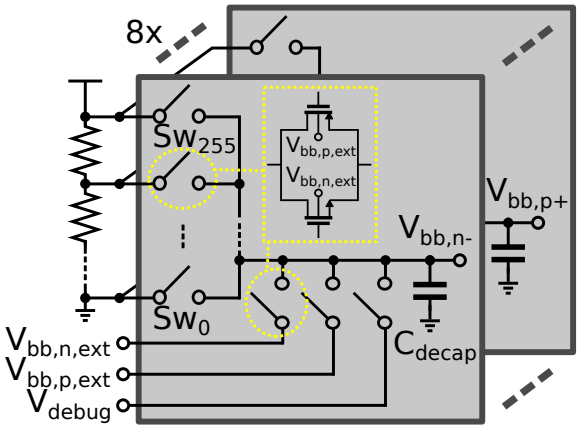

Figure 4.10: Body-bias DAC providing all 8 back-bias voltages used in the amplifier in figure 4.7.

degeneration.

The design uses back-biasing for all core transistors to enhance operation at cryogenic temperatures. Most importantly, the input transistors M_1-M_4 need to be back-biased at cryogenic temperatures if using DC coupling, as discussed in section 4.2.2. By biasing the body of the input transistors separately ($V_{bb,n+/-}$ and $V_{bb,p+/-}$), we also allow for input offset cancellation. We target an offset of 1 LSB to avoid significant SNDR degradation, which dictates the body-bias DAC resolution. For an expected gain of 7, an LSB of ≈5 mV at the ADC input, a body-bias factor $\zeta = 0.25$, and a total DAC range of 1.1 V we require approximately 8b resolution to cover the expected mismatch range when applying the body-bias to one of the four input transistors M_1-M_4. To get a reliable pass-gate operation, the complementary transistors in $SW_{+/-,A/B}$ need to be back-biased, as discussed in section 4.2.1. And finally, back-bias can also be applied to the cascode transistors $M_{5,A/B\text{-}8,A/B}$ for additional swing, avoiding the cascode transistors driving the input pairs towards triode. According to RT simulation, we would be able to adjust for the expected increase in V_{th} at cryogenic temperatures and recover the target driver linearity of >50 dB. In addition to enabling cryogenic operation, the adjustable body-bias also allows for compensation of the process spread affecting open-loop amplifiers, as the spread in the threshold can now be compensated in the field.

As discussed in section 4.2, FBB can cause leakage by forward-biasing the device diodes. To identify possible sources of leakage, we show a sketch of the amplifiers' well layout in figure 4.9, which does not differ from the usual layout in a triple-well process. The problematic diodes in this context are formed by the source/drain diffusion of transistors (labeled $D_{S/D}$ in figure 4.9). All well-to-well diodes ($D_{DNW,1}$, $D_{DNW,2}$ and D_{NW}) are never forward biased for FBB within the supply rails. Among the $D_{S/D}$ diodes, the worst-case for leakage is found at the source of the cascodes $M_{5,A/B}$-$M_{8,A/B}$ when a full nominal supply is applied as FBB. During reset, M_1-M_4 are in triode and the supply is reset to the nominal ground/V_{dd} rails. Hence, the forward voltage for the source-bulk diodes of the cascode is a full V_{dd}. This leads to an estimated leakage of 58 nA from the PMOS cascode onto a node in reset, see discussion of figure 4.19, causing only additional power dissipation at a negligible magnitude in the context of our application. All other diodes carry less FBB, specifically the ones connecting to the ADC top-plate, and are therefore not expected to contribute measurable effects.

The static body-bias DAC uses a simple resistive ladder between ground and V_{dd}, which is tapped by a set of switches addressed by binary decoders, see figure 4.10. For compactness, the DAC uses the surrounding DNW to contain all PMOS circuitry, as explained in section 4.2.3. As the DAC is fully passive, decoupling is added at the output to isolate the resistive ladder from kickback. The small-size pass-gates implementing these switches must be operational for switching mid-rail voltages at cryogenic temperatures. To ensure that, in this prototype chip the switches themselves are also back-biased by externally supplied voltages $V_{bb,n/p,ext}$. In a future iteration, these voltages can be generated with a low-resolution and low accuracy DAC, as the necessary body-bias for guaranteeing full functionality (around 0.4 V for NMOS, 0.7 V for PMOS) are easily switchable by switches without body-bias and low-precision is acceptable for these biases. The DAC also allows for using the external voltages $V_{bb,n/p,ext}$ instead of the resistive ladder, as well as read-back of the control voltages to detect abnormalities via V_{debug}, connected to a pad.

4.3.2 Timing generation

The timing-generation block, see figure 4.5, produces all pulses shown in figure 4.6 from the full-rate input clock. The output of the pseudo-differential clock receiver is divided and aligned on the negative clock edge, while the primary pulse is initiated at each positive edge. The entire timing calibration block, except for the clock divider, is implemented with open-loop delays and combinational logic. This saves power compared to using the high-frequency clock required to produce all the phases and fine-grain adjustments necessary here. A DLL-based alternative would improve the robustness but at the cost of increased power consumption and design complexity. Care was therefore taken to make the delay-based logic robust to PVT variations by only using relative delays and carefully matching driving capabilities of parallel paths, thus achieving reliable operation from RT to 4.2 K.

The primary pulse generator (figure 4.11) is shared between both slices to avoid the additional calibration necessary to generate the control pulses via separate blocks. The produced pulses are multiplexed in the timing generator, see figure 4.5. Both T_R and T_S are ideally kept short to allow more conversion time for the ADC, and are adjustable from 120 ps to 400 ps. For applications requiring the FIA gain to be robust against extended

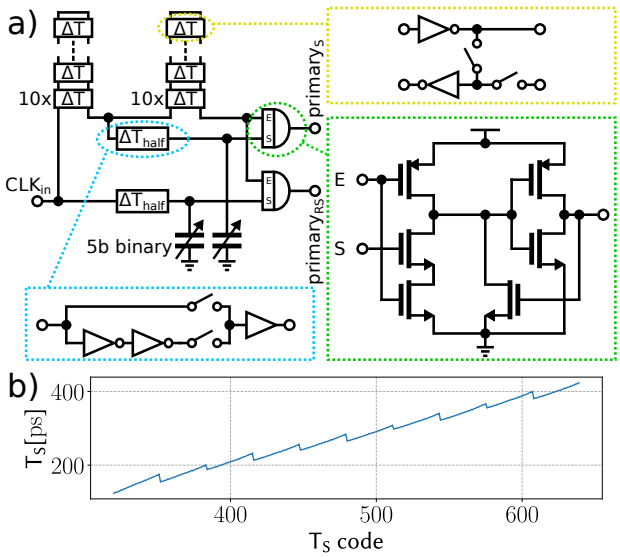

Figure 4.11: a) Master pulse generator b) RT simulation of pulse-width control.

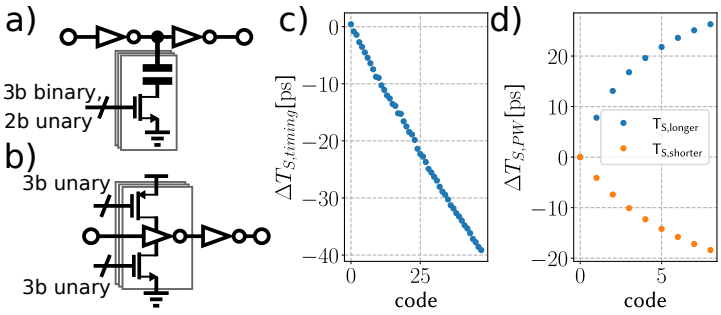

Figure 4.12: a) Timing calibration, b) pulse-width calibration, c) RT simulation of timing calibration, d) RT simulation of pulse-width calibration.

PVT variations, circuit techniques as proposed in [106] can be employed. To generate this range, three functions are used: a full delay step (ΔT) defined by the combined delays of two inverter delays and a pass-gate, a half step (ΔT_{half}) corresponding to two inverter delays, and a 5b binary weighed capacitor array for fine steps. While the main effect of adjusting T_S is varying the amplifier gain, the duration of T_S also affects the inherent filtering introduced by the windowed integration [94]. As the windowed integration corresponds to a *sinc* response, this could allow, for example, adjustment of the notch to reject a spurious out-of-band tone like mixer LO feed-through.

While most of the timing blocks are shared between the slices, the non-shared sections cause inter-slice mismatch, among which, the relative pulse timing mismatch $\Delta T_{S,timing}$ and gain mismatch caused by pulse-width mismatch $\Delta T_{S,PW}$. $\Delta T_{S,timing}$ is calibrated by a

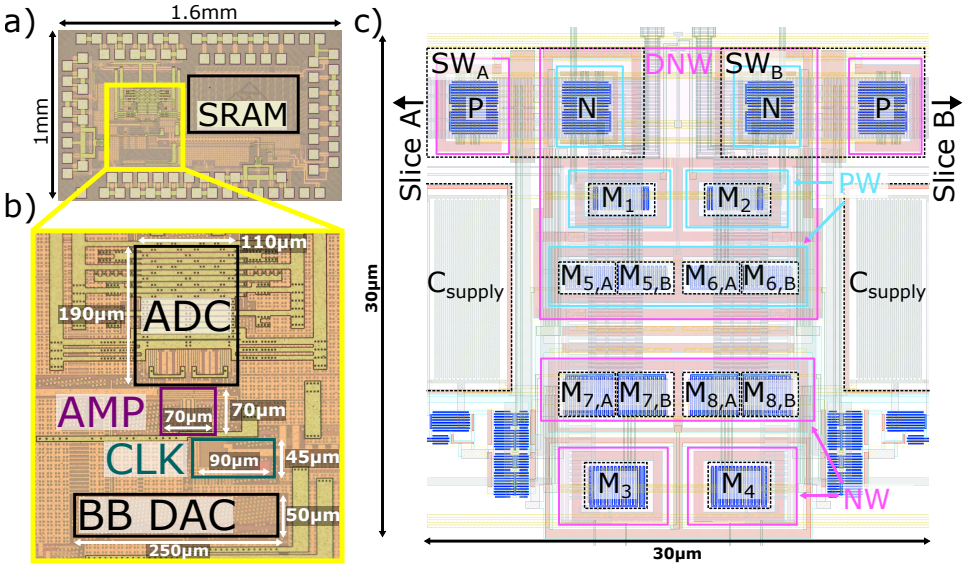

Figure 4.13: a) Micrograph of the test chip; b) Micrograph of core analog blocks, c) Layout details of the amplifier core.

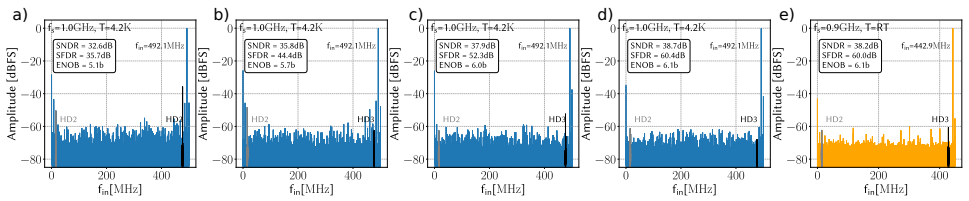

Figure 4.14: Measured spectrum at 4.2 K a) without FBB, b) with FBB on the input pair, c) with FBB on the input pair and the reset switches, d) with FBB on the input pair, the reset switches and the cascodes. e) Spectrum at RT

capacitor array with a 3b binary and 2b unary control, see figure 4.12 a), which allows for delay adjustments for each slice up to 40 ps in ≈1 ps steps. The $\Delta T_{S,PW}$ calibration allows for gain calibration by adjusting unary-coded inverter weights, see figure 4.12b), allowing to calibrate up to ±20 ps of mismatch per slice. The described calibration circuits are sufficient to reduce the interleaving spurs below 60 dBc in RT simulations. The total jitter contributed by the timing generation is about 0.5 ps/0.6 ps for the rising/falling edge of $S_{A/B}$ in extracted simulation, and therefore not limiting the amplifiers SNR [94]. Typical quantum computing systems require spectral purity significantly beyond this level [107], resulting in no additional system constraints due to the amplifier.

4.4 Measurement results

A micrograph of the test chip, implemented in a 40 nm LP bulk technology, is shown in figure 4.13 a), with more details of the analog core in b). In the amplifier layout (figure 4.13 c),

we minimized the distance of the active devices to the body contacts, keeping it below 1 µm for most of the circuit. The triple-well layout shown in figure 4.9 consumes more area than minimally required (approximately 4×), but such an increase is insignificant compared to the size of the floating capacitor C_{supply} or the ADC slice.

The chip was tested in a dip-stick setup with chip-on-board assembly, similar to the test-setup in [98]. We concentrate on testing the ADC with the driver, as the stand-alone ADC achieves performance similar to [83] thanks to the minor changes in slice design. Both input and clock signals are provided by a single signal generator (SMA100B) and converted to differential signals by on-board baluns (BAL-3SMG). The conversion result is recorded at full rate in the on-chip SRAM and then read back via a low-speed opto-coupled serial link to an RT FPGA for analysis. The timing and ADC slice calibrations are performed in the foreground via loop-back through the RT equipment. The calibration decks differ between RT and cryogenic temperature due to the drastic changes in transistor characteristics. For applications requiring background calibration due to higher expected PVT variations than in the target use case, calibration techniques as proposed in [108] may be applied. All reported measurements have been performed under the following conditions unless otherwise noted: all supplies are kept at the nominal value of 1.1 V, the amplifier input and output common mode are set to 550 mV, the gain is set to 6.6/8.9 at RT/4.2 K, corresponding to the same timing calibration setting (T_S *code* in Fig. 11b)). This gain was chosen as a representative value in the mid-range of available gain settings, see figure 4.17 b). SFDR/SNDR values are always excluding the spurs at DC and Nyquist, as these are outside the band of interest for the target application.

In figure 4.14, we activate the body-biasing for different parts of the circuit in succession to observe their influence on the amplifier performance. With no body-bias applied (figure 4.14 a), the circuit is still operational at 4.2 K but shows numerous spurious tones, with the 3rd harmonic dominating at 35.7 dB due to the input inverters entering weak inversion in the middle of the amplifier input voltage range, as discussed in section 4.2.2. In figure 4.14 b), turning on the FBB on the input pair (with 539/−669 mV for NMOS/PMOS similar to the expectation from section 4.2.2) leaves the 2nd harmonic as the dominating spur. This is attributed to the incomplete reset via $SW_{+/-,A/B}$ leaving significant differences in starting condition between the two slices. In figure 4.14 c), now activating a full V_{bb} of FBB on the reset switches, the performance reaches the design target (SFDR > 50 dB), with the 3rd harmonic dominating, This is expected to be caused by the input pair being compressed by the cascodes towards the edge of triode for part of the swing. In figure 4.14 d), a full V_{bb} of FBB is also applied to the cascodes, achieving the optimal SFDR performance. It is important to note that in the body-bias calibration, the 8b body-bias DAC is necessary only for the input pair's offset cancellation and threshold compensation, as the switches and cascodes are operated at the inverted supply. The spectrum also demonstrates that the timing calibration reduces the gain and timing mismatch spurs sufficiently not to limit the amplifiers performance. figure 4.14 e) shows that RT performance is similar, but at slightly lower sampling speed, as discussed next.

In figure 4.15 a), we show the flexibility of the proposed amplifier over a wide range of sampling frequencies. The circuit has a speed advantage when operating at lower temperature, thanks to the speedup of the ADC logic [72]. As required, the SFDR stays over 50 dB at maximum sampling speed over the entire bandwidth (figure 4.15 b), in accordance

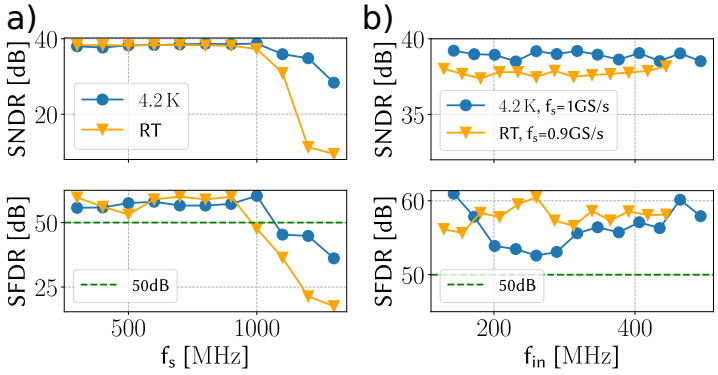

Figure 4.15: Measured a) SNDR/SFDR vs. f_s @ Nyquist b) SNDR/SFDR vs f_{in}.

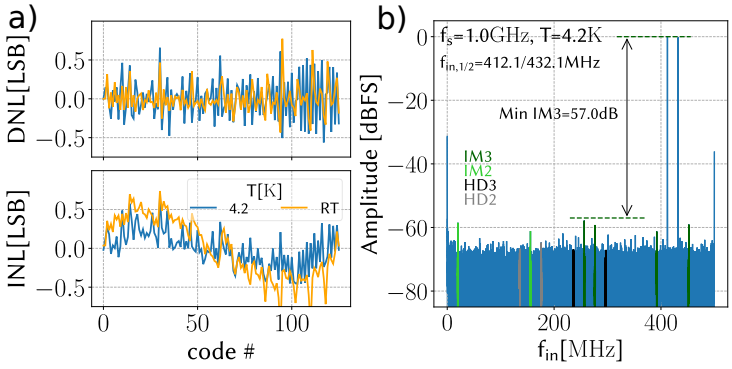

Figure 4.16: Measured a) INL/DNL, b) two-tone test.

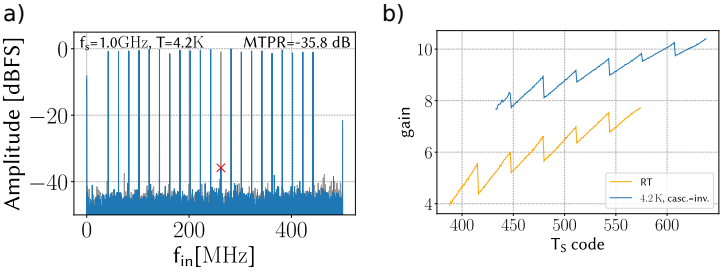

Figure 4.17: a) Multi-tone power ratio, b) Measured gain at RT/4.2 K; the x-axis is the same as in Fig. 11b.

with RT simulation expectation.

In figure 4.16 a) we show the circuits DNL and INL that is on the order of half an LSB, likely limited by the comparator calibration accuracy. In figure 4.16 b), a two-tone test, performed with two signal generators (both SMA100B) and a passive combiner, yields a maximum IM3 spur at 57dB. If exciting the input with a comb of continuous wave (CW)

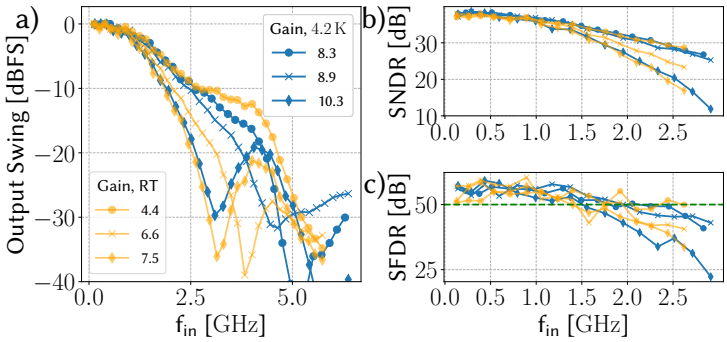

Figure 4.18: a) Measured magnitude response beyond first Nyquist band, b) SNDR vs f_{in}, c) SFDR vs f_{in}, 4.2 K@1 GS/s, RT@0.9 GS/s

Table 4.2: Gain setting overview

T_S setting	min	nominal	max
Gain@RT	4.4	6.6	7.5
T_S@RT [ps]	172	256	317
Gain@4.2 K	8.3	8.9	10.3
T_S@4.2 K [ps]	192	225	323

tones produced by a VSG (SMW200A) and removing one tone, we observe a multi-tone power ratio of 35.8dB (figure 4.17 a)), indicating sufficient isolation between the channels.

The VGA functionality is demonstrated in figure 4.17 b) by setting the gain (4.0-7.9/7.7-10.4 at RT/4.2 K) by using the pulse-width control (T_S) showcased in figure 4.11 b). The deviation from a linear scaling with T_S *code* is caused by the limited intrinsic gain of the amplifying transistors. The lower bound in the gain range is due to the impossibility of reliably triggering the ADC conversion with a very short T_S pulse. In a future redesign, this could be easily solved by using the wider $primary_{RS}$ pulse, see figure 4.6, to trigger the ADC. At RT, the maximum gain setting is limited by the time T_S reducing the conversion time available to the ADC slice, while this is not a limitation at 4.2 K thanks to the ADC slice being faster. The higher gain at cryogenic temperature can be traced to the increased g_m at cryogenic temperature [44].

In figure 4.18, we explore the circuit behavior beyond the first Nyquist zone for various gain settings, while keeping a constant input signal amplitude and calibration. The amplifiers' output swing, normalized to the swing at low frequency, shows approximately the expected *sinc* shape of equation (4.2) and allows for estimation of T_S in the circuit, shown in table 4.2. The deviations from the ideal *sinc* shape are caused by the parasitic capacitance at the cascode node [94]. We can observe sustained SFDR performance >50 dB in the 3rd Nyquist zone, as SFDR tracks the driver output swing. The SNDR performance drops as the swing reaching the ADC input is reduced due to the *sinc* shaped transfer characteristic reducing the circuit gain beyond the 2nd Nyquist zone significantly, limiting sub-sampling operation to this zone.

Table 4.3: Comparison table

	This work		Kiene, ESSCIRC 2022 [21]		Kiene, JSSC 2023 [9]		Kull, ISSCC 2013	D. Li, JSSC 2020	Duan, JSSC 2014 [31]	Malki, JSSC 2014 [14]	Jiang, TCAS 2021 [16]	Miki, JSSC 2017 [11]
Temperature [K]	RT	4.2	RT	4.2	RT	4.2	RT	RT	RT	RT	RT	RT
Architecture	TI SAR		TI SAR		TI SAR		SAR	SAR	TI SAR	SAR	TI SAR	TI SAR
Max f_s[MS/s]	900	1000	1000		900	1000	1300	900	12800	80	2000	2000
Resolution [bit]	7		7		6-8		8	7	7	10	7	8
Technology [nm]	40		40		40		32	40	65	40	40	40
Supplies [V]	1.1		1.1		1.1, 2.5		1	1.1	1.2	1.1	1.2/0.9	
SNDR@Nyquist [dB]	38.2	38.7	38.2	41.1	33.4	36.2	39.3	39.7	26.4	52.3	36.4	39.4
SFDR@Nyquist [dBc]	>50	>50	>50	>50	48.4	48.5	49.6	54.8	32.4	55[3]	47.8	55
Power [mW]	1.87[1,2]	1.79[1,2]	1.94[1]		10.3[1]	10.6[1]	3.1	2.6	162	2.86 (7.16[2])	7.62	54.2
FoM$_w$[fJ/c.step]	31.3[1,2]	25.4[1,2]	29.2[1]	20.9[1]	260[1]	200[1]	28	36.6	740	106 (267[2])	70.8	355
Core area [mm^2]	0.042[1,2]		0.042[1]		0.045[1]		0.0015	0.014	0.23	0.065	0.0082	0.54
ADC driver	✓		×		×		×	×	✓	✓	✓	✓

[1]Full ADC, clock receiver and all timing circuitry [2]including amplifier [3]estimated from plot

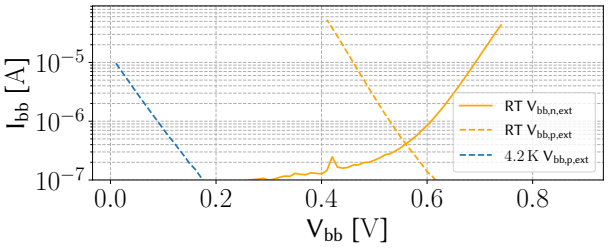

Figure 4.19: Back-bias leakage for the full DAC; the 4.8 K. The leakage on $V_{bb,n,ext}$ was below the measurement floor ($\approx 1 \times 10^{-7}$ A) and therefore not plotted.

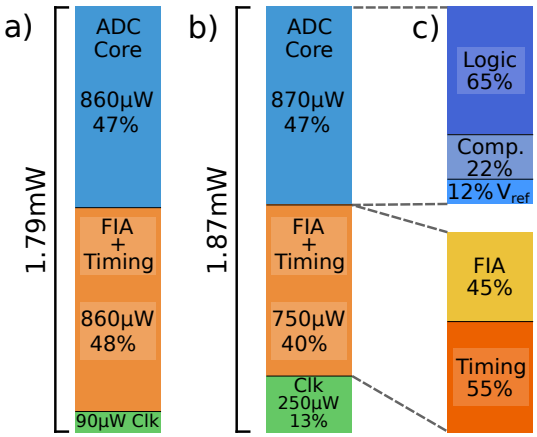

Figure 4.20: Measured power: a) 4.2 K for f_S =1 GS/s, b) RT for f_S =0.9 GS/s, b) simulated breakdown at RT.

To estimate the junction leakage due to FBB, figure 4.19 reports the back-bias leakage of the full body-bias DAC, measured via $V_{bb,n,ext}/V_{bb,p,ext}$ in figure 4.10. For the measurement, the biases are swept individually with the other terminal fixed to its nominal supply. The measured leakage is produced by a total width of 4.7 mm contributed by each NMOS/PMOS

transistor in the decoders and switches of the body-bias DAC. As mentioned in section 4.2.3, this does include the DNW to PW diode leakage thus making the leakage in figure 4.19 an upper bound for a layout avoiding these diodes as in figure 4.9. At RT we measure significant leakage upon reaching the diode thresholds. At 4.2 K, we observe no leakage above the measurement noise floor of 100 nA for NMOS and up to 10 µA for the PMOS. Normalized to the total transistor width, this corresponds to a leakage below 20 pA/µm (2 nA/µm) for NMOS (PMOS) over the full FBB range, and below the measurement noise floor within the bias ranges used for the DAC in cryogenic measurements, i.e., 0.7/0.4 V for NMOS/PMOS in all measurements shown here. We did not find any signature possibly caused by diode leakage in any measurement, also when observing the DAC output voltages via V_{debug}. This demonstrates that, except for extremely leakage-sensitive circuits, cryogenic-aware FBB opens new design options for bulk-CMOS circuits.

Looking at the power breakdown in figure 4.20, the power is approximately equally split between the ADC core and the FIA (including the timing generation for both the FIA and the ADC). Simulations at RT in figure 4.20 c) show that the core ADC power is dominated by the logic, while the FIA core and the timing generation use approximately the same power, about half of the combined power. Despite the need for timing circuitry for the dynamic amplifier, this performance results in a FOM_W of 31.3 and 25.4 pJ/conv.-step at RT and 4.2 K, respectively, with the FIA core power (excluding the timing) degrading the FOM_W by only 6.6/5.5 pJ/conv.-step at RT/4.2 K, thus demonstrating the efficiency of the proposed driver.

When compared to prior ADCs with similar sample rate and resolution at RT and 4.2 K in section 3.3.2, the proposed ADC achieves comparable FOM_W while also including the driving amplifier. Among the ADCs including a driver, we improve the FOM_W by 2× over the state-of-the-art at RT and report the first ADC with a dynamic driver at 4.2 K.

4.5 Conclusion

In this chapter, we have presented an FIA amplifier driving a time-interleaved SAR ADC at RT and 4.2 K. The driver pioneers the extensive usage of FBB in bulk technologies in cryo-CMOS analog circuit design, thus enabling cryo-CMOS designers to use techniques and topologies that were usually confined to RT applications. The proposed driver uses an effective combination of dynamic amplification, floating supply, cascode sampling and cryogenic-aware FBB to efficiently drive interleaved SAR ADCs. The design also shows the reliable performance of a dynamic amplifier under an extreme temperature variation, irrespective of the drastic changes in all transistor parameters. To the authors' knowledge, this is the first reported dynamic ADC driver operating at cryogenic temperatures. Furthermore, the proposed circuit achieves the best FOM among state-of-the-art RT ADC with a driver and comparable FOM among cryogenic and RT ADC operating at similar sampling speeds and resolution while also including the driver. As a result, the combination of the dynamic driver introduced in this chapters and the SAR ADC presented in this and the previous chapters represents a power-efficient compact solution for the mixed-signal section of the cryo-CMOS SoC for the RF frequency-multiplexed readout of spin qubits.

