



Delft University of Technology

Cryogenic electronics for the read-out of quantum processors

Homulle, Harald

DOI

[10.4233/uuid:e833f394-c8b1-46e2-86b8-da0c71559538](https://doi.org/10.4233/uuid:e833f394-c8b1-46e2-86b8-da0c71559538)

Publication date

2019

Document Version

Final published version

Citation (APA)

Homulle, H. (2019). *Cryogenic electronics for the read-out of quantum processors*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:e833f394-c8b1-46e2-86b8-da0c71559538>

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.

CRYOGENIC ELECTRONICS

FOR THE READ-OUT OF QUANTUM PROCESSORS

CRYOGENIC ELECTRONICS
FOR THE READ-OUT OF QUANTUM PROCESSORS

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus Prof. dr. ir. T. H. J. J. van der Hagen;
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op
woensdag 1 mei 2019 om 10:00 uur

door

Harald Arjan Robert HOMULLE

Elektrotechnisch ingenieur,
Technische Universiteit Delft, Nederland,
geboren te 's Gravenhage, Nederland.

Dit proefschrift is goedgekeurd door de promotor:

Prof. dr. ir. E. Charbon

Samenstelling promotiecommissie:

Rector Magnificus
Prof. dr. ir. E. Charbon

Voorzitter
Technische Universiteit Delft

Onafhankelijke leden:

Prof. dr. ir. S. Hamdioui

Technische Universiteit Delft

Prof. dr. D. P. DiVincenzo

Technische Universiteit Delft

Prof. dr. E. Prati

Forschungszentrum Jülich, Germany

Prof. dr. H. Bluhm

Politecnico di Milano, Italy

Dr. F. Gonzalez-Zalba


RWTH Aachen, Germany

University of Cambridge, United Kingdom

Overig lid:

Dr. F. Sebastiano

Technische Universiteit Delft



Keywords: cryogenic electronics, FPGA, CMOS, transistor,
qubit, quantum processor

Printed by: Ipskamp Printing, Enschede

Cover designed by: Harald Homulle

Copyright © 2019 by H. A. R. Homulle

ISBN 978-94-6384-029-3

An electronic version of this dissertation is available
in the TU Delft Repository.

0	0	1	1	1	1	1	1	0	0	0	1	0	1	0	0	1	0	0	1	0	1	1	0	1	1	0	0	0	0
0	1	0	0	0	1	0	1	0	0	1	1	0	1	0	1	0	0	1	1	0	0	1	0	0	1	0	0	1	1
0	1	1	1	0	0	0	0	1	0	0	0	1	1	0	0	0	1	1	0	0	1	1	0	1	1	0	1	1	
0	0	0	0	0	1	0	1	0	1	1	1	0	0	0	0	1	1	0	0	0	0	0	1	0	0	0	1	0	
1	1	0	1	0	1	0	0	0	0	1	1	0	0	1	1	0	0	0	1	1	1	0	0	1	1	0	0	1	
1	1	0	0	0	1	1	1	0	1	1	1	0	1	0	1	0	1	1	1	1	1	1	1	0	1	1	0	0	
1	0	0	1	0	1	0	0	0	1	0	1	0	1	1	1	1	1	0	0	0	0	0	1	0	1	1	0	0	
1	0	0	0	1	1	1	1	0	1	1	1	0	0	1	0	0	1	1	0	0	1	0	0	0	1	1	1	0	
1	0	1	1	0	1	0	1	0	1	0	1	0	0	0	1	0	0	1	1	1	1	0	1	1	1	1	1	0	
1	0	1	0	1	0	0	1	1	1	0	1	1	1	1	0	1	0	1	0	0	1	1	0	0	1	1	1	0	
1	1	1	0	1	1	1	0	0	1	0	1	0	0	0	1	1	1	1	1	1	0	1	1	1	0	1	1	1	
0	0	0	0	0	1	0	1	1	1	1	0	0	1	1	0	1	1	0	0	1	0	0	1	1	0	1	1	1	
1	0	0	0	0	1	0	1	1	1	0	1	1	0	0	0	1	1	0	0	0	1	1	0	1	1	1	0	0	
0	1>	0	1	1	1	1	0	0>	0>	1>	1	0	1>	1	0	1	0	1	0	1	0	1>	0>	0	0	1>	1	1	
>	0	1>	1	1	0	1	1	1	0	1>	0	0	0>	1	0	1	1	0	1>	1	1	0	1>	0	1	1>	1>	0>	
0	0	1	1>	0>	1	0>	0	1	0	1	0	0	0	0	0	1	0	0	1	1>	0	1	1	1>	1	0	0	1	
1	0	1	0	0	1	1	0>	0	0	0	1	0>	0>	1>	1	0	0>	0	0>	1	1	0	0>	0	1	0	0	1	
0	1	1	1>	0	1	0>	0	0>	1	1>	0>	1>	1	0	1>	1>	1	0	0>	0	1>	0>	0	0	1	1>			
>	1	0	0	1	0	1>	0>	0	0>	0	1	1																	
>>	1	1	1	0	1	0	0	0	1>	0	0	0																	
>	1	0	1	1	1>	0	1	1	0	1	1	0	0>	1															
1	1	1	0	0>	0>	0	0>	0	1	1>	1	1	1	1	1	1>													
>>	1>	1	1	0>	0	1>	1	0	1	1	1	0	0	1>	0														
0	1	0	0	0>	0	0	1	0	0	1	1>																		

PREFACE

It is little over four years ago, when I started to work on the topic of cryogenic electronics. A lot has emerged since then, most of it covered in this dissertation. Before diving into the technical content, I would like to thank several people, without whom this work would not have been possible.

Of course there is Edoardo Charbon, together with Fabio Sebastiano, the founding fathers of the group. I remember well, shortly after my master defense, Edoardo asked me for a Ph.D. in this topic, which was completely novel for our group at that time. We basically started all the way from a white paper on the desk, I believe that we have made a gigantic step in those past years towards a real cryogenic electronic interface. At that time, we started our first collaborations with QuTech and later Intel Corp., learning in a hard way that co-operation is not always easy and straightforward, especially when many disciplines are involved. I would like to thank the numerous people that joined the group along the way, most notably Bishnu Patra, Rosario Incandela and Stefan Visser, with whom I shared a lot of the first cryogenic experiences. For those cryogenic measurements, I am grateful that we could learn so much from Enrico Prati and Giorgio Ferrari in Milan, where we measured our CMOS chips and FPGAs for the first time at 4 K. We kind of replicated and improved their setup to suit best our planned measurements, and those cryogenic probes, so-called dip-sticks, worked out pretty well. I spent many hours in our basement lab in TNW (F-16), doing numerous cool downs with various chips, FPGAs and other components. Of course, I also wasted a lot of time there, when something once again didn't work out as expected. Those hours paid off in many publications, all listed at the end of this dissertation, and a lot of joy when most of the stuff turned out to be working at such low temperatures. I am grateful to all secretaries (Minaksie Ramsoekh, Lidwina Tromp, Joyce van Velzen, Marja Plas) of our group, changing year by year, and I appreciate all the help I got from the technicians (Zu-Yao Chang, Lukasz Pakula, Hans van der Does, Remco Roeleveld, Siebe Visser) in both EWI, TNW and DEMO, for bonding my chips, for setting up the lab, making sure there is always helium available, for PCB related matters, and numerous other small things that I most likely forgot about. Recent advances towards the completion of this work led to measurements in a dilution fridge; a novel experience for me, but thanks to Andrea Corna a pleasant new adventure. Although it turned out to be one of the most difficult parts, with several things not working as expected compared to our liquid helium measurements, we learned a lot on the thermalization and other aspects of measuring in a fridge compared to a dip-stick.

Besides all cryogenic work, I remained a bit in the single-photon imaging group, and would like to thank all office mates for sharing a wonderful time with them (Augusto Carimatto, Esteban Venialgo and Ting Gong), and above all Claudio Bruschini, with whom I had many meetings and discussions, enjoying most of them, which resulted in some other papers. Although I was fond of spending time in both fields, it was not always easy to manage, getting me lost in between and doing nothing for a day (at least we had a nice view from the office).

I am grateful to my family, who always supported me during this four year long Ph.D. and the preceding five years of my studies at TU Delft. I hope I was not too much of

a burden, when still living at home. Cycling to Delft (25 km round trip) turned out to be an interesting ride over the past nine years, since so much has changed along the road. Nine years ago, I was still driving through farmlands and in between green houses, where complete suburbs have now been erected. Also in Delft, so much has changed, with the new train tunnel and station, and the amazing tram tracks on the TU Delft campus. It is sad somehow, taking away the good old stuff and replacing it with something that is hopefully better, but mostly isn't. Let's look at how this process turned out on the TU Delft: new coffee machines, worse tea, hot chocolate and probably worse coffee; a new caterer, less variation, higher prices and even more waste, because too lazy to wash dishes; new printers, we have to walk back and forth, because direct printing no longer works; a new purchase management system, less user friendly and using an obsolete sil-verlight interface; a new hour registration system, three clicks per day were not enough, so lets make the people have to click even more; a new EWI building for WI only, now we have doubled the cost and not solved anything. Extrapolating this trend, well, maybe better not to.

All of that brings me back to the topic of today, cryogenic electronics, which could find an application in many fields, such as deep-space, high-energy physics and astronomy. However, in this work, supported by the fact that our group became part of QuTech, is focussed on quantum systems and their quantum–classical control interface. We explored at depth the operation of many parts of such a system, vital for its operation at cryogenic temperatures, and, although this work is by no extend complete, it is a first step. I hope that you, the reader, will find (parts of) this work relevant and that it can be a guide for future cryogenic designers. Enjoy this cryogenic journey.

*Harald Homulle
Delft, April 2019*

CONTENTS

Preface	vii
Contents	viii
1 Introduction	1
I Discrete electronics	7
2 Cryogenic capacitors, resistors, power transistors & op-amps	9
2.1 Capacitors	10
2.2 Resistors	12
2.3 Power transistors	12
2.4 Operational amplifiers	14
3 Commercial & discrete cryogenic voltage regulators	15
3.1 Commercially available voltage regulators	15
3.2 Design of a discrete voltage regulator & measurement setup	17
3.3 Results	18
3.4 Summary	23
4 Crystal & MEMS oscillators for cryogenic clock generation	25
4.1 Implementation & measurement setup	26
4.2 Results	27
4.3 Summary	33
5 Cryogenic characterization of Altera & Xilinx 28 nm FPGAs	35
5.1 Comparison & measurement setup	36
5.2 Results	38
5.3 Summary	46
6 FPGA-based analog-to-digital conversion	49
6.1 System architecture	50
6.2 System calibration	53
6.3 Results	56
6.4 Summary	68
7 FPGA-based voltage stabilization	71
7.1 System architecture	71
7.2 Results	76
7.3 Summary	79

II	Integrated electronics	81
8	Cryogenic MOS & bipolar transistors in standard CMOS processes	83
8.1	Cryogenic behaviour of MOS transistors	85
8.2	MOS transistors in 40 & 160 nm	87
8.3	Weak inversion MOS & bipolar transistors	90
8.4	Summary	99
9	Basics for a cryogenic voltage reference	101
9.1	Bandgap cores	102
9.2	Voltage-reference circuits	105
9.3	Summary	109
10	Operation of digital circuits under cryogenic conditions	111
10.1	Inverter in 160 nm	112
10.2	Ring oscillators in 40 & 160 nm	115
10.3	Shift register in 40 nm	117
10.4	Summary	119
III	A quantum read-out platform	121
11	QuRO – Cryogenic quantum read-out	123
11.1	Spin qubits	123
11.2	Qubit reflectometry read-out	125
11.3	Amplifiers	126
11.4	Directional coupler	131
11.5	Summary	132
12	The read-out chain	133
12.1	Integration	133
12.2	Towards a dilution fridge	140
12.3	Summary	142
13	Conclusion	145
	Abbreviations	149
	Bibliography	151
	Publications	161
	CryoCMOS gallery	165
	CryoPCB gallery	166
	CryoProbe	167
	Summary	169
	Samenvatting	171
	About the author	173

1

INTRODUCTION

Our society is completely dependent on the internet and the availability of seemingly infinite computing power. We can easily search through the apparently endless world wide web, stream millions of videos and safely transfer money around the globe. Computing power roughly doubles every two years, still following Moore's law since its inception in 1965 [1]. However, Moore's law is soon predicted to become obsolete, as the growth in computing capacity stagnates [2]. Various problems, unsolvable today, require significantly more processing power than even the largest supercomputers can provide.

Quantum computers promise an exponential speed-up over conventional computing, thus providing the resources to solve complex problems such as climate models, to breach conventional security algorithms and to search more quickly in extremely large databases. Entanglement and superposition are the key quantum phenomena that empower these computers. Instead of bits being either 0 or 1, they can be in a superposition of these two states, denoted $|0\rangle$ and $|1\rangle$. The superposition can be envisioned as an arrow pointing to any position on a sphere, instead of only to $|0\rangle$ or $|1\rangle$, as shown in Figure 1.1 with $\alpha|0\rangle + \beta|1\rangle$. When quantum bits (qubits) are correlated to one another and no longer independent, they are said to be entangled. Acting on a single qubit influences all entangled qubits in the system, making powerful interactions possible. As a qubit can be 0 and 1 at the same time, it can represent two classical bits, i.e. n qubits encode 2^n classical bits. A million qubits can be compared to more classical bits than ever produced. Electrons, ions and photons are the most common elements to act as a quantum bit. They can be implemented in various technologies, such as transmons, quantum dots and nitrogen-vacancy centres [3]–[14].

Unfortunately, these quantum phenomena only become 'visible' at extremely low temperatures. Moreover, the qubits are very sensitive to disturbance and their coherence time is usually limited to several microseconds for the best qubits currently available [15]–[17]. Therefore, the quantum device has to be operated close to absolute zero, ideally below 100 mK, and quantum error-correction is required to track and repair qubit states. Both requirements impose severe constraints on the electronics employed in the quantum-classical interface. This interface controls the qubits, reads the quantum information and processes this data to track and correct errors. It should be able to produce very accurate signals for qubit operations and to read tiny signals from the quantum device. Furthermore, the error-correction cycle should be faster than the qubit coherence time to enable fault-tolerant quantum computing.




Figure 1.1: The Bloch sphere represents all possible quantum states as a superposition of $\alpha |0\rangle + \beta |1\rangle$.

Currently, this electronic system is implemented at room temperature, which is about 300°C warmer than the qubits. To interface these environments, interconnects are implemented between cryogenic and room temperature. The number of interconnects is limited, as heat injection into the qubits will disturb the quantum states. For now, this is not a problem per se, but millions of qubits are required for large-scale quantum operations and all of them need to be interfaced with electronics. To reduce heat injection, one solution is to operate qubits at higher temperatures [18], [19], but their performance is currently limited due to thermal noise. A second solution is to operate the electronics at lower temperatures [20], [21], [4]. The ultimate solution is to operate both quantum device and electronics at the same temperature, but this is not predicted to happen in the very near future. Another problem for these large-scale systems is the vast amount of data that needs to be processed in each error-correction cycle, demanding a local and highly parallel electronic infrastructure. Even when assuming that we can frequency-multiplex 100 qubits in a 1 GHz bandwidth and that we digitize their signals at 2 GSa/s with 10 bits of resolution, we already need to process 1 TB/s of data for a system consisting of 5000 qubits.

Therefore, we propose implementing a tiled cryogenic read-out and control interface. One such tile is schematically depicted in Figure 1.2 and only operates on a select number of qubits. The quantum device is situated at the lowest temperature level. To simplify the connection to the electronics at a higher temperature, we can reduce the number of interconnects between electronics and qubits by implementing (de)multiplexers. The main part of the electronics will be situated at liquid helium temperature, i.e. 4.2 K (we will refer to it as 4 K), as it is the best trade-off between available cooling power and electronic requirements. Even with the 1–2 W of cooling power at 4 K, the electronics already needs to be so highly efficient, in the order of a few mW/qubit for the complete read-out and control process, that the operation at even lower temperatures with exponentially lower cooling power will be extremely challenging. The tiny signals from the quantum device first need to be amplified, after which the quantum information can be read with the analog-to-digital converters (ADCs). The qubit control in contrast is realized with digital-to-analog converters (DACs). Digital processing




Figure 1.2: Quantum tile with electronics for qubit control and read-out.

can be implemented at higher temperatures and in a digital controller, such as a field-programmable gate array (FPGA). Using FPGAs is advantageous as they can be highly parallelized and reprogrammed at any time. This is also the place for the execution of the quantum algorithms (QEX) and the correction of qubit errors (QEC) [22].

Implementing electronics at temperatures as low as 4 K or -269°C is not trivial, as most electronic components are not supposed to be operated outside the industrial temperature range down to -55°C . Various technologies have been investigated to implement cryogenic electronics, such as GaAs high-electron-mobility transistors (HEMTs) [23], SiGe heterojunction bipolar transistors (HBTs) [24], rapid single flux quantum (RSFQ) devices [25], [26], and custom semiconductors [27]–[30]. However, the most reliable and common technology for fabricating integrated circuits is CMOS (complementary metal-oxide-semiconductor). It is the only technology that allows the integration of billions of transistors and at the same time ensures a low power consumption and an extremely high reliability. Although CMOS transistors have been shown to operate at temperatures as low as 100 mK [31], [32], [C12], little is known about the integration of large-scale systems, such as the proposed quantum–classical interface.

Therefore, we will thoroughly study the operation of electronics over a wide range of temperatures and ultimately realize an initial version of such a quantum–classical interface at 4 K. To do so, the remainder of this dissertation is split into three parts.

In Part I, the focus is on the characterization and design of circuits from commercially available components. We start by investigating basic elements: capacitors, resistors and power transistors in Chapter 2. These elements are required for almost any circuit and are used throughout the following chapters. As we are working towards an FPGA-based system, it is worthwhile to first study some auxiliary components at low temperatures. Voltage regulators to stabilize supply voltages are designed in Chapter 3, and oscillators

for the generation of a clock signal are investigated in Chapter 4. Finally, two FPGAs from Altera and Xilinx are characterized at cryogenic temperatures in Chapter 5. The observed changes can be related to those observed later on in Part II, as FPGAs are basically very-large-scale integrated circuits. To demonstrate the usability of FPGAs over 200°C below their normal temperature range, several firmware designs were developed and studied in Chapters 6 and 7. The most practical design is that of an analog-to-digital converter, requiring only the FPGA and some external resistors.

Since commercially available components are (usually) not optimized for operation at cryogenic temperatures, we strive for a quantum–classical interface completely designed in a CMOS process, with the goal of slowly replacing all discrete elements from Part I. Therefore, in Part II, we focus on the characterization and design of integrated circuits. We first study the behaviour of the basic element of any electronic circuit, namely the transistor, and characterize it from 300 K down to 4 K in Chapter 8. With the characterized transistors, we can build circuits that are optimized for operation at 4 K and we demonstrate this with the creation of both analog and digital building blocks. The creation of a voltage reference is described in Chapter 9, for the generation of stable cryogenic voltages. The behaviour of digital elements, such as inverters and oscillators, is investigated in Chapter 10. These results demonstrate that CMOS is a viable choice for the integration of the quantum–classical interface, and that it indeed allows to replace most discrete electronic components.

Both discrete and integrated electronics are combined in the final Part III. We realized a complete quantum read-out chain, with the components highlighted in purple in Figure 1.3. In order to read the tiny signals from the qubits, they are first amplified with two stages of amplification. This amplification is realized with a CMOS low-noise amplifier (LNA) followed by a discrete SiGe amplifier. The signal is then digitized in our FPGA-based ADC, and the data is locally processed. A cryogenic oscillator provides the FPGA's clock, and the supplies are stabilized by the voltage regulators. We first discuss the implementation of the system in Chapter 11, followed by results on the complete read-out chain in Chapter 12. Conclusions, suggestions for further improvements and an outlook into the future are drawn in Chapter 13.




Figure 1.3: The quantum-classical interface implemented at cryogenic temperatures, close to the quantum processor. Highlighted in green are the components of the read-out chain touched upon in this work.

I

DISCRETE ELECTRONICS



2

CRYOGENIC CAPACITORS, RESISTORS, POWER TRANSISTORS & OPERATIONAL AMPLIFIERS

In this first part, we focus on commercially available or off-the-shelf electronics. These elements are readily available from various distributors and allow for rapid prototyping of circuits at cryogenic temperatures. Because FPGAs are a very versatile solution to easily built custom logic, we are working towards a platform based on FPGA technology in Chapter 5, where we demonstrate the functionality of these devices. However, we first have to study auxiliary components required for such a system, shown in this and the following chapters. In this chapter, we consider basic passive and active elements, such as capacitors, resistors and power transistors.

Any electronic system requires a power distribution network, usually implemented with regulated voltages and decoupled supplies. The regulation of voltages is discussed in Chapter 3. To decouple, capacitors play a key role. They stabilize the voltage, even when the current through the electronic system or load fluctuates. Several studies show the effects of temperature on capacitors [33]–[37], [J6], revealing that some materials behave properly and some fail completely. Generally, capacitance goes down, whereas the effective series resistance (ESR) goes up. Luckily, there are some types, such as NP0/C0G, PPS and silicon, that are very stable over the complete temperature range down to 4 K and change by less than 2%. We present an overview of tested capacitors in Section 2.1.

Other components that are used in many systems, such as in the designed voltage regulator (Chapter 3), are resistors, power transistors and operational amplifiers (op-amps). Various resistive materials are available, and again not all behave properly at low temperatures. [36] reveals that metal-film resistors are among the most stable, with less than 1% spread. For completeness, we briefly feature resistors in Section 2.2. Power transistors are active instead of passive like the capacitors and resistors, and can be compared to the transistors in Chapter 8, with the exception of much higher ($\gg 100\times$) current levels. There are some curious effects occurring in these devices at low temperatures, as we show in Section 2.3. Finally, op-amps are required to create a high-gain feedback-loop. A brief literature overview of them is given in Section 2.4.

Parts of this chapter have been published in Review of Scientific Instruments [J4], [J6] and Cryogenics [J9].

2.1 CAPACITORS

Passive components form an important part of any printed circuit board (PCB) design, whether the PCB houses a commercial IC (such as an FPGA) or an application specific integrated circuit (ASIC) as the main embodiment. Especially capacitors, which are not only used for the decoupling networks of power supplies, but also for analog filters, are important. At cryogenic conditions, the material properties can differ significantly, altering the dielectric values and thus the resulting junction capacitance and/or resistance. As this is such an important part of any system, we acquired the frequency response of several capacitors in the range from 4.7 nF to 330 μ F that cover a wide range of dielectric materials. Several studies [33]–[37] already revealed that various dielectric types can be excluded, such as high-k dielectric constant materials (X5R or X7R).

An example of an ill-behaving capacitor is shown in Figure 2.1. This 4.7 μ F solid tantalum capacitor shows a significant reduction of its capacitance at especially higher frequencies, and a steady increase of its effective series resistance (ESR), indicating a reduced quality of the capacitor at lower temperatures.

A complete overview of tested capacitors can be found in Table 2.1, in which the best capacitors per value are highlighted. Capacitors were measured at 300 K and 4 K with a Bode-100 (Omicron Lab) impedance analyzer. Test results are reported for low frequencies (around 100 Hz). Since we covered such a wide capacitance range (4.7 nF to 330 μ F), it is impossible to rely on a single dielectric (due to size and availability). Below 1 μ F, one can resort to the use of NP0/C0G, acrylic film and silicon capacitors, all of which perform fairly stable over temperature, i.e. < 10% capacitance loss and stable or improving ESR at 4 K. We recommend silicon capacitors only in specific cases (cost-effective for RF and small area); we generally advise either NP0/C0G or acrylic film capacitors. For capacitances above 1 μ F, only tantalum proved to be relatively stable. Notably, tantalum polymer outperformed solid tantalum capacitors in all test cases. Still, the capacitance is reduced by 10–30% and ESR increases for some values (4.7 and 47 μ F). There are also tantalum capacitors (EPPL2 [38]) optimized for low temperatures, but they only outperformed standard tantalum above 77 K.




Figure 2.1: Example of an ill-behaving capacitor (TR3A475K010C1000) at cryogenic temperatures, (a) impedance, (b) capacitance and (c) effective series resistance versus frequency.

Table 2.1: The cryogenic performance of several capacitors. All values are measured at low frequencies (100 Hz range).

Specified capacitance [μF]	Temperature		300 K		4 K			
	Type	Part number	Capacitance [μF]	ESR [Ω]	Capacitance [μF]	ESR [Ω]	ΔC [%]	ΔESR [%]
0.0047	NP0/C0G	C0603C472K1GECAUTO	0.0066	0.1	0.0061	0.1	-8	0
0.047	Silicon	935.131.424.547	0.043	0.7	0.041	0.4	-3	-46
0.047	PPS	LDBAA2470JC5N0	0.047	0.2	0.049	0.1	5	-20
0.047	NP0/C0G	C0805C473J3G	0.049	0.3	0.051	0.1	5	-64
0.47	Tantalum solid	TAJR474K035RNJ	0.500	1.5	0.432	112.3	-13	7479
0.47	Tantalum polymer	TCJB474M063R0400	0.474	0.5	0.404	1.5	-14	202
0.47	PPS	SMDIC03470TB00JQ00	0.475	0.1	0.478	0.1	1	0
0.47	NP0/C0G	C2220C474J5GACTU	0.482	0.1	0.497	0.1	3	-30
0.47	Acrylic	FCA1206C474M-H3	0.513	0.4	0.427	0.1	-17	-90
1	Tantalum polymer	TCJP105M025R0500	0.95	0.6	0.76	6.0	-20	846
1	PPS	SMDIC04100TB00KQ00	1.01	0.1	1.02	0.1	1	0
1	Silicon	935.121.427.710	0.95	0.4	0.96	0.3	1	-42
4.7	X8L	C1206C475J8NACTU	4.75	6.0	0.32	189.0	-93	3050
4.7	Tantalum solid	TR3A475K010C1000	4.76	0.8	4.27	26.5	-10	3244
4.7	Tantalum polymer	T55A475M010C0200	4.51	0.2	3.95	0.5	-12	128
47	Tantalum solid	TR3A476K6R3C0800	48.1	0.9	21.8	742.8	-54	86606
47	Niobium oxide	NLJA476M006R1600	48.3	0.8	10.0	85.9	-79	10731
47	Tantalum polymer	T58A0476M010C0100	43.1	0.2	29.3	15.6	-32	6403
100	Tantalum solid	T495D107K010ATE050	102.8	0.2	91.9	1.8	-11	800
100	Tantalum solid	T495X107K016ATE080	100.3	0.2	91.8	1.0	-9	350
100	Tantalum polymer	T520A107M06ATE070	105.1	0.6	73.0	0.2	-31	-65
330	Tantalum solid	T495X337K010ATE035	331.2	0.2	265.9	1.4	-20	685
330	Tantalum polymer	T520C337M004ATE025	319.0	0.2	266.9	0.1	-16	-42

The frequency responses of the best performing capacitors (highlighted in Table 2.1) are shown in Figure 2.2 at both 300 K and 4 K. In (a, b), the impedance of only the higher valued caps deviates significantly from 300 K down to 4 K. The same effect is visible in the extracted capacitance plot (c, d), especially for the 47 μF cap.

Although most capacitors show a degradation at cryogenic temperatures, we believe that this set of characterized caps gives a valuable foundation for the selection of decoupling and filter capacitors used at low temperatures. We recommend the use of NP0/C0G and acrylic film below 1 μF and tantalum polymer above 1 μF .




Figure 2.2: Impedance, ideally $\frac{1}{\omega C}$, and extracted capacitance versus frequency at (a, c) 300 K and (b, d) 4 K for the highlighted capacitors in Table 2.1. The performance of especially the 4.7 and 47 μF capacitors deteriorates at 4 K.

2.2 RESISTORS

Although not as important as capacitors, resistors are commonly used for filters, protection circuits, termination etc. As for capacitors, several resistive materials have been studied in the past [36]. Metal film resistors are found to be most stable over temperature, with fairly limited change in resistance ($< 1\%$). Figure 2.3(a) shows the resistance of a Vishay 50 Ω (FC0603E50R0BST1) and a Panasonic 499 Ω SMD resistor (ERA-3AEB4990V) over temperature. As comparison, a thick film equivalent is shown in Figure 2.3(b).