5

Low-Frequency Noise in Cryo-CMOS Devices

This chapter presents an extensive characterization of the low-frequency noise (LFN) at room temperature (RT) and cryogenic temperature (4.2 K) of 40-nm bulk-CMOS transistors. The noise is measured over a wide range of bias conditions and geometries to generate a comprehensive overview of LFN in this technology. While the RT results are in line with the literature and the foundry models, the cryogenic behavior diverges in many aspects. These deviations include changes with respect to RT in magnitude and bias dependence that are conditional on transistor type and geometry, and even an additional systematic Lorentzian feature that is common among individual devices. Furthermore, we find the scaling of the average LFN with the area and its variability to be similar between RT and 4.2 K, with the cryogenic scaling reported systematically for the first time. The findings suggest that, as no consistent decrease of LFN at lower temperatures is observed while the white noise is reduced, the impact of LFN for precision analog design at cryogenic temperatures gains a more predominant role.

This chapter has been published as: *Gerd Kiene, Sadik Ilik, Luigi Mastrodomenico, Masoud Babaie, and Fabio Sebastiano, Cryogenic Characterization of Low-Frequency Noise in 40-nm CMOS, IEEE Journal of the Electron Devices Society (2024) [109].*

5.1 Introduction

Cryo-CMOS design in general and the design of readout circuits considered in this thesis in particular face significant challenges due to the incomplete understanding of cryogenic transistor behavior, as discussed in e.g. chapter 2. These challenges frequently cause over-design or even design failures. To address the situation, extensive characterization and modeling of DC and RF behavior of transistors have been performed, e.g. in [51, 52].

Within the cryogenic behavior, noise is generally less explored as a topic, but prior work on both broad-band noise and low-frequency noise has been reported. In broad-band noise characterization, the noise is scaling less than expected from a purely thermal origin, thus suggesting that shot noise dominates at cryogenic temperatures [53], also discussed and explored in comparator noise measurements in chapter 3. For designing systems at cryogenic temperature, LFN also plays a crucial role, e.g., for the phase noise of phase-locked loops (PLL) [107] and the input-referred noise of transimpedance amplifiers (TIA) [37]. As discussed in detail in chapter 6, LFN may represent the main limitation in the performance of DC-readout circuits for spin qubits. Hence, this chapter experimentally investigates the LFN of cryo-CMOS devices, so that fundamental limitations for spin-qubit DC readout can be identified.

Early work on cryogenic LFN provided evidence pointing to both carrier number and mobility fluctuations as possible origins of the noise at low temperatures [110, 111]. In [112], measurements of a modern SOI process over a range of bias points are shown to be compatible with the model including carrier number fluctuations and correlated mobility fluctuations [113]. In [54], the lack of the expected decrease of LFN at cryogenic temperatures is attributed to band-tail states acting as traps. This is corroborated by the same band-tail states possibly causing the saturation of the MOSFET sub-threshold slope at cryogenic temperatures [114, 115]. In [116], characterization and modeling of large-area devices manufactured on different crystal orientations suggest carrier number fluctuation caused by interface traps as an alternative explanation for cryogenic LNF. While these characterizations, typically only of one or few devices, have added significantly to the understanding of LFN at cryogenic temperature, they can only partially inform design decisions.

Missing in current literature is a more extensive characterization that can both guide the circuit designers with accurate predictions and help understand the physical origin of LFN. To fill this gap, we measure several different geometries and characterize them over a wide range of bias voltages, as commonly required for design space explorations. Multiple individual devices with the same geometry are measured at the same bias, since statistical characterization is crucial for accurate analysis due to the large variations between individual devices, even for large-area devices [117], which are already present at RT and hence expected also at cryogenic temperatures. Such an extensive characterization enables us to also perform a systematic analysis of the scaling of LFN and its variability with device area, for the first time at cryogenic temperature. Furthermore, the analysis of multiple equal devices enabled the discrimination of a Lorentzian feature that is unexpectedly systematic among different devices. This feature is dependent on device geometry, bias, and temperature, and causes increased device noise that is detrimental to cryogenic circuits.

This chapter is organized as follows. We begin by introducing the experimental methods

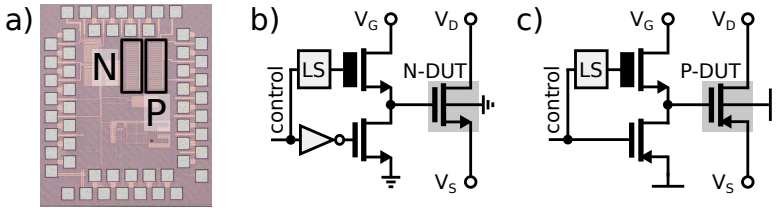

Figure 5.1: a) Micrograph of the test chip; b) Schematic of each NMOS; c) Schematic of each PMOS. The thick-oxide selection transistors are shown with a thicker-gate symbol.

in section 5.2, then present the measurement results and analyze them in section 5.3, discuss the results in the light of their impact on precision analog design in section 5.4 and draw the conclusions in section 5.5.

5.2 Methods

The measurement setup was designed to facilitate the characterization of many devices at cryogenic temperature, as necessary to generate the statistics for LFN. For this, a chip was fabricated in 40 nm LP CMOS bulk technology (figure 5.1 a), which carried 8 individual transistors for each of the several geometries, for a total of 400 heavily multiplexed devices under test (DUT) per chip. In the chip, the devices' source and drain (V_d, V_s) are connected in parallel. Multiplexing between them is implemented, similar to [118], by connecting a single device gate at a time to the gate bias (V_g), see figure 5.1 b) and c) for the NMOS and PMOS schematic. For reliable rail-to-rail operation, multiplexing of V_g is implemented with thick-oxide transistors, supplied by 2.5 V. All deselected transistors are biased with V_{gs}=0, and the source is kept at ground/V_{dd} for all NMOS/PMOS measurements. Unless otherwise noted, all measured transistors are of low-threshold-voltage (LVT) flavor.

Assembled in a DIP package, the chip is mounted on a PCB placed in a dipstick setup and submerged in liquid helium with a temperature sensor mounted close to the sample. Such a liquid cooling at 4.2 K results in a stable and accurate definition of the ambient temperature, unlike prior characterization in the vacuum environment of the typical cryogenic probe stations, which may suffer from thermalization issues. For the RT measurements, the ambient temperature was controlled to within ±2.5 K around 295 K. This setup allows for up to 5 days of continuous measurement at cryogenic temperatures using a single 100 L helium dewar, without labor-intensive thermal cycles. The sample was kept mounted in the same dipstick for both RT and cryogenic measurements. Measurements are performed via a resistive bias-T followed by a transimpedance amplifier (TIA) and digitized by an acquisition card. The resulting data is accurate over the frequency range from 1 Hz (or 5 Hz, depending on settings of the bias-T) up to 50 kHz, limited by the TIA bandwidth. The reported sweeps over V_{gs} were performed in two bias configurations: one in the linear region with fixed V_{ds}=50 mV and one in an effective "diode" connection with V_{ds}=V_{gs}. The analyzed data is filtered with a rolling median filter for better visibility. Further details about the measurement setup and the data processing can be found in appendix A.

All data and analysis code used for the figures in this paper is available in [119].

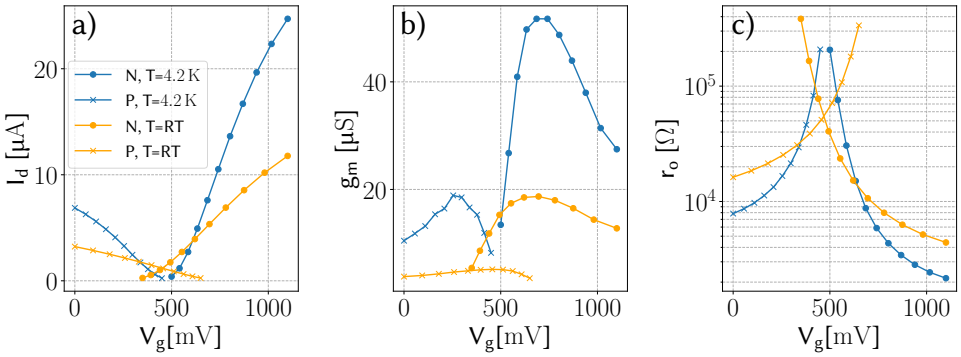

Figure 5.2: Example N/PMOS (1 µm×1 µm) DC-characteristics at V_{ds}=50 mV: a) drain current I_d, b) transconductance g_m, c) output resistance r_o. Source of NMOS/PMOS at 0/1.1 V.

5.3 Characterization results

An example DC characterization of an N/PMOS with W×L=1 µm×1 µm is shown in figure 5.2. The characterization shows the typical changes in behavior when moving to cryogenic temperature: the increased threshold voltage (V_{th}) and the increased transconductance (g_m) as well as the reduced output resistance (r_o) in strong inversion [44].

Example noise spectra of eight 1 µm×1 µm devices at a single bias point V_{gs}=V_{ds}=1.1 V for both N/PMOS are shown in figure 5.3. The plots also include a conservative noise-floor estimation based on the measured noise floor of the test equipment scaled by the expected additional noise due to the transistor r_o, further discussed in appendix A. Alongside the individual spectra, a logarithmic mean curve is shown, around which individual devices differ significantly at RT, as also reported in [117], due to the random distribution of single traps. This effect is here, for the first time, shown to persist also at cryogenic temperatures. The RT mean shows typical 1/f behavior and matches the foundry device model well. The cryogenic results differ, with the NMOS devices showing increased output-referred noise and the PMOS devices showing reduced noise. While the variability at lower frequencies is qualitatively similar to the RT results, a systematic deviation from the 1/f shape appears in the spectra at higher frequencies.

To substantiate this, we plot in figure 5.4 the mean spectra for W×L=1 µm×1 µm and W×L=1 µm×40 nm NMOS and PMOS devices for 11 logarithmically spaced overdrive-voltage biases (V_{od}=V_{gs}-V_{th}), which are color-coded in the plot. We concentrate on the mean and do not show the individual devices for better visibility of the effect. The device geometries were chosen to provide examples of typical devices preferred in analog (high intrinsic gain, long channels) and RF (high speed, short channels) circuits. For the RT results, we see again that the mean noise approximates a 1/f slope and matches well the foundry device model over a wide range of bias conditions. The 1/f slope is also visible in the cryogenic results, but a systematic Lorentzian feature appears in the mean curves at higher frequencies. Signatures of this feature are visible in all the devices (NMOS and PMOS, long and short channels) with the effect appearing at lower frequency in longer devices. As this effect appears in all the individual devices, it does not average out to a 1/f behavior as otherwise common for random traps. To understand its possible origin, we

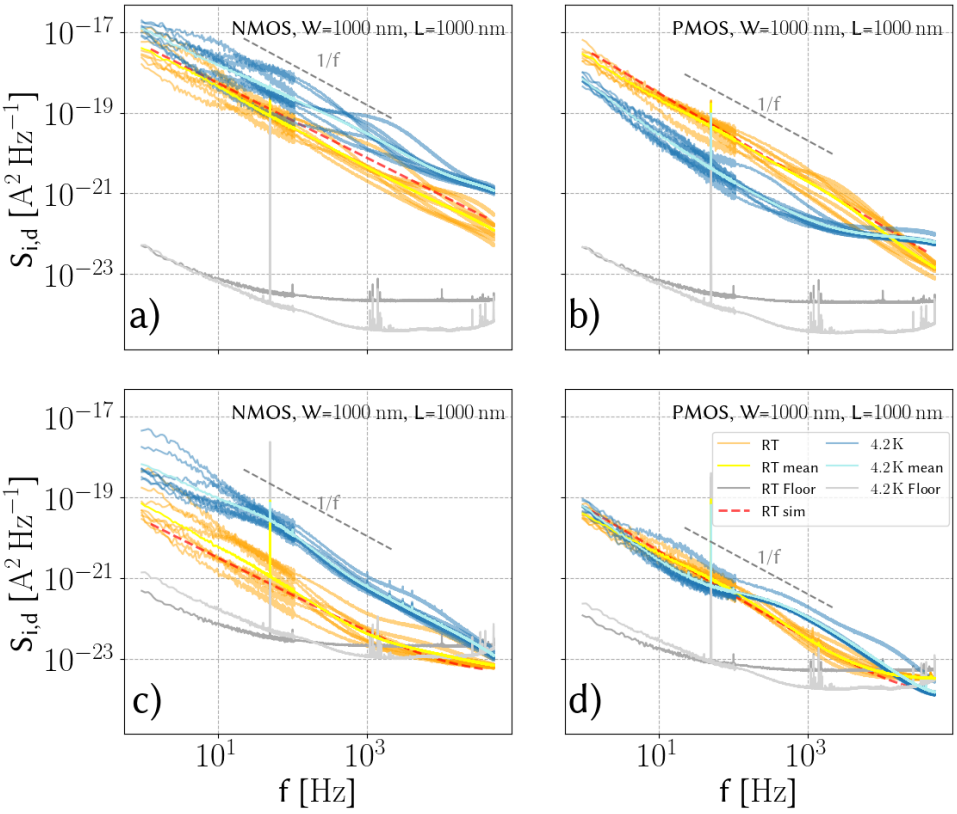

Figure 5.3: Noise spectra for NMOS (a, c) and PMOS (b, d) with W×L =1 µm×1 µm at V_{gs}=1.1 V and V_{ds}=1.1 V (a,b) and V_{ds}=50 mV (c,d), respectively. 8 devices at the same bias points are shown.

analyze this Lorentzian in more detail in the following subsection.

5.3.1 Systematic Lorentzian

The systematic Lorentzian feature common to multiple devices and mentioned above has not been reported in prior works. To ensure that it is not an artifact of the measurements, we performed additional verifications. As the equipment and the cabling adopted in the cryogenic setup are identical to the RT setup, for which the Lorentzian does not appear, we specifically investigate the components cooled to cryogenic temperature, including the chip itself and the associated passive components (resistors, capacitors) on the measurement board. Leakage currents from the deselected transistors are excluded by recording spectra with all transistors deselected and sweeping the drain bias. The on-chip multiplexing was tested by reducing the thick-oxide switch supply by 500 mV without measurable effects on either the DUT's DC behavior or its noise spectrum. The on-board tantalum capacitors have been excluded by measuring the DUT on a board without populating those, resulting in increased interference but the same effect visible. For testing the thin-film resistors, we

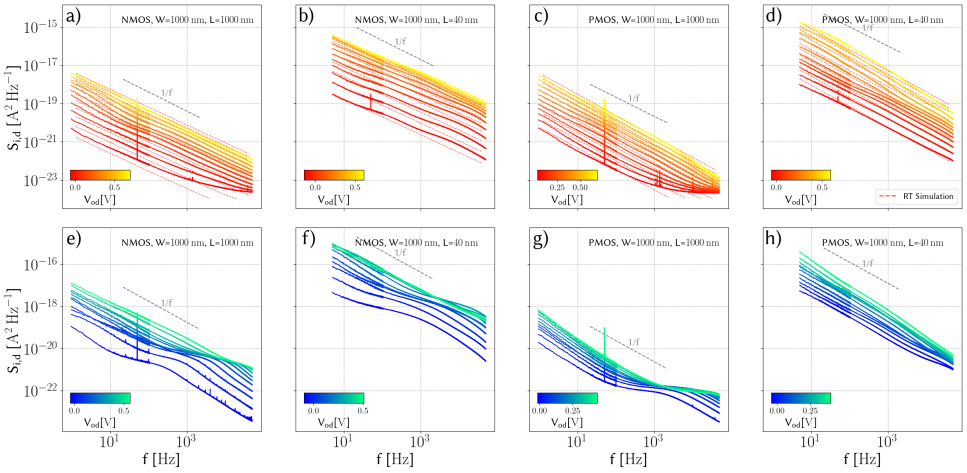

Figure 5.4: Mean noise spectra for PMOS and NMOS devices with W×L=1 µm×1 µm and W×L=1 µm×40 nm with V_{gs}=V_{ds} for logarithmically spaced overdrive-voltage (V_{od}=V_{gs}-V_{th}) at RT (a-d) and 4.2 K (e-h). Each curve represents the mean of 8 individual devices. The RT plots also show the corresponding simulations of the foundry model with dashed lines.

5

replaced the chip with resistors and ran noise acquisitions under various bias conditions, without observing measurable LFN contributed by the resistors. The additional verifications are discussed in more detail in appendix A.

To gain a deeper understanding of the Lorentzian effect, we characterized large-area (4 µm×4 µm) NMOS and PMOS devices (figure 5.5), since they show a relatively low mismatch and the Lorentzian feature appears more distinctly compared to small devices, as visible in figure 5.4. Figure 5.5 shows that the Lorentzian's corner moves to higher frequencies for higher overdrive. As we observe negligible individual differences in these large devices concerning the Lorentzian, we performed, in the following, a more detailed series of measurements on a single device, which is shown in figure 5.6.

Figure 5.6 a) shows a fine sweep of V_{gs} of a single device, from which the corner frequency $f_{C,Lorentzian}$ as function of the overdrive voltage has been extracted and plotted in figure 5.6 b). The strong dependence of the corner frequency with the bias can be attributed to barrier lowering or barrier thinning effects [120, 121].

To further investigate the nature of the effect, the temperature dependence of the trap at a given voltage bias has been measured by varying the position of the DUT mounted in the dipstick with respect to the liquid-helium surface. It is important to note that, while the temperature sensor was glued as close as possible to the sample on the ceramic package for this sweep, the accuracy in the thermalization of the DUT cannot be precisely quantified for the adopted setup above 4.2 K. When sweeping the temperature from 4.2 K to 25 K, the Lorentzian moves to higher frequencies, see figure 5.6c). Above 12 K, the corner frequency exceeds the upper bandwidth limit of the setup.

Assuming the presence of a trap, the trap time-constant $\tau_{C,Lorentzian} = 1/(2\pi f_{C,Lorentzian})$ is extracted and plotted in figure 5.6 d) versus the inverse of the thermal energy. The saturation of the time constant at low temperatures suggests the existence of a temperature-

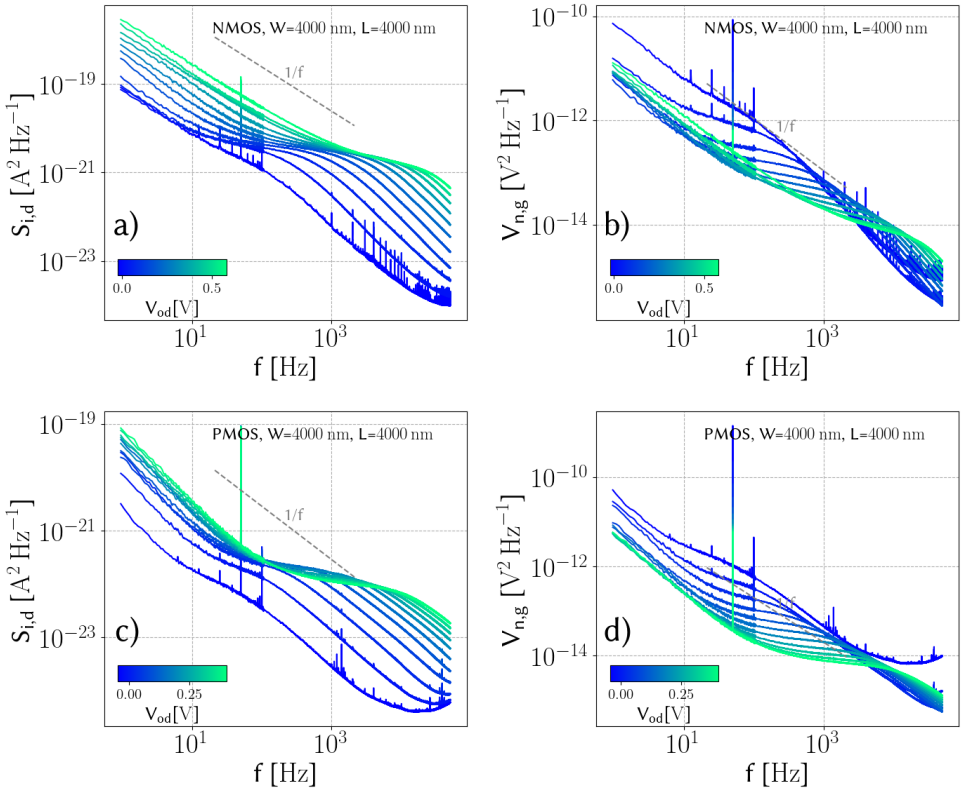

Figure 5.5: Mean noise spectra at 4.2 K. The output (a, c) and input-referred (b, d) of eight N/PMOS devices with WxL=4 µmx4 µm are shown, to demonstrate the bias dependence of the systematic Lorentzian.

independent contribution due to elastic tunneling. As in figure 5.6 b), a decrease in the time-constant can be observed with increasing overdrive voltage. For thermally activated inelastic tunneling, [122] suggests a temperature dependence of the trap time-constant:

$$\tau_{thermal} = \tau_0 e^{\frac{E_b}{k_b T}} \tag{5.1}$$

where $\tau_0 = \frac{1}{\sigma_0 v n}$ is depending on the cross-section pre-factor σ_0, the carrier velocity v and channel carrier concentration n, E_b is the energy barrier and k_b the Boltzmann constant. Assuming a temperature-independent contribution of elastic tunneling, with time constant τ_{el} [123], the combined effect can be described as [120]:

$$\frac{1}{\tau_{C,Lorentzian}} = \frac{1}{\tau_{thermal}} + \frac{1}{\tau_{el}} \tag{5.2}$$

Fitting equation (5.2) to the measured data in figure 5.6 d) we can determine the corresponding energy barrier E_b. The extracted E_b shifts from 5.9 meV to 2.3 meV for a V_{gs} bias from 600 mV to 1100 mV. By assuming a linear field dependence of E_b [120, 121], E_b can be extrapolated to lie on the order of 10 meV when all terminals are at 0 V.

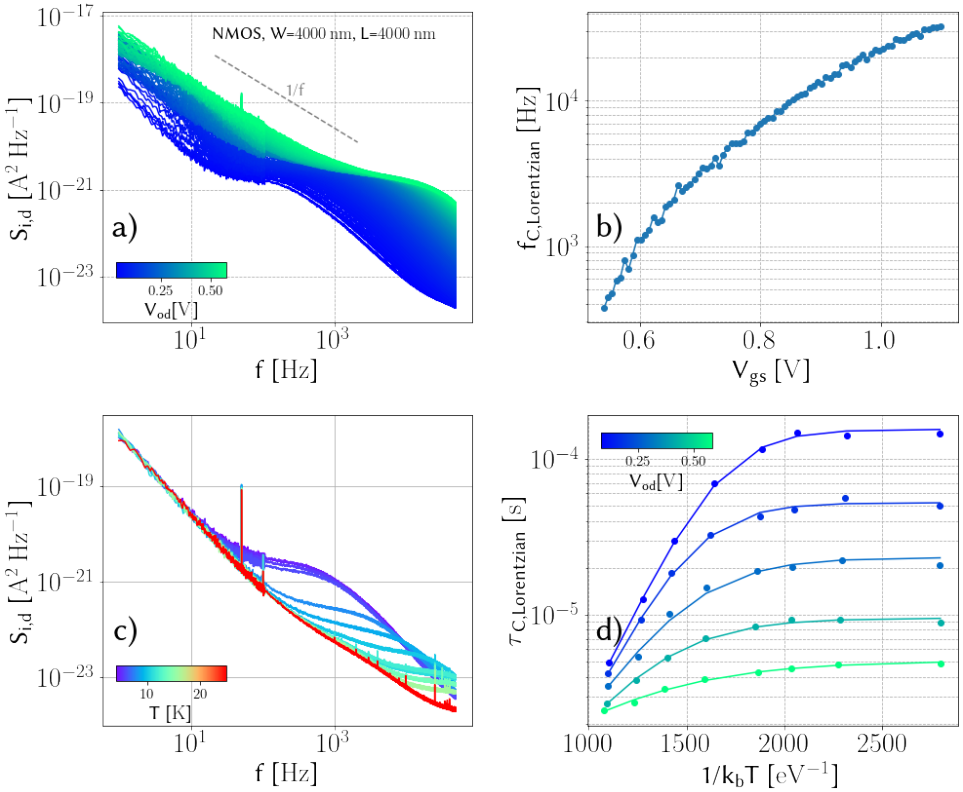

Figure 5.6: Detailed measurements on a W×L=4 µm×4 µm NMOS device: a) fine bias sweep at 4.2 K, b) extracted corner frequency of the Lorentzian at 4.2 K, c) temperature sweep for $V_{od} = V_{gs} - V_{th} = 80$ mV, d) extracted time constant of the Lorentzian.

As the observed effect is systematic over multiple individual devices of the same geometry, randomly distributed oxide traps are an unlikely explanation for the results. A possible physical cause for this phenomenon could be found in the band-tail: depending on doping, the density of states (DoS) might not only show an exponential tail [115], but also an additional distinct Gaussian peak below the conduction band at a doping-dependent activation energy [124]. Normally, states in this tail would be thermally ionized, but at cryogenic temperatures they might act as traps, possibly explaining the substantiation of noise at these temperatures [54]. A doping density in the order of 10^{18} cm^{-3} to 10^{19} cm^{-3}, which is a typical doping for the lightly doped drain regions of a 40 nm bulk CMOS technology, would lead to a barrier energy in the order of 10 meV [124]. This hypothesis is compatible with the observed behavior, including the commonality among different devices, the extracted barrier potential, and the increased magnitude of the effect in saturation compared to triode.

However, to gain a more comprehensive understanding of this phenomenon, further research is required, for instance by testing devices with several different precisely known

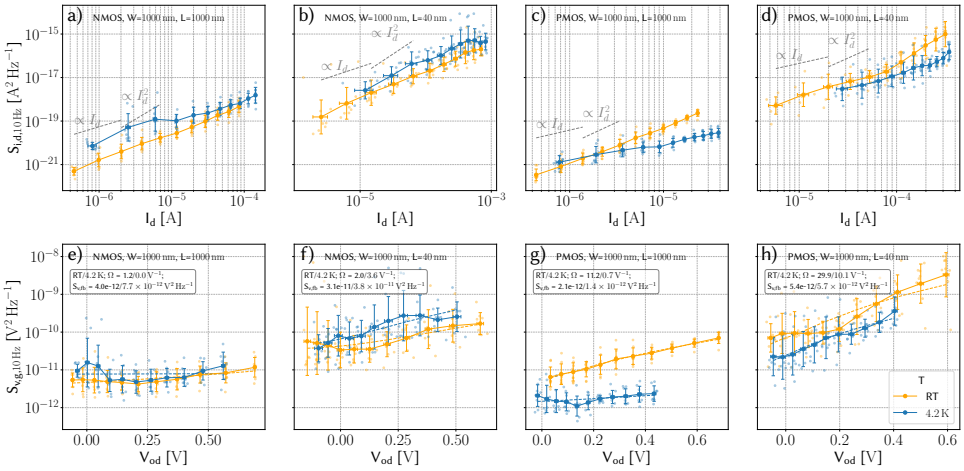

Figure 5.7: Output-referred (a-d) and input-referred (e-h) noise power spectral density at 10 Hz for N/PMOS at V_{gs}=V_{ds} over different bias. Each point represents the measurement from an individual device, while solid lines indicate the mean over the device for each voltage bias.

5

doping levels and by characterizing other bulk technologies under well-controlled thermal conditions, e.g. by employing liquid helium cooling, and over significant device statistics.