2.3 POWER TRANSISTORS

In Chapter 8, we will study the behaviour of IC transistors in detail, here we show the properties of some discrete power transistors (current levels $\gg 100\times$ compared to integrated transistors). Since integrated MOS transistors are verified to work at 4 K, in contrast to bipolar transistors [J7], we selected a range of power MOSFETs as well.

In general, the technology or implementation of these transistors is not revealed by the manufacturers, so the changes observed at cryogenic temperatures are hard to attribute to design or technology choices. Twelve MOSFETs from APEC, Diodes Inc., Infineon, NXP, ON, TSM and Vishay were selected and immersed in liquid helium for cryo-




Figure 2.3: Relative resistance for different materials, (a) metal and (b) thick film, over temperature.

Table 2.2: Summary of the main MOSFET performance characteristics at 300 K and 4 K for several devices.

Temperature	300 K				4 K			
Device	V_{th} [V]	$V_{th(DS)}$ [V]	$R_{DS(on)}$ [Ω]	SS [mV/dec.]	V_{th} [V]	$V_{th(DS)}$ [V]	$R_{DS(on)}$ [Ω]	SS [mV/dec.]
BSH105	0.8	0	2.0	107	1.2	1.0	0.2	53
BSS806	0.8	0	1.6	101	1.2	0.3	1.0	34
FDN337	0.8	0	1.7	103	1.2	1.4	0.1	37
IRLML2402	1.2	0	3.2	122	1.7	0.4	18.2	58
IRLML6344	1.0	0	1.6	121	1.5	0.2	1.8	38
PMV40	0.7	0	1.6	116	1.4	0.3	1.0	34
SI2312	0.6	0	1.5	102	1.1	0.3	0.8	23
SI2336	0.6	0	1.5	101	1.0	0.4	0.8	31
SI2374	0.6	0	1.5	102	1.1	0.2	1.0	22
TSM2312	0.7	0	1.6	110	1.3	0.7	0.2	33
TSM2314	0.7	0	1.6	109	1.2	0.7	0.2	35
ZXMN2A01	1.0	0	2.3	120	1.5	0.4	4.3	29

genic testing. All transistors were tested in an SOT-23 package to avoid deviations caused by different packaging. An overview of the main performance characteristics of the various devices is presented in Table 2.2 at both operating temperatures. As expected, the majority of devices exhibit elevated current levels at cryogenic temperatures, thanks to an increase in carrier mobility and shown in Figure 2.4 for the TSM2314. Furthermore, the threshold voltage V_{th} is increased for all devices by roughly 0.5 V. More surprisingly, all devices exhibit a non-zero turn-on voltage on the drain–source junction, indicated with $V_{th(DS)}$. This effect has been shown in literature [39], [40], and is attributed to potential hills in the source–drain channel, which can be caused by zero gate overlap on the source and drain, or by increased oxide thickness on the source and drain overlaps. The subthreshold slope decreases by a factor of 3 to 4, less than expected from theory, but in line with literature [J8]. Finally, the kink effect is not present at cryogenic temperatures in any device, even at high drain–source voltage levels and currents, which might be expected from literature [J8].

Although devices are immersed in liquid helium, self-heating is non-negligible at the higher current levels in these power MOSFETs, and significant hysteresis is observed in




Figure 2.4: I_{DS} versus (a) V_{DS} and (b) V_{GS} for the TSM2314 at 4 K and 300 K, indicating the various performance measures from Table 2.2.

all devices. The results shown in Figure 2.4 are typical for a forward V_{DS} and V_{GS} sweep obtained with a short sweep time to limit self-heating effects.

2.4 OPERATIONAL AMPLIFIERS

Several works have addressed the operability of commercial operational amplifiers at cryogenic temperatures, either 77 or 4 K [41]–[46]. Characterization down to 77 K performed by [41] on the InterSil ICL7622, Motorola MC14575, and Texas Instruments TLC251 and TLC271 shows a general trend of increasing performance and power consumption peaking at roughly 120 K. Down to 77 K, all devices are still functional. [42] reports that the Texas Instruments OPA350 is fully functional at least down to 115 K. In [43], the first attempts to operate op-amps at 4 K are undertaken. A Microchip MIC860, and an Analog Devices AD8605 and AD8651 are shown to work, even at 4 K. However, different power supply settings are required for proper operation, and power consumption generally increases, accompanied by a slightly improved performance in terms of gain-bandwidth. The offset voltage of both ADs is fairly stable (± 1 mV), whereas the MIC860 shows a deviation of almost 30 mV from 300 K down to 4 K. [44] identifies the InterSil ICL7611 to be non-working at 4 K within its default supply voltage range, but confirms it to be working with higher supplies and self-heating to at least 10 K. [45] shows that the National Semiconductor LMC660 is working at 4 K, despite a 10 fold increase in its power consumption. Finally, in [46] both an Analog Devices OP27, and Texas Instruments TLC271 are tested. The OP27 works down to 77 K. Within a 4 K environment, the TLC271 needs to be heated to 130 K for proper operation.

3

COMMERCIAL & DISCRETE CRYOGENIC VOLTAGE REGULATORS

For large-scale electronic systems, a stable supply voltage is crucial for performance and reliability. Therefore, voltage regulators are available that cover a wide range of requirements, such as noise, stability and power consumption. They can be divided into two categories: low-dropout and switching regulators. The former features a high stability, low noise and a high power consumption, while the latter stresses a low power consumption at the expense of noise and stability. Unfortunately, both types of commercial regulators tend not to operate at extremely low temperatures, as we will show in Section 3.1. The best tested device operates down to 60 K; it becomes unstable at lower temperatures. The main reasons for them not to operate correctly are the use of bipolar devices (confirmed not to work in Chapter 8), and the use of thermal protection circuitries. Therefore, we require another solution to stabilize voltages at low temperature.

In the past, a few attempts have been made to implement cryogenic voltage regulators, but none were shown to operate at deep-cryogenic temperatures. [47] demonstrates that simple shunt regulators based on Zener diodes operate stably down to 100 K. Furthermore, low-dropout (LDO) regulators are shown to operate down to 77 K by [48].

Since in the previous chapter we showed that various commercially available components operate at cryogenic temperatures, we can build a low-dropout voltage regulator from them. We show the design methodology, requiring only resistors, power transistors and operational amplifiers, in Section 3.2, followed by its performance over temperature down to 4 K in Section 3.3. The devices exhibit excellent temperature stability and load regulation, and they feature a low power consumption. The main drawback is an increase in noise, which is unexpected as noise should go down with temperature. To conclude, we briefly discuss regulator performance and future improvements in Section 3.4.

3.1 COMMERCIALLY AVAILABLE VOLTAGE REGULATORS

Several commercially available voltage regulators have been tested in the temperature range from 4 K to 300 K and their output voltage relative to 300 K is plotted versus temperature in Figure 3.1(a). Measurements were done at the specified V_{IN} and the regula-

Parts of this chapter have been published in Review of Scientific Instruments [J6] and Cryogenics [J9].




Figure 3.1: (a) Output voltage and (b) consumed power of several DC-DC converters / voltage regulators over temperature. All tested regulators fail operation at temperatures roughly below 90 K when assuming a 2.5% margin on the voltage at 300 K. More complete characterization can be found in the supplementary material of [J6].

tors were loaded with 100 Ω . As can be seen, none of the regulators are stable in voltage over the complete temperature range. Most regulators simply switch off or clip (proportionally) to their input voltage. Switching off behaviour can be explained by the internal protection circuitry, which shuts down the regulator when a certain temperature or internal voltage limit has been reached. To concretize functionality, a 2.5% margin was allowed on the output voltage with respect to room temperature. The corresponding output voltage plots for all regulators are presented in the supplementary material of [J6]. The best performing regulator was found to be an LDO from Texas Instruments

(TPS7A7002), at least in terms of output voltage. The device is extremely stable down to 90 K, at which temperature it completely turns off.

Besides a stable output, the power consumption is extremely important at cryogenic temperatures, due to limited power budgets in any cryogenic system. The power consumed in the same set of regulators is shown in Figure 3.1(b). It is well known that switching voltage supplies outperform other regulator types in terms of power consumption, as is also shown here. The most stable converter (TPS7A7002) has a power consumption of roughly 60 mW, roughly an order of magnitude higher than most switching converters.

The best performing switching converter is the ISL80031 from InterSil, which is stable in voltage until roughly 60 K (although with a larger variation compared to the TPS7A7002), consuming less than 4 mW. While these regulators introduce more noise, this is in general not a severe complication for FPGA-based systems, that are digital in nature.

3.2 DESIGN OF A DISCRETE VOLTAGE REGULATOR & MEASUREMENT SETUP

A discrete solution, with components that properly operate at low temperatures, can be designed as an alternative to the non-functional commercial regulators. With a fairly limited number of components, namely resistors, capacitors, a transistor and an operational amplifier, a discrete LDO can be built. We implemented a basic LDO circuit on a PCB, such as the one shown in Figure 3.2, with the component values and voltages as presented in Table 3.1. The op-amp drives both inputs to equal each other through the feedback loop, thus making both terminals equal to V_{REF} . This reference voltage is currently provided from room temperature, although this is suboptimal due to the long supply leads (ground fluctuations). A better solution is to place a voltage reference at cryogenic temperatures, such as the one presented in Chapter 9. Thanks to the resistive divider R_1/R_2 , each LDO can be configured with a different output voltage V_{OUT} from a single reference. In our design, V_{OUT} was set to 1.0 V, the logic supply of a Xilinx Artix 7 FPGA (Chapter 5). At the input and output, we placed large decoupling caps of 47 μF , mainly for stability concerns. In a practical application, such as regulating the FPGA




Figure 3.2: Schematic view of the implemented discrete cryogenic LDO, the blue area is situated at 4 K, while the supplies are placed at room temperature. Component values are presented in Table 3.1.

supplies, a wider range of decoupling capacitors (100 nF, 1 μ F etc.) is required to smooth fast transients in the 100+ MHz range. Although the majority of commercially available op-amps operate over a wide range of supply voltages V_{DD} , it should exceed the required pass transistor's gate voltage, which depends on the output voltage and current requirements.

3

3.3 RESULTS

We have implemented and tested several op-amp variations for the LDO circuit presented in Figure 3.2. The PCB, shown in Figure 3.3(b), contains the various LDO circuits and MOSFETs for testing at cryogenic temperatures. Op-amps were selected to be purely MOS-based, as per their data sheets, and the following 5 models were chosen: Analog Devices AD8601, AD8605, Microchip MCP6001, MCP6291 and ON Semiconductor TLV271. Testing, the setup is shown in Figure 3.3(a), was done with a dip-stick, where the sample is immersed in liquid helium or helium vapours to reach temperatures between 4 K and 300 K (see Figures A.3 and A.4). DC characterization was performed with 2636B (Keithley) source measure units (SMUs), whereas spectral (noise) analysis was performed with an SR770 (Stanford Research) low-frequency spectrum analyzer in combination with B2962A (KeySight) low noise supplies. Power supply rejection was measured with a Bode-100 (Omicron Lab) vector network analyzer in combination with a J2120A (Picotest) power line injector. Ambient temperature was measured with a DT-670 (LakeShore Cryotronics) temperature diode in close proximity to the circuits. The circuit temperature, especially of the dropout transistor, will be slightly elevated due to self-heating.

3.3.1 INPUT VOLTAGE

As indicated in the MOSFET characteristics, a certain threshold voltage in the drain at cryogenic temperatures has to be surpassed before conduction occurs. This means V_{IN} must at least be larger than $V_{OUT} + V_{th(DS)}$ to reach a stable output. To demonstrate this limitation, V_{IN} was swept from 1.1 to 1.5 V and over temperature as shown in Fig-

Table 3.1: Component selection and supply voltages for the discrete cryogenic LDO.

Component	Value	Remark
Op-amp	AD8601, AD8605, MCP6001, MCP6291, TLV271	CMOS-only amplifiers
MOSFET	TSM2314	
R_1	8.2 k Ω	Metal film resistor
R_2	12 k Ω	Metal film resistor
C_1	47 μ F	Tantalum capacitor
C_2	47 μ F	Tantalum capacitor
C_{DD}	4.7 μ F	Tantalum capacitor
Z_L	100 Ω	Unless otherwise indicated
V_{REF}	0.6 V	
V_{IN}	1.5 V	Unless otherwise indicated
V_{OUT}	1.0 V	
V_{DD}	3.3 V	Unless otherwise indicated
V_{drop}	$V_{IN} - V_{OUT}$	




Figure 3.3: (a) Schematic of the test setup employed for LDO characterization. (b) PCB with the implemented LDOs and power MOSFETs for cryogenic testing.




Figure 3.4: LDO output voltage V_{OUT} versus (a) input voltage V_{IN} and (b) op-amp supply voltage V_{DD} . The exact temperature dependence varies with the chosen op-amp (a) AD8601, (b) TLV271, but a similar trend is observed in all variations. Dashed lines in (b) indicate a 0.2% margin on the output voltage at room temperature.

ure 3.4(a). The output voltage is stable for all input voltages (fixed load) down to roughly 21 K, at which temperature 1.1 V at the input is no longer sufficient to drive the output to 1 V. 1.2 V fails at 14 K and 1.3 V at 7 K, indicating a semi-linear temperature dependence. Hence, for the remainder of the results presented here, V_{IN} was fixed at 1.5 V.

3.3.2 SUPPLY VOLTAGE

Another important aspect is the supply voltage of the operational amplifier. It influences not only its power consumption, but also the maximum drive strength on the gate, and thus indirectly the maximum output voltage. As a low output voltage of 1.0 V is set, the latter does not impact LDO performance; however, for higher required output voltages, for example for chip I/Os at 2.5 V or 3.3 V, a higher op-amp supply voltage needs to be considered. The output characteristics versus supply voltage and temperature are plotted in Figure 3.4(b) for the implementation with the TLV271. Down to 150 K, the output voltage is extremely stable with respect to room temperature. At lower temperatures, the output is still well within the 0.2% deviation indicated by the dashed lines. This shows

that, although a performance deviation is noticeable, most likely in the op-amp, a stable output voltage can be reached over an extremely wide temperature range from 4 K to 300 K. For the remainder of this chapter, the supply voltage V_{DD} was set to 3.3 V. The output voltage plot relative to room temperature is shown for all configurations in Figure 3.5(a); all circuits are well within a 0.2% margin from their room temperature output at 4 K.

3

3.3.3 POWER CONSUMPTION

Both the dropout in the MOSFET and the op-amp contribute to the overall LDO power consumption. Especially at cryogenic temperatures, the power consumption of several op-amps increases dramatically, as shown in Figure 3.5(b). The MCP6291 starts to dissipate several tens of milliwatt, after it has reached the minimum at 40 K. Only the AD8601 and TLV271 have a reduced power requirement at cryogenic temperatures. The TLV shows a power reduction of over 90%, from 1.7 mW to 0.1 mW.

3.3.4 LOAD REGULATION

One of the most important aspects of the LDO is the load regulation. Ideally, the output voltage is independent of the load, even with fast transients. These transients are normally captured by a decoupling capacitor network, but the changes in load require a regulator for compensation. We tested the regulators with a DC load ranging from 0 to 300 mA, equivalent to an ohmic load changing from high-impedance down to 3.3 Ω . The regulated load curves are shown in Figure 3.6 for the various configurations; the MCP6001 fails to regulate non-negligible loads at 4 K. Furthermore, the room temperature regulation is generally more stable than at cryogenic temperatures, except for the MCP6291 and TLV271, which perform slightly better at 4 K. The MCP6291 is the most stable with 0.25 mV/A load stability at 4 K compared to 1.47 mV/A at 300 K.




Figure 3.5: LDO (a) relative output voltage V_{OUT} with respect to room temperature and (b) power consumption of the op-amp V_{DD} in the five tested configurations.




Figure 3.6: LDO load regulation at 300 K and 4 K for the different op-amps, the MCP6001 clearly ceases proper operation for non-negligible loads.




Figure 3.7: LDO power supply rejection ratio (a) for all tested devices at 4 K and 300 K and (b) for the TLV271 over a wide range of temperatures.

3.3.5 POWER SUPPLY REJECTION

Besides generating a stable output voltage, the LDO should be capable of rejecting any ripples on the input voltage, caused by either noise or load changes. The power supply rejection ratio (PSRR) of the LDO circuits was acquired over temperature, and is shown in Figure 3.7(a) for all tested devices at 300 K and 4 K. Although the PSRR improves at lower frequencies reaching cryogenic temperatures (for the AD8605, MCP6291, and TLV271), the high-frequency behaviour is dominated by a self-resonance point with negligible supply rejection. This self-resonance at roughly 500 kHz for the TLV271 (Figure 3.7(b)) only appears at the lowest temperatures, below 100 K. At low frequencies, the PSRR improves significantly from 60 to 76 dB.

3.3.6 NOISE

Finally, we measured the noise of the regulators over temperature. The noise density (PSD) is shown in Figure 3.8(a) for all regulators at both 4 K and 300 K, and is, in con-




Figure 3.8: LDO noise spectral density (a) for all regulators at 4 K and 300 K and (b) for the TLV271, showing increased noise levels at low temperatures.

trast to expectations, higher at low temperatures. Previous work mainly attributes this to an increase in flicker noise [43], which rises at lower temperatures. Moreover, both AD8601 and AD8605 show a significant number of spurs in the spectrum, whose magnitude grows at temperatures below 100 K.

The TLV271-based implementation is shown in more detail in Figure 3.8(b) at various temperatures. The noise at higher frequencies systematically increases over temperature, whereas the noise at lower frequencies initially decreases, it again increases at the lowest temperatures. This results in a considerably higher noise density at 4 K than at 300 K. Although the MCP6291 has the lowest noise at 4 K, we do not recommend it because of its extreme power consumption.

Table 3.2: Summary of the voltage regulators at 300 K and 4 K. Specifications are given for $V_{IN} = 1.5$ V, $V_{REF} = 0.6$ V, $V_{DD} = 3.3$ V and $Z_L = 100$ Ω .

Temperature	AD8601		AD8605		MCP6001		MCP6291		TLV271	
	300 K	4 K	300 K	4 K	300 K	4 K	300 K	4 K	300 K	4 K
V_{OUT} [V]	1.0099	1.0096	1.0098	1.0086	1.0129	1.0137	1.0097	1.0086	1.0095	1.0106
P_{DD} [mW]	2.37	0.66	3.33	4.46	0.41	5.78	3.09	39.2	1.67	0.10
Load regulation [mV/A]	0.11	3.53	0.08	1.59	0.30	81	1.47	0.25	4.24	0.81
PSD (1 kHz) [nV/Hz]	65	1143	38	191	69	1855	44	852	80	688
PSD (10 kHz) [nV/Hz]	56	1375	35	135	138	1023	34	71	62	302
PSRR (1 kHz) [dB]	65.4	53.4	59.5	75.1	59.5	27.1	38.2	69.2	60.7	76.4
PSRR (100 kHz) [dB]	45.7	4.0	47.0	51.5	45.7	25.9	23.9	56.9	44.8	30.8

3.4 SUMMARY

All implementations are compared in Table 3.2, although the MCP6001 and MCP6291 are not recommended, since they have an extremely high power consumption or do not even properly regulate a load at 4 K. The AD8605 consumes 4.4 mW at 4 K, over 40 $\times$ the consumption of the TLV271 of only 0.1 mW. The load regulation is best in the AD8605 at room temperature, but regulation is better in the TLV271 at 4 K. Finally, in terms of noise, all levels increase, some by over a decade. The TLV271 performs best considering its low power consumption, and does not show the spurs in the noise spectrum that both ADs exhibit. Other important aspects not addressed in this manuscript are the LDOs transient behaviour and sample-to-sample variation. Limited tests on a select number of samples with the TLV271 used to power a cryogenic FPGA revealed no significant spread between boards.

Cryogenic voltage regulators offer a solution to distribute voltages at low temperatures, and to avoid large and fluctuating voltage drops in the cabling towards cryogenic environments. For future designs, the main challenges are the high input voltage required to overcome the drain–source threshold $V_{th(DS)}$ and the high output noise levels at low temperatures. A high V_{IN} leads to a higher dropout voltage than ideally required and thus to a relatively large power consumption; other discrete MOSFETs with a lower $V_{th(DS)}$ may overcome this limitation. To reduce the noise, an op-amp has to be found or designed that features lower noise levels at low temperatures. Nonetheless, to the best of our knowledge, this is the first time voltage regulators are implemented and demonstrated at deep-cryogenic temperatures as low as 4 K, offering a reliable solution for stable low temperature electronics.

4

CRYSTAL & MEMS OSCILLATORS FOR CRYOGENIC CLOCK GENERATION

Now that we have obtained stable voltage supplies for our FPGA-based system in the previous chapter, we require just one other component to make it ‘run’: a clock. A clock signal can be provided from room temperature at the cost of one or two high-frequency interconnects for single-ended or differential signalling, respectively. A more elegant solution is the integration of a clock source in the cryogenic environment, eliminating the interconnects with room temperature.

Although several works have addressed the issue of generating microwave frequency signals at low temperatures, the generation of clock signals in the range of interest, 50–100 MHz, has not been covered. For example, in the field of atomic clocks, crystals at low temperatures are generally used for their better and longer stability. [49] shows the implementation of such a system, capable of generating 5 MHz, 100 MHz and 10 GHz signals at 6 K with a cryogenic sapphire oscillator, but still requiring a 100 MHz reference to be provided from room temperature. [50] designs a voltage-controlled oscillator (VCO) in a deep-submicron SOI CMOS process for the tuning of a single-electron transistor (SET) integrated on the same die. The VCO operates stably down to 1.1 K and can be tuned from 300 kHz to 1.8 GHz. In [J5], a microwave signal generator is fabricated based on a transformer oscillator tunable around 6 GHz, and is operated down to 4 K. [51] demonstrates a 30 GHz VCO, implemented in a SiGe BiCMOS process, to operate properly with a frequency shift of only 3% at 4 K. The most relevant research shows that clock generators for FPGAs (20 and 50 MHz) do not work below roughly 120 K [52].

We tested several devices, capable of providing a single-ended clock signal to a cryogenic FPGA, that generate 50 or 100 MHz. Several commercial oscillators, based on crystal or MEMS technology, were characterized over the temperature range from 300 K down to 4 K. Eight devices from eight manufacturers were selected, as we discuss in Section 4.1, together with the measurement setup. The characterization results are presented in Section 4.2, followed by a brief recap in Section 4.3.

Parts of this chapter have been published in the IEEE International Frequency Control Symposium [C14].

4.1 IMPLEMENTATION & MEASUREMENT SETUP

For the characterization of commercial clock generators, a selection of eight devices that generate 50 or 100 MHz was made. All of them are suitable for our cryogenic FPGA platform (Chapter 5). A comprehensive list of the devices under test is given in Table 4.1, it comprises two MEMS-based and six crystal-based oscillators.

All devices were placed on a PCB with standard FR-4 dielectric, together with a single 0.1 μF decoupling capacitor (as recommended on most datasheets) per oscillator. This capacitor has an NP0 dielectric as it is one of the most stable capacitive materials over the temperature range of interest. Separate SMP connectors were mounted at the output of each device, and coaxially connected to either a WaveMaster 813Zi-B (LeCroy) oscilloscope or an FSW43 (Rohde & Schwartz) spectrum analyzer for, respectively, time domain and phase noise measurements. The setup is sketched in Figure 4.1(a), and an actual PCB is shown in (b), mounted on our dip-stick (see Figures A.3 and A.4). All devices were powered independently with a B2962A (KeySight) low-noise power source.

Although these devices should ideally be CMOS loaded, i.e. high-impedance with 15 pF maximum capacitance, their outputs were AC coupled to a 50 Ω load in the oscilloscope or spectrum analyzer to avoid reflections in the coaxial interconnects. This poten-

Table 4.1: Summary of tested commercial clock generators at low temperatures featuring both MEMS and crystals at 50 and 100 MHz.

	Manufacturer	Type	Frequency [Mhz]	Voltage [V]
ASFLMPC-50.000MHZ-LR-T	Abracon	MEMS	50	2.2–3.6
S53305-50.000-X	Aker	Crystal	50	3.3
SG5032CAN	Epson	Crystal	50	1.6–3.6
QX533A50.00000 B15M	Qantek	Crystal	50	3.3
FXO-HC536R-100	Fox	Crystal	100	3.3
LFSPXO009438	IQD	Crystal	100	3.3
501JCA100M000CAG	Silicon Labs	MEMS	100	1.7–3.6
7W-100.000MBB-T	TXC	Crystal	100	1.8–5.0




Figure 4.1: (a) Schematic of the test setup employed for oscillator characterization of cryogenic cooled devices. (b) PCB showing four oscillators coupled to individual SMP connectors for cryogenic testing.

tially limits their performance, compared to datasheet specifications, but the main interest of this study is the operability and performance deviation between deep-cryogenic and room temperature. The PCB was immersed in liquid helium or helium vapours to reach a temperature between 4 K and 300 K.

4.2 RESULTS

Initial testing was done with an oscilloscope to confirm operability at cryogenic temperatures. After 10 to 20 cooling cycles, the devices were tested a second time with the spectrum analyzer to confirm operability and assess phase noise performance at cryogenic temperatures. Although the MEMS-based devices were initially working with relatively high stability, even at 4 K, their performance degraded significantly, even at room temperature, after several cooling cycles. The crystal-based oscillators on the contrary are more robust and do not suffer even after several cycles.

4

4.2.1 MEMS-BASED OSCILLATORS

Since the MEMS devices are apparently ageing rather fast at cryogenic temperatures, we will highlight both initial results and those achieved after several cooling cycles. In Figure 4.2, the oscillation frequency is shown as a function of supply voltage for the two




Figure 4.2: (a, c) oscillation frequency versus supply voltage in the first cool down cycle. (b, d) same after 10 to 20 cooling cycles.

MEMS devices in the first cooling cycle and after several cycles. The first MEMS device, the Abracon (50 MHz), operates at first glance (Figure 4.2(a)) over the complete temperature range, but with a significant change in frequency for the 75–150 K range. The frequency drops by 3 MHz, but it recovers around 25 K to 50.1 MHz. The power consumption at 3.3 V has increased from 85.3 mW at 300 K to 142 mW at 4 K, about 66% higher. The period jitter doubles from roughly 5 to almost 10 ps. After several cooling cycles (b), the device can no longer operate in the desired voltage range and a higher voltage is required to start-up. At lower temperatures, the device no longer operates despite the higher supply voltage.

The Silicon Labs device, with a nominal frequency of 100 MHz, completely fails at 150 K, but surprisingly recovers at around 25 K. It is the most complex from all tested devices (according to available information), and includes an on-chip low-dropout voltage regulator (LDO), a frequency-locked loop (FLL), the MEMS resonator and a voltage-controlled oscillator (VCO). The circuit thus recovers around 25 K, with a slightly higher oscillation frequency at 4 K. The period jitter at 4 K is close to the result obtained at 300 K, around 5 ps. Although a stable clock signal was generated over the complete voltage range during the first cooling cycle (c), the usable voltage range reduced after several cycles (d). The power consumption at 3.3 V is slightly higher at 4 K, 37.3 mW versus 33.6 mW.

4.2.2 CRYSTAL-BASED OSCILLATORS IN THE TIME DOMAIN

The crystal-based oscillators do not experience the ageing, apparent for MEMS devices; even after several cooling cycles, their performance is stable. Therefore, we only report their performance after several cooling cycles. The oscillation frequency (Figure 4.2) reduces over temperature by roughly 0.1% (300 K down to 4 K), in contrast to the MEMS devices that both show a higher operating frequency at 4 K. All 50 MHz crystal oscillators are functional over the complete temperature range. The 100 MHz devices are stable, except for the Fox. The Fox required a higher supply voltage after several cooling cycles. Furthermore, the frequency deviated by almost 3 MHz (at 5 V) from 100 MHz at 300 K. Therefore, we only focus on the five functional devices, for which the obtained specifications at both 4 K and 300 K are summarized in Table 4.2.