5.3.2 Bias dependence

At low frequencies, the noise behavior is largely unaffected by the systematic Lorenzian feature and lends itself to a more traditional analysis. To this end, we sample the data in figure 5.4 at 10 Hz and plot the current noise versus the bias drain current (figure 5.7 a-d)) and the input-referred noise over V_{od} ((figure 5.7 e-h) in figure 5.7. As expected, we observe a significant variation in the noise of the individual devices around the mean and a drop in the input-referred noise voltage with larger device size, which are analyzed in details in the next subsection. The output current noise scales with current, with proportionalities ranging from I_d to I_d^2. Generally, the NMOS results are found to scale roughly $\propto I_d$, while the PMOS results are better described with $\propto I_d^2$. The input-referred data in figure 5.7 can be qualitatively well-described by

$$S_{v,g} = \left(1 + \Omega \frac{I_d}{g_m}\right)^2 S_{fb} \tag{5.3}$$

with S_{fb} the flat-band voltage spectral density and Ω the coefficient of the correlated mobility fluctuation [113]. This qualitative correspondence is in-line with the observations in [112]. The extent of the correlated mobility fluctuations changes with device type and geometry. For the large-area NMOS in e) there are little correlated mobility fluctuations both at RT and cryogenic temperatures, while for the small-area PMOS in h) the effect of correlated mobility is significant at both temperatures. Furthermore, it is interesting to note the significant differences between the RT and cryogenic results for the large-area PMOS in g): the RT results show a strong correlated mobility contribution, which decreases

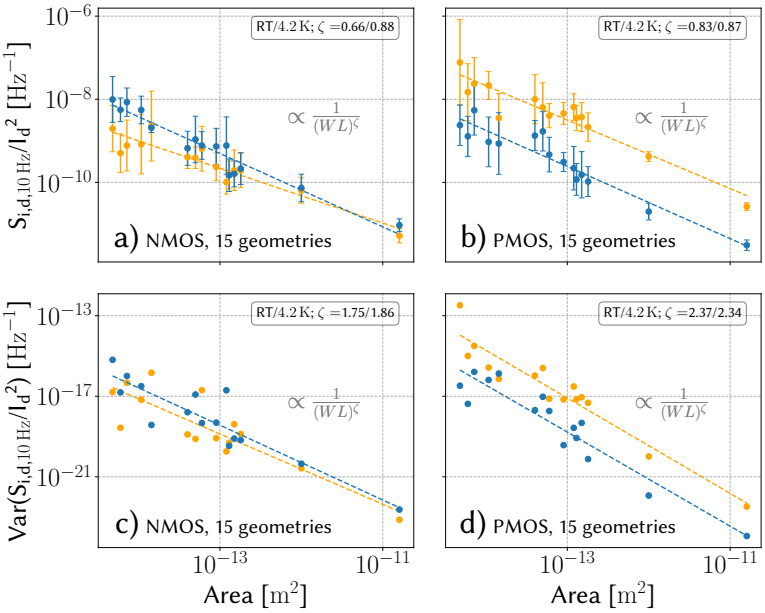

Figure 5.8: a,b) Normalized current noise versus device area for (a) NMOS and (b) PMOS at V_{ds}=V_{gs}=1.1 V for 15 different geometries. c,d) Variance of normalized drain current for NMOS (c) and PMOS (d) at same bias point as in a,b.

significantly in the cryogenic results. The noise magnitude is similar at RT and 4.2 K for plots e), f) and h), with only the large-area PMOS in g) showing a significant drop in input-referred noise at the same overdrive. The differences in scaling with bias between NMOS and PMOS suggest that the noise is generated via different, possibly doping-dependent, noise mechanisms. Explaining the generally small difference between the RT and the cryogenic behavior with an increase in interface defect states would demand unlikely high values for the defect density. This again suggests band-tail states as a possible cause of the substantiated noise as discussed in [54]. A possible further explanation can be found in [125], with the conclusion that Boltzmann statistics can not be applied, and therefore an effective temperature needs to be introduced, that partially explains the reduced scaling with temperature.

In subthreshold (for negative V_{od}), we would expect a constant characteristic in e)-h), assuming the flat-band voltage spectral density is constant. Although we only probed the edge of this regime, we can observe a tendency for increased input-referred noise in the triode measurements, see appendix A. This diverging input-referred voltage in subthreshold can be attributed to mobility fluctuations [113], which could be confirmed by further characterization in the deep-subthreshold regime.

5.3.3 Area dependence

At RT, it was shown in [117] that significant variability is present even in larger devices for the cases of a non-uniform channel. At cryogenic temperature, this area dependence

of LFN noise has not been systematically explored before. Here we report measurements over area from 8 devices per each of 15 different geometries for both PMOS and NMOS, showing both the normalized current noise (figure 5.8 a)/b) and its variance (figure 5.8 c)/d) as a function of device area. We observe that the noise scaling of NMOS devices shows generally little behavior differences at RT and cryogenic temperature in both expected value and the variance. While we do observe a less than proportional scaling with area for the normalized noise at RT in figure 5.8 a), at 4.2 K that scaling is closer to the expected inverse proportionality. However, this difference is not highly significant in light of the device variability. In contrast to this, the PMOS devices in figure 5.8 b) show significant reduction in normalized noise and variance compared to RT. For PMOS, scaling of the noise is very similar over temperature and close to $1/WL$. For both NMOS and PMOS, the variance scaling in saturation is comparatively weak compared to the cubic dependence expected in the linear region, as explained in [117] with effects caused by the halo-implants. The variance scaling exponent with area is larger for PMOS than for NMOS. At RT, the measurements in the linear region, shown in appendix A, are closer to the expected scaling of mean and variance with $1/(WL)$ and $1/(WL)^3$ [117], but the cryogenic results differ significantly due to influence of the systematic Lorentzian. As shown in appendix A, different layout effects were also investigated, but neither metallization on top of the DUTs nor the presence of dummy transistors affected the results shown here.

5.4 Discussion

The observations made here paint a rich picture of differences for transistor LFN between RT and cryogenic temperatures. The differences in behavior are strongly dependent on transistor type, geometry, and bias.

For analog designers adopting transistors in saturation, often in moderate inversion, the changes are however smaller than expected given the radical change in temperature, with the input-referred noise almost unaffected by the temperature for most geometries. The expected scaling with temperature predicted in [113] is not observed. A minor increase in input-referred noise can be observed for the NMOS devices, while the minimum-length PMOS devices show a slight decrease. Only the long PMOS devices significantly benefit from cooling to cryogenic temperatures. However, the significant advantage in lower LFN for longer PMOS devices needs to be traded against the difficulties associated with a larger increase in threshold voltage for PMOS devices, compared to their NMOS counterparts [44], also visible in figure 5.2. As LFN is broadly invariant to temperature, the decrease of the white noise [53] results in an expected large increase of the flicker noise corner. This may require the adoption of techniques for LFN mitigation [126] for a wider range of frequencies than at RT.

The systematic Lorentzian further complicates the picture. As we expect the effect to be doping-dependent, it might differ significantly between technologies. Where it influences the results, it constitutes a significant amount of additional noise, on top of the background 1/f noise. The different spectral shape of the noise may also be critical: the plateau caused by the Lorentzian gives additional noise power at moderate frequency but falls off steeply for higher frequencies.

Given the ample set of effects, it would be challenging to generate a unified compact model, and a binned model that discriminates based on device size could be more suitable.

5.5 Conclusion

The extensive study of LFN in MOS transistors presented in this chapter, comprising a wide range of bias points and device geometries at both RT and 4.2 K, uncovers a diverse set of effects and dependencies. Most prominently, we observed and described a systematic Lorentzian spectrum appearing at cryogenic temperatures. Apart from this, we observed remarkably constant behavior with temperature for the bias-dependence of the NMOS, and reductions in input-referred noise for the PMOS. Furthermore, the area scaling was shown to be intact at cryogenic temperature. Since the observed changes have a significant impact on analog design at cryogenic temperatures, the reported characterization will aid the designers of the next-generation cryogenic systems.

In particular, the data and the analysis presented in this chapter, in combination with the design guidelines presented in chapter 6, can be the basis for the design of a scalable power-efficient cryo-CMOS DC readout for spin qubit, which, as discussed in chapter 1 could be a valid alternative to the RF readout treated in 2, chapter 3, and chapter 4.

6

SET DC READOUT: OPPORTUNITIES AND LIMITS

This chapter presents extensive guidelines for the design of an integrated DC-readout interface for semiconductor spin qubits. Since the focus is on the readout via a single electron transistor (SET), the SET behavior and performance are first described and modeled, showing that the signal-to-noise ratio (SNR) theoretically achievable by a SET-based DC-readout is significantly beyond the state-of-the-art. Practical circuit architectures for implementing a DC-readout, such as the voltage amplifier, the transimpedance amplifier, the charge sampling, and the current pre-amplifier, are then analyzed by deriving their design equations and trade-offs. As a result, the practical performances of those different solutions are evaluated and compared, thus presenting clear selection criteria for the readout architecture and its design equations given the specific parameters of the SET sensor.

This chapter is in preparation to be published as *Gerd Kiene, et al., SET DC readout: opportunities and limits (preliminary title).*

6.1 Introduction

Dc-readout, especially when compared to RF-readout that has been the focus of the ADC designs in chapters 2 to 4, follows a very simple concept. In it, the sensor is directly interfaced, and a signal current or voltage is recorded and compared to a threshold to detect the measured bit. In this chapter, we detail what is necessary to optimize a DC-readout circuit for an SET and what limits exist for such a readout. To set the context, we re-visit the discussion about the benefits and challenges of DC-readout compared to RF-readout here to give a more detailed overview that sets the context for the following derivations.

RF-readout provides the fastest single-shot readouts to date - a performance close to the shot noise limit [127]. Additionally, it allows for potential frequency multiplexing [42], which can reduce the number of lines needed to interface the quantum device. This reduction is especially beneficial when operating the readout circuit at a different temperature stage. Further, the high frequencies used in the RF-readout allow for amplification far from the low-frequency noise of amplifiers in the electronic interface.

But RF-readout also carries heavy disadvantages in the perspective of scaling: it relies on bulky on-chip inductors for matching the sensor impedance [42], requires high-frequency sources and amplifiers for stimulation and amplification, as well as often relying on big off-chip passive components like directional couplers [18]. These factors contribute to a large scaling-unfriendly footprint compared to the qubit dimensions and to a significant power consumption for operating the RF interface. The same RF-readout strategy can also be adopted for gate-based readout [34], which can avoid the use of the SET, still sharing the same (dis)advantages of the SET-based RF-readout.

In contrast, DC-readout has offered a much slower speed, mostly constrained by the large parasitic capacitance of the interconnect between the SET and the readout electronics. By using cryogenic amplification, however, this capacitance can be reduced, resulting in fast amplifiers being demonstrated using SiGe [48] and HEMT [49] transistors operating close to the quantum sample. An inherent downside of the DC-readout is the necessity of an individual electrical connection for each charge sensor to be read. If the sensor and the readout reside on different chips, this requires a bond-pad per sensor, a large offset to the otherwise compact readout. At the same time, no components are necessary for DC-readout that can not be integrated on the chip. This allows for a potentially very dense co-integration of a CMOS-based DC-readout.

Thanks to those appealing features, several implementations of an integrated DC-readout circuit have been proposed: current comparators [45, 46], a trans-impedance amplifier [37], and a charge-sampling integrator [47]. While these constitute significant practical advances towards a scalable DC-readout, a systematic benchmarking of the different possible architectures and their trade-offs is still missing. Also, it is still unclear where the limits of DC-readout lie and how to approach them in practice. To aid the development of scalable amplifiers for DC-readout, we fill this gap here by deriving the limits and comparing practical readout interfaces. For such interfaces, we set the design goal of reaching a readout fidelity above 99.9%, i.e., the quantum-error-correction threshold [128], within a readout period of $T_b = 1\,\mu s$, i.e., a time compatible with spin qubit coherence time and gate duration [129], while minimizing the necessary power. This corresponds to bit-error-rates (BER) of 0.01 and, therefore, requires a signal-to-noise ratio (SNR) of 10 dB,

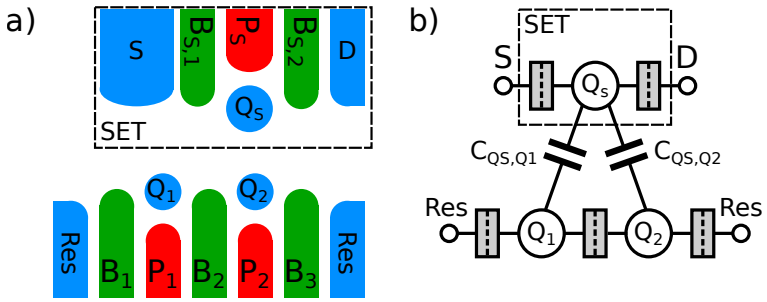

Figure 6.1: a) Layout sketch of a typical double quantum dot with SET readout, b) schematic representation of setup.

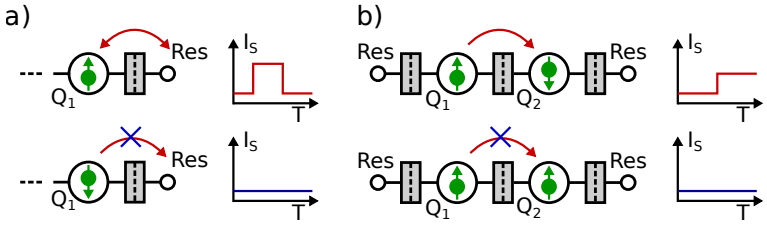

Figure 6.2: a) Elzermann readout, b) Pauli-Spin-blockade readout. Repeated here from the introduction, figure 1.3, for the convenience of the reader.

as shown in the appendix. We try to achieve this SNR within a readout period of $T_b = 1\,\mu s$, while minimizing the necessary power.

The chapter is organized as follows: First, we review the SET as a charge sensor in section 6.2 to provide the understanding and modelling necessary to build the readout interface. In section 6.3, a brief review of the CMOS and BiCMOS technologies available for the implementation of cryogenic amplifiers is given. Then, the fundamental limits for reading the SET in a DC-setup are derived in section 6.4, followed by an extensive architecture exploration and comparison for practical amplifier topologies in section 6.5. Finally, a conclusion is drawn in section 6.6.

6.2 SET MODEL

Figure 6.1a) illustrates a typical planar layout of a double-quantum dot with an SET. The data-dots denoted as $Q_{1/2}$ are electrostatically defined by the potential of metal gates and can be filled with single electrons to form a two-qubit system. To measure the charge arrangement of this double-dot system, we utilize an SET. The SET itself is, in fact, a quantum dot with a high electron occupancy represented as Q_S. In this section, we describe how this example system is operated for the purpose of reading the spin information in the data-dots.

6.2.1 Spin-to-charge conversion

To read the spin state of the electrons in the data-dots $Q_{1/2}$, first, a spin-to-charge conversion needs to be performed. This manuscript focuses on electronics for the charge-state detection step. However, the choice of protocol for spin-to-charge conversion results in significant differences in the generated signals, which we discuss here briefly.

The main two protocols for performing a spin-to-charge conversion are energy-selective readout [27] and Pauli-spin-blockade readout [30, 31], each pictured in figure 6.2. Energy-selective readout (or Elzerman readout) places the energy level of a neighbouring electron reservoir in-between the energy levels of the qubit, such that only one (the higher energy state, generally spin-up) of the spin-states can tunnel out of the quantum dot. This removal of an electron in the spin-up state leaves a vacant energy level, that is subsequently filled with a new electron in the ground state coming from the reservoir, leading to a transient "blip" in quantum dot charge. On the other hand, if the spin is in the ground state the electron will stay in the quantum dot and no exchange with the reservoir takes places, thus the charge signal is expected to be flat. Since the processes of tunneling in and out of the reservoir are probabilistic in nature, the duration of the "blip" can be statistically modified by adjusting the tunneling rate, yet the exact duration and position remain unpredictable. This method has been demonstrated to yield high fidelities [130], but it presents technical difficulties for scaling to large quantum processors, as it requires sub-K operating temperatures and complex interfacing electronics. First, this protocol is not robust against temperature: the energy level of the electron reservoir smears out rapidly with temperature, increasing the probability of tunneling of spin-down and thus increasing false-positives detection, decreasing the fidelity of the spin-to-charge conversion. Second, the uncertainty in the blip's timing implies the need for both a high sampling rate of the electrometer and performing detection algorithms on the resulting data.

Alternatively, in the PSB protocol, two electrons residing in two different quantum dots are used. The spin-to-charge conversion is obtained thanks to the inability of electrons to occupy the exact same energy level, i.e., the same spin-state in a single quantum dot. Thus, when one tries to combine two electrons of different spin-state in the same quantum dot, tunnelling happens and both electrons end up together. If the electrons are in same spin states, however, they do not tunnel and remain in their own dots. The transition is "blockaded," as there is no quantum level at which the electrons can coexist in the same dot. Two flavors of PSB exist, known as Singlet-Triplet readout and Parity readout [131], both leading to the same pattern of charge signal. As during PSB, no electrons need to be exchanged with the reservoir, the process is more robust against temperature [132], which is a crucial advantage. Moreover, because the tunnel rates between dots is generally faster than to/from reservoirs, and the lifetime of the blockaded states is orders of magnitude longer than readout times (a milliseconds to a few microseconds), this results in a much more predictable signal with respect to the Elzerman readout: PSB generates a charge step response at a known time for a known duration. A drawback is that, since no electrons are removed from the system (only reconfigured), the effect on the electrometer is generally weaker, and the detection, therefore, is more difficult.

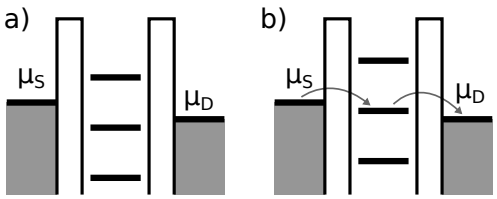

Figure 6.3: SET operating principle (adapted from [28]): a) Coulomb blockade, b) single-electron transport through the dot.

6.2.2 SET DC CHARACTERISTICS

The quantum chip's overall design is such that significant capacitive coupling between the double-dot and the SET is present, depicted in the schematic of Figure 6.1b). The number of trapped electrons on the SET and the conductance through it can be manipulated by five control voltages: the plunger voltage (P_S) controlling the electrochemical potential of the dot and, therefore, the electron occupation number; the two barrier voltages ($B_{s,1/2}$) controlling the tunnel-barrier resistance in and out of the island; and the reservoir voltages (S, D). Due to close proximity, significant capacitive crosstalk exists between the different control voltages and the barriers/dots. Relating the SET to an NFET, we refer to the reservoir contact at lower/higher potential as source/drain, respectively, and to the plunger as a gate.

For a sufficiently small charge island Q_S, the available energy levels on the island are spaced by the charging energy $E_C = e^2/C$, where C is the island's capacitance [32]. When a small bias window is opened between the source and drain terminals of the SET, if none of the quantized energy levels is positioned between the lead potentials, no current can flow, the SET is in "Coulomb blockade" [figure 6.3 a)] and the number of bound electrons remains fixed. If there is a level available, however, single electrons can flow through the quantum structure via the available energy state, resulting in a measurable current [figure 6.3 b)]. A typical measurement to characterize the SET is shown in figure 6.4 a). Clearly visible are the parallelograms with blockaded transport, termed "Coulomb diamonds" [28]. If the bias window becomes large enough, such that more than one quantized energy level is available for transport, the SET cannot be found in a blockaded state anymore. In figure 6.4 b), we show a cut at fixed bias through the measured characteristic, that clearly presents Coulomb blockades and peaks (figure 6.3 a) and b) respectively, as noted in the figure). To optimize the performance of SETs as electrometers, they are usually operated at a low-bias regime where bias-broadening effects are minimized and the steepness of the peaks maximized. The SET is biased around the most responsive region of the Coulomb peak, indicated by the bias current, I_B in Figure 6.4b). In this specific configuration, the movement of a single electron between $Q_{1/2}$ is enough to cause a change in the SET's current levels $I_{s,0/1} = I_B \pm I_s$, that can be associated with the corresponding spin-states via the previously discussed spin-to-charge conversion. For a detailed review of the SET operation and spins in quantum dots, the reader is referred to [28, 32].

6.2.3 SET READOUT PULSE SEQUENCE

As briefly explained above, during PSB readout, the quantum dots are pulsed into a spin-dependent electron configuration for a finite duration of readout time, T_b. The SET current

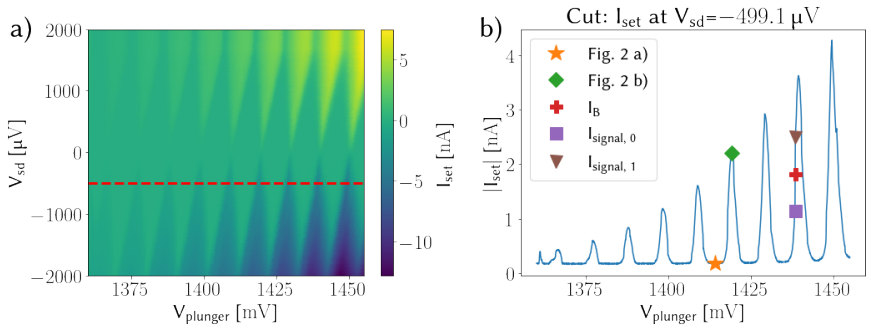

Figure 6.4: a) Measured SET large signal characteristics, b) slice of characteristic for specific V_{sd}, indicating the operating modes in figure 6.3 and an example for the bias points in figure 6.5. Device is cooled to ≈15 mK and measures an effective electron temperature of ≈80 mK.

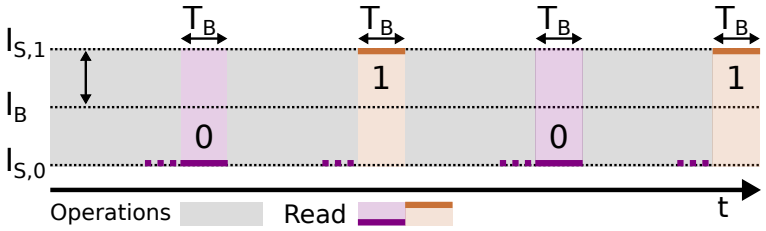

Figure 6.5: SET current evolution in the readout pulse sequence.

is measured during this time and registers the charge configuration change thanks to the capacitive coupling to $C_{QS,Q1/2}$ by assuming the value $I_{s,0/1}$, respectively. We assume the ideal PSB in the following, which means no decay of the blockaded states occurs during the SET acquisition time, such that the readout signal appears as a step.

Before and after the readout times, information processing in the quantum dots is executed. This processing encompasses a combination of fast bias-pulses on the gates defining the qubits, combined with microwave pulses. The pulses used to perform operations are generally not targeting single electrodes, but "virtual gates" combining the effect of multiple electrodes [133]. The SET's plunger is often included in this virtualization due to its large capacitive coupling to the quantum dots. Due to the movement of electrons in the vicinity of the sensor and due to transient bias changes, the sensor current can, therefore, fluctuate prior to the readout phase, indicated by the grey areas in figure 6.5. This is important for the interfacing read-out circuit as either the circuit needs to be able to cope with these variations prior to readout, or the SET needs to be configured in the blockaded regime during qubit operation, where it is relatively insensitive to changes in its electrostatic surroundings. Depending on the choice of operation sequence, the sensor might also put out a well-defined zero level before the pulse, indicated by a dotted zero line before T_B in figure 6.5. While this increases the necessary measurement time, it provides some rejection for the low-frequency noise discussed in the following section.

6.2.4 Noise in SET devices

The SET intrinsically generates shot noise as discussed in [134] and tested experimentally [135]. In this work, we operate significantly below the transition frequency [134] (≈2 GHz for the lowest typical bias of interest for our target application I_B≈100 pA), such that a white shot noise power spectral density (PSD) can be assumed:

$$S_N^2 = 2qFI_{SET} \tag{6.1}$$

where q is the elementary charge, F is the Fano factor [136] and I_{SET} is the SET current. The value of the Fano factor F, also called the shot-noise suppression factor, was measured to be between 0.5 - 1 [135]; thus, for the rest of this work, we adopt a value of 1 as a conservative estimate. The intrinsic, unavoidable shot noise sets the lower limit on the SET noise.

In addition to the broadband noise, SETs are also affected by low-frequency charge noise with 1/f shape [137, 138]. This charge noise is dependent on temperature [139], with higher temperatures corresponding to higher charge noise. The charge noise is converted to a current by the SET *IV* characteristic. Furthermore, single traps, which can be described as two-level systems, can be strongly coupled to the sensor and cause random telegraph noise (RTN). Such strongly coupled traps cause Lorentzian shapes in the noise spectrum to rise above the 1/f background. In the following analysis, we assume the noise to be dominated by 1/f, but we comment on the possible effects of RTN in the design section.

Summarizing, the intrinsic SET single-sided power spectral density can be modeled as a combination of shot and 1/f noise:

$$S_{N,SET}^2(f) = 2qFI_{SET} + \frac{S_{LFN,f=1\,\mathrm{Hz}}}{f^{\alpha}} \tag{6.2}$$

with $S_{LFN,1\,\mathrm{Hz}}$ the low-frequency noise at $f = 1$ Hz and $\alpha \approx 1$ the low-frequency noise slope. $S_{LFN,1\,\mathrm{Hz}}$ is typically measured at the maximum conductance point, and, assuming a first-order model, can be scaled by the device transconductance in other operating points: $S_{LFN,1\,\mathrm{Hz}}(V) = S_{LFN,1\,\mathrm{Hz},max}\frac{g_m(V)}{g_{m,max}}$. Typical values for $S_{LFN,1\,\mathrm{Hz}}$ lie between $10^{-23}\,\mathrm{A^2\,Hz^{-1}}$ to $10^{-24}\,\mathrm{A^2\,Hz^{-1}}$.

6.2.5 SET model for design

For verification of the proposed circuits, we develop a large-signal Verilog-A model that fits the data presented in figure 6.4 over the target operating region of the SET.

For developing the electrical interface, we model the SET setup as shown in figure 6.6, with typical values for the parameters reported in table 6.1. As detailed above, the SET produces a qubit-state-dependent current swing I_s around a bias I_B. The signal amplitude I_s is highly dependent on the coupling of the quantum dot to the sensor, the temperature and the bias conditions. For a practical readout, we assume a typical value of 300 pA, as commonly reached in experimental practice. In parallel with the qubit-dependent current, the setup model includes the setup capacitance C_S and resistance R_S. $C_S = C_p + C_{SET}$ is typically dominated by the parasitic capacitance C_p of the interconnect between the SET and the electrical readout. The intrinsic C_{SET} is extremely small, and corresponds to a diffusion capacitance in a MOSFET, setting a lower bound of C_S in the order of fF, which is only achievable if co-integrating the qubits and the readout. If the SET and the amplifier

Table 6.1: Model parameters

Parameter	Typical	Range
I_s	300 pA	100 pA to 500 pA
I_B	1 nA	0.3 nA to 5 nA
R_S	133 kΩ	30 kΩ to 1000 kΩ
C_S	2 pF	10 fF to 100 pF
F	1	0.5 to 1
$I_{LFN,f=1\,\mathrm{Hz}}$	-	$10^{-23}\,\mathrm{A^2\,Hz^{-1}}$ to $10^{-24}\,\mathrm{A^2\,Hz^{-1}}$

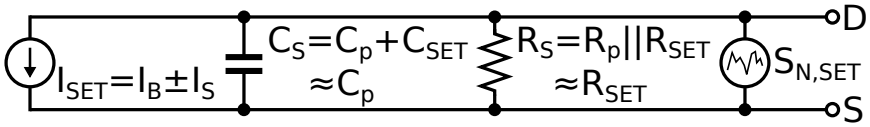

Figure 6.6: Simplified SET electrical model.

are mounted on different temperature plates of a dilution refrigerator or even at room temperature as in [140], the long (≈1 m) coaxial interconnection can result in large capacitance (≈100 pF). The quoted typical C_S=2 pF is estimated based on a two-chip configuration with ESD protection on both sides and direct chip-to-chip bonding. In some practical amplifier structures, the amplifier input capacitance will also add significant additional capacitance on the input node. The SET output resistance R_{SET} typically dominates the setup resistance R_S and can be extracted from the SET DC characterization.

6.3 Cryogenic electronics

Figure 6.7 shows the measured device characteristics of a typical 40 nm CMOS process that forms the basis of the numerical estimations in this paper. A higher threshold voltage, a steeper sub-threshold slope, and a significantly increased g_m/I_d are observed, like reported in [44]. Specifically, the increased g_m/I_d in weak inversion offers significant benefits for power-efficient low-noise design, however, limited by changes in the subthreshold slope leading to large variability in this regime [141]. The threshold voltage can, also in bulk technologies, be influenced by a wide range of body-bias voltages [99]. The gate-tunneling-current is, to first order, unchanged [142]. The low-frequency noise (LFN) changes when moving to cryogenic temperatures, as reported in [112], but the general amplitude of the change is often small compared to the device-to-device variability of the LFN. White noise in MOS transistors has been reported to not scale as thermal noise, but rather seems to approach a lower limit set by shot noise [53]. Possible causes for the increased noise might also be found in self-heating in CMOS [82]. As limited white noise modelling and characterization are available, we here adopt an effective thermal temperature of the MOSFET of 150 K for the noise in active devices for the reference 40 nm technology, in line with the improvement observed in comparator measurements done in the same technology [83]. This reduced scaling, however, is so far only shown in devices in saturation. If the transistors are operated in triode, the resistive channel noise is expected to scale with

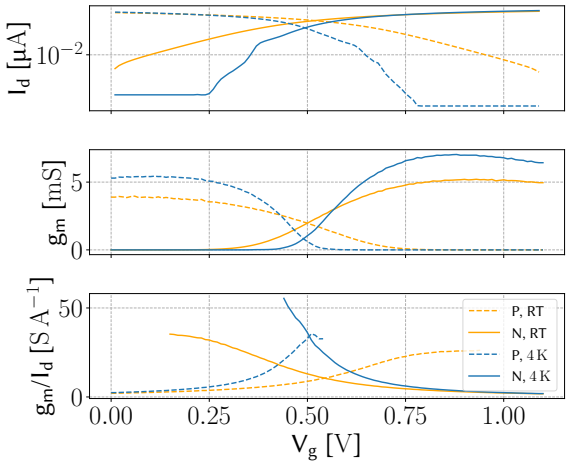

Figure 6.7: Measured DC characteristics at room temperature and 4.2 K of W/L=6×1.2 µm/100 nm NMOS and W/L=6×2.4 µm/100 nm PMOS: a) I_d, b) g_m and c) g_m/I_d. All measured with V_{ds}=550 mV, NMOS/PMOS source voltage at 0/1.1 V.