The three 50 MHz devices drop in frequency by roughly 0.1%, and show very similar frequency degradation. Their power consumption (Figure 4.4(a–c)) shows an overall increasing trend at lower temperatures, e.g. 25% for the Qantek device. The RMS phase jitter, the phase noise integrated between 100 Hz and 10 MHz offset frequencies from the main carrier, increases at lower temperatures as shown in Figure 4.5. According to literature this increase can be mainly attributed to an increase in flicker noise at cryogenic temperatures [53]. The Aker increases its phase jitter from 520 fs to over 14 ps; the Qantek on the contrary shows a significantly lower increase (3×) from 320 to 910 fs. The driving capabilities of the devices can be derived from the output amplitude over the 50 Ω load. The amplitude of the main carrier (not considering the higher frequency tones visible in the square wave) shows little deviation in the Epson and Qantek devices (Figure 4.6(a)). However, the Aker decreases its amplitude with over 30%.

The two 100 MHz oscillators, the IQD and TXC in Figure 4.3(e,f), show somewhat opposite frequency deviation from RT in the 75–150 K range after which both devices




Figure 4.3: Oscillation frequency versus supply voltage for the six tested crystal-based oscillators. All are fully operational over the complete temperature range with the Fox as exception.

converge to a similar frequency of 99.88 MHz. Both devices consume less power, -14% for the IQD and -5% for the TXC (Figure 4.4(d, e)), whereas the 50 MHz devices consume more power at 4 K. The phase jitter increase in the TXC is the lowest of all tested devices (Figure 4.5(e)), from 350 to 620 fs. The IQD, although best at room temperature with 240 fs phase jitter, shows an increase to 2.3 ps at 4 K. Finally, the amplitude of the TXC (Figure 4.6(b)) shows a particular trend; it first drops significantly in the 75–150 K range to less than 250 mV, after which it recovers around 25 K to its final value around 500 mV.




Figure 4.4: Power consumption of the five functional crystal oscillators at various temperatures. The 50 MHz devices show an increasing trend, the 100 MHz show lower power consumption at 4 K.

4.2.3 CRYSTAL-BASED OSCILLATORS IN THE FREQUENCY DOMAIN

The phase noise of the five tested devices was acquired over an offset range from 10 Hz to 10 MHz from the main carrier (either 50 or 100 MHz). As the phase jitter, shown in the previous section, increases for all devices at cryogenic temperatures, the phase noise has to be elevated at those temperatures as well. The phase noise of the two devices with lowest jitter, i.e. the 50 MHz Quantek and the 100 MHz TXC, is shown in Figure 4.7 for 300 K and 4 K. The phase noise is overall higher at cryogenic temperatures, but increases mostly in the lower offset frequency range due to the increase of flicker noise at lower temperatures. On the higher offset frequencies, the phase noise at 4 K is close to the result obtained at room temperature; however, even lower phase noise would be expected from the decrease of thermal noise at lower temperatures.

4.2.4 CRYSTAL-BASED OSCILLATOR COUPLED TO AN FPGA

To demonstrate the suitability of implementing an oscillator directly at cryogenic temperatures, the TXC crystal was mounted on an FPGA-board for cryogenic testing. The board, hosting an Artix 7 FPGA from Xilinx, was specifically designed to operate at cryogenic temperatures (Chapter 5). The output of the TXC, supplied with the same voltage




Figure 4.5: Phase jitter in picoseconds derived from the integrated phase noise over the offset frequency range from 100 Hz to 10 MHz of the five functional crystal oscillators at various temperatures.





Figure 4.6: Carrier amplitude for the 50 MHz Quantek and 100 MHz TXC oscillators.

(2.5 V) used for the IO banks of the FPGA, was routed towards a clock capable input pin of the FPGA. The signal was multiplied by an FPGA clock manager by a factor of 12 to 1.2 GHz and divided to obtain a broad set of useful clock frequencies from 25 to 400 MHz. Although the frequency still shifts by 0.1% at 4 K, the clock managers were able to lock properly, and to generate the requested output frequencies. The period jit-



4

Figure 4.7: Phase noise over 10 Hz to 10 MHz offset from the main carrier (a) 50 MHz Quantek and (b) 100 MHz TXC at 300 K and 4 K.

Table 4.2: Obtained specifications for the 5 functional crystal devices. The phase noise (PN) is reported at 100 Hz, 10 kHz and 1 MHz offsets from the main carrier, either 50 or 100 MHz.

Temperature	Aker S53305-50.000-X		Epson SG5032CAN		Quantek QX533A50.0B15M		IQD LFSPXO009438		TXC 7W-100.0MBB-T	
	300 K	4 K	300 K	4 K	300 K	4 K	300 K	4 K	300 K	4 K
Frequency [MHz]	49.999	49.940	50.000	49.944	49.999	49.948	99.998	99.882	99.999	99.884
Phase jitter _{RMS} [ps]	0.52	14.06	0.32	2.52	0.32	0.91	0.24	2.25	0.35	0.62
Power [mW]	18.3	20.6	14.5	15.4	11.0	13.8	28.0	24.0	22.5	21.3
Carrier [mV]	515	351	591	617	527	562	617	638	620	524
PN @ 100 Hz [dBc/Hz]	-114	-89	-119	-82	-113	-90	-110	-81	-104	-99
PN @ 10 kHz [dBc/Hz]	-136	-107	-142	-133	-140	-140	-136	-113	-131	-122
PN @ 1 MHz [dBc/Hz]	-147	-117	-153	-146	-152	-147	-146	-135	-143	-142




Figure 4.8: Period jitter of the FPGA clock signal derived from the TXC crystal operating at 300 K and 4 K. The dashed line shows the jitter of the crystal routed directly through the FPGA fabric without passing through an PLL or MMCM clock manager.

ter, as obtained from the FPGA output, is shown in Figure 4.8 for a clock signal merely passing through the FPGA fabric, or as obtained after routing through either a phase-locked loop (PLL) or mixed-mode clock manager (MMCM). A stable output, even with

reduced jitter, could be generated at cryogenic temperatures. The jitter is dominated by the FPGA, i.e. IO jitter, routing jitter and the jitter of the PLL or MMCM itself.

4.3 SUMMARY

Eight commercial crystal and MEMS oscillators were characterized at deep-cryogenic temperatures. The MEMS devices were not able to sustain several cooling cycles, whereas the majority of crystal devices were fully functional without degradation after several cooling cycles. The crystal devices were fairly stable over the complete temperature range (4–300 K), and could be operated at their nominal supply voltage. The tested devices showed the same trend of decreasing frequency with cooling of about 0.1%, whereas their power consumption was generally higher, combined with a larger signal swing at low temperatures. The phase noise and corresponding phase jitter were elevated, mainly due to an increase in flicker noise. The best performing oscillator is the TXC with 620 fs phase jitter at 4 K, compared to 350 fs at room temperature. These results show that standard commercial oscillators, at least crystal-based, can be operated with fairly high stability over a significantly wider temperature range compared to the standard industrial temperature range down to -55°C . The oscillators enable frequency generation at low temperatures, mitigating the need to route high frequency signals into a cryogenic environment from room temperature, and can well complement our cryogenic FPGA system, discussed in detail in the next chapter.

5

CRYOGENIC CHARACTERIZATION OF ALTERA & XILINX 28 NM FPGAs

With the functional components from the previous chapters, we are able to implement an FPGA-based system that operates even at deep-cryogenic temperatures. We can properly decouple its supplies (Chapter 2), regulate its voltages (Chapter 3) and provide it with a stable clock (Chapter 4), thus enabling us to characterize the performance of the FPGAs themselves.

Several FPGAs have already been tested at cryogenic temperatures, focussing on applications such as deep-space exploration and cryogenic digital data links for physics experiments. So far, Altera (Stratix II), Microsemi (IGLOO series) and Xilinx FPGAs (Virtex 2, Virtex 5, Spartan 3) have been shown to be functional down to 77 K [52], [54], [55]. However, for the quantum controller, we require electronics to operate at even lower temperatures. Only a few devices from Xilinx have been demonstrated below 77 K, namely a Spartan 3 and an Artix 7 operating at 4 K [20], [56], [J4].

We once again stress the significance of FPGAs in the quantum–classical interface, as a shorter feedback time can significantly improve the performance of qubit systems. [57]–[61] are all using FPGAs in their measurements setups to benefit from parallelism (faster decision times), smaller system size and better integration. For example, [61] shows increased performance when using FPGA systems with a feedback calculation time of 200 ns. Unfortunately, the propagation through the cabling in and out of the cryogenic environment adds an additional 100 ns to the overall feedback time, limiting the benefits. Thus, placing the electronics physically closer to the qubits will reduce latencies and enhance qubit fidelity. Therefore, we explore the behaviour of two commercially available FPGAs, a Cyclone V from Altera and an Artix 7 from Xilinx, about 200°C below their intended temperature range.

In Section 5.1, we first compare the two FPGAs in terms of specifications and clarify the selection of just these two. As both FPGAs are fabricated in a similar 28 nm CMOS process, we expect similar deviations at lower temperatures. Although this is partially true, the Cyclone V fails operation below 30 K, whereas the Artix 7 is fully functional even at 4 K. These results are discussed in depth in Section 5.2. We conclude with some suggestions for further work in Section 5.3.

Parts of this chapter have been published in Review of Scientific Instruments [J4] and the IEEE International Conference on Field-Programmable Technology [C13].

5.1 COMPARISON & MEASUREMENT SETUP

Two 28 nm FPGAs were selected, as highlighted in Table 5.1, namely an Altera Cyclone V and a Xilinx Artix 7, both available in a 17×17 mm BGA package with 256 pins. The largest available models in terms of logic resources in this package were selected, delivering a Cyclone V 5CEBA4F17I7 with roughly half the available resources compared to the Artix 7 XC7A100T. Both FPGAs are industrial grade, which guarantees a functional device down to −40°C, still decades higher than the intended operating temperature.

Although both FPGAs are fabricated in a 28 nm process from TSMC, Altera uses the LP process flavor with a conventional polysilicon gate and embedded SiGe implants for increased PMOS performance. Xilinx on the other hand selected the HPL process flavor with a high-k metal gate. In theory, the HPL process delivers a higher performance at the same operating voltage thanks to the metal gate, which would result in lower dynamic power consumption, however the pros and cons of both processes are more complex for a simple comparison. In the end, the performance is mainly determined by the architecture of the two FPGAs, which is fundamentally different for Altera and Xilinx.

Both FPGAs were placed on dedicated PCBs, together with the before mentioned low-dropout voltage regulators, a crystal oscillator, decoupling networks and the necessary interface connectors. For the decoupling, silicon capacitors were selected for 0.1 and 1 μ F, whereas tantalum was used for all higher values up to 330 μ F (Chapter 2). Both PCBs are shown in Figure 5.1 and were placed in liquid helium or helium vapours to reach cryogenic temperatures as low as 4 K (see Figures A.3 and A.4 for a drawing of the setup employing a dip-stick for cryogenic measurements). The test setup is schematically drawn in Figure 5.2, showing the cooled PCB with the FPGA and its interface connections. For operating both FPGAs, the logic supply VCC_INT was set to 1 V (Artix 7) or 1.1 V (Cyclone V), the auxiliary/PLL supply VCC_AUX to 1.8 V (Artix 7) or 2.5 V (Cyclone V) and the IO banks VCC_IO to 2.5 V.

Table 5.1: Comparison of Altera Cyclone V and Xilinx Artix 7 FPGAs tested at cryogenic temperatures in terms of manufacturer specifications, available device resources and operating voltages.

		Altera	Xilinx
Specifications	Family	Cyclone V	Artix 7
	Model	5CEBA4F17I7	XC7A100T
	Package	BGA 256 pins	BGA 256 pins
	Size [mm]	17×17	17×17
	Technology [nm]	28	28
	Process	TSMC 28LP	TSMC 28HPL
	Gate material	Polysilicon	High-k metal
Resources	ALMs/Slices	18,480	15,850
	Registers	73,920	126,800
	LUTs	18,480	63,400
	Carries	36,960	63,400
	PLL + MMCM	4 + 0	6 + 6
	Temperature diode	-	1
Supply	VCC_INT [V]	1.1	1
	VCC_AUX [V]	2.5	1.8
	VCC_IO [V]	2.5	2.5




Figure 5.1: Altera Cyclone V and Xilinx Artix 7 on two dedicated PCBs for cryogenic testing.




Figure 5.2: Test setup with the FPGA, oscillators and LDOs placed on a dedicated PCB at cryogenic temperatures.

Both FPGAs were powered with the onboard LDOs, whereas LDO voltages were supplied from room temperature with Keithley 2636B source measure units (SMUs). Programming was performed through the JTAG interface, either with a XUP-USB-JTAG controller (Digilent for Xilinx) or an USB-blaster II (Altera). For stable programming at cryogenic temperatures, the programming speed had to be reduced to 3 MHz. Communication with the FPGAs was handled with an UART-to-USB converter (FTDI). For the majority of experiments, a 100 MHz LVDS clock was supplied from room temperature with an LMK61 (Texas Instruments) low jitter clock generator. In experiments requiring a variable clock frequency, an LVDS clock signal was generated with a CG635 (Stanford Research). Signals from the FPGA were captured and measured with a WaveMaster 813Zi-B (LeCroy).

Similar tests were executed on both FPGAs, focussing on the combinatorial logic, i.e. the look-up tables (LUTs) and fast adders (carry chains); the clock managers, i.e. phased-locked loops (PLLs) and the mixed-mode clock manager (MMCM); the slice registers (flip-flops) and the IOs. In principle the same HDL was used, except for the primitives

specific to either Altera or Xilinx. All tests were executed on three Altera and three Xilinx FPGAs, revealing no significant differences between devices.

5.2 RESULTS

5.2.1 IDLE CONDITION

Both FPGAs were cooled down and the idle power consumption, without anything programmed into the fabric, was extracted in Figure 5.3. Device currents were measured at room temperature with the Keithley SMUs and multiplied with the LDO output voltages of 1, 1.8 or 2.5 V. The Cyclone V was not able to boot-up below roughly 120 K, but once switched on, it could be cooled down further to around 30 K. At 30 K, there is a steep increase in the auxiliary supply current and most likely some internal over-current protection circuit disables the FPGA, as can be seen in Figure 5.3(b). Below 30 K the FPGA loses the programmed state and can not be reprogrammed, in contrast to the Artix 7, which is capable of booting up and being programmed even at 4 K. However, an even steeper increase in the auxiliary supply of the Artix 7 is noticeable, which significantly increases power consumption at cryogenic temperatures. The most likely cause of the increased auxiliary power consumption is the analog biasing, for the PLLs and other analog circuitries, operating far out of their normal intended operating range.

5.2.2 CLOCK MANAGERS

The clock managers, PLL and MMCM, form a crucial part in almost any HDL program. We implemented different configurations in both FPGAs, multiplying the 100 MHz clock with a factor 12 to 1.2 GHz and dividing from there to derive clock frequencies in the range 25 to 400 MHz, covering the majority of use-cases. For each configuration, a separate programming file was generated and sequentially programmed into the fabric, implementing a single LVDS input buffer, the clock manager and a single-ended output pin.




Figure 5.3: Idle power consumption, without anything programmed into the FPGA fabric, for the three different supplies shows a significant increase in auxiliary power consumption, whereas the internal logic and IO supplies remain more stable.




Figure 5.4: Clock managers (PLL and MMCM) jitter and power consumption. As a reference, the jitter of a clock passing through the FPGA fabric without PLL or MMCM is shown. In the Artix 7, the jitter is reduced by almost 50%, whereas the Cyclone V PLL is not operating properly any more at 30 K.

The output jitter, mainly limited by the jitter of the single-ended IO into a $50\ \Omega$ load, is plotted for various output frequencies in Figure 5.4. Most interestingly, the PLL in the Cyclone V doesn't properly divide the clock any more at 30 K, as shown Figure 5.4(b). In the Artix 7, the jitter at 30 K is reduced by almost 50% compared to 300 K for both PLL and MMCM. Finally, the auxiliary power consumption is significantly larger at both 30 K and 4 K compared to that at room temperature (Figure 5.4(c)), but it is dominated by the increase in idle power consumption as shown before in Figure 5.3(a).

5.2.3 COMBINATORIAL LOGIC

Together with the registers, the combinatorial logic forms one of the main building blocks for any firmware implemented on the FPGA and thus occupies most of the FPGA fabric, implemented in slices (Xilinx) or logic array blocks (LABs, Altera). Xilinx's architecture houses four look-up tables (LUTs) and four carry elements in each slice. Altera's architecture on the contrary places 10 adaptive logic modules (ALMs), each with one LUT and two carry elements, in a single LAB. The slices or LABs are placed in columns, in which the carry has a dedicated routing path between elements for optimal speed of the





Figure 5.5: Cell delay for LUTs, implemented both as inverters and buffers, and carry elements over the FPGA array for each of the 64, respectively 42, full columns in the Artix 7 and Cyclone V.

5


adders. To characterize the performance, oscillators with 779 LUTs or carry elements were placed in each of the 64, respectively 42, full columns in the Artix 7 and Cyclone V. The oscillators span almost the complete column to cover the majority of logic blocks in the FPGA. For each oscillator, a separate programming file was generated and programmed into the FPGA, consisting of the oscillator, constrained to one of the columns, with an enable that is activated after a count down of the 100 MHz clock. From each oscillator, the unit cell delay for both the LUTs and carries is extracted, as shown in Figure 5.5.

The cell delay of the carry elements in the Cyclone V, 8.4 ± 0.1 ps at 300 K, reduces to 7.5 ± 0.1 ps at 30 K. This cell propagation is unexpectedly fast compared to the carry delay in the Artix 7, 20 ± 0.3 ps at 300 K and 19.5 ± 0.3 ps at 4 K and would potentially allow the implementation of higher resolution time-to-digital converters [62]. Furthermore, the delay in the Cyclone V can be seen to reduce more (almost 11%), compared to the speed-up achieved in the Artix 7 (3%), this phenomenon can be explained by the higher supply voltage of the Cyclone V of 1.1 V compared to 1 V and is clearly shown in Figure 5.6. In both FPGAs, the logic at cryogenic temperatures is only faster at higher voltages, as we will show later in Chapter 10 for other integrated oscillators. Due to the increase in transistor threshold voltage at cryogenic temperatures, the delay is significantly longer at lower supply voltages [63]. As the two 28 nm processes are similar, the behaviour with respect to the supply voltage can be expected to be similar as well. In both devices, the logic operates equally fast at 300 K and 30 K at a supply voltage of around 1.02 V, which is close to the supply voltage of the Artix 7.

The LUTs experience a similar speed-up as the carry elements, summarized in Table 5.2 for the LUTs and carry elements in both FPGAs operating at nominal supply voltage. The LUTs in the Artix 7 configured as buffers are faster than the LUTs configured as inverters, in the Cyclone V, the opposite is true. The LUTs in the even columns of the Artix 7 are systematically slower than in the odd columns, an effect that is not visible in the Cyclone V and seems unrelated to the slice type. However, in the Cyclone V overall more spread is present in the average cell delay for LUTs compared to that in the Artix 7,



(a) Artix 7




(b) Cyclone V


Figure 5.6: Cell delay for LUTs and carry elements versus logic supply voltage V_{CC_INT} for a nominal oscillator from Figure 5.5.

5

especially at cryogenic temperatures. Overall, although changes are expected and clearly visible, changes in the order of 10% are accounted for in the compilers to compensate for process spread, voltage and temperature variations and would therefore not affect operation in the majority of HDL programs.



(a) Artix 7



(b) Cyclone V

Figure 5.7: Power consumption of the core logic P_{INT} for the oscillators from Figure 5.6.

Table 5.2: Summary of the combinatorial logic performance. Average cell delay, standard deviation (SD) over all columns and delay decrease are summarized for the LUTs and carries in the Cyclone V and Artix 7 operating at their nominal logic supply voltage. Values derived from the column oscillators, of which both average frequency and jitter are indicated.

Temperature		Altera Cyclone V		Xilinx Artix 7		
		300 K	30 K	300 K	30 K	4 K
Buffer delay	[ps]	239.5	210.3	238.7	235.3	234.7
↳ decrease	[%]	-	12.2	-	1.4	1.7
↳ SD (σ)	[ps]	5.9	7.1	5.7	5.3	4.8
Buffer oscillator frequency	[MHz]	2.69	3.07	2.68	2.72	2.73
Buffer oscillator jitter	[ps]	125	171	151	286	264
Inverter delay	[ps]	232.7	201.8	244.4	243	243
↳ decrease	[%]	-	13.3	-	0.6	0.6
↳ SD (σ)	[ps]	7	8.3	5.5	5.4	5.1
Inverter oscillator frequency	[MHz]	2.77	3.22	2.62	2.63	2.63
Inverter oscillator jitter	[ps]	119	162	152	257	279
Carry delay	[ps]	8.4	7.5	20	19.5	19.5
↳ decrease	[%]	-	10.8	-	2.4	2.4
↳ SD (σ)	[ps]	0.1	0.1	0.3	0.3	0.3
Carry oscillator frequency	[MHz]	74.2	83	31.6	32.4	32.5
Carry oscillator jitter	[ps]	9.9	9.4	18.6	21.7	21.5

The power consumption of the core logic, with only one of the oscillators running in it, is shown in Figure 5.7. Clearly, the power consumption reduces at lower temperatures. This can be attributed to more efficient logic on one hand (see Subsection 5.2.6), and to a reduction in leakage of the transistors on the other hand. The main power reduction is obtained by lowering the temperature from 300 K to 30 K, while the step from 30 K to 4 K is fairly small. The Artix 7 shows a linear dependence on supply voltage at 30 K and 4 K, whereas a quadratic behaviour is visible at 300 K. This leads to a significant power reduction at especially the higher supply voltages. Although the voltage dependence in the Cyclone V is more quadratic than linear, its behaviour is the same at both room and cryogenic temperatures.

5.2.4 REGISTERS

To benchmark the slice or ALM registers, firmware was designed with two counters and a latched comparator. Although the implementation is not comparable to an actual design, it gives an indication of logic changes in the flip-flops at low temperatures. The two n -bit counters are constantly incremented and compared against each other. In the second counter, the two most significant bits are tied to zero, resulting in equal counters and thus a positive comparator output in 25% of the clock cycles. The comparators duty cycle is measured together with a single bit of the counter, while increasing the frequency of the clock routed to the circuit. At a certain frequency $f_{\max}$ the circuit fails in the comparator, implemented with LUTs. The two counters, implemented in the fast carry elements, can be systematically operated at 10 to 30 MHz higher clock rates above $f_{\max}$. The maximum operating frequency $f_{\max}$ is shown in Figure 5.8 for an increasing number of bits n in the circuit. The shorter counters and comparators are expected to




Figure 5.8: Maximum operating frequency $f_{\max}$ for a circuit with two counters and a comparator for an increasing number of bits. The Artix 7 clearly outperforms the Cyclone V which is operating stably around 600 MHz.

5

run at a higher maximum frequency as fewer bits are propagating, which is indeed visible at all temperatures. The $f_{\max}$ in both FPGAs actually reduces with temperature by 30 to 60 MHz.

In contrast to the combinatorial logic, which becomes faster at cryogenic temperatures, the registers evidently require a longer setup and/or hold time for proper operation. These clock rates are still far beyond the nominal specified operating frequency of around 400 to 500 MHz and no problems were observed in larger systems with clock frequencies up to 400 MHz, such as the ADC of Chapter 6.

5.2.5 IO

Without doubt the IOs of both FPGAs are functional, otherwise none of the previous test could have been completed. However, it is noteworthy to check the variation in output




Figure 5.9: 2.5 V LVCMOS driver output amplitude over a 50 Ω load at room temperature versus signal frequency. The Cyclone V is more stable in terms of output amplitude over the tested frequency range, but is limited at cryogenic temperature to 830 MHz, whereas the Artix 7 is capable of operating up to 1160 MHz at 4 K.

drive strength of the IOs over frequency, since the IOs are fundamental for communication in and out of a cryogenic environment. In both FPGAs, the outputs were configured to a single-ended 2.5 V LVCMOS driver and terminated at room temperature with a $50\ \Omega$ load. Although an LVCMOS driver is generally not terminated with $50\ \Omega$, due to the long coaxial lines a $50\ \Omega$ load had to be applied to avoid reflections. The most significant difference between the IOs in the Altera and Xilinx FPGAs is the pull-up or pull-down. By default, the IOs in the Cyclone V are operating in pull-up fashion, whereas the Artix 7 IOs operate as pull-downs.

The output amplitude over the $50\ \Omega$ load is plotted in Figure 5.9 versus the signal frequency at which the IO is switching. The Cyclone V has a more stable output swing over the complete frequency range and fails at 960 MHz at 300 K and 830 MHz at 30 K, at which frequency the IO is no longer capable of pulling down from the positive supply. The Artix 7 deteriorates faster in output swing, but is capable of operating up to 970 MHz at 300 K, similar to the Cyclone V. At cryogenic temperatures, the performance is enhanced and the IOs are operational up to 1160 MHz at 4 K.

5

5.2.6 POWER CONSUMPTION & DIE TEMPERATURE

The actual temperature of the FPGA die or substrate will always be higher than the ambient temperature outside the FPGA package. Thanks to the internal diode in the Artix 7, the actual FPGA temperature can be accurately reconstructed while the FPGA is operating in an environment temperature of roughly 4, 30 or 300 K as shown in Figure 5.10. In order to reconstruct the die temperature, the internal diode was calibrated against a DT-670 (LakeShore Cryotronics) reference diode with the FPGA switched off, as shown in Figure 5.11. With the FPGA turned on, the reference diode remained stable at the ambient temperature, while the internal diode temperature increases linearly with the power consumed by the FPGA. Burning 0.5 W inside the FPGA, the temperature has risen by 13 K, 11 K and 9 K at respectively 300, 30 and 4 K, indicating the cooling power of the surrounding medium has steadily increased from normal air towards liquid helium. Although no temperature diode is available in the Cyclone V for the same calibration pro-




Figure 5.10: Xilinx Artix 7 power consumption versus die temperature. The FPGA's internal temperature diode was calibrated against a reference diode while powered off to extract internal temperatures (Figure 5.11).

cedure, we expect a similar temperature rise as the same package is used to house the FPGA.

To sweep the power consumption, an array with 2048 oscillators was implemented, utilizing over 90% of the resources in the Cyclone V and just under 50% in the Artix 7. The oscillators, each made up of 13 LUTs, were not constrained to particular locations and freely placed by the compiler, resulting in increased overhead from longer interconnects. This effect is especially visible for the Artix 7 with an oscillation period of only 117 MHz, versus 165 MHz in the Cyclone V. The extracted unit delay of the LUTs in the Artix 7 of 0.33 ns is significantly higher than the 0.24 ns indicated in Table 5.2, confirming the overhead caused by the compiler.

The power consumption was corrected for changes in the oscillation frequency of the main oscillator due to the increase in temperature and increase in voltage drop from the regulator to the FPGA supply pins as more oscillators are switched on. The power consumption versus the number of active oscillators is shown in Figure 5.12 for both FPGAs. The transient power consumption per LUT in the Artix 7 was found to be $20.5 \mu\text{W}$ at 300 K with an oscillation frequency of 117 MHz, resulting in a consumed energy per transition




Figure 5.11: FPGA diode calibration against external reference diode, both diodes biased with $10 \mu\text{A}$. Over 3,000 measurement points were used in order to recreate the diode curve. As hysteresis of $< 0.5 \text{ K}$ was present in the measurement data, the curve is averaged out between temperature decrease and increase.




Figure 5.12: FPGA power consumption versus number of active oscillators. Static power consumption is increased at lower temperatures, while dynamic or switching power is reduced.

of 0.17 pJ. While correcting for different frequencies at cryogenic temperatures, the transient power consumption decreases to 19.6 μW at 30 K and 18.1 μW at 4 K at, a decrease of almost 12%.

The transient power consumption in the Cyclone V was 24 μW with the oscillators running at 165 MHz, equivalent to 0.14 pJ per transition, which is more efficient than the LUTs in the Artix 7. At 30 K, the power is reduced to 20.5 μW , i.e. 15% less energy was required to run the oscillators.