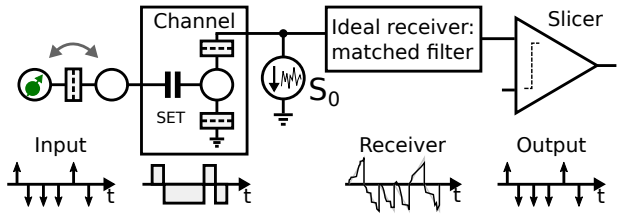

Figure 6.8: Sketch of the idealized readout sequence: upon an impulse input, the channel produces current pulses based on the bit polarity alongside white noise with spectral density S_0. A matched-filter receiver integrates the signal current plus noise and a slicer recovers the bitstream.

temperature. In the absence of current, also no shot noise is produced. As a result, sampling kT/C is expected to scale with temperature, suppressing it sufficiently to make it negligible for many of the architectures considered in the following.

Another promising technology for the implementation of cryogenic amplification is the SiGe BiCMOS process, as the SiGe BJTs show very high current gains at cryogenic temperature [143], and the associated g_m/I_d is much larger than achievable in CMOS. Discrete SiGe BJTs have been used in readout designs [48], enabling a very low power consumption, that is unlikely to be achievable with CMOS. While SiGe does offer superior noise performance, it is also significantly more difficult to co-integrate technologically. For this reason, this work focuses on a CMOS design.

6.4 DC-readout at the shot noise limit

In this section, we explore the intrinsic SNR limits when reading an SET used as a charge sensor for the case of a PSB-type readout, see figure 6.8. The information about the qubit spin states $b_i = \pm 1$ is modeled as a train of impulses $W_b(t)$ each separated by the bit period T_b:

$$W_b(t) = \sum_{i=0}^{\infty} b_i \delta(t - iT_b) \tag{6.3}$$

This neglects the spacing of readout pulses due to operations, as shown in figure 6.5, but without loss of generality in describing the readout itself. As discussed in section 6.2.3, the sensor reacts to this information with current pulses for each bit, describable by a communication channel with impulse response $s(t)$:

$$s(t) = I_s(u(t) - u(t - T_b)) \tag{6.4}$$

with I_s the SET signal current swing and $u(t)$ the Heaviside step function. In an idealized case of manufacturing SETs with little oxide defects, the low-frequency noise is negligible, and only white noise needs to be considered. Under this assumption, the ideal filter for extracting the information from this channel is a matched filter [144]. The matched filter impulse response $h(t)$ is obtained from the channel impulse $s(t)$ by time-reversal and shifting by the bit period T_B:

$$h(t) = s(T_B - t) = u(T_b - t) - u(T_B - t - T_B) = u(t) - u(t - T_b) \tag{6.5}$$

This matched filter represents an ideal integrator with reset. Without loss of generality, we can assume that the integration is done on a capacitor C_P. The output signal provided to the bit slicer at the moment of making the ith decision can now be written as:

$$V_{out}(T_i) = \int_{(i-1)T_b}^{iT_b} \frac{\pm I_s}{C_P} dt = \pm \frac{I_s T_b}{C_p} \tag{6.6}$$

Noise analysis can now be performed with the methodology in [86], given the receiver impulse response and the single-sided noise spectral density $S^2_{N,SET}(f) = S_0$. The integrated noise at the time of slicing can then be calculated as:

$$n^2(t = T_i) = \frac{1}{2} S_0 \int_0^{T_b} |h(t)|^2 dt \tag{6.7}$$

$$= \frac{1}{2} \frac{S_0}{C_P^2} \int_0^{T_b} (u(t) - u(t - T_b))\, dt \tag{6.8}$$

$$= \frac{S_0}{2C_P^2} T_b \tag{6.9}$$

Without loss of generality, for a memory-less system, we assume $i = 0$ in the following, such that evaluation happens at time $t = T_b$. The ideal white-noise limited SNR of DC-readout is then:

$$SNR_{limit} = \frac{V_{out}(T_b)^2}{n^2(T_b)} = \frac{\frac{I_s^2 T_b^2}{C_P^2}}{\frac{S_0}{2C_P^2} T_b} = \frac{2 I_s^2 T_b}{S_0} \tag{6.10}$$

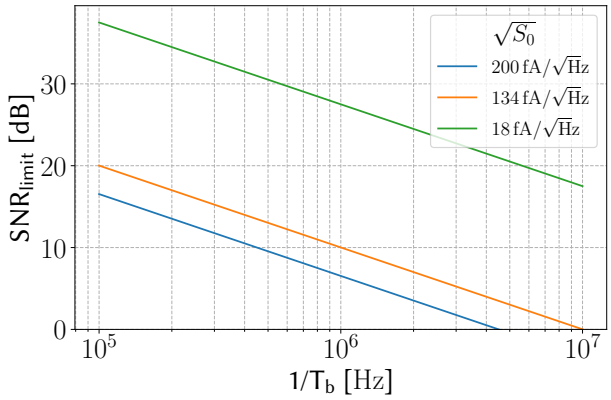

Figure 6.9: Ideal receiver assuming a signal amplitude of $I_{SET} = 300\,\text{pA}$ and a range of noise densities S_0.

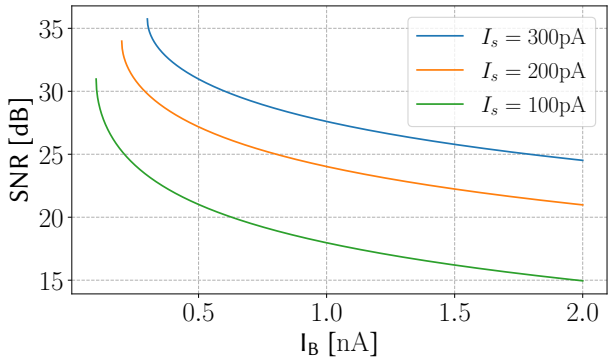

Figure 6.10: SNR limit for $T_b = 1\,\mu\text{s}$, slicing level for equal SNR on both bits.

Solving for the current-noise density, we gain a specification for a given SNR in the case of an ideal receiver:

$$\sqrt{S_0} = \sqrt{\frac{2I_s^2 T_b}{SNR}} \tag{6.11}$$

If, for example, assuming the typical values given in table 6.1, this results in a requirement of better than $134\,\text{fA}/\sqrt{Hz}$ for the input-referred noise of the integrator if targeting an SNR of 10 dB with $T_b = 1\,\mu\text{s}$. For the values in table 6.1, we plot equation (6.10) in figure 6.9, assuming a selection of noise densities.

A limiting case is reached if we consider that the *only* contribution to S_0 is the shot noise generated by the SET itself, assuming an otherwise noiseless readout. The signal swing I_s is around bias point I_B, such that the total current $I_{SET,\pm} = I_B \pm I_s$. For shot noise, we now have $S_0 = 2qI_{tot}$, which would result in a symbol-dependent SNR, as $I_{SET,\pm}$ varies with symbol. For the requirement of equal SNR for both symbols, the slicing level needs to be closer to the signal level with lower noise. This results in the limit on the average

equivalent SNR of:

$$SNR_{shot} = \frac{4I_s^2 T_b}{q\left(\sqrt{I_B - I_s} + \sqrt{I_B + I_s}\right)^2} \tag{6.12}$$

For the typical bias case in table 6.1, this limit is equivalent to the $\sqrt{S_0} = 18\,\mathrm{fA}/\sqrt{\mathrm{Hz}}$ case plotted in figure 6.9, forming the fundamental limit under this signal level. To explore the effect of bias on the SNR limit, we evaluate SNR_{shot} over I_B for various I_s in figure 6.10, keeping $T_b = 1\,\mu\mathrm{s}$. A maximum occurs for a bias equal to the signal level, corresponding to no current when transmitting one of the symbols. This has a physical interpretation: if for one of the symbols the SET is in Coulomb blockade, only minimal second-order currents flow. When minimal current flows, also the intrinsic noise is minimal, allowing a slicing level very close to this bit value maximizing the overall SNR.

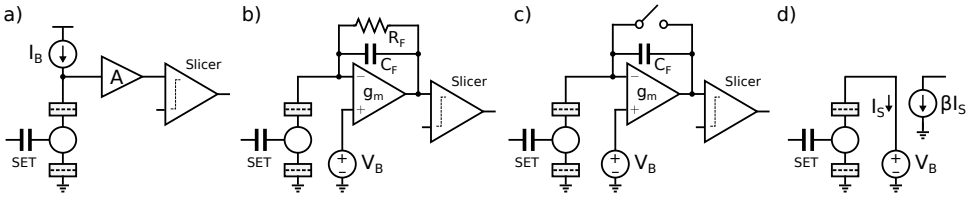

Figure 6.11: Amplifiers considered here: a) voltage amplifier, b) transimpedance amplifier, c) charge-based amplifier, d) current pre-amplifier. The SET symbol is for all following analyses replaced by the small-signal model in figure 6.6.

6.5 Practical front-end requirements

To approach the limits set in the previous section in practice, we analyze and benchmark four readout strategies in the following sections: voltage-mode, current-mode with a TIA, charge-sampling, and current pre-amplification. In voltage-mode readout (figure 6.11 a)), the SET is biased with a current, in current-mode (figure 6.11 b)) with a voltage, while the other quantity is read, respectively. As a special case of current-mode readout, charge sampling (figure 6.11 c)) is discussed separately, as it is functionally close to the matched filter in section 6.4. Finally, the current pre-amplifier (figure 6.11 d)) grants a current gain with low input and high output impedance, which in practice must be followed up by a voltage- or current-mode detection circuit.

In the discussion of these readout methods, emphasis is placed on the challenges arising in their practical implementation. Design examples are carried out with the target of the first stage, limiting the output SNR to 10 dB for a 1 µs integration time. Of course, to get a 10-dB SNR for the overall detection chain, the noise of the first stage must be decreased, but the same design equations can be adopted for the specific choice of follow-up stage without loss of generality. Further, we assume negligible low-frequency noise, requiring either sufficiently large sizing of the input transistors for this to hold, or cascading the readout with a correlated double sampling (CDS) stage [126]. If a CDS stage is used, however, the requirements on the broad-band noise become more stringent, as this noise is now sampled twice. For the remainder of this work, we assume that the amplifier is operated at 4 K.

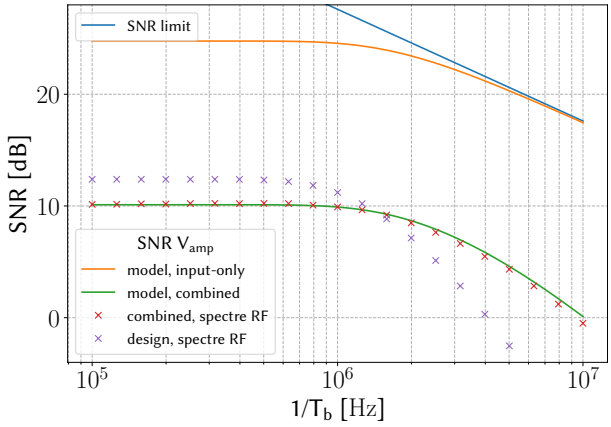

Figure 6.12: SNR of the voltage amplifier in model calculations and spectre RF simulations, compared to the SNR limit.

6.5.1 Voltage amplifier

At the core of a voltage-mode readout, see figure 6.11 a), the SET is biased with a current, the signal current integration is done on the setup capacitance C_S and is sliced after amplification. This core setup deviates from the ideal case discussed above as now the amplifier has a high input impedance. Therefore, all impedances connected to the SET node, indicated in figure 6.6, need to be considered: both the lumped parasitic capacitance C_S and resistance R_S.

For the purpose of this analysis, we also assume an inter-symbol reset or settling during the operations, resulting in no channel memory. The input system constitutes a leaky integrator, with impulse response:

$$h(t) = \frac{1}{C_S} e^{-\frac{t}{\tau_s}} u(t) \tag{6.13}$$

with the time constant $\tau_s = R_S C_S$. For the signal, we then get period T_b:

$$s * h|_{t=T_b} = R_s I_s \left(1 - e^{-\frac{T_b}{\tau_s}}\right) \tag{6.14}$$

The noise power from the SET is now filtered by the leaky integrator, and can be calculated with the same method as applied in the derivation of equation (6.9):

$$n^2(t = T_b) = \frac{\tau S_{N,SET}}{4C_S^2} \left(1 - e^{-\frac{2T_b}{\tau_s}}\right) \tag{6.15}$$

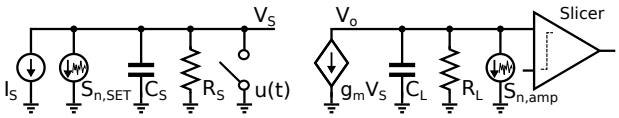

Figure 6.13: Voltage-based amplifier small signal model

This noise needs to be combined with the RMS noise of the voltage amplifier $V^2_{n,a,rms}$:

$$SNR_V = \frac{(s * h(t = T_b))^2}{n^2(t = T_b) + V^2_{n,amp,RMS}} \tag{6.16}$$

$$= \frac{R_S^2 I_s^2 \left(1 - e^{-\frac{T_b}{\tau_s}}\right)^2}{\frac{R_s S_{N,SET}}{4C_S}\left(1 - e^{-\frac{2T_b}{\tau_s}}\right) + V^2_{n,amp,RMS}} \tag{6.17}$$

In figure 6.12, we plot the achievable SNR over $1/T_b$ for the typical parameters from table 6.1 and compare two special cases: $V^2_{n,a,rms} = 0$ and $V^2_{n,a,rms} = 1.5 \times 10^{-10}\,\mathrm{V}^2$. First, without amplifier noise, the SNR saturates for low readout speeds as the resistive input leakage limits the integrated signal, while for high speeds, the ideal limit is approached (model, input-only). In the second case, the amplifier noise amplitude has been chosen such that the SNR target of 10 dB can be reached for a 1 µs integration time (model, combined). The calculation is benchmarked against a simulation using ideal components in Spectre RF, demonstrating a close match (combined, spectre RF). In the following, we replace the abstracted amplifier with a more realistic model.

The amplifiers current noise $S_{i,n,a}$ can be modeled by:

$$S_{n,i,a} = 4kT\gamma NEF g_m \tag{6.18}$$

where k is the Boltzmann constant, γ a technology parameter, NEF the amplifiers noise efficiency factor [145] and g_m the amplifier transconductance. Assuming the small signal noise model of figure 6.13, the output voltage noise density at DC due only to the amplifier is:

$$V^2_{n,o} = 4kT\gamma NEF g_m R_L^2 \tag{6.19}$$

With the equivalent noise bandwith $NBW_V = \frac{1}{4\tau_a} = \frac{1}{4R_L C_L}$, the amplifier noise is the classic kT/C result:

$$V^2_{n,rms,o} = \frac{kT\gamma NEF g_m R_L}{C_L} \tag{6.20}$$

It is important to note that this expression only holds for the settling case in which $T_b \gg \tau_a/2$, with $\tau_a = R_L C_L$ [86]. Referring this to the input by approximating the gain with $A_{DC} = g_m R_L$ and using equation (6.17) under the assumption that the SNR is dominated by the amplifier noise, we gain a specification for the necessary capacitance:

$$C_L = \frac{kT\gamma NEF}{A_{DC} R_S^2 I_S^2 (1 - e^{-\frac{T_b}{\tau_s}})^2} SNR_V \tag{6.21}$$

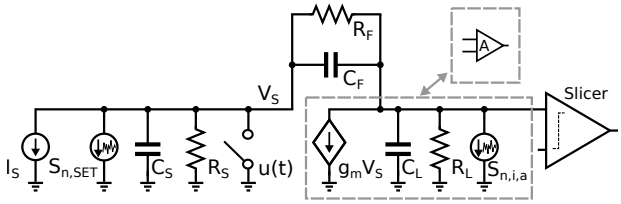

Figure 6.14: Transimpedance amplifier small signal model

Now the g_m is set by the gain requirements at the symbol rate:

$$g_m = \frac{A\sqrt{1+\omega^2 C_L^2 R_L^2}}{R_L} \approx A\omega C_L \tag{6.22}$$

With $\omega = 2\pi/T_b$. For minimizing the necessary capacitance, the gain A needs to be maximized, while for minimizing the g_m in turn the output impedance needs to be maximized.

Assuming a target gain of $A_{DC} = 100$, a noise efficiency factor of $NEF = 2$, a target SNR of 10dB, a bit period of $T_b = 1\,\mu s$ and a MOSFET noise temperature of 150 K we gain: $C_L = 272\,fF$, $g_m = 170\,\mu S$. This design is then simulated alongside the ideal model in figure 6.12 (design, Spectre RF), showing a slight overdesign caused by the simplifying assumptions. The load resistance in the simulation $R_L = 1\,M\Omega$ is chosen to be implementable with MOSFETs.

For a practical design also the input bias current needs to be generated. This bias current generation could be implemented by a bias voltage supplying a polysilicon resistor, that is expected to fully benefit from the noise scaling with temperature, making its contributions negligible.

This discussion of the voltage amplifier further assumed that the signal current is not majorly affected by the voltage signal build-up. In general, this can be violated: in our measurement data, the SET would significantly change its output characteristics if experiencing just a 50 µV swing (the numerical example above would correspond to a settled swing of ±40 µV). An improved sensor design, demonstrated in [146], can increase the sensor's output impedance significantly and increase the voltage range over which a stable current can be supplied to 3 mV.

6.5.2 Transimpedance amplifier

The current-mode front-end in form of a TIA is shown in figure 6.11 b). If $R_F \to \infty$, $C_F \to 0$, the TIA model reduces to the voltage amplifier in figure 6.11 a). Here we concentrate on cases far from this limit, where full treatment of the TIA transfer function is necessary.

In a settled TIA design, the low-frequency transconductance is set by the feedback resistor R_F. The settling also implies that the signal is not integrated, deviating functionally from the ideal receiver in section 6.4.

To quantitatively evaluate the TIA, we analyze the small-signal model shown in figure 6.14. If attempting a full direct solution, the corresponding equations are unsuitable for analytical estimations. Therefore, we divide the problem into two separate parts: First, we examine the effect of SET noise in the TIA; then we look at the noise of the amplifier itself.

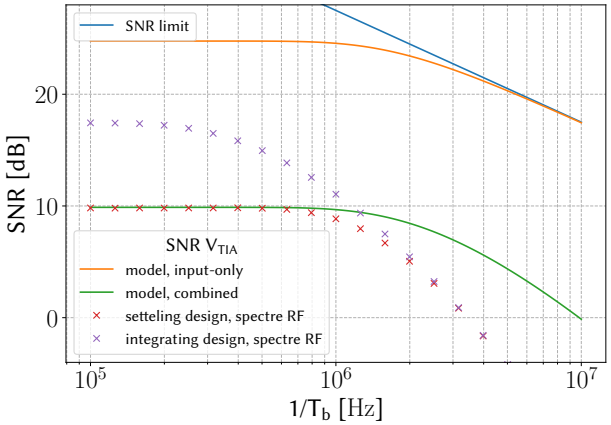

Figure 6.15: SNR of the TIA in model calculations and spectre RF simulations, compared to the SNR limit.

For analyzing the effect of SET noise, we replace the amplifier model in figure 6.14 with a voltage amplifier with gain A. The impulse response of the TIA now becomes:

$$h(t) = \frac{A}{(1+A)C_F + C_S} e^{-\frac{R_F+(1+A)R_S}{((1+A)C_F+C_P)R_F R_S}t} u(t) \tag{6.23}$$

The signal amplitude at the time of slicing ($t = T_b$) is then:

$$s * h|_{t=T_b} = \frac{I_s R_F R_S A}{R_F + (1+A)R_S} \left(1 - e^{-\frac{R_F+(1+A)R_S}{((1+A)C_F+C_S)R_F R_S}T_b}\right) \tag{6.24}$$

Following the approach previously detailed for the voltage amplifier, the output noise due to the SET noise is:

$$\begin{aligned} n^2(t=T_b) = &\frac{A^2 R_F R_S}{2((1+A)C_F + C_S)(R_F + (1+A)R_S)} \\ &\cdot \left(1 - e^{-\frac{R_F+(1+A)R_S}{((1+A)C_F+C_S)R_F R_S}T_b}\right) S_{N,SET} \end{aligned} \tag{6.25}$$

The SNR is therefore:

$$SNR_{TIA} = \frac{s * h(t=T_b)^2}{n^2(t=T_b) + V^2_{n,rms,TIA}} \tag{6.26}$$

where $V^2_{n,rms,TIA}$ is the RMS noise of the TIA structure. We conservatively assume that the TIA is compensated for a flat response, with $C_F = \frac{R_S}{R_F} C_S$. The value of the feedback resistor R_F is mostly defined by the swing requirements of the following stage, as the resistor noise benefits fully from the temperature scaling. To reach for example an output peak-to-peak swing of 1 mV, we require a resistance of at least 1.7 MΩ, significantly exceeding the minimum resistor necessary due to the resistor current noise (≈49 kΩ@4 K for half the noise). Evaluating equation (6.25) for the case of no amplifier noise in figure 6.15, we

see that the TIA, if settled, drops significantly below the limiting ideal performance, but approaches it in the high-frequency limit in which the input current is integrated on the feedback capacitance (model, input-only). The influence of the amplifier gain A on SNR performance is small, for evaluation it was assumed to be 100. Any practical amplifier, however, shows noise of its own, we show the combined performance with an amplifier noise of $V^2_{n,rms,TIA} = 1.2\times10^{-7}\,\mathrm{V}^2$, which would limit the SNR to the 10dB target at $T_b = 1\,\mu\mathrm{s}$ (model, combined). In the following we design an amplifier for the settled TIA case.

For tractable design equations, we again assume $C_F = \frac{R_S}{R_F}C_S$. For this case we have for the noise bandwidth of the TIA:

$$NBW_{TIA} \approx \frac{g_m}{4(1+\alpha)C_L} \tag{6.27}$$

where $\alpha = \frac{R_F}{R_S}$ and we assumed $\frac{1+\alpha}{g_m} \gg R_F + R_S$. The corresponding DC voltage noise density is, with the same amplifier current noise model as in equation (6.18) for $S_{n,i,a}$:

$$V_n = \left(\frac{1+\alpha}{g_m}\right)^2 4kT\gamma NEF g_m \tag{6.28}$$

Combining this with equation (6.27) results in:

$$V^2_{n,rms,TIA} = \frac{(1+\alpha)kT\gamma NEF}{C_L} \tag{6.29}$$

Now assuming a settled output signal $V_{sig} = R_F I_s$, the output SNR is:

$$SNR_o = \frac{V^2_{sig}}{V^2_{n,rms,TIA}} = \frac{R_F^2 I_s^2}{V^2_{n,rms,TIA}} \tag{6.30}$$

The necessary output capacitance is then derived as

$$C_L = \frac{(1+\alpha)kTNEF}{R_F^2 I_s^2} SNR_o \tag{6.31}$$

To calculate the necessary bandwidth, we approximate the input transfer characteristic with its dominant pole:

$$\frac{V_o}{I_{in}} \approx \frac{R_F}{1 + \frac{(C_L+C_S)R_S + C_L R_F}{g_m R_S} s} \tag{6.32}$$

The finite bandwidth causes a settling error E, setting a requirement on g_m:

$$g_m \approx \frac{(C_L + C_S)R_S + C_L R_F}{T_b R_S} \ln\frac{1}{E} \tag{6.33}$$

This estimation slightly underestimates the necessary conductance, as it neglects the effects of the input pole.

For a $R_F = 1.7\,\mathrm{M\Omega}$ and assuming a MOSFET noise temperature of $T = 150\,\mathrm{K}$, a $NEF = 4$ for the amplifier and a settling error of $E = 2\%$ to be acceptable, we obtain $C_L = 4.4\,\mathrm{pF}$ and $g_m = 244\,\mu\mathrm{S}$. The simulated behavior with the typical parameters in table 6.1 is shown in

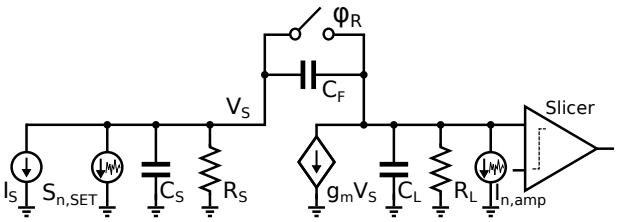

Figure 6.16: Charge-based front-end.

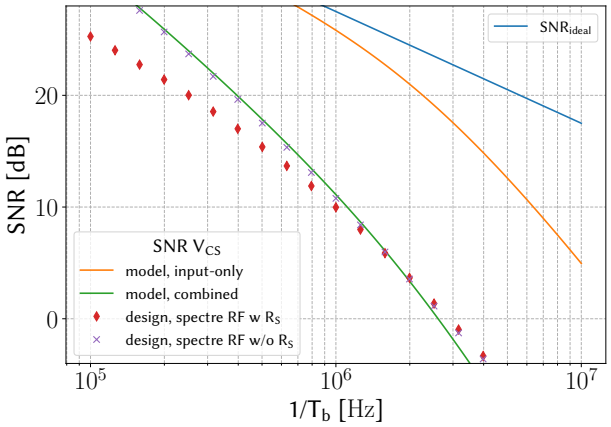

Figure 6.17: SNR of the charge-based amplifier in model calculations and spectre RF simulations, compared to the SNR limit.

figure 6.15, demonstrating that a settling TIA falls just short of the design target, due to the influence of the neglected second pole. However, if changing to a non-settling design by increasing the feedback resistance to $R_F = 10\,\mathrm{M\Omega}$ and keeping the same g_m and C_L parameters, the integrating design exceeds the target specifications as seen in figure 6.15.

There is an important caveat of the non-settling TIA design: if resetting the input capacitance $C_S = 2\,\mathrm{pF}$ via $R_F = 10\,\mathrm{M\Omega}$, the reset time constant is 20 µs, leading to significant inter-symbol interference. This issue could be addressed either via an adaptive time constant for resetting or by replacing the resistor by a switch. The latter approach leads to the charge-based amplification outlined in the following section.

6.5.3 Charge-based amplifier

To come close to the ideal matched-filter-based readout strategy, we can use a charge-integration structure, as shown in figure 6.11 c). For the purpose of making this analysis analytically tractable, we will for now assume both R_L and R_S to be very large in the circuit in figure 6.16. Under this assumption, the impulse response $h(t)$ of the circuit in figure 6.16

is:

$$h(t) = \left(\frac{1}{C_F} - \frac{(C_F + C_L)(C_F + C_S)}{C_L C_S C_F + C_F^2 (C_L + C_S)} \cdot e^{-\frac{C_F g_m}{C_L C_S + C_F (C_L + C_S)} t} \right) u(t) \tag{6.34}$$

By convolution with the input step, we gain the output signal:

$$s * h(t = T_b) = \frac{I_s}{C_F} \left(T_b - \frac{(C_F + C_L)(C_F + C_S)}{C_F g_m} \cdot \left(1 - e^{-\frac{C_F g_m}{C_L C_S + C_F (C_L + C_S)} T_b} \right) \right) \approx \frac{I_s T_b}{C_F} \tag{6.35}$$

The approximation holds approximately for sufficient settling of the integration loop, i.e., for a sufficiently large g_m. The SET input noise integration can be computed from the impulse response as done in the previous section, but in the interest of space, we approximate it here by assuming full integration of the noise.

In this noise analysis, we follow the methodology described in [147] and divide the operation into two phases: the reset phase (R) and the integration phase (INT). In the reset phase we can simplify the circuit by assuming C_F to be shorted, assuming that the switch conductance is sufficiently large. In this case, the noise transfer function during the R-phase is:

$$H_{n,R}(s) = \frac{\frac{1}{g_m}}{1 + s \frac{1}{g_m} (C_S + C_L)} \tag{6.36}$$

As equivalent noise bandwidth is $NBW = \frac{g_m}{4(C_p + C_L)}$, the total noise during the R-phase is:

$$V_{n,R,RMS}^2 = S_n NBW = \frac{1}{4 g_m (C_S + C_L)} (I_{n,amp}^2 + I_{n,set}^2) \tag{6.37}$$

This noise is re-distributed to the feedback capacitor during the integration phase with the ratio $\alpha = \frac{C_S}{C_F}$. During the INT-phase the switch in figure 6.16 is an open, resulting in the noise transfer function:

$$H_{n,INT} = \frac{\frac{1+\alpha}{g_m}}{1 + s \frac{1+\alpha}{g_m} \left(\frac{\alpha}{1+\alpha} C_F + C_L \right)} \tag{6.38}$$

Therefore the settled noise in the INT-phase is:

$$V_{n,INT,RMS}^2 = \frac{(1+\alpha) I_{n,amp}^2}{4 \left(\frac{\alpha}{1+\alpha} C_F + C_L \right) g_m} + \frac{S_0^2}{2 C_F^2} T_b \tag{6.39}$$

where we use equation (6.9) for the SET noise integration. Using equation (6.35) for the signal amplitude, the SNR becomes:

$$SNR_{CS} = (s * h(t = T_b))^2 \Bigg/ \left(\frac{S_0^2}{2 C_F^2} T_b + \frac{\alpha^2}{4 g_m (C_S + C_L)} (I_{n,amp}^2 + S_0^2) + \frac{(1+\alpha) I_{n,amp}^2}{4 g_m \left(\frac{\alpha}{1+\alpha} C_F + C_L \right)} \right) \tag{6.40}$$

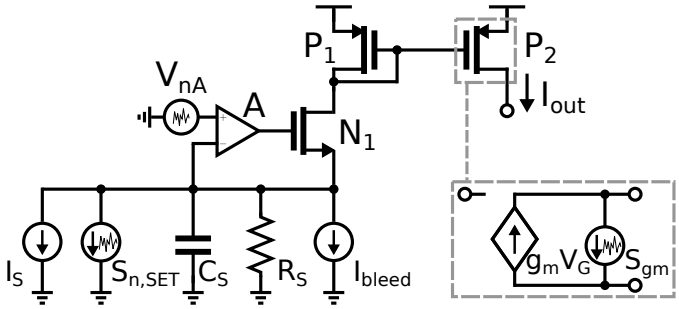

Figure 6.18: Current pre-amplifier CMOS implementation

As a starting point for the design, we assume a MOSFET noise temperature of 150 K, a $NEF = 4$,$C_F = 150\,\text{fF}$, $C_L = 6\,\text{pF}$ and $g_m = 500\,\mu\text{S}$. This corresponds to a slight overdesign, accounting for the neglected R_S. The resulting model and simulated characteristics are plotted in figure 6.17, showing that the design meets the 10 dB SNR target.

6.5.4 Current pre-amplifier

All previously analyzed receiver topologies benefit significantly from an increase in signal current. Here, we evaluate if preceding them with a continuous-time current amplifier, as sketched in figure 6.11 d), could be used to relax their specifications. In such an amplifier, we would prefer a large current gain β, a low input-referred noise current as well as low input impedance. Given the sensor output impedance of R_S, the input impedance needs to be below this by $R_{in} < \frac{R_S}{\zeta}$, we set a target of $\zeta \approx 10$, to avoid excessive signal loss at the input.

To realize a low-power current-amplifier in a CMOS technology in practice, the current mirror, see figure 6.18, is a natural choice. A current mirror with a current gain of $I_{out}/I_{in} = g_{m,P1}/g_{m,P2}$ has an input referred noise of:

$$I_{n,in}^2 = 4kT\gamma g_{m,P1}\left(1 + \frac{g_{m,P1}}{g_{m,P2}}\right) \tag{6.41}$$

For an input-referred noise of $134\,\text{fA}/\sqrt{Hz}$ (equivalent to a 10-dB SNR for $T_b = 1\mu s$ as mentioned in section 6.4) and calculating the acceptable $g_{m,P1}$ based on the above equation and assuming a ratio of $g_{m,P2}/g_{m,P1} = 10$ as well as a MOS noise temperature of 150 K results in an upper bound for the conductance of $g_{m,P1} < 1.8\,\mu\text{S}$. As the corresponding bias current is significantly larger than the SET current, a current-bleeding path (I_{bleed} in figure 6.18) in parallel with the SET needs to be used. Following this, the input resistance of the structure is $R_{in} > 540\,\text{k}\Omega$, violating the input impedance specification. To avoid this loss, a common-gate (CG) stage can be employed. Even if biased in weak inversion with the same current, the input resistance of the CG-stage is still only ≈3× smaller than the one of the current-mirror stage in strong inversion (assuming $\frac{g_{m,N1}}{I_{d,N1}}/\frac{g_{m,P1}}{I_{d,P1}} \approx 3$), thus requiring a boosted CG stage, as shown in figure 6.18 and as used in the design in [45]. The noise

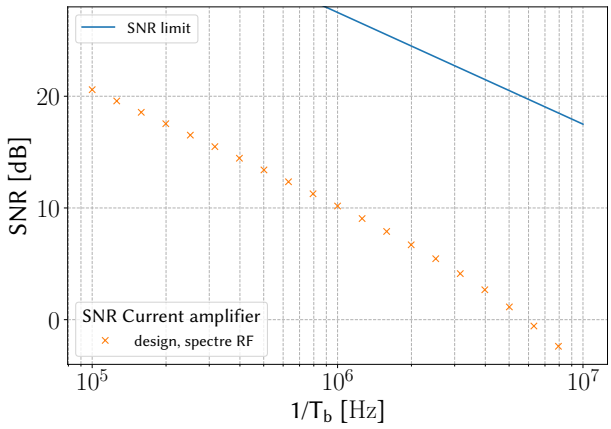

Figure 6.19: SNR of current amplifier in spectre RF simulations.

contributed by the boosted CG stage is approximately:

$$I_{n,CG} \approx \frac{g_{m,N1} A V_{n,A} + I_{n,N1}}{1 + g_{m,N1} R_S A + \frac{R_S}{r_{o,N1}} (1 + g_{m,N1} r_{o,N1})^2} \tag{6.42}$$

If assuming no boosting ($A = 0$), there is only weak suppression of the CG noise due to the comparatively low sensor impedance, since $R_S g_{m,N1} \approx 133\,\text{k}\Omega \cdot 5.5\,\mu\text{S} \approx 0.7$. Also, the resulting $R_{in} \approx 180\,\text{k}\Omega$ input impedance still leads to a $\approx 2\times$ signal loss. If employing boosting, the amplifier noise appears across the input impedance as $I^2_{n,boost} = V^2_{n,A}/R^2_S$ (assuming a large gain A), resulting in a specification for the amplifier input stage $g_{m,boost}$ of:

$$g_{m,boost} = \frac{4kT\gamma NEF}{R_S^2 I^2_{n,boost}} \tag{6.43}$$

Budgeting half the maximum input-referred noise from equation (6.11) to the amplifier and assuming $NEF = 4$ for a differential OTA leads to a requirement of $g_m = 440\,\mu\text{S}$. The gain requirements for the amplifier are relatively moderate if targeting $R_{in} < \frac{R_S}{10}$, leading to a required gain $A \approx 35$ (also considering halving the current in P1 to make room for the amplifier noise). The boosted CG-stage has the added advantage of a well-defined input bias voltage, limited however by the pole formed by the amplifiers g_m and C_S.

We evaluate this design in simulation for $g_{m_{P1}} = g_{m_{N1}} = 1\,\mu\text{S}$, $g_{m_{P2}} = 10 g_{m_{P1}}$, $A = 35$ and budgeting half the available noise to the amplifier at a MOSFET noise temperature of 150 K. For the purpose of this verification, we assume the circuit is followed by an ideal integrator as the matched receiver. The result in figure 6.19 achieves sufficient performance. The drop of SNR at higher frequencies corresponds to the signal current being lost in C_S, which can be counteracted by a larger boosting via A.

An advantage of using the current pre-amplifier is the reduction of kickback from the receiver circuit. The effects of sampling or resetting in the receiver circuit will be reduced by the reverse-isolation. In figure 6.18 this isolation can be further increased by adding cascodes at the output of the mirror.

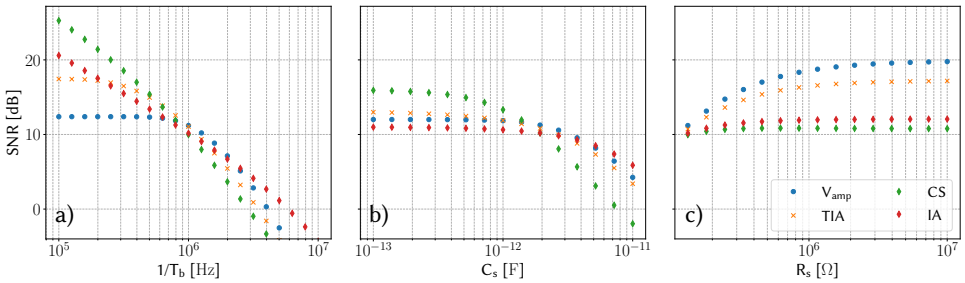

Figure 6.20: SNR comparison of voltage amplifier (V_{amp}), integrating TIA and, sampler (CS) and current amplifier (IA) in response to a) symbol period, b) input capacitance, c) input resistance.

Table 6.2: Design comparison

Metric	VA	TIA	CS	IA
C_L [pF]	0.27	4.4	6	✗
g_m [µS]	170	240	500	440
Power [µW]	19	26	55	48
Fixed input bias	✗	✓	✓	✓
Reset-free	✗	(✓)	✗	✓

6.5.5 Benchmark

The four front-end architectures discussed above each have substantial benefits and challenges, summarized in table 6.2. The voltage amplifier provides the lowest-power solution, but its input voltage definition relies on external resetting and is not stable during integration. The integrating transimpedance amplifier needs a significantly more complex amplifier topology that results in a worse NEF and therefore more power. It however does fix the sensor voltage at the input and can work without a dedicated reset if allowing enough time for settling. The charge integration uses the most power and requires the largest output capacitance, but includes both a practical input-voltage definition and a fast reset. Finally, the current amplifier relies on one of the preceding front-ends before slicing and uses considerable power, but at the same time isolates the sensor from any switching and allows for fast settling of the input node.

In figure 6.20, we compare spectre-RF simulations of the designs presented in this work. For this, we use the models derived in the previous sections and values reported in table 6.2). In a), the input-signal settling observed for the voltage amplifier and TIA limits their SNR at lower frequencies, while the charge sampling architecture and the current amplifier continue to benefit due to integration. While the extent of this integration is in theory limited by the dynamic range, this is likely not problematic given the minute signal levels. For the case of the current amplifier, this improvement only holds if it is followed up by an integrating stage as detector (like the charge sampler). In b), the charge sampling again benefits most from reducing the input capacitance, as this reduces the feedback factor α and leads to both a more complete signal integration and lower noise gain during integration. If sweeping the input resistance up from the default value of 133 kΩ,

the voltage amplifier benefits dramatically, as this increases the signal amplitude that can be integrated on the parasitic input capacitance. This benefit can only be taken advantage of if the sensor can sustain the signal under the larger voltage swings.

We judge from this analysis that the charge-sampling amplifier is best-suited for a potential co-integration. It requires no significant innovations on the sensor side, and benefits significantly from reducing the parasitic capacitance, allowing to scale the power of this architecture in a co-integration environment. The charge-sampler could be implemented similarly to the one in [47], but a targeted co-integration design could save much power. For a two-chip solution, on the other hand, without further sensor innovations, the transimpedance amplifier or current pre-amplifier are well suited, due to the comparatively lower power and good SET bias definition. If however improving the sensor design as suggested in [146], resulting in a significant increase in SET output impedance and therefore R_S, the voltage amplifier promises to perform exceptionally well for a two-chip setup, especially when implementing the g_m with a high-mobility transistor, such as, for example, a SiGe BJT. While the long-term prospects of co-integration of SiGe BJTs are more difficult, such a solution could offer very low power dissipation, significantly beyond analyzed CMOS implementations.

6.6 Conclusion

In this chapter, we have modeled the electronic interface of a SET charge sensor for spin-qubit readout and derived its intrinsic limit in signal-to-noise ratio. In order to practically approach such limits, several circuit architectures are proposed and analyzed analytically and numerically, showing the advantages and disadvantages in terms of their power dissipation, kickback towards the sensor, and capability in accurately biasing the SET. While the numerical results are representative of a typical case, the approach can be adjusted for the specific SET sensors, as required given the significant spread of their parameters over reported experiments. The design equations derived here represent the foundation to build future low-frequency spin-readout interfaces with minimal footprint and power dissipation, as required for large-scale quantum computers.

7

Conclusion

7.1 Research overview

Quantum computers promise significant advantages for addressing problems of acute societal relevance, like drug discovery [148]. However, today's quantum computers, while on the verge of the ability to produce useful results [149], lack significantly in scale to be able to provide the full expected benefits. One bottleneck in scaling up quantum computers is the interconnection between cryogenic quantum processors and the controlling room-temperature electronics. To avoid this bottleneck, cryogenic electronic controllers that operate closer to the qubits have been proposed [17, 19]. Broadly, these controllers need to be able *manipulate* the quantum processor for executing the quantum algorithm and then finally *read* the result of the computation.

The research presented in this thesis can be divided into three general areas, all relevant to the development of cryogenic electronics performing this final *read* operation on quantum computers:

1. Cryogenic high-speed analog-to-digital conversion
2. Low-frequency noise characterization at cryogenic temperatures
3. DC-readout specifications for silicon-spin-based quantum computers

In the following sections, these areas are revisited to provide an overview of the specific contributions made in this thesis to each strand of research. In each area, we come back to the research questions posed in the introduction in section 1.6.

7.1.1 Cryogenic analog-to-digital conversion

The first research question in section 1.6 asked for an efficient ADC at cryogenic temperature for RF-readout of spin qubits. In chapter 2, first the lack of suitable specifications for designing an ADC targeted at frequency-multiplexed readout was addressed. For implementing the high-speed ADC required by the found specifications, the absence of dependable transistor models for simulation was a hurdle in the design process. This was addressed by choosing the SAR topology, which is both power-efficient and comparatively

robust against device-model changes. Additionally, the changes in device characteristics were anticipated where important. For example, a variable common mode was adopted to counteract the transistor threshold voltage increase, which otherwise would cause the comparator to slow down. To reach the target sampling speed in the available 40 nm technology, the SAR conversion loop was further unrolled. The resulting ADC was the first reported high-speed ADC at cryogenic temperatures.

In moderate-resolution SAR ADC design, the comparator noise determines the converter's noise performance. This noise is caused by broadband transistor noise, which is commonly associated with thermal noise. To direct future designs, the noise of a strong-ARM comparator was characterized in chapter 3. The noise was found to improve by a factor of approximately 2 when cooling from room temperature to 4.2 K, which constitutes a relatively minor change compared to the dramatic reduction in temperature by a factor of 70. This adds to the evidence that, in addition to thermal noise, other mechanisms like shot noise may play an important role at cryogenic temperatures. These results can be used to direct future cryogenic comparator designs. Additionally, significant enhancements to the energy efficiency of the ADC have been made in chapter 3, especially by adopting a low-voltage clock receiver. Related to question one in section 1.6, this converter is an effective solution for RF readout of frequency multiplexed spin qubits.

In chapter 4, an efficient dynamic driver for the high-speed ADC was designed to address the second research question from section 1.6. The integrating driver features cascode sampling and also implements the time-interleaving. Its operation at cryogenic temperatures was enabled by the extensive usage of forward body-biasing, which was also evaluated with respect to its impact on analog design for cryogenic temperature in general. Cryogenic-aware forward body bias was demonstrated to be an effective technique for mitigating the increase in cryogenic thresholds. Thanks to this, body-biasing enables a wide range of circuit techniques that were previously restricted to room temperature. Combined with power-effective clock handling, this resulted in the best-reported figure-of-merit that includes the driver at room temperature and the first-reported one at cryogenic temperature.

7.1.2 Low-frequency noise characterization at cryogenic temperatures

To assess and understand the cryogenic behavior of low-frequency noise in CMOS devices(question three in section 1.6), chapter 5 describes a specialized setup developed to facilitate the comprehensive low-frequency noise characterization of a 40 nm CMOS process at cryogenic temperature with the goal of enabling more directed analog design. Specifically, for the first time at cryogenic temperature, the setup allowed for many devices of the same size to be characterized, forming the statistics necessary to inform design decisions. In this characterization, a diverse set of differences between noise behavior at RT and 4.2 K was found, broadly summarized by NMOS devices displaying a small increased noise and PMOS devices a mild decrease. Also, a systematic Lorentzian feature was found to be present in the devices. The Lorentzian was observed to be comparatively stable between devices of the same geometry and flavor. The detection of this feature was enabled by the statistics that were available in this work for the first time. Further detailed characterization of the feature across temperature and bias suggested band-tail states that

act as traps at cryogenic temperature as a possible cause for this effect. An analysis of the noise behavior over area revealed that noise scaling with area continues to be applicable at cryogenic temperatures. The large open dataset associated with this study provides a valuable resource for further modeling and analysis.

7.1.3 DC-readout specifications

Finally, chapter 6 addresses the fourth and last research question posed in section 1.6, which asks for investigating the limits for DC-readout of SETs and asks for design guidelines. A set of specifications for the design of DC-readout for silicon spin qubits were developed. First, the charge sensor under investigation, an SET, was described and modeled. Then, limits in SNR for the operation of the sensor under ideal readout conditions were established, forming an upper bound for any practical effort. These limits were found to be far from the performance currently achieved in DC-readout interfaces, demonstrating that significant gains in either SNR or speed are possible if an approximately ideal readout is implemented. Furthermore, a number of practical integratable implementations for DC-readout were evaluated by deriving design equations for a voltage amplifier, TIA, charge sampling, and current amplification. The TIA seems most suitable for two-chip solutions, while the charge sampling amplifier holds good scaling capabilities for a co-integration.

7.2 Main contributions

Here, a brief summary of the main contributions of this thesis is given:

- The first high-speed ADC at cryogenic temperature was demonstrated [55].
- By pioneering the usage of back-biasing in bulk for many cryogenic analog circuits, the limitations previously posed by the cryogenic threshold increase have been overcome. As a result, analog design in bulk CMOS will change significantly, as many previously employed techniques to work around the increase become obsolete (chapter 4).
- The first dynamic ADC driver at cryogenic temperature was shown. The driver's power consumption did not compromise the ADC power-efficiency, leading to a competitive combined figure of merit (chapter 4).
- The most extensive low-frequency noise characterization performed on cryogenic CMOS devices to date was presented, providing valuable information to designers of precision analog circuits (chapter 5).
- Fundamental limits for DC readout of SETs have been established and ways to approach these limits in practice have been analyzed (chapter 6).

7.3 Future work

This work contributed to the development of a scalable cryogenic readout interface for spin-qubit quantum computers. Suggestions for further work include:

- To make use of the efficient ADC and driver developed here, they need to be integrated with the rest of the readout chain and a digital backend. Developing this full system

however while accounting for the highly variable sensor specifications carries the risk of an ineffective design due to wide ranges, e.g. in operating frequency that need to be covered. To arrive at a stable and precise specification for design, more detailed design work on the direct interface of the RF readout chain with the target sensor is necessary. While this is difficult due to the exploratory nature of the field, narrowly targeted well-performing demonstrators will grant better reference points for progress than developing replacements for general purpose instrumentation.

- While sampling noise (kT/C) is not an important issue in the moderate-resolution ADCs designed here, it is a central concern for many high-resolution architectures with regard to driving the ADC and area consumption. While it is well established, and also discussed in this thesis, that the wide-band current noise of CMOS transistors does not scale with temperature, the sampling noise might, as it corresponds to the noise produced by the resistive channel - opening opportunities for novel approaches to high-resolution ADCs specific to cryogenic temperatures. In particular, a high-resolution ADC with a much smaller input capacitance purely determined by matching requirements could ease both driving and area considerations beyond what is possible at RT.

- If necessary, to improve the linearity of the presented dynamic amplifier further, a direct option could be to significantly reduce the swing of the ADC. This, in turn, would mean more difficult comparator noise specifications, but these are much easier to manipulate than the linearity constraints of the amplifier. This concept potentially integrates well with the suggested minimized input capacitance in the preceding comment.

- As back-biasing effectively allows for resolving the threshold increase at cryogenic temperatures for bulk CMOS technologies, many previously discarded circuit techniques become possible at cryogenic temperature. This includes e.g. the usage of cascodes in amplifiers and pass-gate switches for analog voltages. Future designs could go beyond this by using the threshold voltage control to enable novel circuits, especially by incorporating the additional gate in a functional way. To validate the usefulness of the body bias for such AC applications, the back-gate frequency response should be characterized.

- While an extensive low-frequency noise characterization has been performed as part of this thesis, an associated compact model needs to be derived from the data to fully harness the associated benefits in analog verification.

- The systematic Lorentzian behavior found in the characterized bulk CMOS technology lends itself to a much more detailed analysis of the devices. Possible directions for this include: As the Fermi-Level at 4.2 K is a very sharp probe into the device dynamics, it might be possible to use noise as a detailed probe into the band-structure of the device. Also, it could be worth investigating the possible connection of these noise measurements with the sub-threshold slope changes at cryogenics.

- While MOS technology has the far-future promise of co-integration, different technologies, like for example SiGe, offer various advantages, for example in the imple-

mentation of LNAs or for DC-readout. For near-term projects, using these better-performing technologies for several key building blocks (like, e.g., LNAs) enables offering competitive cryogenic electronics to be actually adopted by users on a daily basis, while CMOS is still used predominantly for demonstration purposes.

- DC-readout faces significant short-term challenges: especially the need for one connection per device for a multi-chip implementation might negate the advantages in terms of size and rejecting low-frequency noise becomes a necessity. However, if it can meet the speed requirements by e.g. using SiGe transistors, the significantly lower complexity compared to RF-readout might make it a highly beneficial solution, even in the short term.
- For the long-term outlook, co-integration of DC-readout might provide a crucial platform advantage of semiconductor spin qubits, especially given the potential of DC readout to approximate the minuscule size of the spin-qubit devices physically. A concrete example of an interesting demonstration in the context of DC-readout could be the implementation of the charge sampler in a CMOS process. While not the optimal solution for a two-chip setup, this could allow for testing possible caveats of the concept, like e.g. kickback issues. Also, a demonstration of a design scaled appropriately for the small parasitic capacitance in a co-integration scenario could be an important demonstrator.

A

Additional low-frequency noise characterization results

This appendix chapter has been published as the supplementary information for the following publication: *Gerd Kiene, Sadik Ilik, Luigi Mastrodomenico, Masoud Babaie, and Fabio Sebastiano, Cryogenic Characterization of Low-Frequency Noise in 40-nm CMOS, IEEE Journal of the Electron Devices Society (2024) [109].*.

A

A.1 Measurement setup

Figure A.1 shows a detailed overview of the measurement setup for the NMOS DUTs. The PMOS measurement is symmetric, but referenced to V_{dd}. Switching between different functions is implemented with latching relays, that are chosen as they do not dissipate any static power at cryogenic temperature. The relays *G6KU-2P-Y* have proven to be reliable at cryogenic temperature.

The setup uses resistive bias-T for biasing, with the resistor (R_{BT}) chosen to be 1 kΩ for the short (high-current, HC) devices, and 10 kΩ for the longer (low-current, LC) devices. The capacitor (C_{BT}) in both cases is a 200 µF low-leakage foil capacitor. This RC combination also defines the setup lower frequency corners at 800 mHz/80 mHz for the HC/LC configuration. After passing the AC-coupling capacitor the signal is amplified by a 1 MV/A TIA and digitized by an acquisition card (M2p.5911-x4). The resistor configuration also has an influence on the setup noise floor: the amplifier used in the TIA has an input referred noise of 1 nV/ $\sqrt{Hz}$, which is converted to a current by the resistance at its input. The upper limit of this resistance is given by the resistor in the bias-T. In practice however the effective resistance is lowered by the devices output resistance (r_o). The effect of the parallel output resistance is especially important for measurements in the linear region, where the device output resistance is low. A simple model for the input-referred noise of the setup can be written as:

$$I^2_{n,in,setup} = \frac{V^2_{n,amp}}{(R_{BT}||r_o)^2} + I^2_{n,R_{BT}} \tag{A.1}$$

with $V_{n,amp}$ the input noise spectral density of the amplifier used in the setup, approximately 1 nV/ $\sqrt{Hz}$, R_{BT} the bias-T resistor, r_o the device output resistance and $I_{n,R_{BT}} = \frac{4k_BT}{R_{BT}}$ the thermal noise of the bias-T resistor. When measuring the setup noise floor without the transistors turned on, we always measure the combined effect of $V_{n,amp}$ and $I_{n,R_{BT}}$. A conservative estimate for the noise that will be produced by the setup alone when the DUT is turned on $I_{n,estimate}$, is to scale the measured current noise floor $I_{n,mean,floor}$ by the additional noise due to r_o:

$$I^2_{n,estimate} = \left(1 + \frac{R_{BT}}{r_o}\right)^2 I^2_{n,meas,floor} \tag{A.2}$$

This overestimates the noise as it assumes all noise originates from the TIA, which is not fully correct as the measured noise floor differs between RT and cryogenic temperature.

Each transistor is implemented, unless otherwise noted, without any metallization directly placed above (excluding also dummy metals for metal-density filling) and with a single dummy transistor on each side. Exception from this rule are the 4 µmx4 µm devices that have no dummies due to layout constraints.

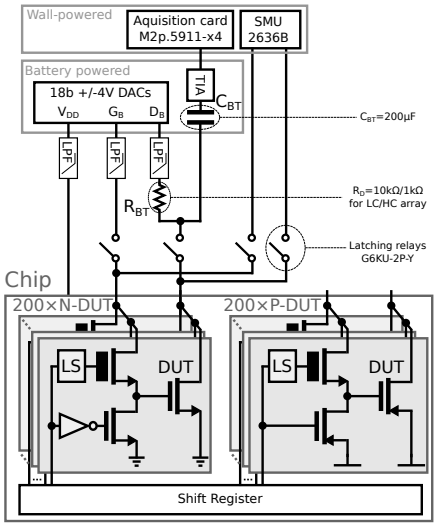

Figure A.1: Simplified noise measurement setup block diagram.

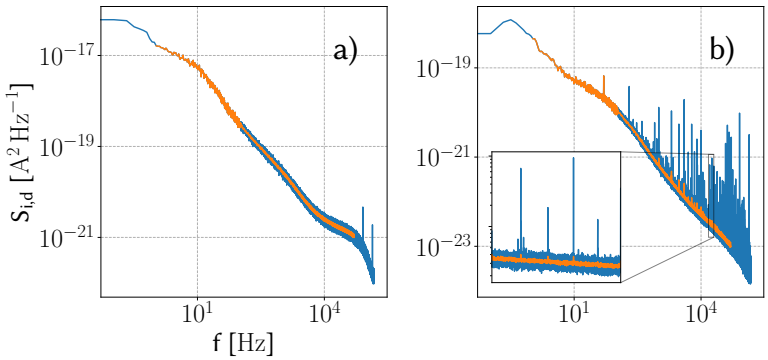

Figure A.2: Data processing: effect on median filter on a) example diode-biased result and b) example triode-biased result.

A.2 Data processing

All plotted data was smoothed with a median filter for better visibility, with the resulting effect on the data showcased in figure A.2. The filter changes the range of the median in the center of the spectrum. Also, the lower frequency bound of the data in any plot is between 5 Hz/1 Hz for HC/LC configurations, to avoid a noticeable effect of the bias-T on the results. At the higher end of the frequency, the data is cut off at 50 kHz to avoid artifacts caused by the setup bandwidth limitations.

A.3 Additional verification

Due to the presence of the unusual systematic Lorentzian feature, we took additional care in setup verification to exclude that any other components being cooled down to cryogenic temperature could be responsible for such a Lorenzian. Here, we discuss these additional verifications in more detail, starting with the one directly related to the chip carrying the transistors and continuing to verify the absence of artifacts from the measurement board components.