5.3 SUMMARY

In this chapter, we combined all components from the previous chapters into a single platform, and demonstrated that commercial FPGAs can withstand temperatures over 200°C below the specified temperature limits. Although the Cyclone V is not capable of operating at 4 K, it does operate at 30 K with fairly similar performance characteristics as room temperature. Still, the Artix 7 outperforms the Cyclone V in most aspects, since it operates at 4 K and can be booted and reprogrammed even at such a low temperature. Its cryogenic performance characteristics are summarized in Table 5.3.

Since this is one of the first broad studies on the performance of cryogenic FPGAs, a vast amount of research can still be dedicated to it. We believe that newer generation FPGAs will operate with even less differences between room and deep-cryogenic temperatures and can be a valuable addition to the quantum–classical interface. New studies should focus not only on new FPGA generations, but also on FPGAs being exploited as co-processor, as is done in the Xilinx Zynq and Altera SoCs on which the FPGA is co-integrated with an ARM processor. To operate such a system, several unknowns

Table 5.3: Overview of the working elements inside the Xilinx Artix 7 operating at 4 K.

Module	Usable	Test	Performance w.r.t. RT
IOs	✓	Output swing versus frequency	Drive strength increases
LVDS	✓	Differential clocks from RT	
LUTs	✓	LUTs connected as oscillator in full columns	Propagation delay decreases < 3%, jitter increases < 78% (151 to 270 ps)
CARRY4	✓	Carry chains connected as oscillator in full columns	Propagation delay decreases < 3%, jitter increases < 16% (18.6 to 21.5 ps)
DFF	✓	n -bit counters + comparators	Maximum operating frequency reduces
BRAM	✓	Transfers of 8 kB (write & read)	No corruption in 100 test sets of 8 kB
MMCM	✓	100 MHz differential input clock multiplied by 12 and divided by various values, single ended output	Jitter decreases on average 52% (13.6 to 6.5 ps)
PLL	✓	100 MHz differential input clock multiplied by 12 and divided by various values, single ended output	Jitter decreases on average 54% (13.7 to 6.4 ps)
IDELAYE2	✓	IDELAYE2 elements connected as tunable oscillator (output frequency variable 13–70 MHz)	Delay decreases < 30%, jitter increases < 50%
DSP48E1	✓	Random calculations (additions, subtractions and multiplications)	No corruption in 400 calculations
Temperature diode	✓	Operating range 4–300 K	

still have to be answered for, such as the operation of DRAM, memories and high-speed links at these low temperatures.

Another recent development that is noteworthy is the co-integration of not only a conventional processor, but also an RF interface with high-speed ADCs and DACs. The Xilinx Zynq UltraScale+ RFSoc [64] hosts not only a large FPGA and an ARM processor, but also 8 4 GSa/s 12-bit ADCs and 8 6.4 GSa/s 14-bit DACs. This could be the basis for any quantum control interface, although its power consumption might be a concern at low temperatures.

6

FPGA-BASED ANALOG-TO-DIGITAL CONVERSION

To read out the qubits, we require an analog-to-digital converter (ADC). In, for example, the quantum–classical interface of superconducting qubits, the ADC has evolved from having a high resolution to having higher sampling rates. In 2012, an ADWIN was used for the read-out, i.e. a 14-bit ADC with 2 MSa/s [65]. In 2015, an 8-bit ADC with 20 MSa/s replaced the slower ADWIN [12]. Currently, those ADCs are being replaced with 200 MSa/s 8-bit variants, namely the ADC08200 from Texas Instruments. Other groups are employing both 500 MSa/s [66] and 1 GSa/s ADCs [67], without stating their resolution. This clearly shows the trend towards higher sampling rates.

Ideally, those ADCs are implemented as an application-specific integrated circuit (ASIC), so that their performance can be trimmed to the exact requirements, resulting in the highest sampling rate with the lowest power consumption. At cryogenic temperatures, there is only a very restricted set of ADCs reported in literature [68]–[75], usually with performance limitations due to a significant change in the analog behaviour of the various components. Cryogenic ADCs reported to date are limited in sampling rate up to a few hundred kHz and in resolution up to 11-bit, thus not supporting the requirements of the quantum–classical interface stated above.

For our cryogenic read-out system, an ad hoc solution is thus required. Therefore, in this chapter, we employ the functional Artix 7 (from the [previous chapter](#)) as a high-speed ADC. Since the FPGA is a completely digital system, it is not straightforward to use it as an ADC. We designed a reconfigurable architecture that can be fully implemented in the FPGA, except for some resistors on its periphery. This not only allows the ADC to be easily interfaced with the remainder of the digital circuitry; it can be scaled to the required sampling rate or resolution, and even enables ADCs with different specifications to co-exist in one system. Above all, it allows calibration to each new environment, in which the system is operating, i.e. it can adapt to changes in voltage, temperature or variations between chips. We show the effectiveness of our calibration techniques by operating the ADC both at room temperature (300 K) and in a cryogenic environment (15 K).

This is not the first attempt to operate the FPGA as ADC, although it is the fastest design to date that can even operate at cryogenic temperatures. Previous designs relied on

Parts of this chapter have been published in the IEEE Transactions on Circuits and Systems I [J3].

conventional architectures. The implementations from [76], [77], based on delta-sigma modulators, operate at a relatively low sampling rate, in the order of tens of kHz, due to the need for a feedback-loop between input and output. The first design stepping away from more conventional ADC architectures is based on time-to-digital converters (TDCs), combined with a reference ramp and comparator. [78] used an RC -filter with 3 resistors and 1 capacitor to create the reference ramp. This reference ramp was then compared with the analog input signal in the FPGAs differential inputs. The time between crossing input and reference ramp was measured in a TDC based on clock phase interpolation, limiting the overall performance and resolution. Both [79] and [C3] improved on this design by implementing the TDC with the carry chain.

The basis of the current design (detailed in Section 6.1) is the use of low-voltage differential signaling (LVDS) buffers as comparators and fast time-to-digital converters to perform the actual conversion. For the start-up calibration (Section 6.2) of the system, we employ the IO delay blocks to equalize the delays to each TDC, and we use a masking circuit to precisely set the length of the TDC to match the reference period. Section 6.3 reports the performance of our ADC and compares it with results reported in literature (both FPGA and ASIC ADCs). A short summary is given in Section 6.4.

6

6.1 SYSTEM ARCHITECTURE

The basic principle of the proposed ADC architecture is shown in Figure 6.1. The system is built around a core consisting of a TDC, a mixed-mode clock manager (MMCM) and an LVDS input buffer. The MMCM receives an input clock of 100 MHz (from an external source) and generates a clock signal V_{OUT} (with 50% duty cycle) that charges and discharges the RC -circuit consisting of a SMD resistor R_{REF} and the parasitic input capacitance C_{INT} of the LVDS differential input buffer. The frequency of V_{OUT} can be tuned in order to increase or reduce the sampling rate. The generated RC -curve is optimal with a time constant in the order of $3 \times$ the reference clock period, optimizing the ratio of curve linearity to input range.




Figure 6.1: Block diagram of the operating principle of the proposed ADC architecture. The architecture comprises a TDC, an LVDS comparator, a MMCM and a resistor R_{REF} . The analog signal V_{IN} is digitized, after conversion of the analog quantity to time, in the TDC.

The LVDS input buffer is used to compare the analog input signal V_{IN} with the RC reference ramp. A '1' at the output of the LVDS buffer is generated as long as the reference ramp is higher than the analog input. The actual conversion to a digital ADC code is done with the TDC.

The TDC needs to be able to measure rising and falling edges. To do this, the TDC decoder is implemented as a bit counter, counting the number of ones in the carry chain. This technique also averages bubbles if they occur [C3].

The sampling rate for a single channel is theoretically maximized to the maximum clock frequency, in the order of 400 MSa/s. To reach both a higher sampling rate and resolution we employ a multi-phase clock interpolation technique to sample the analog signal multiple times in one clock period. Therefore, instead of a single V_{OUT} , we make use of 6 V_{OUT} channels (thus 6 LVDS inputs), each with the same period, but increasingly shifted in phase by $\frac{1}{6}$ of that period.

The waveforms corresponding to a 1.2 GSa/s ADC architecture are depicted in Figure 6.2. The MMCM generates a 100 MHz base clock and the 6 shifted phases. Each clock is routed to a dedicated resistor on the PCB and further towards the LVDS input. The input signal is routed towards the 6 LVDS buffers and can therefore be compared with the 6 phases of the clock. 12 measurements are generated in the 100 MHz clock period leading to a sampling rate of 1.2 GSa/s.

6

The complete implementation is detailed in Figure 6.3. The principle from Figure 6.1 is scaled to 6 channels, each channel is driven by its own clock phase (generated in the MMCM), which creates the 6 phase shifted RC -ramps. Those ramps are compared with the input signal in the LVDS buffer and its output is routed towards the TDC.

The input stage of the TDC consists of an IO delay element (capable of delaying the input signal by various nanoseconds in steps of roughly 50 ps), which is used for the calibration procedure to align the input signal's period with the carry chain (see Section 6.2 for more details on the calibration). Furthermore, a multiplexer is implemented to be






Figure 6.2: Waveforms for a 1.2 GSa/s ADC using 6 phases of a 100 MHz base clock. In each 100 MHz clock period, 12 measurements (6 on the LVDS rising and 6 on the falling edge) are executed. The output clock for the first LVDS channel and its corresponding LVDS output are drawn as reference.



(a) The 6 separate measurement channels



(b) The measurement block: the TDC

Figure 6.3: Detailed block diagram of the implementation of a 1.2 GSa/s ADC consisting of 6 output channels, 6 input buffers and 6 consecutive TDCs.

able to switch between the external input and an internal reference signal. The internal reference signal is the clock routed through another IO delay and is used to calibrate the carry chain's required length. The length is trimmed to the reference clock period and a masking circuit is employed to disable the abundant part of the TDC, i.e. the part of the TDC that extends beyond the clock period. After the completion of the various TDC calibration steps, all TDCs (one for each channel) are covering exactly the TDC clock period (2.5 ns) and are properly synchronised to one another.

The TDC operates at 400 MHz (2.5 ns period), enabling a relatively short carry chain length of 200 carry blocks (the maximum length that fits in one clock domain of the Artix 7). This is done to minimize non-linearities and to reduce the overall FPGA logic utilization. With an average resolution of 16 ps per carry block, the carry chain covers a total of 3.2 ns.

The carry chain's output is double-latched and decoded to a binary value in the thermometer decoder (implemented as a counter-based tree structure). As the decoder is

also running at 400 MHz, a rate reducer to 100 MHz is required to integrate the TDC with the remainder of the circuit.

The communication is implemented with a serial UART protocol, this protocol is chosen as the number of wires towards the cryogenic environment is limited and UART can be implemented with only 2 to 4 wires. However, the data needs to be buffered inside the FPGA as the transfer speed of UART is limited to roughly 1 Mb/s in our setup. The amount of data generated by the ADC is in the order of $1.2 \text{ GSa/s} \times 8 \text{ bits} \approx 10 \text{ Gb/s}$.

The majority of the post-processing is done after read-out on the host. For the interface an UART-to-USB controller was connected to the system to enable a simple interface between PC and FPGA. The FPGA was programmed using JTAG.

The PCB used for testing the FPGA ADC in room and cryogenic temperatures is detailed in [J4]. Due to size and connection limitations the number of components on the PCB is reduced to the bare minimum, i.e. the FPGA is the sole active component present on the PCB. Note that this board is not housing any of the LDOs and oscillators previously discussed.

In the calibration section we discuss the steps required to properly operate the ADC, such as the sorting of the codes from six different phases and the use of time stamps for further interpolation below the 830 ps ADC periods.

6.2 SYSTEM CALIBRATION

In contrast to ASIC ADCs, that only need a small auto calibration or no calibration at all, our ADC needs an extensive calibration process to reach optimal performance. Although calibration takes time and resources, this extensive calibration can compensate for logic changes over a wide temperature range and allows the use of components with higher tolerances. The calibration consists of 4 steps and needs a total of 3 different input signals on the ADC. Two sawtooth and one sinusoidal signal, of which the frequencies are proportional to the sampling rate, are applied to the ADCs input during the calibration.

6.2.1 TDC LENGTH CALIBRATION

The wide temperature range brings significant changes to the timing inside the FPGA. To deal with timing differences, the used carry chain has been made longer than the sample period, which is 2.5 ns. The implemented carry chain consists of 200 delay elements and has a range of 3.2 ns at room temperature. Figure 6.4 shows the problems that are created when the carry chain has more or less range than the sample period. In the case of a chain covering more time than the sample period, the last part of the chain will be a representation of data from the previous period. Data will be double-sampled, leading to an overestimation of the correct time. A chain that is too short will lose a part of the required measurable range, leading to an underestimation of the correct time in this lost range.

To equalize these values, the TDC has an option to disable elements at the end of the chain. The calibration of the TDC length is done by phase shifting a clock signal through the TDC, as can be seen in Figure 6.5, only if the length of the TDC exactly matches the clock period, the output values, while shifting a clock through the line, will be the same.




Figure 6.4: Representations of carry chains that are either longer or shorter than the sample period. A line covering more time than the sample period leads to data being double processed. A too short line leads to data loss.

6




Figure 6.5: Measured clock duty cycle with a too long, too short and a calibrated TDC, while phase shifting the clock input with 360 degrees through the carry chain. Only in the calibrated TDC, the average is constant at 50%.

Too long TDCs will have an overestimation of the time in a certain range; too short TDCs will have an underestimation of the time in a certain spot.

By repeating the same measurement with increasingly more TDC blocks disabled, we can find the number of blocks for which the output of the TDC is always the same, i.e. the smallest standard deviation over one complete 360 degrees rotation of the clock.

To cover a reference period of more than 2.5 ns, multiple consecutive measurements are summed together to form one time stamp per reference period. E.g. for the 1.2 GSa/s ADC, the reference period is 10 ns (however there is one measurement in both the LVDS rising and falling edge), thus 2 consecutive TDC measurements need to be combined.

6.2.2 TDC ALIGNMENT CALIBRATION

After calibration of the TDC length, the TDC needs to be aligned with the reference period. We have to take into account the fact that one LVDS generates two outcomes in one reference period, one for the rising and one for the falling edge. Hereafter, we will refer to these as even, respectively odd, parity.

Figure 6.6 shows on top the RC -curve and a DC input signal in the bottom of the range. These two signals compared in the LVDS buffer lead to the result shown below for an unaligned TDC. The measured values are 0 and 20 for the first and second half of the clock period, which is not aligned to the reference period. Consequently one of the parities can not measure the current input voltage. After alignment of TDC clock and reference period, the values are properly measured to be 5 respectively 15. Both parities can measure input voltages over the complete analog range.

By applying a slow ramp that exceeds the maximum and minimum input voltage, we can find the best alignment of the TDC, i.e. when the TDC parities have a maximum span. At a certain time after starting the ramp, the TDC values of the parities go from 0 to 1, indicating the start of the TDC span. The span ends as soon as the parities go from $max - 1$ to max . By calculating the time between these events for both parities we can calculate the span. When doing this span measurement for all possible alignments we can find the alignment where both parities have the biggest span. That is the configuration in which the alignment is optimal.

6

6.2.3 TDC ADC CALIBRATION

After completing the calibration of the TDCs, the voltage characteristics of the ADCs are calibrated. The transfer of input voltage to digital output is measured by applying a sawtooth on the input of the ADC. As there is a direct relation between input voltage and time after starting the sawtooth, the sample time can be converted to a voltage. For every LVDS input this calibration generates two look-up tables that convert TDC values to ADC values, one for the falling and one for the rising RC -curve.

An actual conversion graph can be seen in Figure 6.7. Rising and falling curves are monotonic. This can be explained by the duty cycle measurement. As shown in Figure 6.8 the outcome of the TDCs is positively correlated to time of measurement for the rising curve and negatively correlated to the falling curve.




Figure 6.6: Schematic view of the TDC alignment process. The measurement period of the TDC does not match with the reference period. The TDC is shifted to match both periods to one another.




Figure 6.7: The conversion graph from TDC to ADC code for the rising and falling edge of the RC -curve. The ramps are monotonically rising and in the calibration corrected for their RC distortion.




Figure 6.8: The timing principle of the TDCs. The principle of counting ones in the thermometer decoder makes the generated time stamps to appear for the rising edge to be between the start of the reference period and the crossing point of the signal and ramp. For the falling edge, the calculated time is between the falling edge crossing point and the end of the reference period.

6.2.4 ADC SYNCHRONIZATION

The entire design of this ADC depends on the accuracy of the ADC synchronization. After all, the results of the 6 interleaved channels can only be combined if the channels are accurately synchronized to each other.

With the previous calibration steps, each LVDS input behaves as a standalone ADC, generating two values in one period of the reference signal. To properly combine the values of six different channels, the time difference of each of the six phases has to be known.

To the inputs of all six channels, a sinusoid with a period covering 32 reference periods is applied. All ADCs measure this signal and a sinusoid is least squares fitted on each ADCs dataset. From each of the 6 fitted sinusoids, the phases is extracted. The phase is then converted to a time and for every ADC, the offset from the mean off all phases is extracted. By taking all these offsets into account, the ADCs timestamps can be properly sorted.

6.3 RESULTS

The proposed ADC design was fully characterized; the shown results were achieved with a combination of our ADC and the voltage stabilizer discussed in the next chapter. Both

its versatility due to reconfigurability and the extensive calibration are discussed. The reconfigurability makes it possible to achieve a wide range of sampling rates and resolutions. The extensive calibration is tested by operating one particular ADC configuration (1.2 GSa/s ADC) at 300 K and 15 K.

6.3.1 PERFORMANCE OF RECONFIGURABLE CONFIGURATIONS

One of the most interesting aspects of this ADC is the possibility to switch between a number of configurations. The only requirements are a small firmware change and, for optimal range, a change in reference resistor for the creation of the RC -ramp.

The performance of ADCs, capable of sampling from 12.5 MSa/s up to 2.4 GSa/s, is shown in Figure 6.9. For each ADC, the best (low input frequency) performance is reported. For a single channel, being only the rising or falling edge of one LVDS comparator, we can reach a performance of over 9 bits (ENOB) with a sampling rate of 12.5 MSa/s. The highest possible sampling rate is achieved with a reference period of 200 MHz on the 6 phase interleaved channels, leading to an impressive 2.4 GSa/s on an FPGA. With each pin added to the ADC, the performance can be improved, however the performance does not scale linearly with the number of pins added to the system. The performance at 2.4 GSa/s is slightly less than 4 bits, especially due to distortion and also due to interference between the different RC s of 200 MHz.

With our ADC we can achieve a wide variety of both sampling rates, a factor of 192 between highest and lowest, and effective resolution, a factor of 5 bits difference. This makes the ADC useful in many applications, in the range of possible configurations. The ADC operating at 1.2 GSa/s is studied in more detail in the following section.




Figure 6.9: ADC performance ENOB versus sampling rate at a 1 MHz input frequency. The ENOB is characterized over a wide range of possible configurations, for each configuration the maximum performance is given. Results are given for one parity (either the rising or falling edge of a single pin) and one to six pins interleaved.

6.3.2 PERFORMANCE OF A 1.2 GSa/s ADC AT 300 K AND 15 K

The 6 interleaved channels ADC is extensively investigated for its performance in terms of range, resolution, single shot precision, non-linearities and signal-to-noise ratio at room temperature (300 K) and immersed in liquid helium (15 K).

FULL RANGE, RESOLUTION AND POWER

The full range of the ADC is dependent on the frequency of the reference signal and to a certain extent the input frequency. The input capacitance is acting as a kind of low pass filter, reducing the input range.

The frequency of the reference signal together with the RC -filter determines how far on the RC -ramp the reference voltage will rise and fall. This is further limited by a finite switching time of the FPGA output pin, especially when switching from 0 to 1. This also implies that there will be a small offset voltage at the input of the ADC.

For the 1.2 GSa/s ADC the range is measured to be [0.9 1.6] V. The least significant bit (LSB) of the ADC represents 3 mV. Thus the full range spans approximately 8 bits. This range can be increased when the reference resistance is reduced. However this will result in more RC deformation and thus a worse DNL and single shot accuracy. The range does not vary significantly over temperature, thanks to the extensive calibration.

The power consumption of the FPGA-based ADC is measured to be 750 mW at room temperature, including clock generation circuitry, data buffering and read-out. In Figure 6.10 the distribution of the dissipated power to the various functions is depicted. It is dominated (around 40%) by the power consumption of the PLL and MMCM circuitry for generation of the high frequency clocks and the different clock phases. The actual conversion power is measured at around 200 mW. Especially at cryogenic temperatures, the clock generation power consumption increases, totalling the systems power consumption at 850 mW.

SINGLE SHOT RESOLUTION AND DC ACCURACY

The single shot resolution (SSR) is the variation over time while measuring a constant/DC input value. It is measured by accumulating samples with a constant input over time. As the accuracy is dependent on the position in the full range, the performance over the complete range has to be studied.

The single shot measurement resolution is shown in Figure 6.11. The resolution is always better than 2 LSB (1σ) and on average 1.1 LSB. This corresponds to an uncertainty in the measured voltage of roughly 3 mV.




Figure 6.10: Distribution of the dissipated ADC power to the various functions of the circuit in the FPGA at 300 K.




Figure 6.11: Single shot measurement obtained by applying a DC voltage to the ADC input. Over the complete measurement range, the precision is 1.1 LSB (1σ) on average at 300 K and decreased to 2.3 LSB at 15 K.

Table 6.1: Worst case jitter measured for different FPGA components at 300 K and 15 K.

Temperature		300 K	15 K
PLL jitter	[ps]	12	10
Carry jitter	[ps]	12	12
IO delay jitter	[ps]	24	45

The ADCs SSR is mainly limited by clock jitter (from the PLL), carry chain jitter (delay line) and IO delay jitter. The jitter measured for these individual blocks is detailed in Table 6.1. The combined jitter, or aperture jitter, can be calculated according to Eq. (6.1), yielding 30 ps and 48 ps, at 300 K and 15 K, respectively.

$$\sigma_a = \sqrt{\sigma_{\text{PLL}}^2 + \sigma_{\text{carries}}^2 + \sigma_{\text{IO delay}}^2}. \quad (6.1)$$

The aperture jitter imposes an upper bound on the ADC performance (signal-to-noise ratio) following:

$$\text{SNR}_{\max} = -20 \log(2\pi f_{\text{in}} \cdot \sigma_a), \quad (6.2)$$

in which f_{in} is the signal input frequency.

Considering an input frequency of 2 MHz, the best achievable SNR is in the order of 68 dB and 64 dB, respectively at 300 K and 15 K. At 40 MHz this is deteriorated to 42 (38) dB.

NON-LINEARITIES

ADC non-linearities are identified using a density test. At the ADC input a slow ramp, exceeding both negative and positive input range, is generated. For each ADC code all occurrences are counted and stored in a histogram, from the histogram the non-linearities can be computed. With an ideal transfer, the counts $C(n)$ in each bin n are equal to the average number of counts $\bar{C} = \frac{\sum_{i=1}^N C(n)}{N}$.

From the histogram the non-linearities can be computed with:

$$\text{DNL}(n) = \frac{C(n)}{\bar{C}} - 1; \quad (6.3a)$$

$$\text{INL}(n) = \sum_{i=1}^n \text{DNL}(i), \quad (6.3b)$$




Figure 6.12: Non-linearities (DNL and INL) extracted from ADC density test. The ADC density is produced by applying a ramp that exceeds both negative and positive input range to the input of the ADC.

6

The DNL and INL for the system (Figure 6.12) are in the range $[-0.75 \ 1.04]$ LSB and $[-0.36 \ 0.52]$ LSB, respectively. At 15 K, the non-linearities are similar, DNL and INL in the range $[-0.85 \ 1.04]$ LSB and $[-0.68 \ 0.77]$ LSB, after calibration.

Main contributors to the ADCs non-linearities are the intrinsic non-linear behaviour of the exponential reference ramp on the analog side and the behaviour of the TDC on the digital side. The non-linearities of the TDC are mainly caused by unevenly spaced delays in the carry chain and a non-perfect clock distribution. Although these phenomena occur everywhere in the device, care should be taken not to place the carry chains in particular bad spots. All these problems can, at least partially, be overcome by the calibration to minimize system non-linearities.

AC PERFORMANCE

The ADCs signal-to-noise ratio (SNR), signal-to-noise-and-distortion ratio (SNDR), harmonic distortion (THD) and effective number of bits (ENOB) are estimated over frequency with a sinusoidal input signal. IM2 and IM3 are measured with a two tone test.

The amplitude of the input signal spans 90% of the ADC input range, in the results presented in this section unless otherwise stated, to avoid large distortion at the ends of the range. While a phase interlaced sampling method is used, the practical Nyquist frequency is the fastest individual sampling frequency (of a single phase), i.e. 100 MHz.

Above this frequency, the signal is faster than the individual phases, resulting in more than one transition in one sampling period, unless the input signal range is reduced to minimize the number of transitions per period to one. More than a single transition per

period cannot be detected, resulting in incorrect time-stamps. In Figure 6.13 this limit is sketched. Figure 6.13(a) shows a 100 MHz sinusoid, which results in one transition per sample period (once on the rising and once on the falling edge), so it can be adequately measured. However the 600 MHz sinusoid, at the theoretical Nyquist rate (b), has four to five transitions per sampling period, leading to inaccurate results. In the case the input amplitude is reduced, again a single transition per period can be reached, i.e. the signal can be properly measured.

The input sinusoidal signals are generated with a DG4202 (Rigol) function generator and low-pass filters are used to suppress out of band noise and harmonics.

In Figure 6.14 digitized sinusoidal signals at both 300 K and 15 K are shown at input frequencies of 2 and 40 MHz. At cryogenic temperatures, the performance can be seen to be degraded especially for higher input frequencies (Figure 6.14(d)). This degradation is caused by the jitter increase of the IO delay, as noted before, and by the significant decrease in decoupling capacitance at cryogenic temperatures. The decrease in total decoupling capacitance is causing more degradation at higher input frequencies as the current consumed by the FPGA is switching more frequently.

The frequency domain plots corresponding to the sinusoids depicted in Figure 6.14 are obtained with a fast Fourier transform (FFT) in Figure 6.15.

The various specifications derived from the frequency domain, such as SNR, SNDR, SFDR, THD and ENOB are summarized in Table 6.2. From the SNR at low and high input frequency, first the aperture time t_a is found to be around 99 ps following Eq. (6.4a).

$$t_a = \frac{\sqrt{(10^{-\text{SNR}_{\text{hf}}/20})^2 - ((1 + \epsilon)/2^N)^2}}{2\pi f_{\text{in,hf}}}; \quad (6.4a)$$

$$\epsilon = 2^N \cdot 10^{-\text{SNR}_{\text{lf}}/20} - 1, \quad (6.4b)$$

in which N is the 8.2 bits of our converter, SNR_{hf} the SNR at high input frequency and SNR_{lf} the SNR at low input frequency. 99 ps aperture time limits the performance, especially at higher frequencies, e.g. maximum SNR at 40 MHz is 32 dB.




Figure 6.13: Nyquist limitation of the ADC. (a) 100 MHz sinusoid that can be properly measured. (b) 600 MHz sinusoid (at theoretical Nyquist frequency), which cannot be measured due to multiple transitions in one sample period, with reduced amplitude, to allow only one transition per period, makes proper conversion possible.




Figure 6.14: Two periods of a 2 and 40 MHz sinusoid digitized at 1.2 GSa/s at 300 K and 15 K. The solid line is a best fitted sinusoid for reference.

Table 6.2: ADC specifications derived from the frequency domain as per Figure 6.15.

f_{in} [MHz]	2		40	
Temperature	300 K	15 K	300 K	15 K
SNR [dB]	41.1	32.2	31.6	28.8
SNDR [dB]	38.8	30.3	29.3	26.5
SFDR [dBc]	57.5	49.6	37.5	35.7
THD [dB]	-42.9	-34.6	-33.2	-30.3
ENOB [bits]	6.2	4.7	4.6	4.1

The signal-to-noise-and-distortion ratio, defined as the power ratio of the signal to the sum of the remainder of the spectrum, was found to be 39 dB at 2 MHz, respectively 29 dB at 40 MHz, at room temperature. Immersed in liquid helium, the SNDR is reduced, especially at higher frequencies, due to the limitations noted before.