On the chip, we identified two possible causes that might influence the measured noise. First, the leakage through the transistors and the ESD diodes on-chip that are nominally "off"; secondly, any effect related to the thick-oxide selection switch. We have excluded the leakage effects by the deselected transistors by recording the spectra with all transistors deselected and sweeping the drain bias, as shown in figure A.3 a) and b). The on-chip multiplexing was tested by reducing the thick-oxide switch supply by 500 mV, without measurable effect on the transistor behavior in figure A.4.

Continuing with the on-board verification, we excluded effects stemming from the resistors and capacitors. In literature, there are some reports of low-frequency noise in thin metal films, e.g. in [150]. For testing if the used thin-film resistors contribute measurable amounts of noise at 4.2 K, we replaced the chip with resistors equal to those in the bias-T (to get realistic current densities) and ran noise acquisitions under various bias conditions, without observing noise contributed by the resistors, see figure A.3 c) and d). Similar reports exist on low-frequency noise in tantalum capacitors, e.g. in [151], motivating the same test for these. To test the possible influence of tantalum capacitors, we tested the samples on a board without populating these capacitors. The resulting measurements show no significant changes in the Lorentzian characteristic due to the removal of the capacitors, see figure A.5.

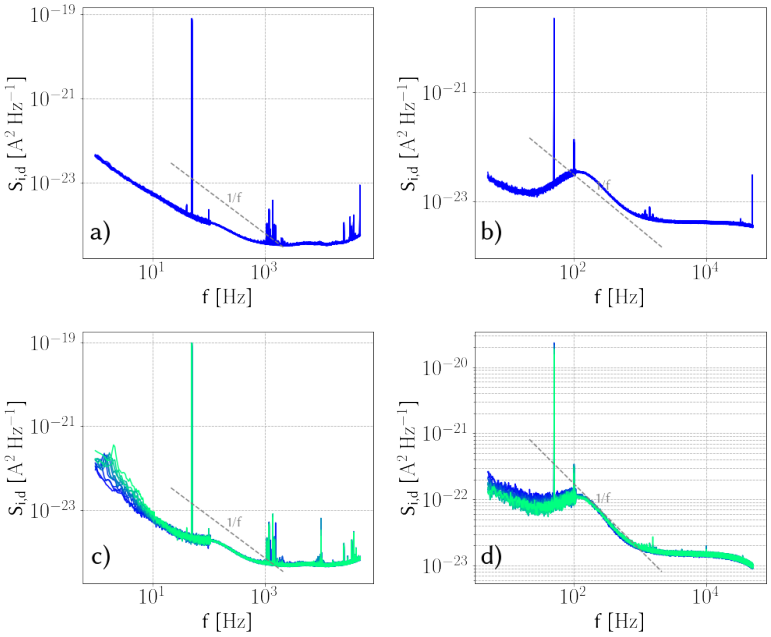

Figure A.3: a) / b) V_{gs}=0 for LC / HC array, respectively. c) and d) DUT replaced by 10 kΩ/1 kΩ resistor in LC / HC array, respectively, with bias swept.

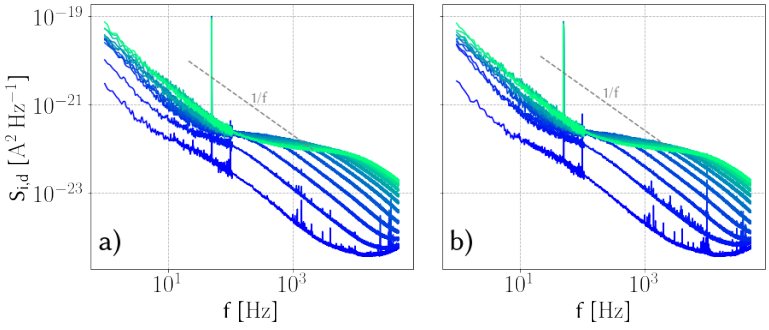

Figure A.4: 4x4u NMOS, a) nominal supply, b) thick-oxide supply decreased by 500 mV

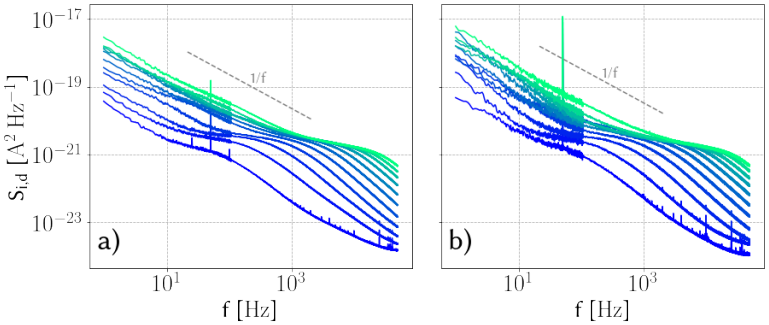

Figure A.5: 4x4u PMOS, a) with tantalum capacitors, b) board without tantalum capacitors (resulting in stronger interference).

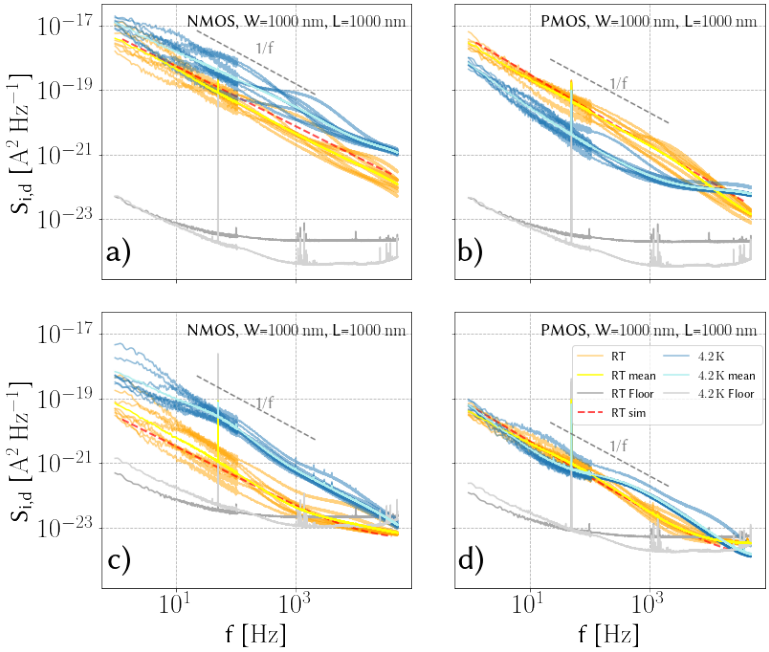

Figure A.6: Input-referred noise spectra for NMOS (a, c) and PMOS (b, d) with W×L =1 µm×1 µm at V_{gs}=1.1 V and V_{ds}=1.1 V (a,b) and V_{ds}=50 mV (c,d), respectively. 8 devices at the same bias points are shown.

A.4 Additional results

In this section, we show additional results from our characterization that can are also part of the associated dataset. The main objective in doing this is to give an overview of the dataset and thereby invite others to use this data for further modeling.

figure A.6 shows example noise spectra for 8 1 µmx1 µm devices for both NMOS and PMOS. These are the same devices as in the main text, but now the input-referred spectra are shown. The input-referring is done using the individual, characterized, g_m of the transistors.

In the main text the diode-spectra with V_{gs}=V_{ds} were discussed. Here we complement this by showing the mean triode spectra with V_{ds}=50 mV in figure A.8. The associated RT results are again close to simulation, and cryogenic results show signs of the described common Lorentzian, especially in the PMOS results. The influence of the additional Lorentzian in the linear region is also appearing at lower frequency, even affecting the results around 10 Hz, in contrast to the observations when biasing in saturation. However, the effects of the common Lorentzian are comparatively less apparent overall in the linear region.

figure A.9 shows the same mean spectra in triode, but now input-referred. At RT the significantly larger bias-dependence of the PMOS input-referred noise appears, suggesting a higher degree of correlated mobility fluctuations. Again also here the bias-dependence of the common bump is preserved.

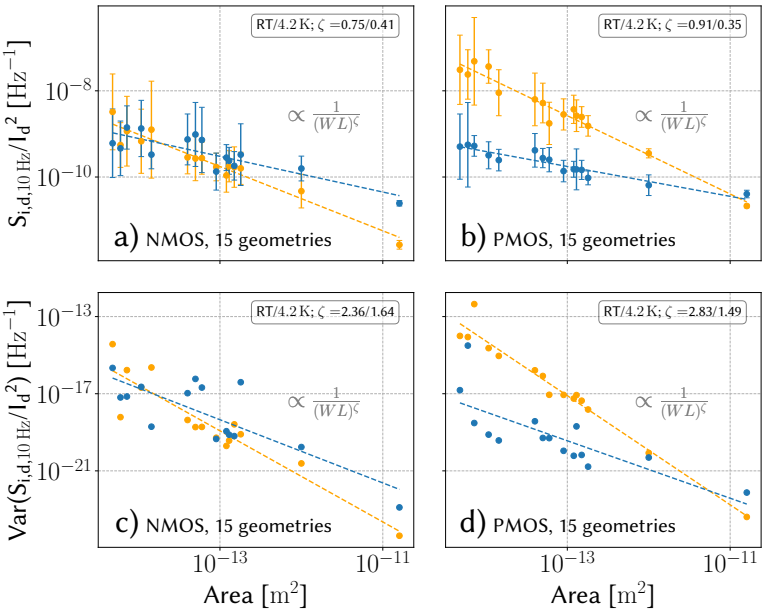

Figure A.7: a,b) Normalized current noise versus device area for (a) NMOS and (b) PMOS at V_{ds}=50 mV and V_{gs}=1.1 V for 15 different geometries. c,d) Variance of normalized drain current for NMOS (c) and PMOS (d) at same bias point as in a,b.

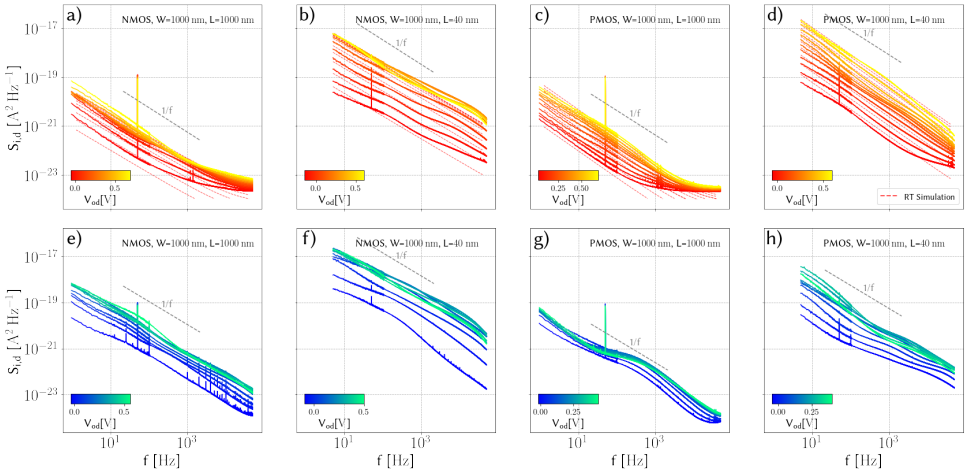

Figure A.8: Mean noise spectra for PMOS and NMOS devices with W×L=1 µm×1 µm and W×L=1 µm×40 nm with V_{ds}=50 mV for logarithmically spaced overdrive-voltage (V_{od}=V_{gs}-V_{th}) at RT (a-d) and 4.2 K (e-h). Each curve represents the mean of 8 individual devices. The RT plots also show the corresponding simulations of the foundry model.

In figure A.10 the input-referred mean NMOS spectra show little bias dependence

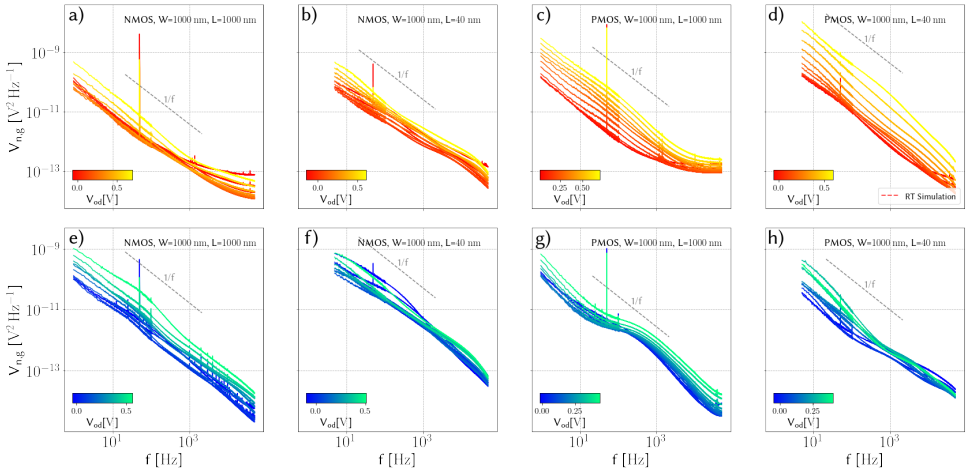

Figure A.9: Mean input-referred noise spectra for PMOS and NMOS devices with W×L=1 µm×1 µm and W×L=1 µm×40 nm with V_{ds}=50 mV for logarithmically spaced overdrive-voltage (V_{od}=V_{gs}-V_{th}) at RT (a-d) and 4.2 K (e-h). Each curve represents the mean of 8 individual devices. The RT plots also show the corresponding simulations of the foundry model.

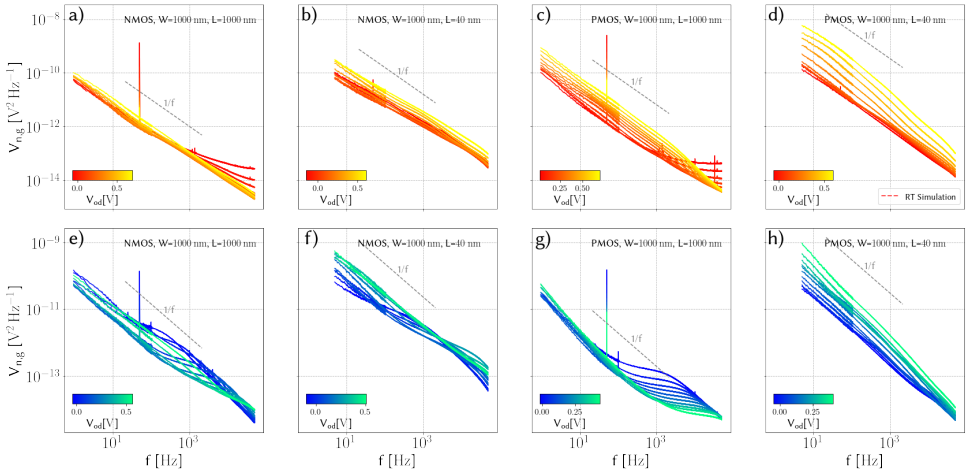

Figure A.10: Mean input-referred noise spectra for PMOS and NMOS devices with W×L=1 µm×1 µm and W×L=1 µm×40 nm with V_{gs}=V_{ds} for logarithmically spaced overdrive-voltage (V_{od}=V_{gs}-V_{th}) at RT (a-d) and 4.2 K (e-h). Each curve represents the mean of 8 individual devices. The RT plots also show the corresponding simulations of the foundry model.

in diode measurements at RT, as discussed in the main text. The bias-dependent effects are again stronger for PMOS devices. The cryogenic measurements show the strong bias-dependence of input-referred Lorentzian, especially in the 1 µmx1 µm PMOS devices.

In figure A.11 the results from figure A.8 and figure A.9 are sliced at 10 Hz. As in

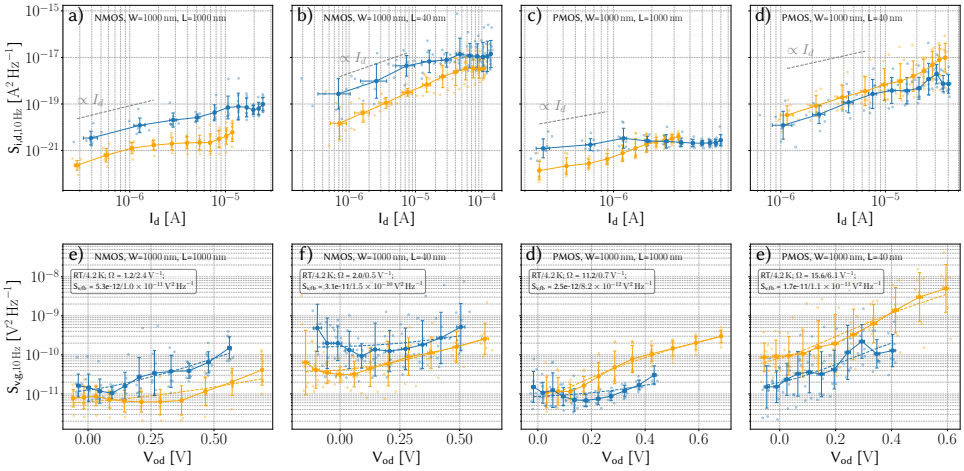

Figure A.11: Output-referred (a-d) and input-referred (e-h) noise power spectral density at 10 Hz for N/PMOS at V_{ds}=50 mV over different V_{gs} bias. Each point represents the measurement from an individual device, while solid lines indicate the mean over the device for each voltage bias.

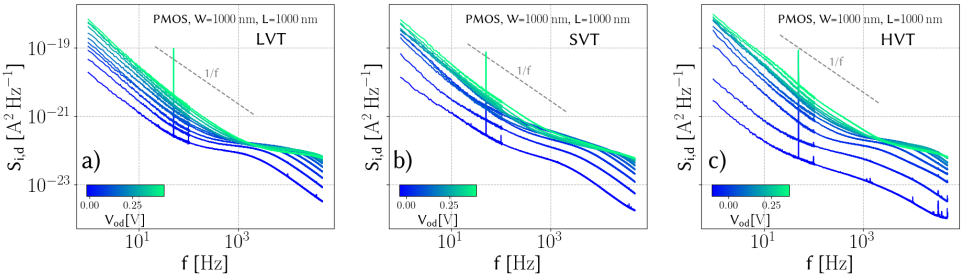

Figure A.12: Several threshold voltage flavors of PMOS 1x1 µm, mean of 8 devices each. All flavors show the common Lorentzian.

the main text the individual dots represent individual devices. As expected, these show significant deviations from the mean also at cryogenic temperatures. However, these triode results are influenced by the common Lorenzian in part, as for triode this lies at comparatively lower frequency. Towards lower bias, we observe an upward trend in the input-referred noise, which would not be expected for number-fluctuations, but rather could be a signature of possible mobility fluctuations [111]. Further investigations in the deep-sub-threshold regime are necessary to substantiate this.

In figure A.12, we compare several threshold-voltage flavors, using the mean of 8 1x1 µm devices for each flavor. This showcases that the bump is present for different threshold-voltage flavors, with all devices in exactly the same setup, but the effect manifests different magnitudes / bias dependencies. This makes a doping-dependent explanation more likely.

In figure A.7, we show the RT results: these are generally closer to standard theory [117], with expected value scaling close to $1/WL$ and variance to $1/(WL)^3$. While being

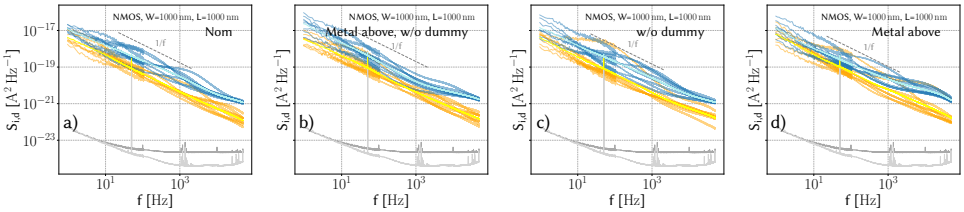

Figure A.13: a) nominal measurement of 8 1x1 µm NMOS devices with metal exclude above and dummies on the sides, b) metal lines placed as a grid on top in layout and dummies removed, c) dummies removed, d) metal lines placed on top in layout.

also closer, the NMOS results continue to be a bit lower in both metrics. The cryogenic results are influenced by the systematic Lorentzian and therefore not easily comparable to room temperature data for the area sweep.

For 1 µmx1 µm devices, we had the following additional devices available: ones with metal fill on-top, ones without dummies and ones with both of these modifications. We expected possible stress-dependent effects to show up in the spectra, but both interventions had no measurable effects on the transistors noise performance, with individual device differences dominating, as seen in figure A.13. Low-frequency noise performance is therefore, within the accuracy of our measurement, stable over those layout variations, resulting in no special layout precautions from the perspective of low-frequency noise.

A

Appendix SET DC readout: opportunities and limits

A.1 BER for PAM

In reading spin qubits, the task is generally to read out spins corresponding to a spin-up or spin-down state, in the following denoted as 0 and 1. This is translated into an SET signal that is a normally distributed with mean μ_0 μ_1 and standard deviation $\sigma_0 = \sigma_1 = \sigma$. The error in this case is given by:

$$P_{error,0} = \frac{1}{2} - \frac{1}{2} erf\left(\frac{x-\mu_0}{\sqrt{2}\sigma}\right) \tag{A.1}$$

The optimal decision threshold for this case is the mid point, such that $x = \frac{\mu_0+\mu_1}{2}$, with this:

$$P_{error,0} = \frac{1}{2} - \frac{1}{2} erf\left(\frac{\mu_1-\mu_0}{2\sqrt{2}\sigma}\right) \tag{A.2}$$

Due to symmetry we have $P_{error,0} = P_{error,1} = P_{error}$. Now the power of this signal is given by:

$$P_a = \frac{1}{2}\left(\mu_0 - \frac{\mu_0+\mu_1}{2}\right)^2 + \frac{1}{2}\left(\mu_1 - \frac{\mu_0+\mu_1}{2}\right)^2 \tag{A.3}$$

$$= \frac{(\mu_0-\mu_1)^2}{4} \tag{A.4}$$

And with the Q-function $Q(x) = \frac{1}{2} - \frac{1}{2} erf\left(\frac{x}{\sqrt{2}}\right)$ and $SNR = \frac{P_a}{\sigma}$:

$$BER = Q\left(\sqrt{SNR}\right) \tag{A.5}$$

This is plotted in figure A.1. Reversed we recover the SNR requirements for a given BER. For example: for a readout fidelity of 99.9%, a common threshold for practical quantum error correction, we require a $BER = 0.001$, so a SNR of better than about 10dB.

A

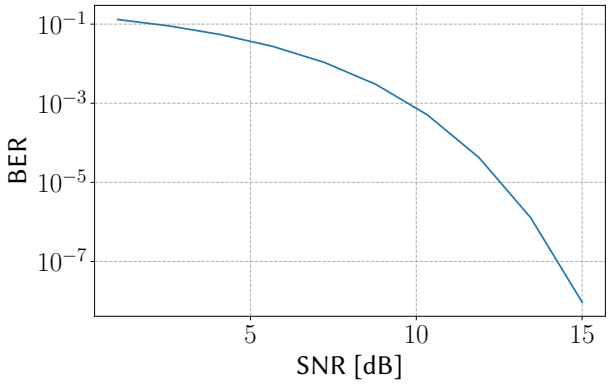

Figure A.1: BER requirements

A.2 Large-signal verilog-A SET model

We list a suggested verilog-A model for the SET in the following, in the interest of space only the core of the model is shown:

```
`define reps 10
module set_model(s, d, p);
electrical s, d, p;
genvar i;
analog begin
  generate i (-`reps, `reps) begin
    I(d, s) <+ i_step
      * (1 / pi * tanh((V(d,s) - V(p,s)
      * ls_p + i * v_charge * 2) / pw)
      + 1/pi * tanh((V(d,s) + V(p,s)
      * la_p - i * v_charge * 2) / pw));
  end
  I(d,s) <+ white_noise(2 * 1.6021e-19
    * I(d, s), "shot")
    + flicker_noise(flicker_noise_1hz,
    1, "flicker");
end
endmodule
```

In the model s, d and p correspond to source drain and plunger of the SET, i_{step} is the conductance per level contributing to the transport, la_p is the leverarm of the plunger, pw the peak width, v_{charge} is the dot charging energy, $flicker_noise_1hz$ is the 1/f noise at 1 Hz measured at the point of maximum dot conductance. Finally $reps$ defines the number of repetitions of the coulomb diamond pattern. A limitation of the model is the 1/f noise representation: this is now constant over the entire range, and therefore only valid at the point of maximum conductance. An extension of the model could scale the $flicker_noise_1hz$ parameter with the g_m.

Bibliography

References

[1] G. E. Moore, "Cramming more components onto integrated circuits," pp. 114–117, 1965.

[2] TSMC, "TSMC Showcases New Technology Developments at 2023 Technology Symposium," *https://pr.tsmc.com/english/news/3021*, 2023.

[3] M. A. Nielsen and I. Chuang, "Quantum computation and quantum information," 2002.

[4] R. P. Feynman, "Simulating physics with computers," in *Feynman and computation*. CRC Press, 2018, pp. 133–153.

[5] I. M. Georgescu, S. Ashhab, and F. Nori, "Quantum simulation," *Reviews of Modern Physics*, vol. 86, no. 1, p. 153, 2014.

[6] I. Kassal, S. P. Jordan, P. J. Love, M. Mohseni, and A. Aspuru-Guzik, "Polynomial-time quantum algorithm for the simulation of chemical dynamics," *Proceedings of the National Academy of Sciences*, vol. 105, no. 48, pp. 18 681–18 686, 2008.

[7] P. W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *Proceedings 35th annual symposium on foundations of computer science*. IEEE, 1994, pp. 124–134.

[8] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, 1996, pp. 212–219.

[9] X. Xue, M. Russ, N. Samkharadze, B. Undseth, A. Sammak, G. Scappucci, and L. M. Vandersypen, "Quantum logic with spin qubits crossing the surface code threshold," *Nature*, vol. 601, no. 7893, pp. 343–347, 2022.

[10] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," *Physical Review A*, vol. 86, no. 3, p. 032324, 2012.

[11] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, "Elucidating reaction mechanisms on quantum computers," *Proceedings of the National Academy of Sciences*, vol. 114, no. 29, pp. 7555–7560, 2017.

[12] J. Gambetta, "Quantum-centric supercomputing: The next wave of computing," *https://research.ibm.com/blog/next-wave-quantum-centric-supercomputing*, 2022.

[13] Bluefors, "XLDsl Dilution Refrigerator System," *https://bluefors.com/products/xldsl-dilution-refrigerator/*, 2023.

[14] Qublox, "Qublox cluster," *https://www.qblox.com/products#cluster*, 2024.

[15] Z. Instruments, "Zurich Instruments Quantum System Hub," *https://www.zhinst.com/europe/en/products/qhub-quantum-system-hub*, 2024.

[16] D. P. Franke, J. S. Clarke, L. M. Vandersypen, and M. Veldhorst, "Rent's rule and extensibility in quantum computing," *Microprocessors and Microsystems*, vol. 67, pp. 1–7, 2019.

[17] B. Patra, R. M. Incandela, J. P. Van Dijk, H. A. Homulle, L. Song, M. Shahmohammadi, R. B. Staszewski, A. Vladimirescu, M. Babaie, F. Sebastiano *et al.*, "Cryo-CMOS circuits and systems for quantum computing applications," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 1, pp. 309–321, 2017.

[18] J. Park, S. Subramanian, L. Lampert, T. Mladenov, I. Klotchkov, D. J. Kurian, E. Juarez-Hernandez, B. P. Esparza, S. R. Kale, A. B. KT *et al.*, "A Fully Integrated Cryo-CMOS SoC for State Manipulation, Readout, and High-Speed Gate Pulsing of Spin Qubits," *IEEE Journal of Solid-State Circuits*, vol. 56, no. 11, pp. 3289–3306, 2021.

[19] L. Vandersypen, H. Bluhm, J. Clarke, A. Dzurak, R. Ishihara, A. Morello, D. Reilly, L. Schreiber, and M. Veldhorst, "Interfacing spin qubits in quantum dots and donors—hot, dense, and coherent," *npj Quantum Information*, vol. 3, no. 1, pp. 1–10, 2017.

[20] L. Petit, H. Eenink, M. Russ, W. Lawrie, N. Hendrickx, S. Philips, J. Clarke, L. Vandersypen, and M. Veldhorst, "Universal quantum logic in hot silicon qubits," *Nature*, vol. 580, no. 7803, pp. 355–359, 2020.

[21] C. H. Yang, R. Leon, J. Hwang, A. Saraiva, T. Tanttu, W. Huang, J. C. Lemyre, K. W. Chan, K. Tan, F. E. Hudson *et al.*, "Operation of a silicon quantum processor unit cell above one kelvin," *Nature*, vol. 580, no. 7803, pp. 350–354, 2020.

[22] J. Y. Huang, R. Y. Su, W. H. Lim, M. Feng, B. van Straaten, B. Severin, W. Gilbert, N. Dumoulin Stuyck, T. Tanttu, S. Serrano *et al.*, "High-fidelity spin qubit operation and algorithmic initialization above 1 K," *Nature*, vol. 627, no. 8005, pp. 772–777, 2024.

[23] A. Zwerver, T. Krähenmann, T. Watson, L. Lampert, H. C. George, R. Pillarisetty, S. Bojarski, P. Amin, S. Amitonov, J. Boter *et al.*, "Qubits made by advanced semiconductor manufacturing," *Nature Electronics*, vol. 5, no. 3, pp. 184–190, 2022.

[24] S. G. Philips, M. T. Mądzik, S. V. Amitonov, S. L. de Snoo, M. Russ, N. Kalhor, C. Volk, W. I. Lawrie, D. Brousse, L. Tryputen *et al.*, "Universal control of a six-qubit quantum processor in silicon," *Nature*, vol. 609, no. 7929, pp. 919–924, 2022.

[25] M. Veldhorst, H. Eenink, C.-H. Yang, and A. S. Dzurak, "Silicon CMOS architecture for a spin-based quantum computer," *Nature communications*, vol. 8, no. 1, p. 1766, 2017.

[26] J. M. Boter, J. P. Dehollain, J. P. Van Dijk, Y. Xu, T. Hensgens, R. Versluis, H. W. Naus, J. S. Clarke, M. Veldhorst, F. Sebastiano *et al.*, "Spiderweb array: a sparse spin-qubit array," *Physical Review Applied*, vol. 18, no. 2, p. 024053, 2022.

[27] J. Elzerman, R. Hanson, L. W. Van Beveren, B. Witkamp, L. Vandersypen, and L. P. Kouwenhoven, "Single-shot read-out of an individual electron spin in a quantum dot," *nature*, vol. 430, no. 6998, pp. 431–435, 2004.

[28] R. Hanson, L. P. Kouwenhoven, J. R. Petta, S. Tarucha, and L. M. Vandersypen, "Spins in few-electron quantum dots," *Reviews of modern physics*, vol. 79, no. 4, p. 1217, 2007.

[29] F. A. Zwanenburg, A. S. Dzurak, A. Morello, M. Y. Simmons, L. C. Hollenberg, G. Klimeck, S. Rogge, S. N. Coppersmith, and M. A. Eriksson, "Silicon quantum electronics," *Reviews of modern physics*, vol. 85, no. 3, p. 961, 2013.

[30] K. Ono, D. Austing, Y. Tokura, and S. Tarucha, "Current rectification by Pauli exclusion in a weakly coupled double quantum dot system," *Science*, vol. 297, no. 5585, pp. 1313–1317, 2002.

[31] A. C. Johnson, J. Petta, C. Marcus, M. Hanson, and A. Gossard, "Singlet-triplet spin blockade and charge sensing in a few-electron double quantum dot," *Physical Review B*, vol. 72, no. 16, p. 165308, 2005.

[32] K. K. Likharev, "Single-electron devices and their applications," *Proceedings of the IEEE*, vol. 87, no. 4, pp. 606–632, 1999.

[33] A. West, B. Hensen, A. Jouan, T. Tanttu, C.-H. Yang, A. Rossi, M. F. Gonzalez-Zalba, F. Hudson, A. Morello, D. J. Reilly *et al.*, "Gate-based single-shot readout of spins in silicon," *Nature nanotechnology*, vol. 14, no. 5, pp. 437–441, 2019.

[34] G. Zheng, N. Samkharadze, M. L. Noordam, N. Kalhor, D. Brousse, A. Sammak, G. Scappucci, and L. M. Vandersypen, "Rapid gate-based spin read-out in silicon using an on-chip resonator," *Nature nanotechnology*, vol. 14, no. 8, pp. 742–746, 2019.

[35] J. P. G. Van Dijk, B. Patra, S. Subramanian, X. Xue, N. Samkharadze, A. Corna, C. Jeon, F. Sheikh, E. Juarez-Hernandez, B. P. Esparza *et al.*, "A Scalable Cryo-CMOS Controller for the Wideband Frequency-Multiplexed Control of Spin Qubits and Transmons," *IEEE Journal of Solid-State Circuits*, vol. 55, no. 11, pp. 2930–2946, 2020.

[36] L. Enthoven, J. Van Staveren, J. Gong, M. Babaie, and F. Sebastiano, "A 3V 15b 157μW Cryo-CMOS DAC for Multiplexed Spin-Qubit Biasing," in *2022 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits)*. IEEE, 2022, pp. 228–229.

[37] L. Le Guevel, G. Billiot, B. Cardoso Paz, M. Tagliaferri, S. De Franceschi, R. Maurand, M. Cassé, M. Zurita, M. Sanquer, M. Vinet *et al.*, "Low-power transimpedance amplifier for cryogenic integration with quantum devices," *Applied Physics Reviews*, vol. 7, no. 4, p. 041407, 2020.

[38] A. Ruffino, Y. Peng, T.-Y. Yang, J. Michniewicz, M. F. Gonzalez-Zalba, and E. Charbon, "A Fully-Integrated 40-nm 5-6.5 GHz Cryo-CMOS System-on-Chip with I/Q Receiver and Frequency Synthesizer for Scalable Multiplexed Readout of Quantum Dots," in *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 64. IEEE, 2021, pp. 210–212.

[39] B. Prabowo, G. Zheng, M. Mehrpoo, B. Patra, P. Harvey-Collard, J. Dijkema, A. Sammak, G. Scappucci, E. Charbon, F. Sebastiano *et al.*, "A 6-to-8GHz 0.17 mw/Qubit Cryo-CMOS Receiver for Multiple Spin Qubit Readout in 40nm CMOS Technology," in *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 64. IEEE, 2021, pp. 212–214.

[40] R. Schoelkopf, P. Wahlgren, A. Kozhevnikov, P. Delsing, and D. Prober, "The radio-frequency single-electron transistor (RF-SET): A fast and ultrasensitive electrometer," *science*, vol. 280, no. 5367, pp. 1238–1242, 1998.

[41] A. Aassime, G. Johansson, G. Wendin, R. Schoelkopf, and P. Delsing, "Radio-frequency single-electron transistor as readout device for qubits: Charge sensitivity and backaction," *Physical Review Letters*, vol. 86, no. 15, p. 3376, 2001.

[42] J. Hornibrook, J. Colless, A. Mahoney, X. Croot, S. Blanvillain, H. Lu, A. Gossard, and D. Reilly, "Frequency multiplexing for readout of spin qubits," *Applied Physics Letters*, vol. 104, no. 10, p. 103108, 2014.

[43] A. Beckers, F. Jazaeri, and C. Enz, "Cryogenic MOS transistor model," *IEEE Transactions on Electron Devices*, vol. 65, no. 9, pp. 3617–3625, 2018.

[44] 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," *IEEE Journal of the Electron Devices Society*, vol. 6, pp. 996–1006, 2018.

[45] K. Das, T. Lehmann, and A. S. Dzurak, "Sub-nanoampere one-shot single electron transistor readout electrometry below 10 kelvin," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 61, no. 10, pp. 2816–2824, 2014.

[46] H. Fuketa, I. Akita, T. Ishikawa, H. Koike, and T. Mori, "A Cryogenic CMOS Current Comparator for Spin Qubit Readout Achieving Fast Readout Time and High Current Resolution," in *2022 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits)*. IEEE, 2022, pp. 234–235.

[47] M. Castriotta, E. Prati, and G. Ferrari, "Fully-integrated Cryo-CMOS spin-to-digital readout for semiconductor qubits," *IEEE Solid-State Circuits Letters*, 2023.

[48] M. J. Curry, M. Rudolph, T. England, A. Mounce, R. Jock, C. Bureau-Oxton, P. Harvey-Collard, P. Sharma, J. Anderson, D. Campbell *et al.*, "Single-shot readout performance of two heterojunction-bipolar-transistor amplification circuits at millikelvin temperatures," *Scientific Reports*, vol. 9, no. 1, p. 16976, 2019.

[49] I. Vink, T. Nooitgedagt, R. Schouten, L. Vandersypen, and W. Wegscheider, "Cryogenic amplifier for fast real-time detection of single-electron tunneling," *Applied Physics Letters*, vol. 91, no. 12, 2007.

[50] J. C. Bardin and S. Weinreb, "A 0.1–5 GHz Cryogenic SiGe MMIC LNA," *IEEE Microwave and Wireless Components Letters*, vol. 19, no. 6, pp. 407–409, 2009.

[51] A. Beckers, F. Jazaeri, A. Grill, S. Narasimhamoorthy, B. Parvais, and C. Enz, "Physical model of low-temperature to cryogenic threshold voltage in MOSFETs," *IEEE Journal of the Electron Devices Society*, vol. 8, pp. 780–788, 2020.

[52] W. Chakraborty, K. A. Aabrar, J. Gomez, R. Saligram, A. Raychowdhury, P. Fay, and S. Datta, "Characterization and modeling of 22 nm FDSOI Cryogenic RF CMOS," *IEEE Journal on Exploratory Solid-State Computational Devices and Circuits*, vol. 7, no. 2, pp. 184–192, 2021.

[53] K. Ohmori and S. Amakawa, "Variable-Temperature Broadband Noise Characterization of MOSFETs for Cryogenic Electronics: From Room Temperature down to 3 K," in *2023 7th IEEE Electron Devices Technology & Manufacturing Conference (EDTM)*. IEEE, 2023, pp. 1–3.

[54] R. Asanovski, A. Grill, J. Franco, P. Palestri, A. Beckers, B. Kaczer, and L. Selmi, "Understanding the Excess 1/f Noise in MOSFETs at Cryogenic Temperatures," *IEEE Transactions on Electron Devices*, 2023.

[55] G. Kiene, R. W. Overwater, A. Catania, A. G. Sreenivasulu, P. Bruschi, E. Charbon, M. Babaie, and F. Sebastiano, "A 1-GS/s 6–8-b Cryo-CMOS SAR ADC for Quantum Computing," *IEEE Journal of Solid-State Circuits*, 2023.

[56] B. Murmann, "ADC Performance Survey 1997-2022," *http://web.stanford.edu/ murmann/adcsurvey.html*, 2022.

[57] O. A. Mukhanov, D. Gupta, A. M. Kadin, and V. K. Semenov, "Superconductor analog-to-digital converters," *Proceedings of the IEEE*, vol. 92, no. 10, pp. 1564–1584, 2004.

[58] J. Wong, R. Dunnegan, D. Gupta, D. Kirichenko, V. Dotsenko, R. Webber, R. Miller, O. Mukhanov, and R. Hitt, "High performance, all digital RF receiver tested at 7.5 GigaHertz," in *MILCOM 2007-IEEE Military Communications Conference.* IEEE, 2007, pp. 1–5.

[59] H. Homulle, S. Visser, and E. Charbon, "A cryogenic 1 GSa/s, soft-core FPGA ADC for quantum computing applications," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 63, no. 11, pp. 1854–1865, 2016.

[60] L. Kull, T. Toifl, M. Schmatz, P. A. Francese, C. Menolfi, M. Braendli, M. Kossel, T. Morf, T. M. Andersen, and Y. Leblebici, "A 3.1 mW 8b 1.2 GS/s single-channel asynchronous SAR ADC with alternate comparators for enhanced speed in 32 nm digital SOI CMOS," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 12, pp. 3049–3058, 2013.

[61] Y. Creten, P. Merken, W. Sansen, R. Mertens, and C. Van Hoof, "A cryogenic ADC operating down to 4.2 K," in *2007 IEEE International Solid-State Circuits Conference. Digest of Technical Papers*. IEEE, 2007, pp. 468–616.

[62] Y. Creten, P. Merken, W. Sansen, R. P. Mertens, and C. Van Hoof, "An 8-bit flash analog-to-digital converter in standard CMOS technology functional from 4.2 K to 300 K," *IEEE journal of solid-state circuits*, vol. 44, no. 7, pp. 2019–2025, 2009.

[63] B. Okcan, P. Merken, G. Gielen, and C. Van Hoof, "A cryogenic analog to digital converter operating from 300 K down to 4.4 K," *Review of Scientific Instruments*, vol. 81, no. 2, p. 024702, 2010.

[64] F. Liu, T. Xue, Z. Deng, and Y. Liu, "Evaluation and lifetime study of a cryogenic SAR ADC for hpge detectors in low-background physics experiments," *Journal of Instrumentation*, vol. 14, no. 07, p. T07005, 2019.

[65] F. Liu, Z. Deng, X. Zhao, and Y. Liu, "A Switched Capacitor Waveform Digitizing ASIC at Cryogenic Temperature for HPGe Detectors," *IEEE Transactions on Nuclear Science*, 2021.

[66] Y.-q. Zhao, M. Yang, and H.-l. Zhao, "A cryogenic 10-bit successive approximation register analog-to-digital converter design with modified device model," *Journal of Shanghai Jiaotong University (Science)*, vol. 18, no. 5, pp. 520–525, 2013.

[67] T. Jiang, W. Liu, F. Y. Zhong, C. Zhong, K. Hu, and P. Y. Chiang, "A single-channel, 1.25-GS/s, 6-bit, 6.08-mW asynchronous successive-approximation ADC with improved feedback delay in 40-nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 47, no. 10, pp. 2444–2453, 2012.

[68] S.-W. M. Chen and R. W. Brodersen, "A 6-bit 600-MS/s 5.3-mW Asynchronous ADC in 0.13 μm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 12, pp. 2669–2680, 2006.

[69] S. Liu, T. Rabuske, J. Paramesh, L. Pileggi, and J. Fernandes, "Analysis and background self-calibration of comparator offset in loop-unrolled SAR ADCs," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 65, no. 2, pp. 458–470, 2017.

[70] P. t'Hart, M. Babaie, E. Charbon, A. Vladimirescu, F. Sebastiano *et al.*, "Characterization and modeling of mismatch in cryo-CMOS," *IEEE Journal of the Electron Devices Society*, vol. 8, pp. 263–273, 2020.

[71] H. Bohuslavskyi, S. Barraud, V. Barral, M. Casse, L. Le Guevel, L. Hutin, B. Bertrand, A. Crippa, X. Jehl, G. Pillonnet *et al.*, "Cryogenic characterization of 28-nm FD-SOI

ring oscillators with energy efficiency optimization," *IEEE Transactions on Electron Devices*, vol. 65, no. 9, pp. 3682–3688, 2018.

[72] J. Van Dijk, G. Kiene, R. Overwater, P. Padalia, J. Van Staveren, M. Babaie, A. Vladimirescu, E. Charbon, F. Sebastiano *et al.*, "Cryo-CMOS for analog/mixed-signal circuits and systems," in *2020 IEEE Custom Integrated Circuits Conference (CICC)*. IEEE, 2020, pp. 1–8.

[73] R. Navid *et al.*, "High-frequency noise in nanoscale metal oxide semiconductor field effect transistors," *Journal of applied physics*, vol. 101, no. 12, p. 124501, 2007.

[74] 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," *IEEE Journal of the Electron Devices Society*, vol. 8, pp. 448–456, 2020.

[75] P. t'Hart, T. Huizinga, M. Babaie, A. Vladimirescu, F. Sebastiano *et al.*, "Integrated Cryo-CMOS Temperature Sensors for Quantum Control ICs," in *2022 IEEE 15th Workshop on Low Temperature Electronics (WOLTE)*. IEEE, 2022, pp. 1–4.

[76] T. Kobayashi, K. Nogami, T. Shirotori, Y. Fujimoto, and O. Watanabe, "A current-mode latch sense amplifier and a static power saving input buffer for low-power architecture," in *1992 Symposium on VLSI Circuits Digest of Technical Papers*. IEEE, 1992, pp. 28–29.

[77] L. Chen, K. Ragab, X. Tang, J. Song, A. Sanyal, and N. Sun, "A 0.95-mW 6-b 700-MS/s single-channel loop-unrolled SAR ADC in 40-nm CMOS," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 64, no. 3, pp. 244–248, 2016.

[78] Y.-H. Chung, W.-S. Rih, and C.-W. Chang, "A 6-bit 1.3-GS/s Ping-Pong Domino-SAR ADC in 55-nm CMOS," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 65, no. 8, pp. 999–1003, 2018.

[79] A. Akkaya, F. Celik, and Y. Leblebici, "An 8-Bit 800 MS/s Loop-Unrolled SAR ADC With Common-Mode Adaptive Background Offset Calibration in 28 nm FDSOI," *IEEE Transactions on Circuits and Systems I: Regular Papers*, 2021.

[80] N. Kurosawa, H. Kobayashi, K. Maruyama, H. Sugawara, and K. Kobayashi, "Explicit analysis of channel mismatch effects in time-interleaved ADC systems," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 48, no. 3, pp. 261–271, 2001.

[81] G. Kiene, "Data for: A 1-gs/s 6-to-8b cryo-cmos sar adc for quantum computing," *doi 10.4121/21089596*, 2021.

[82] P. T. Hart, M. Babaie, A. Vladimirescu, and F. Sebastiano, "Characterization and Modeling of Self-Heating in Nanometer Bulk-CMOS at Cryogenic Temperatures," *IEEE Journal of the Electron Devices Society*, vol. 9, pp. 891–901, 2021.

[83] G. Kiene, A. G. Sreenivasulu, R. W. Overwater, M. Babaie, and F. Sebastiano, "Cryogenic Comparator Characterization and Modeling for a Cryo-CMOS 7b 1-GSa/s SAR ADC," in *ESSCIRC 2022-IEEE 48th European Solid State Circuits Conference (ESSCIRC)*. IEEE, 2022, pp. 53–56.

[84] G. Kiene, A. Catania, R. Overwater, P. Bruschi, E. Charbon, M. Babaie, and F. Sebastiano, "13.4 A 1GS/s 6-to-8b 0.5 mW/Qubit Cryo-CMOS SAR ADC for Quantum Computing in 40nm CMOS," in *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 64. IEEE, 2021, pp. 214–216.

[85] P. Nuzzo *et al.*, "Noise analysis of regenerative comparators for reconfigurable ADC architectures," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 55, no. 6, pp. 1441–1454, 2008.

[86] T. Sepke *et al.*, "Noise analysis for comparator-based circuits," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 56, no. 3, pp. 541–553, 2008.

[87] S. Das and J. C. Bardin, "Characterization of Shot Noise Suppression in Nanometer MOSFETs," in *2021 IEEE MTT-S International Microwave Symposium (IMS)*. IEEE, 2021, pp. 892–895.

[88] S. Van Winckelt, A. Caglar, B. Gys, S. Brebels, A. Potocnik, B. Parvais, P. Wambacq, and J. Craninckx, "A 28nm 6.5-8.1 GHz 1.16 mW/qubit Cryo-CMOS System-an-Chip for Superconducting Qubit Readout," in *ESSCIRC 2022-IEEE 48th European Solid State Circuits Conference (ESSCIRC)*. IEEE, 2022, pp. 61–64.

[89] T. Miki, T. Ozeki, and J.-i. Naka, "A 2-gs/s 8-bit time-interleaved SAR ADC for millimeter-wave pulsed radar baseband SoC," *IEEE Journal of Solid-State Circuits*, vol. 52, no. 10, pp. 2712–2720, 2017.

[90] K. Bult, M. S. Akter, and R. Sehgal, "High-efficiency residue amplifiers," *Low-Power Analog Techniques, Sensors for Mobile Devices, and Energy Efficient Amplifiers: Advances in Analog Circuit Design 2018*, pp. 253–296, 2019.

[91] L. R. Carley and T. Mukherjee, "High-speed low-power integrating CMOS sample-and-hold amplifier architecture," in *Proceedings of the IEEE 1995 Custom Integrated Circuits Conference*. IEEE, 1995, pp. 543–546.

[92] B. Malki, T. Yamamoto, B. Verbruggen, P. Wambacq, and J. Craninckx, "A 70 dB DR 10 b 0-to-80 MS/s current-integrating SAR ADC with adaptive dynamic range," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 5, pp. 1173–1183, 2014.

[93] B. Malki, B. Verbruggen, E. Martens, P. Wambacq, and J. Craninckx, "A 150 kHz–80 MHz BW discrete-time analog baseband for software-defined-radio receivers using a 5th-order IIR LPF, active FIR and a 10 bit 300 MS/s ADC in 28 nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 7, pp. 1593–1606, 2016.

[94] W. Jiang, Y. Zhu, C.-H. Chan, B. Murmann, and R. P. Martins, "A 7-bit 2 GS/s time-interleaved SAR ADC with timing skew calibration based on current integrating

sampler," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 68, no. 2, pp. 557–568, 2020.

[95] M. S. Akter, K. A. Makinwa, and K. Bult, "A capacitively degenerated 100-dB linear 20–150 MS/s dynamic amplifier," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 4, pp. 1115–1126, 2017.

[96] L. Shen, Y. Shen, Z. Li, W. Shi, X. Tang, S. Li, W. Zhao, M. Zhang, Z. Zhu, and N. Sun, "A two-step ADC with a continuous-time SAR-based first stage," *IEEE Journal of Solid-State Circuits*, vol. 54, no. 12, pp. 3375–3385, 2019.

[97] P. T. Hart, J. Van Dijk, M. Babaie, E. Charbon, A. Vladimircscu, F. Sebastiano *et al.*, "Characterization and model validation of mismatch in nanometer cmos at cryogenic temperatures," in *2018 48th European Solid-State Device Research Conference (ESSDERC)*. IEEE, 2018, pp. 246–249.

[98] G. Kiene, R. W. Overwater, A. Catania, A. G. Sreenivasulu, P. Bruschi, E. Charbon, M. Babaie, and F. Sebastiano, "A 1-GS/s 6–8-b Cryo-CMOS SAR ADC for Quantum Computing," *IEEE Journal of Solid-State Circuits*, 2023.

[99] R. W. Overwater, M. Babaie, and F. Sebastiano, "Cryogenic-aware forward body biasing in bulk CMOS," *IEEE Electron Device Letters*, 2023.

[100] E. Alpman, H. Lakdawala, L. R. Carley, and K. Soumyanath, "A 1.1 v 50mW 2.5 Gs/s 7b time-interleaved c-2c SAR ADC in 45nm LP digital CMOS," in *2009 IEEE International Solid-State Circuits Conference-Digest of Technical Papers*. IEEE, 2009, pp. 76–77.

[101] X. Huang, H. Londoño-Ramírez, M. Ballini, C. Van Hoof, J. Genoe, S. Haesler, G. Gielen, N. Van Helleputte, and C. M. Lopez, "Actively Multiplexed μECoG Brain Implant System with Incremental-$\delta\sigma$ ADCs Employing Bulk-DACs," *IEEE Journal of Solid-State Circuits*, vol. 57, no. 11, pp. 3312–3323, 2022.

[102] D. P. Foty, "Impurity ionization in MOSFETs at very low temperatures," *Cryogenics*, vol. 30, no. 12, pp. 1056–1063, 1990.

[103] J. Zhang, Z. Wang, R. Wang, Z. Sun, and R. Huang, "Body Bias Dependence of Bias Temperature Instability (BTI) in Bulk FinFET Technology," *Energy & Environmental Materials*, vol. 5, no. 4, pp. 1200–1203, 2022.

[104] L. Kull, T. Toifl, M. Schmatz, P. A. Francese, C. Menolfi, M. Braendli, M. Kossel, T. Morf, T. M. Andersen, and Y. Leblebici, "22.1 a 90gs/s 8b 667mw 64× interleaved sar adc in 32nm digital soi cmos," in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. IEEE, 2014, pp. 378–379.

[105] Y. Duan and E. Alon, "A 12.8 GS/s time-interleaved ADC with 25 GHz effective resolution bandwidth and 4.6 ENOB," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 8, pp. 1725–1738, 2014.

[106] H. Yoon, C. Lee, T. Kim, Y. Kwon, and Y. Chae, "A 65-dB-SNDR Pipelined SAR ADC Using PVT-Robust Capacitively Degenerated Dynamic Amplifier," *IEEE Journal of Solid-State Circuits*, vol. 58, no. 4, pp. 961–971, 2023.

[107] J. Gong, E. Charbon, F. Sebastiano, and M. Babaie, "A Cryo-CMOS PLL for Quantum Computing Applications," *IEEE Journal of Solid-State Circuits*, 2022.

[108] A. M. Ali, H. Dinc, P. Bhoraskar, S. Bardsley, C. Dillon, M. McShea, J. P. Periathambi, and S. Puckett, "A 12-b 18-GS/s RF sampling ADC with an integrated wideband track-and-hold amplifier and background calibration," *IEEE Journal of Solid-State Circuits*, vol. 55, no. 12, pp. 3210–3224, 2020.

[109] G. Kiene, S. İlik, L. Mastrodomenico, M. Babaie, and F. Sebastiano, "Cryogenic Characterization of Low-Frequency Noise in 40-nm CMOS," *IEEE Journal of the Electron Devices Society*, 2024.

[110] I. Hafez, G. Ghibaudo, and F. Balestra, "Flicker noise in metal-oxide-semiconductor transistors from liquid helium to room temperature," *Journal of applied physics*, vol. 66, no. 5, pp. 2211–2213, 1989.

[111] M. Aoki, H. Katto, and E. Yamada, "Low-frequency 1/f noise in MOSFET's at low current levels," *Journal of Applied Physics*, vol. 48, no. 12, pp. 5135–5140, 1977.

[112] B. C. Paz, M. Cassé, C. Theodorou, G. Ghibaudo, T. Kammler, L. Pirro, M. Vinet, S. de Franceschi, T. Meunier, and F. Gaillard, "Performance and low-frequency noise of 22-nm FDSOI down to 4.2 K for cryogenic applications," *IEEE Transactions on Electron Devices*, vol. 67, no. 11, pp. 4563–4567, 2020.

[113] G. Ghibaudo, O. Roux, C. Nguyen-Duc, F. Balestra, and J. Brini, "Improved analysis of low frequency noise in field-effect MOS transistors," *physica status solidi (a)*, vol. 124, no. 2, pp. 571–581, 1991.

[114] I. Hafez, G. Ghibaudo, and F. Balestra, "Assessment of interface state density in silicon metal-oxide-semiconductor transistors at room, liquid-nitrogen, and liquid-helium temperatures," *Journal of applied physics*, vol. 67, no. 4, pp. 1950–1952, 1990.

[115] H. Bohuslavskyi, A. Jansen, S. Barraud, V. Barral, M. Cassé, L. Le Guevel, X. Jehl, L. Hutin, B. Bertrand, G. Billiot *et al.*, "Cryogenic subthreshold swing saturation in FD-SOI MOSFETs described with band broadening," *IEEE Electron Device Letters*, vol. 40, no. 5, pp. 784–787, 2019.

[116] H. Oka, T. Matsukawa, K. Kato, S. Iizuka, W. Mizubayashi, K. Endo, T. Yasuda, and T. Mori, "Toward long-coherence-time Si spin Qubit: The origin of low-frequency noise in cryo-CMOS," in *2020 IEEE Symposium on VLSI Technology*. IEEE, 2020, pp. 1–2.

[117] M. B. da Silva, H. P. Tuinhout, A. Zegers-van Duijnhoven, G. I. Wirth, and A. J. Scholten, "A physics-based statistical RTN model for the low frequency noise in MOSFETs," *IEEE Transactions on Electron Devices*, vol. 63, no. 9, pp. 3683–3692, 2016.

[118] P. t Hart, M. Babaie, E. Charbon, A. Vladimirescu, and F. Sebastiano, "Subthreshold mismatch in nanometer CMOS at cryogenic temperatures," in *ESSDERC 2019-49th European Solid-State Device Research Conference (ESSDERC)*. IEEE, 2019, pp. 98–101.

[119] G. Kiene and S. Ilik, "Data for: Low frequency noise characterization in 40nm Cryo-CMOS," *doi 10.4121/c41c26f7-586f-48ca-8abe-9127b2d97c60*, 2023.

[120] C. Rogers and R. Buhrman, "Nature of single-localized-electron states derived from tunneling measurements," *Physical review letters*, vol. 55, no. 8, p. 859, 1985.

[121] J. Michl, A. Grill, D. Claes, G. Rzepa, B. Kaczer, D. Linten, I. Radu, T. Grasser, and M. Waltl, "Quantum mechanical charge trap modeling to explain BTI at cryogenic temperatures," in *2020 IEEE International Reliability Physics Symposium (IRPS)*. IEEE, 2020, pp. 1–6.