The SNR is comparable to the SNDR indicating little harmonic distortion, as also indicated by the THD, defined as the ratio of the sum of all signal harmonics by the power of the signal frequency itself. Significantly more harmonic distortion was present at high frequencies, moreover there is clear interference with the 100 MHz sampling frequency, either due to coupling into the analog signal or imperfect synchronization of the different phases. At cryogenic temperatures, the performance has dropped, the noise floor has increased and the harmonic distortion is more apparent.

From the SNDR, the ENOB is calculated with Eq. (6.5) to be 6.2, respectively 4.6 bits, at room temperature for low and high input frequencies. At cryogenic temperatures, the ENOB is reduced by 0.5 to 1.5 bit.




Figure 6.15: Frequency domain representation of the sinusoids depicted in Figure 6.14. Frequency spectra were obtained from 48,000 samples, while using a Blackman-Harris window of length 2^{14} shifted over the samples. Furthermore, the spurious-free dynamic range (SFDR) is indicated.

$$\text{ENOB} = \frac{\text{SNDR} - 1.76}{6.02} \quad (6.5)$$

The ENOB over input frequency and input voltage range is shown in Figure 6.16(a) at room temperature, averaging 10 measurement results per point. In (b) the ENOB is corrected for full scale ENOB_{FS} according to Eq. (6.6).

$$\text{ENOB}_{FS} = \frac{\text{SNDR} - 1.76 + 20 \log\left(\frac{FS}{IR}\right)}{6.02}, \quad (6.6)$$

in which FS is the full scale amplitude and IR the actual input range.

Finally, the ENOB is shown in Figure 6.17 at both temperatures. The trend lines are drawn as a guideline, while at each input frequency 10 measurements are plotted with the standard deviation in the ENOB obtained over these measurements.

The effective resolution bandwidth is the input frequency at which the SNDR drops by 3 dB or the ENOB by 0.5 bit. The ERBW is in the order of 15 MHz at room temperature, but in the cryogenic environment it is roughly 25 MHz, due to the overall lower performance. Note that the ERBW is larger for smaller input amplitudes, due to less distortion (second contour line in Figure 6.16(b)), and up to 60-70 MHz. With an effective bandwidth of 15 MHz, the effective Nyquist sampling rate is roughly 30 MHz, instead of 1.2 GSa/s. This can be used for oversampling (with an oversampling ratio of 40 $\times$) to increase resolution and reduce the noise.

The intermodulation distortion, IM2 and IM3, is acquired with a two tone test, one at 40 MHz and a second at 33 MHz. The result is depicted in Figure 6.18 in time and frequency domain at both 300 K and 15 K. Besides significant power in IM2 and IM3 components, more cross-coupled harmonics are seen, especially at cryogenic temperatures. In room temperature however, second order harmonics from the two base tones are more apparent than at cryogenic.

CROSS DEVICE PERFORMANCE

As the application is so close to the maximum capability of an FPGA and PCB, it is likely that there will be changes in performance from board to board. We had 3 boards at our disposal and tested the 1.2 GSa/s ADC on all the boards at room temperature. The performance of the ADCs on all three boards is very similar thanks to the advanced calibration of our ADC. The calibration can smooth out the performance over different devices making this ADC usable as a soft-core of which the performance can be well regulated.




Figure 6.16: ENOB and ENOB_{FS} over input frequency and input range at 300 K. The result in (a) shows the actual ENOB measured, whereas (b) shows ENOB corrected for full scale as to show performance over the full range expressed in full scale ENOB. Contour lines are drawn at 6, 5.5, 4, 4.5 and 4 bits of ENOB.




Figure 6.17: ENOB over input frequency at both 300 K and 15 K. The resolution at 15 K is roughly reduced by 1.5 bit, due to higher IO delay jitter and reduced decoupling capacitance at low temperature.




Figure 6.18: IM2 and IM3 derived with a two tone test at 33 and 40 MHz.

6.3.3 COMPARISON OF A 1.2 GSA/s FPGA ADC WITH LITERATURE

Our FPGA ADC is compared with other FPGA implementations [76]–[78], [C3] and the hard-core ADC, integrated as XADC, in Xilinx 7 series [80] at room temperature in Table 6.3. Unfortunately, most general specifications are missing for the other reconfigurable ADC designs, nonetheless we tried to give a comprehensive overview.

The best performing ADC so far was found to be based on the carry chain TDC, however with our design we are capable of reaching both a higher sampling rate ($6\times$ faster) and, thanks to calibration, a higher resolution. However the ERBW of both ADCs is in the same order of magnitude. This most likely indicates an aperture time limitation of the LVDS buffers, which is in the order of 100 ps.

To compare directly with implementations as clock phase TDC, delta-sigma or especially ASICs is not fair. Nevertheless we can extrapolate the performance of our design, assuming it runs at a speed comparable to the other implementations.

Compared to clock phase TDC, we show (see Subsection 6.3.1) that a resolution of 9 bits can be achieved with a 25 MSa/s in our system, thus outperforming the clock phase TDC implementation with 3 bit higher resolution.

Comparing to the delta-sigma modulator, our ADC reaches a 9 bit resolution with a 25 MSa/s system, outperforming it in sampling speed with a factor of roughly $50\times$.


As a result, to the best of our knowledge, our proposed ADC outperforms the sampling rate of all previously reported ADCs implemented using reconfigurable logic, and although limited due to aperture limitations, also the widest effective resolution bandwidth.

To extend the comparison, we also compared our design with ASIC ADC implementations. To complete the comparison with ADCs found inside FPGAs, we compared our ADC to the ASIC XADC, found in Xilinx 7 series. We can extrapolate our performance to reach 10 bits of resolution at a rate of roughly 5 MSa/s. This shows that even with a soft-core ADC, a performance similar to that of an ASIC implementation can be reached, although not in terms of power consumption. Furthermore, the XADC is not reconfigurable, implying it is always sampling with 1 MSa/s. There are only two ADC channels available, whereas our design is only limited by the number of slices required to implement one channel. Although the XADC does not require any external components, as the reconfigurable ADCs do, it comes at the cost of occupying additional space on the FPGA die.


To compare more extensively with other ASIC ADCs, we used the ADC overview tables collected by Murmann [81]. The results of this extensive comparison can be found in the Figure of Merit plots of Figure 6.19. The Figures of Merit are intended to give an

Table 6.3: Comparison between our architecture and previous ADCs implemented in reconfigurable hardware. Additionally we compare with the XADC available in Xilinx 7 series FPGAs. *n.a.* is used to indicate that the corresponding data is not available.

ADC based on:	Multi-phase interleaving & carry chain TDC [this work]	Carry chain TDC [C3]	Clock phase TDC [78]	Delta-sigma modulator [76], [77]	XADC [80]
Temperature range [K]	15–300	300	300	300	300
Clock speed [MHz]	400	200	360	100	
Conversion rate [MSa/s]	1200	200	22.5	0.5–0.05	1
Voltage range [V]	0.9–1.6	0–2.5	0–3.3	0–3.3	0–1.0
Digital range [bits]	8.2	7.2	6	10–16	12
ENOB [bits]	6 (@ 1 MHz)	6 (@ 1 MHz)	<i>n.a.</i>	9	10
Resolution (LSB) [mV]	3	17	52	3	0.25
DNL [LSB]	[-0.75 1.04]	[-0.9 1.4]	<i>n.a.</i>	<i>n.a.</i>	± 1
INL [LSB]	[-0.36 0.52]	[-1.1 1.6]	<i>n.a.</i>	<i>n.a.</i>	± 2
Single shot resolution [LSB]	1.1	2	<i>n.a.</i>	<i>n.a.</i>	1
Power consumption [mW]	750	410	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>
FPGA slices	< 3555	< 400	<i>n.a.</i>	< 80	dedicated
External components	7	1	4	2 - 3	0



(a) Energy vs. SNDR



(b) Aperture vs. SNDR

Figure 6.19: Graphical comparison of the ADC presented in this work (square for 300 K, hexagram for 15 K) with different ASIC implementations of ADCs, using the following Figures of Merit:

(a) energy per conversion versus SNDR,

(b) frequency of the measured signal versus SNDR. The aperture error sets a limitation on the maximum frequency of the input signal that can be distinguished due to jitter.

The comparison is done for both designs functioning at room temperature as well as cryogenic temperatures (featuring only a small set of ASIC ADCs [68]–[75]). Especially at cryogenic temperatures, our design compares favourably with ASIC implementations. The comparison is based on the survey maintained by Murmann [81].

immediate ranking of the ADC in terms of speed, resolution and power consumption. In the figures our design is marked with a square (for 300 K) and a hexagram (for 15 K).

This comparison is not fair, as our design is implemented in a reconfigurable device. Even state-of-the-art FPGAs cannot reach the performance generally achieved with recent ASIC designs. In contrast, we can see that, although the design is implemented in a reconfigurable device, it can match the performance of ASIC designs in somewhat older technologies.

Another interesting fact that can be seen is the trend of ASIC designs over the years. Thanks to Moore's law, the performance keeps improving over the years, due to lower feature size, resulting in lower power consumption and higher speed. For FPGAs we see a similar trend, smaller node sizes will lead to a lower power consumption and a higher speed. This will again lead to a better performance of reconfigurable analog-to-digital converters.

A comparison of our ADC at cryogenic temperatures is more complicated. There is only a small set of ASIC ADCs operating at cryogenic temperatures [68]–[75], but still in a relatively wide temperature range, from 4 K up to 77 K. As analog behaviour of devices changes, especially at the lower temperatures (< 10 K), it is not entirely fair to compare all of them. However, we tried to give a comprehensive overview by including all cryogenic implementations in the Murmann plots of Figure 6.19 and compare them with our design.

The first thing to notice is the significant reduction in performance of the ASIC ADCs operating at cold, a reduction of over $100\times$ compared to the best performing ADC at room temperature. These implementations are generally limited in sampling rate up to a few hundred kHz, whereas our ADC was capable of sampling at 1.2 GSa/s at cryogenic temperatures, although with an ERBW in tens of MHz range. In terms of resolution, a general trend shows lower ENOB at lower temperatures, as does our implementation (0.5–1.5 bit reduction). The main limitation is found in the power consumed by our ADC. Most of the cryogenic designs burn less than 1 mW, whereas the FPGA consumes over 850 mW, due to an increase in static power consumption of the FPGA in the cryogenic environment and a large overhead [J4].

6

6.4 SUMMARY

In this chapter, we showed a completely reconfigurable ADC design, employing an Artix 7 FPGA. The design was demonstrated to operate from 12.5 MSa/s up to 2.4 GSa/s and can obtain a resolution as high as 9 bits (ENOB) up to 25 MSa/s. A full characterization has been done for the 1.2 GSa/s variant. Achieving an ENOB of 6 bits over an ERBW of 15 MHz makes the converter a non-Nyquist sampling ADC with 30 MHz Nyquist conversion rate and an oversampling ratio of $40\times$. It has a single shot resolution of 1.1 LSB and a high linearity (DNL $[-0.75 \ 1.04]$ LSB and INL $[-0.36 \ 0.52]$ LSB).

Besides changes in process and voltage from PCB to PCB, the PVT compensation has been demonstrated at cryogenic temperatures as low as 15 K. The ADC is performing in a wide temperature range, from 15 K to 300 K, but a small drop in performance is noticed at deep-cryogenic temperatures due to higher jitter of IO delay components and lower decoupling capacitance in the cryogenic environment, which is not due to or dependent from our architecture. The system is completely soft-core and can be implemented in an

FPGA, only 7 additional resistors need to be placed beside the FPGA (6 for the RC -curves and one $50\ \Omega$ matching resistor).

To the best of our knowledge, this is the fastest fully reconfigurable FPGA ADC reported to date, operating from deep-cryogenic to room temperature. Moreover, we believe this is the fastest cryogenic ADC (ASIC or FPGA) presented to date. Thanks to the reconfigurability and calibration features, this ADC can be used in many applications in a wide set of operating environments with specifications ranging from low sampling speed and high accuracy to very high sampling rates as 1.2 GSa/s.

7

FPGA-BASED VOLTAGE STABILIZATION

To enhance the FPGA-based ADC's performance; we can apply a voltage stabilization technique. It is similar to the use of voltage regulators (Chapter 3), but differs as it is purely firmware-based, thus requiring no extra dedicated hardware. Through firmware design, we compensate for fluctuations in the power dissipation. This design is based on real-time measurements of the variation of cell delay in the carry chain, which is mainly caused by voltage drop on the logic supply, but also by temperature fluctuations. A small farm of oscillators is used to flatten out the power consumption in real-time, in order to automatically stabilize the supply voltage as well.

Similar techniques have been shown before, but mainly for the purpose of avoiding security attacks. In those systems, the power consumption or heat is flattened in order to prevent intruders from reading data through power or heat changes (side-channel attacks). For instance, [82] proposes the use of a distributed oscillator farm to equalize the heat map from the FPGA as seen from outside. Its control is based on the difference in frequency shift between two differently sized ring oscillators, which is proportional to temperature.

We will first explain our implementation and measurement setup in Section 7.1. The results with and without the stabilizer are then highlighted in Section 7.2. Finally, a short summary can be found in Section 7.3.

7.1 SYSTEM ARCHITECTURE

7.1.1 MEASUREMENT SETUP

The problematic voltage drop is mainly caused by powering an FPGA over extremely long wires (compared to regulating the voltage on the PCB itself). The wire length is needed, because the power has to be supplied from outside the cryogenic environment, as the power supplies don't work at these temperatures (shown in Chapter 3). The measurement setup is schematically drawn in Figure 7.1. As can be seen, there are roughly 2 meters of wires from supply towards the cryogenic setup and roughly 1.5 m within the cryogenic setup. While different cable types have to be used, low ohmic wires do not represent a problem at room temperature, but cause significant heat injection into the

Parts of this chapter have been published in Review of Scientific Instruments [J6].




Figure 7.1: Schematic drawing of the FPGA operating at cryogenic temperatures while supplying its power through long wires coming from room temperature.

cryogenic environment. The main resistive path is therefore the cryogenic cable, which is estimated to be roughly 0.2Ω for R_+ and 0.1Ω for R_- . Besides the problems arising from these resistances, as will be discussed in the next section, wire inductance and latency are also important. The latency over several meters of coaxial cable accounts for 10 to 15 ns of delay, combined with the finite power supply reaction time (limited supply bandwidth) and the long wires with a few hundred micro Henri of inductance, the effect of the fluctuating voltage drop is significant for several milliseconds after each large current burst due to switching. Placing decoupling capacitors reduces this effect, however the reduced performance of capacitors at cryogenic temperatures limits the inductance compensation and/or may not be possible at all.

7

7.1.2 OPERATING PRINCIPLE

To regulate the voltage from within the FPGA, the power consumed by the FPGA has to be made constant, so as to keep internal delays also constant. To achieve constant power consumption, an architecture was developed to monitor the FPGA's internal delays and act upon detected variations.

The operation and its results are simulated in Figure 7.2, the room temperature power supply was set to 1.1 V, to arrive at the nominal FPGA supply voltage of 1 V in the cryogenic environment. In (a) the operation without regulation is shown. In this example, the system is triggered at a certain time to start with an operation. At that time, the power consumption is suddenly increased, in this example from 0.25 to 0.4 W. The resulting voltage on the FPGA drops by 5%, when assuming a resistive drop in the cable of 0.3Ω from 1.02 to 0.97 V. As a result, the propagation delay of the carry chain is increased by approximately 12% or 3.3 ps when operating at 300 K. At 4 K, this effect is severely worsened, the delay increases by roughly 20% or 4 ps. This change is significantly more than the 4% drop in voltage due to their semi-exponential relationship as shown in Figure 7.3. The resolution of a time-to-digital converter (TDC), implemented in the carry chain of the FPGA, is shown as a function of the supply voltage. The TDC resolution equates to the delay of a single carry element and approximates to 21 ps at the nominal supply voltage of 1 V. At cryogenic temperatures, the delay is increased at lower supply voltages due to the increase in the transistors threshold voltage, thus resulting in a stronger influence of the supply voltage on the carry delay.




Figure 7.2: Simulated impact of the power stabilizer on the FPGA voltage and the resulting delay of the carry elements, i.e. the resolution of the TDC derived from Figure 7.3. The resistive loss in the cables is approximately 0.3Ω , while the power dissipated by the system is 0.25, 0.4 and 0.5 W (Figure 7.5), respectively in steady-state, after the trigger and with the stabilizer active. The voltage at the room temperature supply is set to 1.1 V, while the voltage on the FPGA is nominal 1 V.




Figure 7.3: TDC resolution or carry delay versus the internal FPGA supply voltage (VCC_INT) at both 300 K and 4 K. Simplified from Figure 5.6.

With the regulation turned on, this effect is effectively smoothed, as shown in Figure 7.2(b). The power consumption, before the system is triggered, is slightly higher than that in the non-regulated system, around 0.5 W; after the trigger, the power consumption remains stable. While this technique can achieve a more stable system, the main drawback is the increased power consumption. This is especially important for designs operating at cryogenic temperatures, for which the power budgets are limited.

These problems can be partially addressed with a supply capable of forcing and sensing the voltage through different terminals. In this way, the voltage is measured closer to the FPGA through high-ohmic contacts and the supply can therefore adjust its internal voltage to compensate for the voltage drop in the cables. In practice, this technique is limited by the bandwidth of the feedback and the supplies capability of acting quickly

and accurately enough, especially with a system switching current frequently and at a high rate.

7.1.3 IMPLEMENTATION

To achieve the stabilization of the supply voltage, a small oscillator farm is employed together with a regulator consisting of a time-to-digital converter, an input-output (IO) delay block and some control logic. The complete circuit is shown in Figure 7.4(a).

The internal clock signal, provided by a mixed-mode clock manager (MMCM), is used as stable reference. It is routed through an IO delay block and a carry chain towards a small TDC. This TDC constantly measures the timing of the clock edge. At system start-




Figure 7.4: FPGA internal stabilizer architecture. (a) The oscillator stabilizer control system is divided in the following parts: a TDC with IO delay and a dummy carry chain in order to calibrate the clock delay into the delay line and to generate extra impact of voltage change on the TDC outcome. The outcome is routed towards a decision maker which detects if the voltage is increased or decreased. Control signals are finally sent to the oscillator farms. (b) Detail of decision maker and farm. The decision maker first integrates and filters the TDC data, leading to increment (inc) or decrement (dec) signals that are sent to four oscillator farms, each consisting of 128 oscillators.

up, the IO delay is used to shift the clock edge to the middle of the TDC range. As a result, the output of the TDC should be constantly indicating half of its range, as the position of the clock edge.

However, as a result of voltage drop and consequent shifts in logic delay, this is not the case. As the resistive voltage drop increases, the delay of the elements increases and the propagation of the clock will slow down, hence a higher TDC output is expected. On the other hand, with lower power consumption, and thus lower voltage drop, the internal logic will be faster and the edge will shift to lower TDC values. As the TDC is one of the most precise elements in the FPGA, the changes can be significant even with slight variations in the power consumption, as discussed in the previous section.

To correct for the changes in power consumption, oscillators in the farm are turned on or off depending on the direction of the shift. With lower measured TDC output, some oscillators are turned on. Higher TDC output will lead to a decrease in the amount of enabled oscillators.

The TDC is operated with a 400 MHz clock, while decisions are made at a rate of 6.25 MHz, i.e. every 160 ns. The decision maker and the oscillator farms are detailed in Figure 7.4(b). First, the decision maker filters out noise and jitter from the TDC by averaging multiple measurements, 64 in total for a decision rate of 6.25 MHz. Four results are summed to cross from the 400 MHz to the 100 MHz clock domain. The remaining 16 results are processed using an infinite impulse response (IIR) filter. The result of the averaged measurements is then taken as a measure for the required increase or decrease in the number of enabled oscillators. An averaged value higher than half of the TDC range will flag the decrement signal, a value lower than half of the TDC range will force an increment. A value of exactly half the TDC range is ideal and no change to the number of running oscillators is performed.

The oscillators are implemented as rings of logic elements in which at least one inverting element is present to enable oscillation. A NAND is used in the ring as enable element to turn the oscillators on and off as decided by the decision maker. Each oscillator consists of 6 logic elements, implemented in the look-up tables of the FPGA, 5 configured as buffers and one as NAND gate for the first stage to implement both the inversion and the enable element.

The number of required oscillators greatly depends on the power consumed in the core circuitry. First, the minimum and maximum power consumption of the core circuitry has to be roughly known. Afterwards, the number of implemented oscillators should be capable of bridging the gap between minimum and maximum power consumption with some margin.

In our design, the core circuit is our analog-to-digital converter (Chapter 6), which dissipates 250 mW in idle mode and about 400 mW at maximum conversion rate. To bridge this gap, four farms were implemented, each containing 128 oscillators. Therefore, the total number of active oscillators out of the 512 oscillators can be incremented or decremented with four at a time, i.e. the same increment and decrement signals are routed to each oscillator farm.

7.2 RESULTS

7.2.1 VOLTAGE REGULATION

The implemented system consists of our 1.2 GSa/s analog-to-digital converter (ADC) (Chapter 6) combined with the voltage regulation circuit as discussed in the last section. The system was tested at 300 K and 15 K as the performance at 4 K was deteriorated compared to 15 K by increased jitter in some components and the reduced decoupling capability. The ADC is triggered at time $t = 0$ to start converting the analog input and store this data in the internal FPGA memory. At time zero, the power consumption is significantly increased for the non-stabilized system, as shown in Figure 7.5. In (a), the ADC starts converting a sinusoidal signal with a 2 MHz frequency, while in (b), the conversion is done with a 40 MHz input signal. Although the difference in power consumption between these two cases is small, about 2.5%, it will still affect the signal-to-noise ratio (SNR). The main problem for our ADC though is the difference between calibration and final conversion. While the calibration is done with signals in the low kHz frequency range, the difference in the power consumption is more significant. In principle, the final conversion system stability is not the same as the system stability at the time of the calibration, leading to significant differences in signal shapes that cannot be easily accounted for.

Therefore stabilizing the power consumption in all these cases, can improve the final ADC conversion results. Our stabilization technique indeed stabilizes the current as shown in both figures, however the current is not as flat at 15 K as it is at 300 K and a clear transient at $t = 0$ is still visible. This transient is mainly caused by inaccurate tuning of the initial settings for the oscillator farms at 15 K, which is somewhat harder due to the more significant impact on both carry delay (Figure 7.3) and temperature compared to tuning at 300 K. Combined with increased jitter of especially the IO delay component, the decision logic is no longer capable of fully correcting the power difference.




Figure 7.5: Measured FPGA current with stabilizer turned on and off at 300 K and 15 K. At time $t = 0$, the ADC starts a measurement with an input sinusoid of (a) 2 MHz and (b) 40 MHz.

7.2.2 CALIBRATION AND CONVERSION

Although we can indeed see a significantly more stable current over time, the question remains whether that improves the performance of our ADC. To study the performance of the ADC, tests were done in the four operating conditions, with and without stabilization at 300 K and 15 K.

First, the ADC is calibrated at both temperatures. As the ADC consists of 12 individual sources, i.e. 6 TDC channels and all measuring rising and falling edges, the calibration is essential. The calibration is discussed in detail in Section 6.2 and can be summarized with the following four steps: tuning of the length of the 6 TDCs to match the clock period of 2.5 ns, aligning of the 6 TDCs with the clock edge, calibrating of the time stamp to the analog value and finally synchronising of the phases of the 6 channels for proper sorting of the analog values.

Using those calibration sets, measurements were performed with an input sinusoidal signal of 2 and 40 MHz. The results for the 2 MHz conversion are shown in Figure 7.6 after merging 12 sources and applying the calibration. The digitized sine waves can be seen to be quasi spur free when the stabilizer is used. However, especially at 15 K, the result is not as clean as at 300 K. This can be partially explained by the non-equalization of the current at 15 K, as shown before in Figure 7.5, but is also caused by the decreased decoupling capabilities of our capacitor network and increased jitter in some of the FPGA components.




Figure 7.6: Three periods of a 2 MHz sinusoid digitized at 1.2 GSa/s at 300 K and 15 K with the stabilized and non stabilized system. The input sine wave spans 90% of the ADC input range to prevent the system from clipping etc. The solid line is a best fitted sinusoid for reference.




Figure 7.7: Frequency domain representation of the sinusoids depicted in Figure 7.6. Frequency spectra were obtained from 48,000 samples while using a Blackman-Harris window of length 2^{14} shifted over the samples. Furthermore, the spurious-free-dynamic-range SFDR is indicated.

7

To quantify the improvement in performance, the sine waves are converted to the frequency domain, as shown in Figure 7.7. Clearly, the stabilizer brings significant reduction in signal distortion, especially at cryogenic temperatures. The harmonics of the 2 MHz input tone are reduced, increasing the spurious-free dynamic range (SFDR) by more than 18 dB at 15 K. Furthermore, the signal-to-noise ratio (SNR) and signal-to-noise and distortion ratio (SNDR) are enhanced by up to 9 dB at 15 K thanks to a reduction in especially high frequency noise when enabling the stabilizer. Stabilizing the supply voltage enables a better calibration effectively visible in a lower noise floor.

The main specifications of the ADC are listed in Table 7.1 for the use-cases described before. At 300 K, the stabilizer doesn't bring significant improvements. The effective number of bits (ENOB) is improved by 0.4 bit at low frequencies, but at higher frequencies, there is only a minor improvement of 0.1 bit (which is in the error-margin). At cryogenic temperatures however, the improvement is more significant. For low frequencies, there is a 1.5 bit better ENOB and at high frequencies, the ENOB still improves with 1.2 bit. These results clearly show that the ADC operation at cryogenic temperatures is challenging, especially when not considering power fluctuations and the significant effect of voltage drop on the FPGA behaviour. Though, there are still many unknown effects of electronics, materials and the corresponding behaviours while operating at temperatures so far out of the normal industrial temperature range, thus prompting more research in many aspects of cryogenic FPGA design.


Table 7.1: Summary of the obtained ADC specifications while converting a 2 and 40 MHz sine at 1.2 GSa/s with the stabilizer turned on or off. The stabilizer improves the performance, especially at cryogenic temperatures, with up to 1.5 bit in terms of ENOB or roughly 9 dB SNDR.

f_{in} [MHz]	2				40			
Temperature	300 K		15 K		300 K		15 K	
Stabilizer	Off	On	Off	On	Off	On	Off	On
SNR [dB]	38.5	41.1	24.7	32.2	30.8	31.6	22.9	28.8
SNDR [dB]	36.8	38.9	21.3	30.3	28.6	29.3	19.4	26.5
SFDR [dBc]	-51.4	-57.5	-31.0	-49.6	-36.7	-37.5	-27.1	-35.7
ENOB [bits]	5.8	6.2	3.2	4.7	4.5	4.6	2.9	4.1

7.3 SUMMARY

In this chapter, we have demonstrated a firmware approach to regulate the FPGA's core voltage. By constantly measuring the carry cell delay and acting upon change by decreasing or increasing the amount of running oscillators in a small oscillator farm, the power consumption and consequently the FPGA voltage can be flattened. This technique has been verified and its performance improvement has been quantified on our FPGA ADC system. The ENOB of the ADC was at most 1.5 bits higher, while using the stabilizer.

However, this technique is only capable of stabilizing the FPGA's core voltage (VCC_INT), and not any other voltage on the PCB. Therefore, a more robust and versatile solution can be found in the use of a cryogenic voltage regulator, such as the ones presented in Chapter 3. Since these LDOs are fairly stable, this firmware approach might be more useful for other purposes, for example to avoid security attacks by smoothing the power and heat inside the FPGA. Our recommendation is the use of before mentioned hardware approach to implement a regulator directly on the PCB at cryogenic temperatures.



II

INTEGRATED ELECTRONICS

8

CRYOGENIC MOS & BIPOLAR TRANSISTORS IN STANDARD CMOS PROCESSES

In this second part, we switch from commercially available devices, to custom integrated circuits. Since we can optimize the behaviour of circuits and ensure a low power consumption, while still maintaining the required specifications, CMOS is a better approach to build a large-scale quantum–classical interface. The basis for any such (integrated) circuit is the transistor. Since its inception in 1959, the MOS transistor has become the main element in modern day technology. We are able to reliably integrate billions of these transistors in only a few square millimeters, enabling more complex and advanced electronics than their inventors could have ever imagined in the 1960s.

As the use of CMOS has grown exponentially, people started to look into its use in unconventional environments, such as deep space, high voltage installations or chemical sensors. These environments are subject to, for example, extreme temperatures, radiation and chemical interactions [113]. In this work, we focus on the lower extreme temperatures, namely cryogenics. Since the 1980s, transistors have been studied at temperatures as low as 4 K [83]–[85], [87]–[89]. In Figure 8.1, we summarize the attempts to characterize CMOS transistors over the past four decades [31], [32], [70], [83]–[112], [C12]. During this period, transistor length has shrunk, from over 1 μm [83]–[89] to less than 7 nm [107]–[112], [C12]. In order to understand the implications of operating integrated electronics at low temperatures, a detailed characterization of the particular CMOS process of interest is thus required.

In this chapter, we will focus solely on the low temperature behaviour of transistors, whereas Chapters 9 and 10 will highlight the effects on, respectively, analog and digital circuits. Therefore, we will cover in depth the deviations in the behaviour of transistors in Section 8.1; the main changes are the kink effect, a higher mobility, a higher threshold voltage and a steeper subthreshold slope. The kink effect, elevated current levels at higher drain–source voltages, appears only in older technologies, as highlighted in Figure 8.1. The other three deviations are present in all technologies, but the effects are

Parts of this chapter have been published in IEEE Sensors [C6], the IEEE European Solid-State Device Research Conference [C12] and the IEEE Journal of the Electron Devices Society [J7], [J8].




Figure 8.1: Overview of characterized MOSFETs at low temperatures.

weaker in modern processes, in which the transistors behave more as they do at room temperature.

After highlighting the main deviations in Section 8.1, we continue by comparing the devices in a 40 and 160 nm CMOS processes in Section 8.2. For a deeper understanding over temperature, we focus on subthreshold characteristics of both MOS and bipolar transistors in Section 8.3. From literature, it is known that bipolar transistors suffer from silicon freeze-out and cease operation at roughly 70–90 K [C6], [J7]. A good alternative is the use of MOS transistors operated in dynamic-threshold configuration. Finally, a reflection on bipolar and MOS transistors is given in Section 8.4.

8




Figure 8.2: Cryogenic behaviour of MOS transistors. (a) drain current versus drain-source voltage shows increased mobility at 4 K. (b) drain current versus gate-source voltage shows a steeper subthreshold slope and higher threshold voltage at 4 K. Thick-oxide devices with $W/L = 400/320$ nm implemented in a 160 nm CMOS process.

8.1 CRYOGENIC BEHAVIOUR OF MOS TRANSISTORS

Then when the temperature changes, it is inevitable that the behaviour of the elementary particles also changes. This effect is even stronger at deep-cryogenic temperatures, since these elements move slower and slower, as their thermal energy diminishes to about zero. These phenomena directly impact the functioning of transistors, required to implement cryogenic circuits. We will shortly describe the main deviations with reference to Figure 8.2, in which we show the drain current versus drain and gate voltage at 300 K and 4 K. In this example, we use large, thick-oxide transistors, fabricated in a 160 nm CMOS process, that clearly feature the effects of cooling.

Let us first consider Figure 8.2(a), depicting the transistor's drain current I_{DS} versus the drain voltage V_{DS} . We can directly see two principal effects, namely an increased current (mobility enhancement μ) at cryogenic temperatures and the kink effect. In (b), an increased threshold voltage V_{th} and a steeper subthreshold slope SS can be seen at 4 K.

8.1.1 KINK EFFECT

One of the most prominent effects visible mainly in older technologies and larger devices is the cryogenic kink effect. The drain current of a MOS transistor is elevated at higher drain-source voltages, as shown for the 160 nm thick oxide NMOS transistor in Figure 8.2(a). This effect has been detailed in the past [89], [114], and can be explained through Figure 8.3. At these higher drain-source voltages (large V_{DS}), the electric field in the channel becomes large enough to create electron-hole pairs through impact ionization, and thus for a larger current to flow from drain to source. However, it is inevitable that a fraction of this current will flow to the bulk instead of the source (Figure 8.3(a)),




Figure 8.3: The main cause of the kink effect at cryogenic temperatures can be explained through (a) impact ionization at higher drain voltages causing a non-negligible bulk current, and (c) the bulk resistance to be increased by orders of magnitude, causing a voltage drop on the bulk R_B (b), locally reducing the transistor's threshold voltage and thus increasing its current I_{DS} .

which would cause no issues at room temperature (since bulk and source are at the same potential). But silicon freeze-out, at cryogenic temperatures, increases the bulk resistance by several orders of magnitude (Figure 8.3(c)), causing a non-negligible voltage drop over the bulk R_B and thus local elevation of the bulk potential near the device. Since an elevated bulk potential lowers the transistor's threshold voltage, a jump in drain current is observed at larger V_{DS} and temperatures below 50–70 K.

Although the kink is such a visually striking effect, we only observe it in older technologies (as can be derived from Figure 8.1), with high supply voltages and large devices. This is mainly related to the effects of CMOS scaling, and its associated reduction in supply voltages and transistor dimensions. Smaller transistors require a thinner gate oxide to keep-up device performance, i.e. to reduce the vertical electric field and the transistor's threshold voltage. Nonetheless, electric fields in the channel have increased over the past decades in nanometer devices, causing for example surface scattering in the channel (mobility degradation), and thus reducing impact ionization and related effects. Furthermore, lower supply voltages, usually below 1.2 V, do not enable electrons and holes to reach critical energies to breach the silicon bandgap, and thus prevent the kink effect from occurring at all times. Hence, the kink effect can be safely ignored for modern CMOS technologies (better than 160 nm).

8.1.2 MOBILITY & THRESHOLD VOLTAGE

Two other low temperature effects, present in all devices, partially cancel each other. On one hand, an increase in mobility allows for larger currents to flow in the transistor's channel; on the other hand, an increase in threshold voltage reduces the maximum current, since the device will turn on later and the maximum $V_{GS} - V_{th}$ will be lower.

The mobility in silicon is mainly limited by scattering effects in the lattice, and is dominated by phonon scattering at high temperatures, and by Coulomb or ionized impurity scattering at lower temperatures [115]. Since phonon scattering is generally domi-




Figure 8.4: (a) Electron mobility μ increase in a large NMOS transistor with a W/L of 200/100 μm for various surface electron densities N_s . Redrawn from [115]. (b) Threshold voltage V_{th} increase for a small and large PMOS transistor with a W/L of 0.232/0.16 μm and 2.32/1.6 μm , respectively.

nant at higher temperatures, the effective mobility thus increases with cooling, as shown for a large NMOS transistor in Figure 8.4(a). The electron mobility of this transistor, with a W/L of 200/100 μm , increases by up to one order of magnitude for larger surface electron densities [115].

The threshold voltage, as measured for a small and large PMOS transistor, increases by, respectively, 100 and 250 mV, while lowering the temperature from 300 K to 4 K. Remarkably, the V_{th} increase in the small transistor rolls off, whereas the large transistor's threshold voltage grows more rapidly at lower temperatures. In the majority of modern processes, we see a threshold voltage increase in the order of 100 to 150 mV.

8.1.3 SUBTHRESHOLD SLOPE & HYSTERESIS

The subthreshold slope is a measure of the drain voltage step required to increase the drain current by one decade ($10\times$) for a device working below its threshold voltage. Below threshold, the device behaviour is similar to that of a diode and follows a semi-exponential voltage-current relationship. The ideal slope can be approximated with

$$SS = \ln(10) \frac{kT}{q}, \quad (8.1)$$

with k being the Boltzmann constant, T the absolute temperature and q the electron charge. The ideal SS should thus be linear with temperature and improve in slope from 60 mV/dec. at 300 K to less than 1 mV/dec. at 4 K. The subthreshold slope for a small and large PMOS transistor is shown in Figure 8.5(a); at 300 K, the SS is indeed close to its ideal value of 60 mV/dec., although the small transistor exhibits a slightly higher slope due to short-channel effects. At temperatures below roughly 80 K, the slope rolls off to a minimum of roughly 15 mV/dec., over one order of magnitude degradation compared to the ideal value. This is in agreement with previous studies, reporting a highly non-linear slope below 77 K [88]. The roll-off in SS can be attributed to the non-ideality factor strongly diverging from unity (shown in Subsection 8.3.3) [116].

Another effect that is mostly noticeable for larger devices that exhibit the kink effect is hysteresis. At cryogenic temperatures, the drain current no longer follows the same characteristics for an increasing or decreasing V_{DS} , as shown in Figure 8.5(b) for a large NMOS transistor. This effect is attributed to the much slower release time of electrons trapped in impurities in the transistor's channel [114].

8.2 MOS TRANSISTORS IN 40 & 160 NM

After having highlighted the main changes at cryogenic temperatures, we quickly review the effects visible in the smallest transistors in 40 and 160 nm CMOS technologies, with a W/L of 232/160 nm, respectively, 120/40 nm. The current-voltage relationships for both PMOS and NMOS transistors are shown in Figures 8.6 and 8.7.

Again, the majority of the effects discussed in Section 8.1 can be seen in these transistors. The current levels increase by up to 60%, the subthreshold slope becomes steeper and the threshold voltage shifts up to 150 mV. Although the current levels in NMOS and PMOS transistors in 40 nm are similar, the current in the NMOS increases by 60%, whereas the PMOS shows only a 30% improvement. Similarly in 160 nm, the current increase is about twice in the NMOS compared to the PMOS transistor. For circuit design,




Figure 8.5: (a) Subthreshold slope SS over temperature for a small and large PMOS transistor with a W/L of $0.232/0.16 \mu\text{m}$ and $2.32/1.6 \mu\text{m}$, respectively. (b) Hysteresis observed at cryogenic temperatures in an NMOS transistor with a W/L of $0.232/1.6 \mu\text{m}$. Arrows indicate the forward and backward V_{DS} sweep.




Figure 8.6: Drain current versus V_{DS} for the smallest transistors in 40 and 160 nm CMOS processes. (a) 160 nm NMOS with $W/L=232/160 \text{ nm}$, (b) 160 nm PMOS with the same W/L , (c) 40 nm NMOS with $W/L=120/40 \text{ nm}$, (d) 40 nm PMOS with the same W/L .




Figure 8.7: Drain current versus V_{GS} for the smallest transistors in 40 and 160 nm CMOS processes. (a) 160 nm NMOS with $W/L=232/160$ nm, (b) 160 nm PMOS with the same W/L , (c) 40 nm NMOS with $W/L=120/40$ nm, (d) 40 nm PMOS with the same W/L .

Table 8.1: Main performance figures for the smallest PMOS and NMOS transistors available in a 160 and 40 nm CMOS process.

Device	Temperature	160 nm				40 nm			
		NMOS		PMOS		NMOS		PMOS	
		300 K	4 K	300 K	4 K	300 K	4 K	300 K	4 K
V_{max}	[V]	1.8		1.8		1.1		1.1	
$I_{DS, max}$	[μA]	209	272	105	121	33	54	25	34
$\hookrightarrow$ increase	[%]	-	30.1	-	15.2	-	63.6	-	36.0
V_{th}	[V]	0.47	0.60	0.37	0.53	0.46	0.58	0.49	0.65
SS	[mV/dec.]	88.2	24.6	88.8	43.4	85.1	33.0	87.6	48.2
n_{MOS}	[-]	1.5	29.6	1.5	52.1	1.4	41.5	1.5	60.7

especially standard digital cells, the NMOS/PMOS ratio has thus to be adapted for cryogenic temperatures, to enable a similar maximum current in both transistors (see also Chapter 10).

Despite the scaling of supply voltage from 1.8 V in 160 nm to 1.1 V in 40 nm, the threshold voltage V_{th} shows no such reduction. This limits the operation at low temperatures in moderate and strong inversion, since the device remains in weak inversion

over a much wider voltage range. Furthermore, the ideality factor n_{MOS} of all devices deteriorates at cryogenic temperatures, an indication of the significant deviation from the ideal subthreshold slope. The main performance metrics for these four transistors are summarized in Table 8.1. In the next section, we will further investigate the behaviour of the 160 nm transistors operating in weak inversion, compared to bipolar transistors in the same technology. A similar characterization is performed on the 40 nm transistors for the implementation of the voltage references in the next chapter.

8.3 WEAK INVERSION MOS & BIPOLAR TRANSISTORS

Bandgap reference circuits, used to create a stable reference voltage over temperature (Chapter 9), and bandgap temperature sensors, used to accurately monitor die or environment temperature [117], can be implemented with either bipolar or MOS transistors in weak inversion as core devices. Bipolar transistors, especially substrate PNP transistors (Figure 8.8(a)), readily available in any CMOS process, are generally preferred for their lower spread compared to MOS transistors (Figure 8.8(b)) [117]–[120]. However, BJTs are not well behaved at cryogenic temperatures (< 70 K) due to a significant decrease in their current gain β and an increase in their base resistance R_B .

As an alternative, standard MOS transistors or MOS in dynamic-threshold MOS (DT-MOS) configuration, i.e. a MOS with short-circuited gate and body terminals, can be used as core devices in bandgap references [119] and temperature sensors [121]. In particular, DTMOS are preferred over MOS for their lower spread, better matching, steeper subthreshold slope and operation at lower supply voltage [119], [121], [122], although they have only been characterized down to 77 K [123]–[125].

8.3.1 MEASUREMENT SETUP

8

In bandgap references and temperature sensors, the voltage drop on a diode-connected device (BJT or MOS in weak inversion) is compared or combined with the voltage difference on two diode-connected devices biased at different current densities, which is expected to be proportional to absolute temperature (PTAT) [117]. Since the accuracy of those voltages directly influences the performance of bandgap references and temperature sensors, they are thoroughly analyzed hereafter.




Figure 8.8: Simplified cross section of a substrate PNP (a) and a PMOS (b) in a standard CMOS process.




Figure 8.9: Micrographs of (a) the PNP and (b) MOSFET test structures, indicating the relevant contacts mentioned in Figure 8.8.

The PNP and MOS test structures fabricated in a standard SSMC 160 nm 1P5M CMOS process comprise a pair of substrate PNP transistors arranged in a cross-coupled common-centroid configuration (each of the 4 unit elements with $5 \times 5 \mu\text{m}^2$ emitter area), employing similar layout and dimensions as used in state-of-the-art temperature sensors [120], [126] (Figure 8.9(a)), and MOS transistor pairs with two dummies on the outer side for improved matching (Figure 8.9(b)). To observe the impact of short-channel effects and channel length modulation, PMOS with the same aspect ratio but different length have been fabricated, namely a pair of transistors with a W/L ratio of $0.232 \mu\text{m}/0.16 \mu\text{m}$, referred to as SS (Small-Small), and a pair with $2.32 \mu\text{m}/1.6 \mu\text{m}$, denoted LL (Large-Large). Furthermore, an n-well resistor (W/L ratio of $4 \mu\text{m}/8 \mu\text{m}$ and nominal resistance of $3.5 \text{ k}\Omega$) was fabricated on the same die to characterize the resistivity of the n-well PNP base.

The PNPs have been characterized in an ST-500 (Janis) cryogenic probe station with reference temperature sensor on the sample holder, while the MOS transistors have been mounted on a PCB immersed in liquid helium or in helium vapours with the reference temperature sensor mounted on the PCB (see Figures A.3 and A.4). Electrical characterization has been performed using 2636B (Keithley) source measure units (SMUs).

As depicted in Figure 8.10, the drop on one of the diode-connected devices (V_{BE} for PNP, V_{GS} for MOS and DTMOS) has been characterized while applying a bias current (I_1) and measuring the base current (I_B) for the PNP, while the difference in the voltage drop over two diode-connected devices ($\Delta V_{BE} = V_{BE1} - V_{BE2}$ for the PNP, $\Delta V_{GS} = V_{GS1} - V_{GS2}$) has been characterized while biasing the pair with a current ratio p ($I_1 = pI_2$). For the standard MOS, the drain and gate were shorted to ground, whereas the body was biased with the standard supply voltage (1.8 V). The current was injected into the source. For the DTMOS, the drain, gate and body were shorted to ground, and again current was injected into the source. To improve measurement accuracy in the ΔV_{BE} measurements, the base has been shorted to ground. Mismatch of the PNP and MOS pairs has been averaged using a dynamic element-matching (DEM) technique by averaging the results for $I_1 = pI_2$ and $I_2 = pI_1$, as shown in [127].




Figure 8.10: (a) Schematic diagram of the measurement setup with details of the device under test (DUT) for (b) PNP, (c) MOS and (d) DTMOS characterization. Current is injected in the emitter or source, the collector or drain and gate are always grounded, whereas the base is grounded in the case of ΔV_{BE} measurements, but connected to a current meter for V_{BE} measurements. For MOS measurements, the transistor's body is connected to the standard supply, i.e. 1.8 V in this technology.

8

8.3.2 PNP RESULTS

The collector current $I_C = I_E - I_B$ was computed from the measured base and emitter currents, and plotted in Figure 8.11 as a function of V_{BE} , showing that the exponential




Figure 8.11: PNP Collector current versus emitter-base voltage for temperatures ranging from 7 K to 294 K.

relation between I_C and V_{BE} holds over a wide range of currents and down to approximately 70 K. From the same data, the current gain $\beta = I_C/I_B$ is computed and found to be relatively current independent over more than three decades of I_E for temperatures above 70 K (Figure 8.12). As common for vertical PNPs in modern CMOS technologies [126], [128], the PNPs show a low current gain ($\beta \approx 6$ at room temperature), which drops below unity for temperatures below 130 K. β degrades quasi-exponentially with cooling, which is expected due to the apparent bandgap narrowing associated with the emitter [129].

Device parameters can be extracted by observing that

$$I_C = I_s \left(e^{\frac{qV_{BE,i}}{n_{\text{BJT}}kT}} - 1 \right), \quad (8.2)$$

where $V_{BE,i}$ is the intrinsic base-emitter voltage (i.e. V_{BE} for $R_B = 0$), n_{BJT} is the effective emission coefficient and I_s is the saturation current. Consequently,

$$V_{BE} = V_{BE,i} + R_B I_B = n_{\text{BJT}} \frac{kT}{q} \ln \left(\frac{I_C}{I_s} + 1 \right) + R_B \frac{I_E}{\beta + 1}, \quad (8.3)$$

where I_E is the emitter current, R_B the base resistance and β the current gain.

Applying Eq. (8.3) to a PNP pair biased at current ratio $p = I_{E1}/I_{E2}$, we get

$$\Delta V_{BE} \cong n_{\text{BJT}} \frac{kT}{q} \ln(p) + R_B \frac{p-1}{\beta+1} I_{E2}, \quad (8.4)$$

under the following assumptions: $I_C \gg I_s$, i.e. biasing far from the low-injection region; R_B is current independent; β is current independent, which is valid for a large range of currents as shown in Figure 8.12; accurate PNP matching, which is achieved by proper device layout and dynamic element matching. Figure 8.13(a) shows the measured ΔV_{BE} as a function of the injected emitter current, indicating a stable ΔV_{BE} over a current




Figure 8.12: PNP current gain β versus (a) emitter current and (b) temperature.




Figure 8.13: $\Delta V_{BE} = V_{BE1} - V_{BE2}$ versus (a) I_{E1} and (b) temperature for $I_{E1} = pI_{E2}$ with $p = 3$.




Figure 8.14: Reciprocal of the extracted effective emission coefficient n_{BJT} for the PNP for temperatures ranging from 7 K to 294 K.

range from approximately 10^{-10} – 10^{-5} A and for temperatures above 70 K. ΔV_{BE} is approximately PTAT above 70 K (Figure 8.13(b)) as expected from Eq. (8.4), especially for emitter currents in the medium-low current regime.

Using the measured current gain to fit Eq. (8.4) to the measured ΔV_{BE} (Figure 8.13(b)), the base resistance R_B and the emission coefficient n_{BJT} have been extracted. As shown in Figure 8.14, the emission coefficient is close to unity at room temperature, as expected in CMOS substrate PNPs [127], and is approximately constant over a wide range of emitter currents at least down to 130 K. At lower temperatures, n_{BJT} shows a steep increase and is no longer current independent. The base resistance R_B increases significantly at lower temperatures, as discussed in more details in [C6].




Figure 8.15: (a) Comparison between measured V_{BE} (markers) and a fit against the model expressed by Eq. (8.5) (dashed lines). (b) Difference between fit and measurements showing a steep increase of the fitting error at lower temperatures.

Finally, in order to study the temperature behaviour of V_{BE} , the following standard model is fitted to the measured data [117], [127]:

$$V_{BE}(T) = V_{g0} \left(1 - \frac{T}{T_r} \right) + \frac{T}{T_r} V_{BE}(T_r) - \eta \frac{kT}{q} \ln \left(\frac{T}{T_r} \right) + \frac{kT}{q} \ln \left(\frac{I_C(T)}{I_C(T_r)} \right), \quad (8.5)$$

where V_{g0} and η are constant fitting parameters and T_r is a reference temperature. Fitting Eq. (8.5) to the measurement data over both the temperature range from 111 K to 294 K with $T_r = 294$ K and the current range from 10 pA to 100 μ A results in an average fitting error below 2 mV over the fitting range (Figure 8.15) and in fitting parameters very close to prior work [127] ($\eta = 4.20 \pm 0.4$, $V_{g0} = 1.13 \text{ V} \pm 5 \text{ mV}$, with 95% confidence intervals).

8.3.3 MOS RESULTS

Thanks to the nature of MOS devices, the gate current I_G can be neglected, thus simplifying the MOS current and voltage relationship in the subthreshold region to:

$$I_{DS} = I_s e^{\frac{q(V_{GS} - V_T)}{n_{MOS} kT}} \left(1 - e^{-\frac{qV_{DS}}{kT}} \right), \quad (8.6)$$

in which n_{MOS} is the non-ideality factor. Figure 8.17 shows the measured I_{DS} – V_{GS} curves for the four MOS combinations. The drain current exhibits an exponential behaviour over a wide range of currents, as expected from Eq. (8.6). Although this range is smaller compared to that of the PNP, the exponential relationship holds even at temperatures as low as 4 K. As shown in Figure 8.16, the subthreshold slope SS decreases at lower temperatures, with the SS of DTMOS being lower with respect to MOS and closer to the SS of bipolar transistors as expected from theory [119], [122]. Interestingly, the large DTMOS shows a steeper slope compared to the other devices above 100 K, whereas the small DTMOS improves beyond that down to 4 K.




Figure 8.16: Subthreshold slope for all tested devices.




Figure 8.17: MOS/DTMOS drain current versus gate-source voltage for (a) all tested devices at 4 K and 297 K and (b) the LL DTMOS over the temperature range from 4 K to 297 K.

ΔV_{GS} was measured with the same current ratio $p = 3$ as the bipolar transistor to enable a fair comparison (Figure 8.18). ΔV_{GS} is slightly less current-dependent for larger devices, which can be attributed to the absence of short-channel effects. Mismatch was evident for the small devices when operated without dynamic element matching (not shown), since the ΔV_{GS} could even become negative at the lowest temperatures. For the large DTMOS, ΔV_{GS} is stable, i.e. current independent, over a current range from roughly 10^{-11} – 10^{-8} A, which is two orders of magnitude less than the BJT.

The large DTMOS pair (LL) shows the lowest current dependency of ΔV_{GS} , as highlighted in Figure 8.19. In a similar way as done for the emission coefficient of the PNP (Figure 8.14), its non-ideality factor n_{MOS} is shown in Figure 8.20. n_{MOS} is close to its ideal value of 1 for low currents and for temperatures down to roughly 130 K. At lower temperatures, n_{MOS} increases similarly as in the bipolar device and deviates significantly from ideality.




Figure 8.18: ΔV_{GS} versus I_{DS2} with $I_{DS1} = pI_{DS2}$ with $p = 3$, for (a) all tested devices at 4 K and 297 K and (b) the LL DTMOS over the temperature range from 4 K to 297 K.




Figure 8.19: ΔV_{GS} versus temperature with $I_{DS1} = pI_{DS2}$ with $p = 3$, DTMOS LL.




Figure 8.20: Reciprocal of the extracted non-ideality factor η_{MOS} for the DTMOS LL.

Similarly to Figure 8.15, V_{GS} is plotted versus temperature in Figure 8.21 for the large DTMOS. To clearly assess the effect of replacing the BJTs with (DT)MOS devices in standard bandgap references and temperature sensors, the same model used for V_{BE} Eq. (8.5) is used to fit the measured V_{GS} . However, fitting over both the temperature range and the bias current range gives significant deviation, as shown in Figure 8.21(c, d), already below 200 K (13-mV average fitting error for the large DTMOS, 26 mV for the large MOS) over the complete temperature range when using $V_{g0} = 0.47$ V and $\eta = 1.37$. On the contrary, fitting the same model against the individual curves results in a lower fitting error (3.4 mV for the large DTMOS, 4.6 mV for the large MOS), but a considerable spread in both V_{g0} and η is observed. V_{GS} extrapolated at 0 K, i.e. V_{g0} in Eq. (8.5), appears to




Figure 8.21: V_{GS} for the DT MOS LL (markers) compared to (a) a fit of Eq. (8.5) over the complete dataset and (b) individual fits of Eq. (8.5) per each I_{DS} value (shown as dashed lines). (c, d) show the error obtained between measurement data and both fits.

8

be roughly exponentially dependent on the biasing current I_{DS} . Over the current range from 10 pA to 100 nA, V_{g0} of the large DT MOS shifts by approximately 100 mV from 0.43 V to 0.53 V. For the large MOS, this shift is even larger (from 0.79 to 0.94 V). η spreads from 3.6 to -0.5 in the DT MOS. Although V_{GS} follows the model down to 4 K fairly well, the apparent current dependence of the extrapolated V_{GS} at 0 K may lead to a lower performance of DT MOS-based references, especially when employing a single point trim and assuming a fixed V_{g0} . This is also confirmed by the larger spread of DT MOS-based circuits when compared to their BJT counterparts. For example, Sourì [120], [121] reports an inaccuracy of $\pm 0.48^\circ\text{C}$ versus $\pm 0.15^\circ\text{C}$ for, respectively, a DT MOS- and BJT-based temperature sensor.

Finally, the non-ideality factor (for the MOS and DT MOS) and the emission coefficient (for the BJT) of all devices is shown in Figure 8.22. The n_{MOS} in DT MOS devices is closer to the BJTs emission coefficient n_{BJT} compared to their MOS counterparts. In all cases, however, n increases rapidly at temperatures below 100 K. Another advantage of DT MOS is its lower threshold voltage, which, especially at cryogenic temperatures, can

be exploited to built sub-1-V ultra-low power circuits with a higher performance than their MOS counterparts.

8.4 SUMMARY

We demonstrated that bipolar transistors freeze-out at temperatures of roughly 70 K, deteriorating their current gain and increasing their base resistance. Therefore, below 90 K, one has to rely on something functional for the implementation of any circuit, and luckily MOS transistors behave fairly well over the complete temperature range of interest (down to 4 K), and even down to 100 mK [31], [32], [100], [C12]. Although the performance of MOSFETs is reduced on one side, as follows from the increased non-ideality factor and higher threshold voltage, it does increase on the other side, with a higher mobility and steeper subthreshold slope at 4 K. In practice, several cryogenic effects, such as hysteresis and the kink, can be easily ignored for modern CMOS processes, allowing device models to be the same as those used for room temperature, with altered parameters. Although not discussed in this work, device modelling is essential to enable more precise simulations for cryogenic circuit design, and several works [94]–[97], [100]–[102], [104], [J8] have addressed this issue by proposing both physical and mathematical models that cover all aspects of cryogenic transistor behaviour. Instead of relying on simulations, we will continue in the next two chapters with an experimental investigation of both analog and digital circuits under cryogenic conditions.




Figure 8.22: Non-ideality factor (for the MOS and DTMOS) or emission coefficient (for the BJT) for all tested devices.

9

BASICS FOR A CRYOGENIC VOLTAGE REFERENCE

With the characterized transistors from the previous chapter, we can design analog and digital circuits, and can predict their behaviour at cryogenic temperatures. As an example, we will design a specific analog building block, namely a (bandgap) voltage reference. For the design of other analog circuitries at cryogenic temperatures, we kindly refer the interested reader to [68]–[75], [130] for analog-to-digital converters, [21], [131], [132] for digital-to-analog converters and [133], [J5] for low-noise amplifiers.