[122] M. Kirton and M. Uren, "Noise in solid-state microstructures: A new perspective on individual defects, interface states and low-frequency (1/f) noise," *Advances in Physics*, vol. 38, no. 4, pp. 367–468, 1989.

[123] J. Michl, A. Grill, B. Stampfer, D. Waldhoer, C. Schleich, T. Knobloch, E. Ioannidis, H. Enichlmair, R. Minixhofer, B. Kaczer *et al.*, "Evidence of Tunneling Driven Random Telegraph Noise in Cryo-CMOS," in *2021 IEEE International Electron Devices Meeting (IEDM)*. IEEE, 2021, pp. 31–3.

[124] A. Schenk, P. P. Altermatt, and B. Schmithusen, "Physical model of incomplete ionization for silicon device simulation," in *2006 International Conference on Simulation of Semiconductor Processes and Devices*. IEEE, 2006, pp. 51–54.

[125] G. Ghibaudo, "On the fluctuation-dissipation of the oxide trapped charge in a MOSFET operated down to deep cryogenic temperatures," *arXiv preprint arXiv:2204.04958*, 2022.

[126] C. C. Enz and G. C. Temes, "Circuit techniques for reducing the effects of op-amp imperfections: autozeroing, correlated double sampling, and chopper stabilization," *Proceedings of the IEEE*, vol. 84, no. 11, pp. 1584–1614, 1996.

[127] D. Keith, M. House, M. Donnelly, T. Watson, B. Weber, and M. Simmons, "Single-shot spin readout in semiconductors near the shot-noise sensitivity limit," *Physical Review X*, vol. 9, no. 4, p. 041003, 2019.

[128] J. M. Martinis, "Qubit metrology for building a fault-tolerant quantum computer," *npj Quantum Information*, vol. 1, no. 1, pp. 1–3, 2015.

[129] M. Veldhorst, C. Yang, J. Hwang, W. Huang, J. Dehollain, J. Muhonen, S. Simmons, A. Laucht, F. Hudson, K. M. Itoh *et al.*, "A two-qubit logic gate in silicon," *Nature*, vol. 526, no. 7573, pp. 410–414, 2015.

[130] D. Keith, S. Gorman, L. Kranz, Y. He, J. Keizer, M. Broome, and M. Simmons, "Benchmarking high fidelity single-shot readout of semiconductor qubits," *New Journal of Physics*, vol. 21, no. 6, p. 063011, 2019.

[131] A. E. Seedhouse, T. Tanttu, R. C. Leon, R. Zhao, K. Y. Tan, B. Hensen, F. E. Hudson, K. M. Itoh, J. Yoneda, C. H. Yang *et al.*, "Pauli blockade in silicon quantum dots with spin-orbit control," *PRX Quantum*, vol. 2, no. 1, p. 010303, 2021.

[132] L. C. Camenzind, S. Geyer, A. Fuhrer, R. J. Warburton, D. M. Zumbühl, and A. V. Kuhlmann, "A hole spin qubit in a fin field-effect transistor above 4 kelvin," *Nature Electronics*, vol. 5, no. 3, pp. 178–183, 2022.

[133] C. Volk, A.-M. J. Zwerver, U. Mukhopadhyay, P. T. Eendebak, C. J. van Diepen, J. P. Dehollain, T. Hensgens, T. Fujita, C. Reichl, W. Wegscheider *et al.*, "Loading a quantum-dot based "Qubyte" register," *npj Quantum Information*, vol. 5, no. 1, pp. 1–8, 2019.

[134] A. Korotkov, "Intrinsic noise of the single-electron transistor," *Physical Review B*, vol. 49, no. 15, p. 10381, 1994.

[135] S. Kafanov and P. Delsing, "Measurement of the shot noise in a single-electron transistor," *Physical Review B*, vol. 80, no. 15, p. 155320, 2009.

[136] U. Fano, "Ionization yield of radiations. ii. the fluctuations of the number of ions," *Physical Review*, vol. 72, no. 1, p. 26, 1947.

[137] E. J. Connors, J. Nelson, H. Qiao, L. F. Edge, and J. M. Nichol, "Low-frequency charge noise in Si/SiGe quantum dots," *Physical Review B*, vol. 100, no. 16, p. 165305, 2019.

[138] G. Zimmerli, T. M. Eiles, R. L. Kautz, and J. M. Martinis, "Noise in the Coulomb blockade electrometer," *Applied Physics Letters*, vol. 61, no. 2, pp. 237–239, 1992.

[139] L. Petit, J. Boter, H. Eenink, G. Droulers, M. Tagliaferri, R. Li, D. Franke, K. Singh, J. Clarke, R. Schouten *et al.*, "Spin lifetime and charge noise in hot silicon quantum dot qubits," *Physical review letters*, vol. 121, no. 7, p. 076801, 2018.

[140] T. Watson, S. Philips, E. Kawakami, D. Ward, P. Scarlino, M. Veldhorst, D. Savage, M. Lagally, M. Friesen, S. Coppersmith *et al.*, "A programmable two-qubit quantum processor in silicon," *nature*, vol. 555, no. 7698, pp. 633–637, 2018.

[141] P. T. Hart, M. Babaie, E. Charbon, A. Vladimirescu, and F. Sebastiano, "Subthreshold mismatch in nanometer CMOS at cryogenic temperatures," *IEEE Journal of the Electron Devices Society*, vol. 8, pp. 797–806, 2020.

[142] R. A. Damsteegt, R. W. Overwater, M. Babaie, and F. Sebastiano, "A Benchmark of Cryo-CMOS 40-nm Embedded SRAM/DRAMs for Quantum Computing," in *ESSCIRC 2023-IEEE 49th European Solid State Circuits Conference (ESSCIRC)*. IEEE, 2023, pp. 165–168.

[143] H. Rücker, J. Korn, and J. Schmidt, "Operation of SiGe HBTs at cryogenic temperatures," in *2017 IEEE Bipolar/BiCMOS Circuits and Technology Meeting (BCTM)*. IEEE, 2017, pp. 17–20.

[144] G. Turin, "An introduction to matched filters," *IRE transactions on Information theory*, vol. 6, no. 3, pp. 311–329, 1960.

[145] M. S. Steyaert and W. M. Sansen, "A micropower low-noise monolithic instrumentation amplifier for medical purposes," *IEEE journal of solid-state circuits*, vol. 22, no. 6, pp. 1163–1168, 1987.

[146] E. Kammerloher, A. Schmidbauer, L. Diebel, I. Seidler, M. Neul, M. Künne, A. Ludwig, J. Ritzmann, A. Wieck, D. Bougeard *et al.*, "Sensing dot with high output swing for scalable baseband readout of spin qubits," *arXiv preprint arXiv:2107.13598*, 2021.

[147] A. Caizzone, A. Boukhayma, and C. Enz, "An accurate kTC noise analysis of CDS circuits," in *2018 16th IEEE International New Circuits and Systems Conference (NEWCAS)*. IEEE, 2018, pp. 22–25.

[148] Y. Cao, J. Romero, and A. Aspuru-Guzik, "Potential of quantum computing for drug discovery," *IBM Journal of Research and Development*, vol. 62, no. 6, pp. 6–1, 2018.

[149] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. Van Den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme *et al.*, "Evidence for the utility of quantum computing before fault tolerance," *Nature*, vol. 618, no. 7965, pp. 500–505, 2023.

[150] R. H. Koch, "Relationship between low-frequency noise and resistance drift in metal films," *Physical Review B*, vol. 48, no. 16, p. 12217, 1993.

[151] J. Pavelka, J. Sikula, P. Vasina, V. Sedlakova, M. Tacano, and S. Hashiguchi, "Noise and transport characterisation of tantalum capacitors," *Microelectronics Reliability*, vol. 42, no. 6, pp. 841–847, 2002.

[152] O. Burkeman, *Four thousand weeks: Time management for mortals*. Farrar, Straus and Giroux, 2021.

Summary

Quantum computers promise large speedups compared to classical computers for specific classes of problems by exploiting quantum phenomena for computation. Despite notable progress towards larger systems, today's quantum computers still lack the necessary size for realizing many of the anticipated benefits. Scaling to these larger numbers is complicated by the fragility of the qubits, which raises the necessity of operating the quantum computing core at deep cryogenic temperatures to limit the influence of noise. However, in the present experimental setups for quantum computing the electronic interfaces are often based on equipment operated at room temperature. This arrangement creates a potential wiring bottleneck towards cryogenic temperature, that is connected to reliability limitations. A proposed solution to this challenge is moving the interfacing electronics close to the qubits into the cryogenic environment. However, implementation of the electronics at cryogenic temperature places a set of system constraints that prevents directly porting the room-temperature electronics. The most prominent of these constraints is the limited cooling power at deep cryogenic temperature, requiring the interfacing electronics to save power wherever possible. To implement the necessary functions of the interface in practice, CMOS technology is uniquely positioned due to its large scale integration capabilities, low-power digital logic and sustained performance at cryogenic temperature.

In this work, contributions to the readout electronics of quantum computers are made, with an emphasis on the readout of semiconductor spin quantum computers. Readout poses a significant challenge for current spin-qubit systems, especially its scalable integration into the electronic interface. This dissertation can be separated into contributions to the readout of quantum computers in three main categories: cryogenic digitization, cryogenic noise characterization and readout benchmarking.

The wide-band signals present in the frequency-multiplexed RF readout interface require corresponding wide-band digitization capabilities of the associated data converters. In this work, we present robust, power-efficient and high-speed time-interleaved SAR ADC designs adopted to operation at cryogenic temperature to serve this application. These include the first high-speed cryo-CMOS SAR ADC suitable for operation in the RF readout interface, further efficiency improvements to this ADC and finally an ADC tightly integrated with an efficient body-bias enabled dynamic pre-amplifier. The body-bias enabled pre-amplifier also demonstrated excellent room-temperature performance, with the best reported combined power efficiency of a high-speed SAR and driver circuit. All designs have been realized in a 40nm CMOS technology and characterized at cryogenic temperature.

For the directed design of many of the functions in the electronic interface, accurate models for simulation are required. While significant work has been presented on DC characterization at cryogenic temperatures, comparatively little work covering low-frequency noise is available. To help further development, an extensive low-frequency noise characterization is performed on a 40nm CMOS technology at cryogenic temperature. The results

of this characterization allow a broad overview over the low frequency noise behavior of the technology. Given the radical change in temperature between room-temperature and cryogenic temperature, a surprisingly small overall change in the input-referred noise is observed. The most prominent difference is the finding of a systematic Lorentzian feature at cryogenic temperature that is for the first time described here.

Lastly, while promising designs have been suggested for the DC-readout of spin qubits, guidelines for directing such designs have been lacking. In order to provide a set of such guidelines, we first demonstrate that the limits of DC-readout lie significantly beyond current state-of-the-art, with much room for improvement on the electronics. Then, we analyze the voltage amplifier, the transimpedance amplifier, the charge sampling, and the current pre-amplifier by deriving their design equations and trade-offs. Finally the various architectures are compared with the fundamental limit and suggestions for specific use-cases are give. The results suggest that DC-readout is a promising path for the development of a scalable spin-qubit readout.

Samenvatting

Kwantumcomputers beloven grote snelheidsverbeteringen ten opzichte van klassieke computers voor specifieke classificaties van problemen, door gebruik te maken van kwantumeffecten in de berekeningen. Ondanks de noemenswaardige vooruitgang naar grotere systemen, ontbreekt het hedendaagse kwantumcomputers nog steeds aan de benodigde schaal om de verwachte voordelen te realiseren. Het groeien naar deze groottes wordt bemoeilijkt door de kwetsbaarheid van de qubits, wat de noodzaak oproept om de kern van de kwantumcomputer te laten opereren op extreem cryogene temperaturen, om zo de invloed van ruis te drukken. Desondanks zijn in hedendaagse experimentele opstellingen voor kwantumcomputers de elektronische interfaces vaak gebaseerd op apparatuur die draait op kamertemperatuur. Deze configuratie creëert een potentieel bedradingsknelpunt naar cryogene temperaturen, wat samenhangt met limitaties van de betrouwbaarheid. Een voorgestelde oplossing voor deze uitdaging is het plaatsen van de elektronica dicht bij de qubits in de cryogene omgeving. Het implementeren van de elektronica op cryogene temperaturen zorgt echter voor een set aan systeembeperkingen die voorkomen dat de kamertemperatuur-elektronica direct gebruikt kan worden. De belangrijkste beperking is het gelimiteerde koelvermogen op extreem cryogene temperaturen, waardoor de elektronica waar mogelijk energiebesparender moet zijn. Om de benodigde functies in de praktijk te implementeren, is CMOS technologie uniek gepositioneerd; door zijn vermogen op grote schaal geïntegreerd te worden, het lage verbruik van digitale logica en de aanhoudende prestaties op cryogene temperaturen.

In dit werk worden bijdragen geleverd aan de uitlees-elektronica van kwantumcomputers, met de nadruk op het uitlezen van halfgeleider spin-qubit kwantumcomputers. Het uitlezen vormt een belangrijke uitdaging voor huidige spin-qubit-systemen, voornamelijk in het schaalbaar maken van de integratie met de elektronica. Dit proefschrift kan in zijn bijdragen aan het uitlezen van kwantumcomputers worden verdeeld in drie hoofdcategorieën: cryogene digitalisering, cryogene ruis-karakterisatie en het maken van maatstaven voor uitlezingen.

De breedband signalen die aanwezig zijn in de frequentie-gemultiplexte RF uitleesinterface vereisen overeenkomstige breedband-digitaliseringsvermogens van de geassocieerde dataomzetters. In dit werk presenteren we robuuste, vermogens efficiënte, hogesnelheids- en time-interleaved SAR ADC ontwerpen, aangepast om op cryogene temperaturen te opereren voor gebruik in deze applicatie. Deze bevatten de eerste hogesnelheids-cryo-CMOS SAR ADC, geschikt om te opereren in de RF uitleesinterface, verdere efficientieverbeteringen aan deze ADC en uiteindelijk een ADC die nauw geïntegreerd is met een efficiënte voorversterker mogelijk gemaakt door substraat-bias. De door substraat-bias mogelijk gemaakte voorversterker demonstreerde ook uitstekende kamertemperatuur prestaties met de best vermelde vermogens efficiëntie van de combinatie van een hogesnelheids-SAR en sturingscircuit. Alle ontwerpen zijn gerealiseerd in een 40nm CMOS technologie en gekarakteriseerd op cryogene temperaturen.

Voor het sturen van het ontwerp van veel van de functies in de elektronica-interface zijn accurate modellen nodig voor simulaties. Hoewel reeds significant werk is gepresenteerd over de DC karakterisatie op cryogene temperaturen, is relatief weinig werk beschikbaar over laagfrequente ruis. Om verdere ontwikkeling te stimuleren, is een uitgebreide laagfrequente ruis-karakterisatie uitgevoerd op een 40nm CMOS technologie op cryogene temperaturen. De resultaten van deze karakterisatie geven een breed overzicht van het gedrag van laagfrequente ruis in deze technologie. Ondanks de radicale verandering in temperatuur tussen kamertemperatuur en cryogene temperaturen is er maar een verrassend kleine verandering zichtbaar in de input-verwezen ruis. Het voornaamste verschil is de ontdekking van een systematisch Lorentzian-kenmerk op cryogene temperaturen, die hier voor het eerst is beschreven.

Tenslotte missen er ondanks veelbelovende ontwerpvoorstellen voor de DC uitlezing van spin qubits richtlijnen voor het sturen van zulke ontwerpen. Om deze regelgeving te kunnen verkrijgen demonstreren we eerst dat de limieten van de DC uitlezing ver voorbij de huidige stand van de techniek zijn, met veel ruimte voor verbetering van de elektronica. Daarna analyseren we de voltage-versterker, de transimpedantie-versterker, de ladingssampling, en de stroomvoorversterker, door middel van het afleiden van hun ontwerpsvergelijkingen en afwegingen. Uiteindelijk vergelijken we de verscheidene architecturen met de fundamentele limiet en geven we suggesties voor specifieke voorbeelden. De resultaten suggereren dat DC uitlezing een veelbelovend pad is voor het ontwikkelen van schaalbare spin-qubit uitlezing.

Acknowledgments

Assuming an average human lifetime of about 4000 weeks [152], this PhD took about 7.1% of my life's total expected time-budget. I found this stretch of time to be filled with a wealth of professional and personal growth. Given the non-refundable nature of life's weeks, that's a lucky coincidence. In the following, I want to express my sincere gratitude towards the humans that made this time so fruitful and/or enjoyable.

First, I want to thank my supervisor **Fabio Sebastiano** for the opportunity to work on this project. Your guidance through the various stages of this effort has been invaluable. I greatly benefited from our technical discussions. Thank you also for your people-centric view of affairs where it was necessary to find non-technical solutions to challenges. I want to also express my gratitude towards my co-supervisor **Masoud Babaie**. Your brilliant attention to technical detail and relentless pursuit of knowledge are inspiring. Thank you especially for your input in design reviews.

As to teachers I want to thank **Alyosha Molnar** for his advanced analog design class that introduced circuits to me as a physicist. Also thanks to **Klaas Bult** for discussions about ADCs in the beginning phase of the project that became central to my PhD. His classes, in which I TA'd, also became a central part of my education.

The work in this thesis was done in a research group that had several fantastic official names over the years, but was constantly referred to as the **coolgroup** by it's members. We spend many hours striving, both (occasionally) failing and (occasionally) succeeding together: **Ramon** we shared a lot of projects and countless technical and non-technical discussions, which I greatly enjoyed. Without the collaboration with you, the ADC projects would never have taken the shape that they did in the end. Thanks to Nina and you for help with the translation of the thesis summary as well as for many fun rounds of boardgames! **Augusto** you were showing me the ropes of the PhD, the insane hours of my first and your last tapeout in the coolgroup will not be forgotten. Thanks also for epic runs and for laying the foundation for *augusto.tcl*. **Luc** you were a great office mate and fellow runner, keep your attention with the breath! **Alessandro** it was awesome that you came to visit us in Delft, the collaboration on the first ADC iteration was great, but also all around Delft and Rotterdam we had a lot of fun. **Job** your dry humor and calm attitude were always very enjoyable. You came to the group as a masters student just about the same time as I joined as a PhD, making you my "longest standing" colleague. Finally thanks a lot for ordering components like no other! **Jeroen** we had a lot of fun as office mates and sports buddies during the first two years. **Sadik** you have great mind for devices, great complementary skillset for our work on LFN, all the best for your next steps into noise characterization! **Bishnu**, RF design always comes across as a bit of wizzardry, and this force runs strong in you. Thanks for dinners and being a fun office mate. **Rob** we had some great times in the lab, your very clear perspective on things was always very helpful. **Pascal** we had a great time on that summer school in China, and you are the most distinguished electro-mechanical hacker I know of. May you always find the pallets necessary to keep

your fireplace warm. **Jiang** you are an absolute circuit design machine, keep on crushing. **Bagas** my fellow read-out warrior - all the best with carrying on the fight. **Niels** we had some great times in the lab and during those gruesome tapeouts. **Harald** we only interacted briefly towards the end of your PhD, but your unique attitude towards work and life is the stuff legends are made of. Thanks a lot to **Pinakin** for the good times we hand when building dipsticks together. **Ali** I'm happy to give the SiGe project into your hands, make it count! **Milad** your calm attitude and design expertise were always looked up to. Also, I kept your plant alive till the very end of my PhD, it was a trusty companion.

I'm greatly indebted to the masters students I supervised. During the supervision of each of you I greatly enjoyed the collaboration and teaching aspects, as well as of course the successes we celebrated together. **Rikadur** your strong practical skills and calmness in face of herculean tasks was greatly appreciated. **Aishwarya** you are a brilliant circuit designer, all the best for your future career! **Luigi** your excellent practical and measurement skills helped a lot with the difficult project at hand, your great resolve helped you prevail over the difficulties in measuring noise.

Also the other master students always added to the progress of the coolgroup. Thanks to **Lizzy**, **David**, **Andrea**, **Ilker**, **Ivor**, **Lennard**, **Rishab**, **Yu-Wen**, **Maoran**, **Deb**, **Tim**, **Zenghui** for countless fruitful and fun discussions.

Also the colleagues in Physics I want to thank: **Andrea** you have taught me what a spin qubit is and how to screw apart a dilution refrigerator. Also your jokes and humor really helped with arriving in Qutech - and turns out we are colleagues now again, abeit in different locations. **Oriol**: it was great fun to share an office during the coolgroup's diaspora (the coolgroup had to move from the EWI tower to Physics for issues of fire-safety for a while) as well as work towards a useful DC readout together. **Xiao** thanks for the great discussions on spin qubits. **Sergey** thanks for the brilliant quantum-device manufacturing help. Thanks also to **Paolo**, **Alberto**, **Brennan** and **Florian** for fun times around lunch and Qutech activities.

Raymond your unmistakable style of explaining and discussing electrical problems was always of great help. I learned a lot from you, especially on interference. **Raymond** thanks a lot as well for many discussions and also for making the chip-test platform integrate with the SPI rack, I hope it finds many users. **Erik** and **Roy**: thank you for the great electronics support. **Jason** thanks for keeping the bonder running and always also carrying a great positive vibe everywhere you go. Which is a lot of places every day as you walk very fast. **Siebe**, **Remco** and **Jelle** thanks a lot for keeping the place physically holding together and the helium running. **Nico** and **Tim** thanks for the mechanical work that enabled all the dipsticks running in our lab and used in this thesis. **Marja** thanks for the administrative support, especially when the going got tough.

Outside of work the time in Delft has been made much richer by: **Tugrul** thanks for wonderful days out in the mountains. I especially appreciate the times you encouraged the more adventurous approach, here is to all the past and future adventures. **Stephan** we had some great times climbing together. Even more than the climbing itself, I enjoyed our conversations while cycling to and from the climbing. **Luka** we share a large enthusiasm in boardgames and bouldering. I got to know many new and fun games from you! **Thomas** and **Etu**, together with Luka we explored a lot of boardgames online during the pandemic and afterwards and had great fun. Especially epic in my memory where all those rounds of

"through the ages". **Douwe** I really enjoyed climbing with you and organizing Prorisc/Safe 2019. Your craftmanship is really something to look up to! **Jelle** thanks a lot for the assistance when my limbs where not coming along!

Thanks also to my friends from Villigst: **Gereon**, **Lioba**, **Paul**, **Elisa**, **Pia**, **Maxim**, **Anjela**, **Chris** for their friendship. Let's make Krehfeld happen again!

Also thanks to the friends from my studies in Heidelberg: **Dima**, we met right at the beginning of our studies and kept the contact over great distances in time and space. Thanks for the discussions! **Johannes** and **Florian** thanks for great hikes, via ferratas and deeptalk on the way. **Benjamin** thanks for your visits to NL, those were great fun!

Thanks to the friends in Munich: **Fabi** for tours in and around munich as well as boardgame evenings and for being a never-ending quell of information on bicycles. **Flo**, **Emi**, **Sascha** and **Elwin** as fellow boardgame ultras we look forward to many more delightful evenings.

Thanks also to the Friends from in and around Holtrup: **Fritt**, **Gerrit**, **Gordon**, **Hannes**, **Jonny**, **Marcel**, **Marten**, **Tjerk**, **Wilken**, thank you for your friendship, to many future trips/adventures at home and around the world.

My deepest gratitude belongs to my Family: **Hella** and **Jan-Dirk** for being the best siblings. **Mom** and **Dad** thank you for your unconditional support. This solid backing enabled me to focus without worries, a circumstance which greatly aided the success of my studies.

Finally, I want to thank **Lovisa**. You joined me at a challenging time during this journey, making it much brighter in an instant. Thank you for proofreading this thesis. Above all, thank you for your love.

Curriculum Vitae

Gerd Kiene

13.03.1991	Born in Nienburg (Weser), Germany
2011-2014	Bachelor in Physics at Universität Heidelberg, Thesis: *Evaluating the Synaptic Input of a Neuromorphic Circuit*
2014-2015	Heidelberg scholar at Cornell University
2015-2017	Master in Physics at Universität Heidelberg, Thesis: *Mixed-signal neuron and readout circuits for a neuromorphic system*
2017-2018	Researcher at Universität Heidelberg
2018-2023	PhD researcher at TU Delft, Thesis: *Cryogenic mixed-signal readout electronics for quantum computers*
2023	Mixed-signal design engineer at Rohde und Schwarz, Munich

List of Publications

Journal Publications

- **Gerd Kiene**, Ramon WJ Overwater, Alessandro Catania, Aishwarya Gunaputi Sreenivasulu, Paolo Bruschi, Edoardo Charbon, Masoud Babaie, and Fabio Sebastiano, *A 1-GS/s 6–8-b Cryo-CMOS SAR ADC for Quantum Computing*. IEEE Journal of Solid-State Circuits (2023).

- **Gerd Kiene**, Ramon WJ Overwater, Masoud Babaie, and Fabio Sebastiano, *A Cryo-CMOS SAR ADC With FIA Sampling Driver Enabled by Cryogenic-Aware Back-Biasing*. IEEE Transactions on Circuits and Systems I: Regular Papers (2023).

- **Gerd Kiene**, Sadik Ilik, Luigi Mastrodomenico, Masoud Babaie, and Fabio Sebastiano, *Cryogenic Characterization of Low-Frequency Noise in 40-nm CMOS*, IEEE Journal of the Electron Devices Society (2024).

- **Gerd Kiene**, et al., *SET DC readout: opportunities and limits (preliminary title)*, in preparation.

Conference Publications

- Mohammadreza Mehrpoo, Bishnu Patra, Jiang Gong, Jeroen van Dijk, Harald Homulle, **Gerd Kiene**, Andrej Vladimirescu, Fabio Sebastiano, Edoardo Charbon, and Masoud Babaie, *Benefits and challenges of designing cryogenic CMOS RF circuits for quantum computers.* In 2019 IEEE International Symposium on Circuits and Systems (ISCAS), pp. 1-5. IEEE, 2019.

- Jeroen van Dijk, **Gerd Kiene**, Ramon Overwater, Pinakin Padalia, Job Van Staveren, Masoud Babaie, Andrei Vladimirescu, Edoardo Charbon, and Fabio Sebastiano. *Cryo-CMOS for analog/mixed-signal circuits and systems.*, In 2020 IEEE Custom Integrated Circuits Conference (CICC), pp. 1-8. IEEE, 2020.

- **Gerd Kiene**, Alessandro Catania, Ramon Overwater, Paolo Bruschi, Edoardo Charbon, Masoud Babaie, and Fabio Sebastiano, *13.4 A 1GS/s 6-to-8b 0.5 mW/Qubit cryo-CMOS SAR ADC for Quantum Computing in 40nm CMOS.* In 2021 IEEE International Solid-State Circuits Conference (ISSCC), vol. 64, pp. 214-216. IEEE, 2021.

- **Gerd Kiene***, Aishwarya Gunaputi Sreenivasulu*, Ramon WJ Overwater, Masoud Babaie, and Fabio Sebastiano, *Cryogenic Comparator Characterization and Modeling for a Cryo-CMOS 7b 1-GSa/s SAR ADC.* In ESSCIRC 2022-IEEE 48th European Solid State Circuits Conference (ESSCIRC), pp. 53-56. IEEE, 2022. *equal contribution

Chip gallery

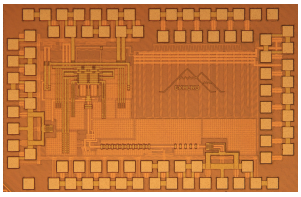

(a) ADC from chapter 2: *"Blisard"*

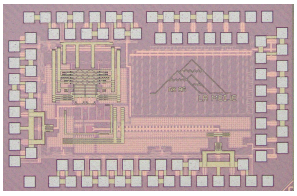

(b) ADC from chapter 3: *"La Meije"*

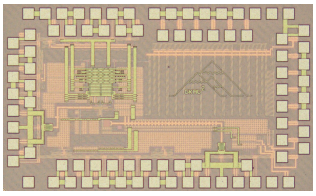

(c) ADC from chapter 4: *"Monte Rosa"*

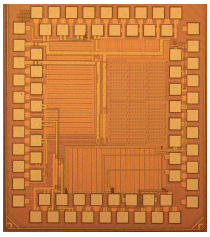

(a) First LFN chip

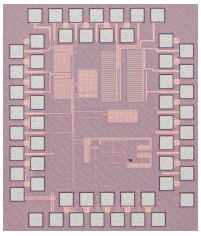

(b) LFN chip from chapter 5