CMOS systems, especially analog circuitries, often require an accurate reference voltage to be generated, preferably on chip. This reference voltage can be generated with a bandgap-reference circuit, comprising either bipolar or MOS transistors in weak inversion as its main embodiment. Although bipolar transistors are preferred for their lower process spread [117]–[120], we revealed in the previous chapter that they are not well-behaved below roughly 90 K. A significant decrease in current gain β and increase in base resistance R_B limit the operation of these devices at lower temperatures. MOS transistors, on the contrary, are operable at 4 K, and show very similar exponential characteristics over a wide temperature range in the subthreshold region.

In the past, several bandgap implementations focussed on the liquid nitrogen temperature range down to 77 K, employing either SiGe HBTs [134], [135] or standard CMOS (without bipolars) [103], [136]. Also, some commercial voltage references were found to be unstable from 130 K and below [137]. The only bandgap voltage reference circuit that operates at 4 K and even down to 700 mK employs SiGe HBTs [138]. The circuit consumes less than 130 μ W at 700 mK and achieves a temperature coefficient better than 0.2 mV/K.

In this chapter, we demonstrate the first cryogenic voltage reference operating down to 4 K in a standard 40 nm CMOS process. We compare the core devices in a similar fashion as for the 160 nm devices of Chapter 8 in Section 9.1. The characterization of the voltage references is presented in Section 9.2, followed by a short recap in Section 9.3. A voltage reference based on DT MOS transistors is fully operational over the 4 K to 300 K temperature range, with a temperature coefficient of 0.85 mV/K.

Parts of this chapter have been published in the IEEE Solid-State Circuits Letters [J10].

9.1 BANDGAP CORES

A typical bandgap circuit exploits the exponential current–voltage relation of diode-connected devices, such as bipolar transistors (Figure 9.1(a)), which can be expressed as

$$I = I_s \left(e^{\frac{V_q}{n k T}} - 1 \right) \quad (9.1)$$

where I and V are the device current and voltage drop and, for a BJT, I_s is the saturation current, $I = I_C$, $V = V_{BE}$, and $n = n_{\text{BJT}}$. Since the voltage drop on such device is approximately complementary-to-absolute-temperature (CTAT), a reference voltage can be obtained by combining it with a proportional-to-absolute-temperature (PTAT) voltage that can be generated as the voltage difference between two diode-connected devices biased at different current densities:

$$\Delta V = n \frac{kT}{q} \ln(p) \quad (9.2)$$

where, for BJTs, $\Delta V = \Delta V_{BE}$, $n = n_{\text{BJT}}$ and p is the emitter-area ratio of the two diodes in Figure 9.1(a), biased with the same current. The reference voltage is then constructed as:

$$V_{REF} = V + \alpha \Delta V \quad (9.3)$$

where α is a constant chosen to ensure a minimum temperature coefficient for V_{REF} .

MOS in dynamic-threshold MOS (DTMOS) configuration (Figure 9.1(b)), i.e. MOS with short-circuited gate and body terminals. For DTMOS in weak inversion, Eq. (9.1), (9.2) and (9.3) hold with $I = I_{DS}$, $V = V_{GS}$, p being the ratio between the aspect ratio of the two diodes and $n = n_{\text{MOS}}$ being the non-ideality factor.

We implemented all structures, including the pairs shown in Figure 9.1 and the circuits discussed in Section 9.2, in a 1P6M TSMC 40 nm process. After bonding on a PCB, the devices were tested in a dip-stick immersed in liquid helium or helium vapours over the temperature range from 4 K to 300 K (see Figures A.3 and A.4). A DT-670 (LakeShore Cryotronics) reference temperature diode was mounted in close proximity to the sample for accurate temperature measurements, and the electrical characterization was performed using 2636B (Keithley) source measure units (SMUs).

9




Figure 9.1: Schematic of the bandgap core devices placed in a 1: p ratio, employing either (a) BJTs or (b) DTMOS. The devices are biased with equal currents, i.e. $I_{E1} = I_{E2} \equiv I_{DS1} = I_{DS2}$.

The BJTs are composed of unit elements with a $5 \times 5 \mu\text{m}^2$ emitter area (comparable to those used in state-of-the-art temperature sensors and bandgap references [120], [126]), and $p = 12$. For the DTMOSs, thick-oxide transistors were used with a $W/L = 4 \mu\text{m}/0.5 \mu\text{m}$ as unit element (since we showed that larger transistors have a lower current dependency on ΔV_{GS}), and $p = 36$.

For the bipolar device, I_C (Figure 9.2(a)) follows the expected exponential relation (following Eq. (9.1)) over a wide range of currents, but only above roughly 100 K, as expected due to increasing base resistance and reduced current gain at lower temperatures. On the contrary, the I_{DS} (Figure 9.2(b)) in the DTMOS follows an exponential relation over the complete temperature range down to 4 K, but only up to $1 \mu\text{A}$, since it enters moderate inversion at higher current levels. Although following an exponential trend, the $V-\ln(I)$ slope of both devices deviates from an ideal diode behaviour, as indicated by n_{BJT} and n_{MOS} deviating from unity and being temperature dependent (Figure 9.3), especially for lower temperatures.

When biasing the BJTs with a constant current, V_{BE} is approximately CTAT down to 100 K over a wide range of currents (10 pA–100 μA), as shown in Figure 9.4(a). Below this temperature, both V_{BE} and ΔV_{BE} (Figure 9.4(b)) show an abrupt change in temperature coefficient. The BJT bandgap voltage V_{REF} is composed according to Eq. (9.3) by numerically combining V_{BE} and ΔV_{BE} with the optimal α , achieving the lowest variation over the 100–300 K temperature range (Figure 9.4(c)). The reference voltage is close to the theoretical silicon bandgap of 1.22 V and exhibits a low temperature coefficient above 100 K (± 40 ppm/K), but is unsuitable for cryogenic applications.

Using the same approach as for the BJTs, the V_{REF} for the DTMOS has been generated in Figure 9.5. With respect to the BJT's case, the reference voltage shows a better stability down to 4 K, but a much larger temperature coefficient (up to ± 830 ppm/K on the 4–300 K range). This is due to the drastic increase of n_{MOS} and the consequent flattening of ΔV_{GS} at lower temperatures. However, the analysis shows that a much flatter reference is obtained when biasing the transistors in strong inversion, i.e. for currents




Figure 9.2: (a) I_C – V_{BE} for a BJT with $5 \times 5 \mu\text{m}^2$ emitter area; (b) I_{DS} – V_{GS} for a DTMOS with $W/L = 4 \mu\text{m}/0.5 \mu\text{m}$.




Figure 9.3: Emission coefficient and non-ideality factor represented as (a) $1/\eta_{BJT}$ and (b) $1/\eta_{MOS}$ versus biasing current.




Figure 9.4: Generating a reference voltage with BJTs: combining (a) the CTAT component V_{BE} with (b) the PTAT component ΔV_{BE} following Eq. (9.3) with (c) the optimal α per each bias current to minimize the variations over the temperature range from 100 K to 300 K.

above approximately $1\ \mu\text{A}$. In this case, the generated reference voltage depends also on the threshold voltage of the devices. The combination of the temperature dependence of several parameters, including the hole mobility and the threshold voltage, cancel out to the first order, as in typical MOS-based references, thus achieving a temperature coefficient as low as $\pm 20\ \text{ppm/K}$ (Figure 9.5). Contrary to intuition, this points to strong inversion references as better candidates for cryogenic operation.

In this analysis, we assumed the current to be constant over temperature. Since in a practical circuit the current is typically PTAT and the temperature coefficients of integrated resistors may be relevant, full voltage references with BJTs and with DTMOS in strong inversion have been implemented and experimentally characterized, as shown in the next section.

9.2 VOLTAGE-REFERENCE CIRCUITS

The complete bandgap circuits are shown in Figure 9.6; the chip micrograph with superimposed layouts is shown in Figure 9.7. A current–voltage mirror using NMOS ($M_{4,5}$) and PMOS ($M_{6,7}$) with a W/L ratio of $3/0.4\ \mu\text{m}$ and $2/0.4\ \mu\text{m}$, respectively, set the voltage across R_1 to ΔV , so that the same current $\Delta V/R_1$ flows through the two core de-




Figure 9.5: Generating a reference voltage with DTMOS: combining (a) the CTAT component V_{GS} with (b) the PTAT component ΔV_{GS} following Eq. (9.3) with (c) the optimal α per each bias current to minimize the variations over the temperature range from 4 K to 300 K.




Figure 9.6: Schematic of the cryogenic bandgap voltage reference employing either (a) BJTs or (b) DTMOS.




Figure 9.7: (a) Chip micrograph of all devices. (b, c) Zoom-in with superimposed layout of the voltage references employing either (b) BJTs or (c) DTMOS (including decoupling capacitors).

vices ($M_{1,2}$ or $Q_{1,2}$). A third branch (M_8 , R_2 , M_3 or Q_3) is used to sum V and ΔV using $\alpha = \frac{R_2}{R_1}$. The resistor values have been chosen large enough for a low power consumption and small enough to ensure a bias current much larger than the sources of inaccuracy, such as noise and leakage currents. We employed polysilicon resistors, since they show

the least spread over temperature; we measured less than 10% resistance variation from 300 K to 4 K for an n-poly resistor in the adopted 40 nm process with a nominal value of 4.5 k Ω at 300 K. Although this variation will affect the biasing current, and hence V_{REF} , this effect is negligible with respect to the steep variation in slope for the core devices at lower temperatures.

For the BJT-based implementation, $R_1 = 3$ k Ω and $R_2 = 7.6$ k Ω , resulting in an $\alpha = 2.53$. Since $\Delta V_{BE} = n_{BJT} kT/q \ln(p) \approx 64$ mV at 300 K, the current in each branch is roughly 21 μ A, reducing to 7 μ A at 100 K. As shown in Figure 9.8, the current through each branch is indeed PTAT down to roughly 180 K, at lower temperatures it slowly rolls off until the sudden increase around 100 K, denoting the clear breaking point of the bipolar transistors. The measured reference voltage is relatively flat above 100 K as shown in Figure 9.9(a). Since cryogenic device models were not available at design time, the chosen α is lower than the optimal $\alpha \approx 7$ from Figure 9.4(c), thus leading to a negative temperature coefficient above 100 K. Behaviour over several devices is presented in Figure 9.9(c, d). The measured power-supply rejection ratio (PSRR, Figure 9.11) is 26 dB at 300 K, and reduces to 22 dB at 100 K. The noise versus temperature is shown in Figure 9.12(a), which does not improve much with cooling. The lowest noise was measured




Figure 9.8: Current through each of the three branches in both implementations over temperature. The sudden increase in current at 100 K of the BJT-based bandgap denotes the breaking point of the bipolar transistor and the clear deviation from the ideal $\Delta V/R_1$ current consumption.




Figure 9.9: Measured reference voltage V_{REF} versus temperature for (a) the BJT and (b) the DT MOS implementations.




Figure 9.10: (a) V_{REF} over temperature for several devices trimmed in post-processing by subtracting the offset to the average reference voltage at 300 K, i.e. the dashed lines in (b) ($V_{DD} = 3$ V). (b) Untrimmed V_{REF} at 300 K ($V_{DD} = 3$ V).

close to 100 K; since the device is not working well below this temperature, the noise increases significantly.

In the DTMOS case, $R_1 = 6.6$ k Ω , $R_2 = 8$ k Ω and thus $\alpha = 1.2$. By comparing V_{REF} in Figure 9.4(c) with the measurements, we estimate $I_{DS} \approx 40$ μ A and $\Delta V_{GS} \approx 260$ mV at 300 K, and $I_{DS} \approx 13$ μ A and $\Delta V_{GS} \approx 86$ mV at 4 K. The measured current in each branch is again shown in Figure 9.8, revealing that the current is roughly PTAT down to 100 K, below which it rolls off, similarly to the roll-off of ΔV_{GS} in this temperature range. The circuit (Figure 9.9(b)) properly operates down to 4 K, albeit V_{REF} still drops by 200 mV from 300 K to 4 K due to α being larger than the optimum value. However, V_{REF} changes by only 10 mV (for a 3 V supply) at cryogenic temperatures between 4 K and 100 K, with the temperature coefficient improving from 1.6 mV/K at 300 K to 0.1 mV/K at 4 K. This behaviour is fairly flat over several devices (Figure 9.10(a, b)); the spread in V_{REF} improves from 32 mV (1σ) at 300 K to 22 mV (1σ) at 4 K. The PSRR (Figure 9.11) is 23 dB and fairly stable over temperature. These results have been achieved thanks to core devices operat-




Figure 9.11: Measured reference voltage V_{REF} versus supply voltage V_{DD} for both implementations at (a) 300 K and (b) 4 K. The dashed lines indicate the slope at the optimal supply voltage of 3 V.




Figure 9.12: Output noise versus temperature for (a) the BJT and (b) the DT MOS implementations at various frequencies. (b) Output noise power spectral density of the DT MOS-based circuit at various temperatures.

ing in strong inversion, while operation in weak inversion would have suffered from the roll-off in ΔV_{GS} shown in Figure 9.5(b). Figure 9.12(c, d) show the noise of the DT MOS-based implementation over both temperature and frequency. There is a striking peak in the noise level around 100 K, whereas the lowest noise value is achieved around 50 K. At lower temperatures, the noise increases, most likely to the increasing flicker noise at deep-cryogenic temperatures. The performance is summarized in Table 9.1.

9.3 SUMMARY

In this chapter, we have shown an example of an analog circuit operating at low temperatures. The DT MOS-based reference is functional over the complete temperature range, with a temperature coefficient better than 0.1%/K between 4 K and 300 K, proving for the first time that CMOS-based references can operate at deep-cryogenic temperatures. Although we cannot yet meet state-of-the-art performance metrics, this is a first step to-

Table 9.1: Summary of the measured performance at 300 K, 100 K and 4 K for both voltage references at $V_{DD} = 3$ V.

Temperature		BJT		MOS		
		300 K	100 K	300 K	100 K	4 K
V_{REF} (avg.)	[V]	0.96	1.18	1.02	0.82	0.81
V_{REF} (1σ)	[mV]	45	97	32	21	22
$V_{REF,trim}$ (1σ)	[mV]	0	68	0	18	17
P_{DD}	[μ W]	189	47	368	153	132
$\Delta V/\Delta T$ (avg.)	[mV/K]	-1.1	-5.5	1.5	0.3	0.1
PSRR	[dB]	26.2	21.9	23.1	27.8	23.4
T range	[K]	100–320		4–320		
$\Delta V/\Delta T^*$ (box)	[mV/K]	-1.15		0.85		

* $\Delta V/\Delta T$ is computed using the box method over the entire temperature range and over all samples.

wards the implementation of such circuits with state-of-the-art performance in the near future.

Many other analog circuits have been demonstrated in literature to operate at low temperatures, but the majority somehow suffers from cooling effects. The absence of optimized cryogenic device models is one of the main limitations in the design of these systems, and should be prioritized for further studies.

10

OPERATION OF DIGITAL CIRCUITS UNDER CRYOGENIC CONDITIONS

Thanks to the nature of digital circuits, a signal is either 0 or 1; the impact of cryogenic temperatures affects its performance, but not its functionality. Take, for example, an inverter, the simplest digital element: its output is 0 when the input is 1, and vice versa. This behaviour is not affected with cooling, but, for example, the speed of the 0–1 transition is altered. Echoing the cryogenic investigation of transistors, digital elements have been studied in the past [41], [86], [90], [98]–[100], [139]–[150], [C12], [J5]. The majority of these investigations addresses the changes observed in an inverter and/or in ring oscillators. Inverter testing reveals mostly DC properties (e.g. gain), whereas ring oscillators indicate changes in AC (e.g. delay). In this case, the latter is more interesting, as it will show us the magnitude of the expected speed-up of digital electronics operating at lower temperatures. In Figure 10.1, we show this speed-up versus the technology node, indicating a clear trend of reduced speed-up for newer technologies. However, in contrast to the abundance of transistor characterizations, there is insufficient knowledge of the operation of digital circuits in, especially, modern technologies.




Figure 10.1: Speed-up with respect to room temperature of CMOS circuits operated at low temperatures at their nominal supply voltage. The speed-up of the 28 nm devices is extracted from the FPGA measurements reported in Chapter 5.

Therefore, in this chapter, we study the behaviour of digital elements in both 40 and 160 nm CMOS processes. The DC characteristics of a 160 nm inverter are presented in Section 10.1. From the inverters, we implement various ring oscillators to obtain the inverter delay and thus the cryogenic speed-up in both 40 and 160 nm technologies (Section 10.2). Similar to the delay-change in the combinatorial logic, the rise- and hold-time of flip-flops are affected. For a better understanding, we implement a shift register, containing just flip-flops, and conduct a short study on the cryogenic effects in Section 10.3.

All devices were bonded on a PCB and measured over the temperature range from 4 K to 300 K with a dip-stick in either liquid helium or the helium vapours. An abstract representation of such a setup is shown in Figures A.3 and A.4. DC characterization was performed with 2636B (Keithley) source measure units, whereas AC measurements were done with a WaveMaster 813Zi-B (LeCroy) oscilloscope. A DT-670 (LakeShore Cryotronics) temperature diode was mounted in close proximity to the samples for accurate temperature registration.

10.1 INVERTER IN 160 NM

The CMOS inverter is the basis for any digital design, but also usable as an analog amplifier. Its output voltage is the inverse of the input with a very strong transition, and hence a high gain. To get a feeling of its cryogenic performance, we first approximated the behaviour of a 160 nm inverter from the single device measurements presented in Figure 8.6. These current–voltage relationships are again shown in Figure 10.3(a, b) with the NMOS and PMOS sized to match those of the standard cell inverter from Figure 10.2(a). Ideally, the NMOS and PMOS transistor are balanced to transition exactly at half the supply voltage $V_{DD}/2$. Therefore, the PMOS is twice the size of the NMOS, to scale the device currents appropriately (Figure 10.3). The intersection points of the overlapping curves, where the I_{DS} in both devices biased at the same V_{GS} is equal, are indicated and mapped on the approximated voltage transfer characteristics of an inverter built with




Figure 10.2: (a, c) Schematics of CMOS inverters in 160 and 40 nm CMOS technologies with the transistor dimensions as used in the standard cell library. (b, d) Corresponding standard cell layouts (not drawn to scale).

these transistors in Figure 10.3(c). From this approximation, we can directly see a steeper and slightly shifted slope, indicating that the gain of the inverter increased and that the NMOS became stronger. Figure 10.3(b) confirms that the NMOS is indeed a bit stronger and shows that the maximum current drawn in the inverter is slightly increased at cryogenic temperatures. The input voltage at maximum gain V_M is a bit higher (around 1 V) compared to the ideal value of $V_{DD}/2$ (0.9 V), an indication that the sizes of the NMOS and PMOS are not properly matched.

The properties of the standard cell inverter, with the layout as shown in Figure 10.2(b), will be slightly different, due to for example differences in layout and edge effects (mismatch), compared to the single device measurements. The voltage transfer characteristics




Figure 10.3: (a, b) Current–voltage relationships of the sized 160 nm NMOS and PMOS transistors (according to Figure 10.2(a)) at both 300 K and 4 K (derived from Figure 8.6). The intersection of drain currents at fixed V_{GS} indicates the allowed operating points of the CMOS inverter implemented with these two transistors. (c) Approximated voltage transfer characteristics of an inverter operating at the nominal supply voltage of 1.8 V.




Figure 10.4: (a) Voltage transfer characteristics and (b) corresponding voltage gain of the 160 nm inverter at 300 K and 4 K and various supply voltages.

tics of the standard cell inverter in this 160 nm CMOS process are shown in Figure 10.4(a) for both 300 K and 4 K. The transition voltage V_M is a bit lower than $V_{DD}/2$, indicating that the NMOS is stronger than the PMOS; this is in contrast to the observation in Figure 10.3(c), where V_M is a bit higher. V_M is shown for various temperatures and voltages in Figure 10.5(a), indicating that V_M shifts to lower voltages at lower temperatures, except for low supply voltages $V_{DD} < 1$ V. This agrees with the observations on transistor level (Figure 10.3(c)). At lower supply voltages, we see another effect kicking in: the increase in threshold voltage of both transistors. A higher gate voltage is required to move the transistors from weak to strong inversion, thus significantly slowing down the transition. Furthermore, the current drawn from the supply at V_M steeply diminishes for these lower supply voltages and lower temperatures (Figure 10.5(b)). In contrast to the increase in maximum current expected from Figure 10.3(b), the current at $V_{DD} = 1.8$ V actually decreases slightly at lower temperatures.

Although not common, the inverter can be used as amplifier when biased at V_M , where the gain is maximum and most linear. Figure 10.4(b) shows the gain for the 160 nm inverter at 300 K and 4 K. At cryogenic temperatures, the gain improves by more than 100% at medium supply voltages (0.9 V). At standard supply voltages (around 1.8 V) the gain still improves by 15–20%. Again we see the shift in V_M to lower voltages for large V_{DD} and to higher voltages for low V_{DD} .

Finally, we quickly note the noise margin improvement of the inverter, defined as the voltage difference at the -1 gain points:

$$NM_H = V_{OH} - V_{IH}; \quad (10.1a)$$

$$NM_L = V_{IL} - V_{OL}, \quad (10.1b)$$

with V_{OH} and V_{OL} being the output high and low voltages and V_{IH} and V_{IL} being the input high and low voltages at the -1 gain points.

The high noise margin NM_H improves from 0.73 to 0.81 V (300 K and 4 K), whereas the NM_L increases by only 20 mV from 0.49 to 0.51 mV at the nominal supply voltage




Figure 10.5: (a) Inverter's voltage at maximum gain V_M , ideally $V_{DD}/2$, and (b) maximum drawn current from the supply during the transition.

(1.8 V). A similar improvement is observed for lower voltages until 0.9 V; further down, the threshold voltage limits the noise margin improvement.

10.2 RING OSCILLATORS IN 40 & 160 NM

In order to assess the dynamic performance of the digital logic, we implemented ring oscillators in both 40 and 160 nm CMOS processes. A schematic view of these oscillators is shown in Figure 10.6. The 160 nm RO consists of 2703 stages (2702 inverters and one NAND gate to enable the oscillator); the 40 nm RO has only 143 stages. Both circuits feature a large output buffer in order to drive a coaxial cable and corresponding 50 Ω




Figure 10.6: Schematic of the (a) 160 and (b) 40 nm ring oscillator. A NAND gate was used to enable the oscillation, and an output buffer was designed to drive a 50 Ω load.




Figure 10.7: Layout of the 40 nm ring oscillator, the oscillation path is indicated in black. Only the first two stages of the output buffer *B* are shown.

load (to avoid reflections). The layout of the 40 nm oscillator is shown in Figure 10.7, on which we can identify the oscillation path and the single inverters previously shown in Figure 10.2(d).

The dynamic performance is evaluated by the frequency change in the oscillator. The nominal frequency at room temperature was measured at 4.8 MHz and 194 MHz for the 160 and 40 nm RO, respectively. We extracted the unit cell delay by dividing the oscillation frequency with twice the number of stages: 38 ps and 18 ps, respectively. The inverter delay is plotted for various temperatures and over its supply voltage in Figure 10.8. At the nominal supply voltage (1.8 V), the inverter delay improves by 20%, from 38 to 31 ps, in the 160 nm process. The 40 nm RO speeds-up by only 11%, from 18 to 16 ps. At the higher supply voltages (> 1 V), this speed-up is fairly much constant. However, around 1 V there is a cross-over point (also visible in the characteristics of the 28 nm FPGA-based ring oscillators in Chapter 5), where the delay is approximately constant over the complete temperature range. Below this voltage, we see the dramatic effect of the threshold voltage increase, raising the delay by orders of magnitudes. As the transistors operate in subthreshold, their delay increases exponentially with lower supply voltages.

We can thus conclude that cryocooling benefits the performance of our circuits, but only above roughly 1 V. For newer technologies, this might be a limitation, since supply voltages are lower in each new process generation (already around 1 V and below for 28 nm and smaller technologies), unless the threshold voltage can be reduced by back-gate modulation as in FD-SOI technologies. Besides the speed-up, also the power consumption increases (Figure 10.9), although not much in the 40 nm process; in 160 nm, it increases with over 25%, to achieve a speed-up of only 20% at 1.8 V.

The performance of both oscillators is summarized in Table 10.1 at their nominal supply voltage. Besides the speed-up of inverters, other dynamic parameters might affect circuit performance, such as the rise and fall times. However, these are hard to char-




Figure 10.8: Inverter delay in (a) 160 nm and (b) 40 nm CMOS technologies versus supply voltage at various temperatures, extracted from the ring oscillators depicted in Figure 10.6.




Figure 10.9: Relative power consumption with respect to room temperature for the (a) 160 nm and (b) 40 nm ring oscillators.

Table 10.1: Ring oscillator performance characteristics at their nominal supply voltage at 300 K and 4 K.

	160 nm		40 nm	
Temperature	300 K	4 K	300 K	4 K
Stages	2703		143	
Nominal V_{DD} [V]	1.8		1.1	
P_{DD} [mW]	97	124	0.86	0.87
↪ increase [%]	-	27.8	-	1.2
f_{OSC} [MHz]	4.8	6.1	194	219
Inverter delay [ps]	38.3	30.6	18	16
↪ decrease [%]	-	20.1	-	11.1

acterize for a single device due to our long coaxial lines, hence its investigation through a shift register.

10.3 SHIFT REGISTER IN 40 NM

Besides the behaviour of the logic cells, the registers play an important part in any circuit. They allow to store information and thus to clock a circuit. One cell that allows this is the data flip-flop (D flip-flop): only at the rising edge of a clock signal, the input is propagated to the output. A simple shift register, allowing to propagate an input signal through a series of D flip-flops, was fabricated in 40 nm CMOS. Its schematic is shown




Figure 10.10: Schematic of the 40 nm shift register consisting of 56 data flip-flops, the output buffer is the same employed in the ring oscillator of Figure 10.6(b)




Figure 10.11: Layout of the 40 nm shift register. An H-tree clock distribution network is employed, indicated in red. The datapath is shown in black.

in Figure 10.10 and consists of 56 stages followed by the same buffer used for the ring oscillator in Figure 10.6(b). The clock network is formed as an H-tree. Although this is not the best solution for a shift register, in which the clock is ideally routed in opposite direction to the data, this is a common structure for large-scale digital circuits, thus better representing the average use-case. The clock network and datapaths are indicated in the circuit layout, shown in Figure 10.11.

The circuit was cooled down and the maximum frequency at which the circuit could be operated without glitches in the output is reported in Figure 10.12(a). Similar to the results for the oscillator, the behaviour at higher voltages (close to the nominal 1.1 V) is very similar at both 300 K and 4 K. The maximum operating frequency $f_{\max}$ improves slightly by 4% (725 MHz versus 700 MHz at 1.1 V). The main deviation is again visible for

10




Figure 10.12: (a) Maximum operating frequency $f_{\max}$ and (b) power consumption P_{DD} of the shift register operated at 300 K and 4 K.


lower voltages below 0.9 V. $f_{\max}$ drops much more rapidly at 4 K compared to 300 K, an effect that should again be attributed to the increased threshold voltage.

10.4 SUMMARY

The performance of digital logic at cryogenic temperatures is very close to room temperature. Following the trend to smaller nodes (newer technologies), we can barely see an improvement. In the old days, the logic speed could increase by over 50%; in recent technologies, 28 and 40 nm, the speed-up is close to non-existent ($< 10\%$). The main performance limitation at cryogenic temperatures is the increase in threshold voltage. Since the devices operate over a wider supply voltage range in the subthreshold region, the switching speed is severely reduced. This, however, gives us the opportunity to sacrifice speed for power in pursuit of ultra-low power digital logic. Especially at cryogenic temperatures, with limited power budgets, this might be a valuable solution for large-scale digital controllers operating at slow speeds and with minimal power consumption.

III

A QUANTUM READ-OUT PLATFORM



11

QURO – CRYOGENIC QUANTUM READ-OUT

In this last part, we will combine all previously obtained knowledge to implement a cryogenic quantum read-out system. This chapter discusses its design and preliminary results, whereas measurements on the complete chain are shown in Chapter 12. The basis for the system is the Artix 7 (Chapter 5), combined with voltage regulators (Chapter 3), the clock generator (Chapter 4) and the ADC (Chapter 6). However, there is a crucial element missing, as we can see from Figure 1.3, namely an amplifier. As the signal, coming from the qubits, has such low levels, the ADC is not capable of digitizing it directly; it has to be amplified first. Therefore, we implemented two amplifiers, one built from discrete components (Part I), the second integrated in a standard 160 nm CMOS process (Part II).

We will first briefly introduce spin qubits and their operation in Section 11.1, followed by the principle of the proposed read-out method and our cryogenic approach in Section 11.2. The final required components for our system, the amplifiers and the directional coupler, are discussed and characterized in Section 11.3 and 11.4, respectively. A short summary is given in Section 11.5, from which we will continue in the next chapter.

11.1 SPIN QUBITS

Since this work is not aimed at describing the physics behind spin qubits in detail, we would like to refer the interested reader to more detailed descriptions in [4]–[6], [151], [152] and constrain ourselves to the basic principles required to understand the read-out procedure. A spin qubit can be constructed in various ways, such as the nuclear spin of atoms, or the spin of electrons. Spin qubits realized in semiconductors are promising, as their fabrication process can benefit from well established technologies, e.g. CMOS. To construct a spin qubit in a semiconductor, such as silicon or gallium arsenide, one has to isolate and manipulate a single electron. To do so, a quantum well is constructed with a very thin layer (around 10 nm) of e.g. silicon in a SiGe-Si layer stack. In this thin layer, a two dimensional electron gas is formed (2DEG), in which electrons are free to move, except for the vertical direction. To isolate a single electron, gate electrodes are deposited on top of the semiconductor and biased with a voltage to deplete the 2DEG below them (Figure 11.1). If the gates are biased properly, a small island is created in the depleted 2DEG, i.e. a quantum dot, which can be populated with a single electron. The




Figure 11.1: Abstract view of a spin qubit in a quantum dot. A 2D electron gas (2DEG) is partially depleted with gate electrodes on top of the semiconductor, creating tiny spots in which a single electron can be trapped, i.e. the quantum dot. A sense quantum dot is placed beside the main dot in order to read its state.

spin of the electron can now be manipulated through the gates, and, in case of multiple quantum dots, multi qubit interactions can be established.

To read the state of the qubit, one usually relies on a sense device, either another quantum dot, or a single electron transistor / quantum point contact, combined with the presence of a strong magnetic field. In this magnetic field, the energy of the electron spin is separated, thanks to the Zeeman effect, i.e. an electron with spin up will have a higher energy than an electron with spin down. This read-out principle is described in detail in [153] and is schematically drawn in Figure 11.2 for both spin up and spin down. The read-out gates are manipulated to lower the tunnel barrier to the electron reservoir and to start the read process. The electron will remain inside the quantum dot if it has a spin down $|\downarrow\rangle$, since it has not enough energy to tunnel back into the electron reservoir. In case of spin up $|\uparrow\rangle$, it will move out to the electron reservoir in the 2DEG, and a low-energy electron with spin down will tunnel back into the quantum dot.

This tunnel effect, which only occurs for an electron with spin up, can be sensed with a sensing device, either by measuring the DC current generated by the tunnelling electron, or by sensing the charge variation of the dot in this short period. Reading the DC current from the sensing device has been proven to be reliable, but it is slow and suffers from $1/f$ noise. Another approach is based on RF-reflectometry. Instead of relying on DC parameters, this technique relies on the reflection of high-frequency signals on the quantum dot gates. Our electronic system is designed to read the qubits through this method, as detailed in the next section.




Figure 11.2: Read-out of a spin qubit in a semiconductor. (a) In case of a spin down, the electron cannot tunnel through the barrier back into the 2DEG electron reservoir. (b) In case of a spin up, the electron with high-energy in the quantum dot will fall back into the electron reservoir and a low-energy electron with spin down will tunnel into the quantum dot. This tunnel effect can be sensed with a sensing device.

11.2 QUBIT REFLECTOMETRY READ-OUT

Reflectometry relies on a signal that is reflected against an unmatched impedance. In this case, the unmatched impedance will be the sense quantum dot during the electron tunnelling event. Since the impedance of the quantum dot, represented by the tunable resistor in Figure 11.3, is in the order of several hundred k Ω , and the remainder of the system is preferably implemented with a standard coaxial impedance of $Z_0 = 50 \Omega$, an LC matching network is implemented to translate the high-impedance quantum dot to 50Ω . A slight change in quantum dot impedance, upon a qubit state change, will be translated to a slight change in the coaxial impedance.

A wave, travelling through a coaxial line with a typical impedance of Z_0 , will be partially reflected at any impedance mismatch, in our case the quantum dot ($Z \neq Z_0$) at the end of the transmission line. The amplitude of the reflected signal can be easily determined as

$$A_R = A_{IN} \frac{Z - Z_0}{Z + Z_0}, \quad (11.1)$$

which gives no reflection for $Z = Z_0$. For example in a typical coaxial line of 50Ω , 1% of the incoming wave is reflected for an impedance mismatch of 1Ω .

Currently, the reflectometry setup is divided into several parts at various temperatures. A directional coupler is placed at mK temperature, a low-noise amplifier (LNA) is situated usually around 4 K for noise reasons, and the remainder of the electronic stack is placed at room temperature: a second amplification stage, a demodulator (to downconvert the RF signal into baseband) and the analog-to-digital converter (ADC). We envision the integration of all these components on a single PCB that can operate at 4 K. This system is schematically drawn in Figure 11.3, with the majority of the components placed at 4 K, connected on one side to the quantum dot at mK and on the other side to the host interface at room temperature.




Figure 11.3: Schematic view of the proposed cryogenic reflectometry setup. At base temperature (mK), the quantum dot (QD) is placed together with the matching network on a single PCB. The remainder of the electronic stack is situated at 4 K on a second PCB: the directional coupler, both amplifiers and the FPGA-based ADC.




Figure 11.4: Quantum Read-Out (QuRO) PCB containing the read-out chain from Figure 11.3: directional coupler, LNA, AMP, ADC and FPGA, combined with the relevant components for proper operation: decoupling, LDOs, crystal oscillator and connectors.

This electronic system integrated on a single PCB is shown in Figure 11.4 with all relevant components highlighted. The read-out chain from Figure 11.3 is identified at the bottom, from the directional coupler towards the FPGA. The performance at 4 K of the directional coupler and the two amplifiers will be demonstrated in the next sections, whereas performance of the complete system is investigated in the next chapter.

11.3 AMPLIFIERS

Our amplification chain consists of a CMOS LNA and a discrete silicon-germanium amplifier (SiGe AMP). The design of the LNA is fully detailed in [J5], therefore we only briefly explain its working principle below. The SiGe AMP is based on [154], [155] with a few modifications, such as an additional amplification transistor.

11.3.1 CMOS LNA

The LNA [J5] is fabricated in a 160 nm CMOS technology and is capable of cancelling the noise of the input matching transistor. Noise cancelling [156] is crucial as the system needs to be matched to a coaxial line (connecting to the qubits) with $50\ \Omega$ characteristic impedance; in other circuits, the noise of the LNA would be dominated by this input matching transistor. A simplified schematic of the LNA is shown in Figure 11.5, the current noise of the input transistor M_1 flows partially through R_f and the drain of M_1 . The first is amplified and inverted by M_2 , whereas the second is only copied through M_3 to the output. Tuning the gain of M_2 and M_3 allows to cancel the inverted and non-inverted noise signals. The input matching of the LNA mainly depends on the transconductance of M_1 : $Z_{IN} = 1/g_{m1}$.

As the gain of only this single stage is limited (around 25 dB), a series of additional gain stages (3 totalling 30 dB) were implemented, together with a $50\ \Omega$ output driver to match again to a coaxial line. From simulations, the expected performance at 4 K was derived, resulting in a 55 dB total gain and a 0.01 dB noise figure, while consuming 75 mW.

11.3.2 SiGe AMP

Following the LNA, a discrete second stage amplifier was constructed on the board, which is based on previous designs [154], [155]. They demonstrated that a circuit, employing two NXP BFU725 SiGe transistors, can deliver 30 dB of gain and a very low noise figure < 0.1 dB at 15 K. Our circuit implementation consists of a cascade of three SiGe amplification transistors (NXP BFU725), shown as $Q_{1,2,3}$ in Figure 11.6, which are AC coupled to one another. The transistors are biased with the resistors on the base and collector to deliver a gain of around 15 dB per stage. Since we used a standard FR-4 dielectric for the PCB, parasitics somewhat limit the bandwidth compared to the implementations of [154], [155], which use a Teflon/Rogers high-frequency dielectric.

Again input and output were matched to $50\ \Omega$ in order to drive coaxial interfaces, at least as far as basic simulations could reveal. As there was no cryogenic model available for the SiGe transistors, we only employed limited room temperature simulations combined with the knowledge from [154], [155] to estimate 4 K performance.




Figure 11.5: Simplified schematic of the noise-cancelling CMOS low-noise amplifier, implemented in a 160 nm technology by R. Incandela [J5].




Figure 11.6: Simplified schematic of the SiGe amplifier, used as second stage after the LNA. Its design is based on [154], [155].

11.3.3 COMPARISON

We have extensively tested both amplifiers separately to demonstrate its performance at 4 K compared to room temperature. Since performance also relies on the design of the PCB (decoupling, parasitics etc.), deviations are visible in the performance of the LNA as presented in [J5], compared to the results shown here. The same measurement principle was used for the analysis of both amplifiers to establish a fair comparison. The setup, employing a dip-stick with relatively long cables, as shown in Figures A.3 and A.4, is not ideal for noise measurements due to de-embedding errors, but it was handled as accurately as possible.

The LNA was programmed (over a simplified SPI interface) from the onboard FPGA in the desired current bias regime, whereas other voltages to both LNA and AMP were supplied from outside with DP832 (Rigol) power supplies. The FPGA was powered with the onboard LDOs connected at room temperature to 2636B (Keithley) source measure units, and its clock was generated with the cryogenic 100 MHz crystal. The firmware of the FPGA is uploaded through JTAG, and FPGA communication occurs through a low-speed UART interface.




Figure 11.7: Gain versus frequency for (a) the CMOS LNA and (b) the SiGe amplifier at 300 K and 4 K. Both gain and bandwidth are enhanced at cryogenic temperatures.

The gain of both amplifiers was measured with a ZNB40 (Rohde & Schwartz) vector network analyzer (VNA), and is shown in Figure 11.7 at both room and cryogenic temperatures. In the LNA (a), both gain and bandwidth are improved at 4 K, from 44 to 58 dB, and 420 to 570 MHz. These values are close to simulation, except for a slightly lower bandwidth due to parasitics. The power consumption (77 mW) is fairly constant over temperature, since the device is programmed to be in a similar current bias condition at both temperatures.

The AMP (b) shows a lower gain improvement compared to the LNA, an indented effect as we reduced the operating voltage from 2.5 to 2 V at 4 K. The bandwidth improves significantly from 490 to 860 MHz, and all of that with a power reduction of > 50% (35 down to 16 mW).

The LNA features many configuration options, all resulting in a slightly different power consumption. All options are shown in Figure 11.8(a) together with the gain in that bias condition. However, the noise cancellation occurs only in a few of those points, rendering the majority of them not to be usable. The AMP has only the supply voltage as tuning parameter, therefore only a single line is visible at each temperature in Figure 11.8(b). The gain flattens out at higher supply voltages, so it is barely worthwhile to spend power on a small gain increase beyond 20–30 mW.

In order to read the tiny signals from the qubits, the noise of the amplification chain should be as low as possible, ideally less than 0.1–0.2 dB. We measured the noise figure of both amplifiers at 300 K and 4 K through the Y-factor method [157], see Figure 11.9, using an FSW43 (Rohde & Schwartz) spectrum analyzer. This method relies on the measurable difference in noise between two points. In order to set those points, we used a noise source with an excess noise ratio (ENR) of roughly 5.5 dB at room temperature, and we amplified it with 30 dB for cryogenic measurements. The principle is as follows. When the noise source is switched off, it acts as a 50 Ω (matching) resistor, generating a noise power density of $P_s = 10 \log_{10}(kT/0.001) = -174$ dBm/Hz at $T = 290$ K. The noise source switched on generates a noise level of -167 dBm/Hz. These values can be transformed into their equivalent noise temperatures of 290 K and 1400 K, respectively. The




Figure 11.8: The maximum gain obtained for (a) the CMOS LNA and (b) the SiGe amplifier at 300 K and 4 K. The LNA is configurable in many current bias conditions, whereas the AMP is only tunable with its supply voltage.




Figure 11.9: Basic principle of the Y-factor method. The output noise at two known input noise temperatures is measured, usually the noise of a 50 Ω matching resistor and the noise of a diode with fixed ENR. The x-axis intercept is equal to the negative noise temperature of the LNA/AMP, before de-embedding.

extrapolated intercept point with the x-axis gives the negative noise temperature of the amplifier $-T_{\text{LNA}}$. However, since we used fairly long cables with non-negligible loss, this value has to be de-embedded by correcting for these losses. At 4 K, 30 dB of attenuation was placed at 4 K, resulting in a cold noise level of -192 dBm/Hz. The hot noise was generated by first amplifying the switched on diode, followed by the attenuation at 4 K, giving a noise level of -171 dBm/Hz.

The extracted noise figure of the two amplifiers, acquired with the Y-factor method and by de-embedding the long cables, is shown in Figure 11.10. The second stage SiGe AMP can be seen to perform slightly better with a lowest noise figure of 1.25 dB at 300 K versus 1.39 dB for the LNA. Both amplifiers exhibit the same improvement of 1.1 dB. The LNA shows a noise degradation at lower frequencies due to a change in input capacitance and the corresponding change in ratio to the AC coupling cap. The SiGe AMP has




Figure 11.10: Noise figure versus frequency for (a) the CMOS LNA and (b) the SiGe amplifier at 300 K and 4 K. A lower noise figure, at 4 K 0.15 versus 0.3 dB, was obtained with the SiGe AMP.

a wide bandwidth in which the noise level is close to the minimum; the bandwidth of the LNA, on the contrary, is much smaller.

We should note that, since the noise of the second stage is actually lower than that in the LNA, it would be ideally used as the first stage. However, at design time, we expected the noise of the LNA to be better (as was derived from simulations), hence it was placed as first stage.

The performance of both amplifiers is summarized in Table 11.1. The second stage SiGe amplifier outperforms the LNA in every aspect: lower noise, higher bandwidth, more gain and lower power consumption. However, the LNA was one of the first to be designed in a standard CMOS process that operates at 4 K, hence its performance was a bit more unpredictable compared to a roughly known discrete SiGe design. Furthermore, the LNA, since it is an integrated circuit, would allow for a tighter integration with the remainder of the read-out chain (ADC, down converter etc.).

11.4 DIRECTIONAL COUPLER

For the reflectometry, a component that splits the microwave signals is required. Conventionally, a directional coupler is used, which couples a room temperature read-out pulse towards the qubits, while the reflected signal is directed towards the read-out chain. For our read-out platform, we opted to co-integrate a directional coupler from Mini-Circuits (ADC-10-4) on the PCB. A simplified schematic view of this coupler is shown in Figure 11.11, with an indication of the relevant ports. The room temperature coaxial line is connected to $RF_{COUPLED}$, which redirects the signal to RF_{IN} (connected to the qubits). The reflected signal on RF_{IN} is extracted at RF_{OUT} and routed towards the LNA. The coupled line has to be terminated on the PCB with a matching resistor $R = Z_0 = 50 \Omega$.

Table 11.1: Performance characteristics of the LNA and the SiGe amplifier at 300 K and 4 K.

Temperature		CMOS LNA		SiGe AMP	
		300 K	4 K	300 K	4 K
Operating voltage	[V]	1.8	1.8	2.5	2
Power consumption	[mW]	76	77	35	16
Gain	[dB]	44	58	56	65
Bandwidth	[MHz]	424	571	490	860
Noise figure	[dB]	1.39	0.29	1.25	0.14




Figure 11.11: Simplified schematic of the Mini-Circuits ADC-10-4 directional coupler. There are three signal ports, RF_{IN} , RF_{OUT} , $RF_{COUPLED}$, and one port that is terminated with $R = Z_0 = 50 \Omega$ (RF_{TERM}).




Figure 11.12: Performance of the Mini-Circuits ADC-10-4 directional coupler over frequency, (a) mainline insertion loss from RF_{IN} to RF_{OUT} , (b) coupling between $\text{RF}_{\text{COUPLED}}$ and RF_{IN} , (c) leakage from $\text{RF}_{\text{COUPLED}}$ into RF_{OUT} .

Since it is a passive three terminal device (ignoring the terminated output), there are three transfer functions, shown in Figure 11.12 at room and liquid helium temperature. The insertion loss (a), i.e. the attenuation of a signal from RF_{IN} to RF_{OUT} , reduces at lower temperatures, although there is a significant increase in loss at lower frequencies. Also the coupling (b) improves at lower temperatures, but the directivity is severely affected by cooling (c). This last effect cannot be explained for, since no information is available on the actual implementation of this commercial device. The directivity is important as it avoids leaking of the strong RF signal directly into the read-out chain.

11.5 SUMMARY

Here and in all previous chapters, we have demonstrated components one-by-one to work at 4 K with relatively stable performance. The amplifiers, all the way up to the digital controller, have been integrated in a single PCB. This single platform would allow to replace the complete conventional electronic stack, and operate at 4 K. However, there are two more steps that need to be taken in order to use this system for quantum experiments: characterization of the full chain and the adaptation to a vacuum environment in a dilution fridge. Although each individual component is operational, the question remains how they influence each other, especially for the amplification chain with extremely high gain (> 100 dB). These aspects will be covered in depth in the next chapter.

12

THE READ-OUT CHAIN

As a final step towards the completion of our read-out system, we have characterized the completed read-out chain from the previous chapter, and transferred it into a dilution refrigerator. In this chapter, we will address all problems that we have faced along the way of operating this small system in a harmonious way. The first step was to operate both amplifiers together. This turned out to be challenging, thanks to the very large gain of both amplifiers together fitted into a small PCB. The next step was to digitize an amplified signal with the FPGA-based ADC and assess the smallest measurable signal amplitude. That is an important measure, since we can compare it directly to the amplitude that we expect to receive from the quantum devices. We summarize our integration findings in Section 12.1.

In order to co-operate our electronic system with a quantum device, the last step after completion of the chain is to move from our dip-stick solution to a dilution refrigerator that can not only produce our wanted temperature of 4 K, but also the mK temperatures required for the qubits. Since our PCB was designed with a dip-stick in mind, it was not at all optimized to function in the vacuum environment of a fridge, demanding an adapted thermalization structure. Combined with the fairly large power consumption of both FPGA, ADC, regulators, oscillator and amplifiers, we would expect an elevated PCB temperature compared to a liquid cooling solution. Section 12.2 lists our initial results of the system operating in a BluForce dilution fridge.

12.1 INTEGRATION

We have shown that all components can operate individually, but the question remains on how they will influence each other. In the following subsections, we will gradually extend the read-out, first with the two amplifiers, later with the FPGA-based ADC.

12.1.1 CMOS LNA + SiGe AMP

From the previous chapter, we can roughly estimate the expected behaviour of the combination of the two amplifiers. The total gain should be the multiplication of the two individual gains in the linear domain, or simply an addition in decibels. The noise figure should be dominated by the first stage, i.e. the CMOS LNA, although negligibly elevated by the second stage noise.

The expected gain is shown in Figure 12.1(a); it is especially high at 4 K with over 120 dB, due to the increase in gain of both amplifiers. As expected, the bandwidth is re-




Figure 12.1: (a) Expected gain versus frequency of both amplifiers connected in cascade, when summing the measured gain values from Figure 11.7. (b) Measured gain versus frequency of both amplifiers connected in cascade, the gain is clearly limited slightly below 90 dB at both temperatures.

duced, but the 550 MHz at 4 K is still close to the 570 MHz of the LNA alone. However, the measured gain (Figure 12.1(b)) shows a significant deviation from our expectations. Both gain and bandwidth are clearly limited; the gain is saturated close to 90 dB and the bandwidth is restricted to 250 MHz. From the fairly similar shape at both temperatures, we presume that a parasitic capacitance between input and output is the main cause of these limitations, forming a feedback loop amplifier configuration. At the design stage, we were indeed expecting these issues, due to the tight space in which we fitted a large gain. A potential solution is an RF-shield, physically separating the various gain stages from each other and thus preventing parasitic couplings. The current design allows for such a shield around the LNA, but this didn't solve any of the before mentioned issues. Coupling is most likely still occurring between the output of the second stage amplifier and the unshielded directional coupler's ports. When we reduced the total gain, by lowering the supply voltage of the second stage amplifier $V_{DD,AMP}$, we found that the discrepancy between measurement and expectation disappears at roughly 1.1 V, giving a total gain of 64 dB (Table 12.1). Apparently, the parasitic coupling is prevented when the total gain does not exceed 65 dB; we indeed found this to be the only useful configuration in combination with our ADC (see Subsection 12.1.4).

Table 12.1: Expected and measured performance characteristics of the cascade of LNA and SiGe amplifier at 300 K and 4 K. We were not able to measure the noise figure of both amplifiers together.

Temperature	Expected				Measured			
	300 K		4 K		300 K		4 K	
$V_{DD,AMP}$ [V]	2.5	1.8	1.1	2.0	2.5	1.8	1.1	2.0
Gain [dB]	96	87	64	123	86	82	64	82
Bandwidth [MHz]	335	325	286	550	250	247	225	256
Noise figure [dB]	1.40		0.29		-		-	

Although the gain plots from Figure 12.1(b) look acceptable, stability is a concern due to the parasitic feedback coupling. Therefore, we again tried to measure the noise figure of the amplifiers together. Unfortunately, we were unable to discriminate between T_{hot} and T_{cold} , thus indicating that there is indeed an issue. The performance metrics of both amplifiers together are summarized in Table 12.1, comparing measurements to our expectations derived from the previous chapter.

12.1.2 ADC

We have discussed the performance of the ADC at length in Chapter 6, but those measurements were not targeted at a quantum system. Therefore, before combining the ADC with the amplifiers, we take another look at its standalone performance. For these measurements, we assumed that we will be reading a signal from the quantum dot at a frequency of $f_c = 133.33$ MHz (in line with the target between 100 and 200 MHz), and we assumed a bandwidth of ± 10 MHz around f_c . Our system is not well suited to sample signals at these high frequencies, due to the six interleaved 100 MHz channels and interference with the clock frequency used on the RC -ramps. These performance limitations can unfortunately not be alleviated in the current setup and would require another type of ADC.

A sinusoidal input signal was generated with a DG4202 (Rigol) and AC-coupled into the ADC from the test point IN_{ADC} in Figure 12.2. Our ADC also requires an external bias of $V_{CC_INT}/2 = 1.25$ V, which we provided together with the sinusoid to the ADC with an onboard bias-tee. This DC-offset is required to shift the input signal to be roughly in the middle of the RC -ramps, which span from 0.9 to 1.7 V.

The ADC was first calibrated at both operating temperatures, following the procedure described in Chapter 6. However, the ADC can only be reached through either the AC-coupled input line (> 10 MHz) or the DC-coupled bias-tee input, thus complicating the calibration routines and limiting their accuracy. The low frequency calibration signals were provided through the bias-tee DC line, filtering their response, whereas our medium frequency signals (1–2 MHz) were damped too much through both inputs, and had thus to be replaced by lower frequency signals.

The 133.33 MHz sinusoid was digitized at both 300 K and 4 K, and the corresponding frequency domain plots were obtained with an FFT as shown in Figure 12.3. (a) shows the full spectrum at 300 K for a signal with 0 dBm input power, revealing the expected harmonics together with the 100 MHz clock signal. Since we know the frequency of operation, in this case at f_c , we can limit our digital bandwidth, for which we choose ± 10 MHz




Figure 12.2: The 4 K part of the read-out chain from Figure 11.3, indicating the various test points and the DC bias connection for our ADC.




Figure 12.3: (a) Frequency spectra of a sampled sinusoid with $f_c = 133.33$ MHz and $P_{IN} = 0$ dBm at 300 K. (b, c) Zoom-in of the spectrum in a ± 10 MHz bandwidth around f_c , indicating the SFDR in this window at both 300 and 4 K. We can see a clear degradation at cryogenic temperatures. Frequency spectra were obtained from 48,000 samples, while using a Blackman-Harris window of length 2^{14} shifted over the samples.

(the expected frequency distance between qubits), as shown in (b, c). The spectrum in this small bandwidth looks fairly clean at 300 K (b) with an SFDR of 53 dBc, but the 4 K spectrum (c) is distorted due to both calibration limitations and the increased instability of some FPGA components at cryogenic temperature. Note that we operated the system at 4 K, in contrast to the 15 K operation used in Chapter 6, resulting in a further performance reduction.

We swept the input power and measured the response from the ADC in this ± 10 MHz bandwidth. As shown in Figure 12.4(a), the measured output power scales as expected with the input at both temperatures, and similar trends can be observed for the SNR and SFDR in (b, c). Both metrics indicate a performance reduction at cryo, in line with Figure 12.3(c), but the SNR is close to its 300 K value for lower input powers. Assuming




Figure 12.4: Performance of the FPGA-based ADC on the QuRO board, operating at 1.2 GSa/s, sampling a 133.33 MHz sinusoid at 300 K and 4 K. (a) Measured output power versus input power of the sine in dBm. (b, c) SFDR and SNR in a ± 10 MHz bandwidth around f_c versus input power. Each point is the average of 10 ADC measurements.

an acceptable SNR is at least 10 dB, the lowest detectable input signal requires a power of -33 dBm at 300 K, versus -30 dBm at 4 K.

12.1.3 SiGe AMP + ADC

We extended the chain by connecting the SiGe amplifier in front of the ADC. The input signal is now passed through IN_{AMP} in Figure 12.2, without applying further changes with respect to the previous subsection and utilizing the same calibration datasets. The amplifier was powered with 2.5 V at 300 K and 2.0 V at 4 K as listed in Table 11.1.

Again, a similar analysis as per Figure 12.4 has been executed and is shown in Figure 12.5. From (a), showing the output power versus the input power, we extracted the gain values in (b). The gain of the ADC alone is close to 0 dB, whereas the gain of the combination of ADC and amplifier is close to 60 dB. The average gain at 300 K is 57 dB, very close to the gain reported in Table 11.1 (56 dB). However, at 4 K, the gain is expected to increase to 65 dB, on average we measured only 53 dB in this amplifier + ADC combination, implying the combination is not performing well. This is confirmed from (c, d),




Figure 12.5: Performance of the FPGA-based ADC combined with the second stage SiGe amplifier on the QuRO board, operating at 1.2 GSa/s, sampling a 133.33 MHz sinusoid at 300 K and 4 K. (a) Measured output power versus input power of the sine in dBm, comparing the ADC only and the ADC+AMP. (b) Calculated gain; the gain of the AMP is close to 60 dB, as expected from Table 11.1 at 300 K, but lower instead of higher at 4 K. (c, d) SFDR and SNR in a ± 10 MHz bandwidth around f_c versus input power for the ADC only and the ADC+AMP.

showing the SNR and SFDR in the ± 10 MHz bandwidth around f_c ; the performance at 4 K is deteriorated compared to measurements of the ADC alone. The SNR is up to 20 dB worse than at room temperature, and large spurious tones inside the ± 10 MHz bandwidth limit the SFDR even more.

Assuming again a minimal required SNR of 10 dB, the lowest measurable input power is -83 dBm at 300 K and -70 dBm at 4 K. The additional amplifier allowed us to measure a signal up to 50 dB lower compared to using the ADC alone, however the performance reduction from room to cryogenic temperature is 13 dB, significantly larger than the 3 dB reduction observed in the previous subsection.

12.1.4 CMOS LNA + SiGe AMP + ADC

The last step in the chain is connecting the LNA, therefore we switched the input signal to the IN_{LNA} port in Figure 12.2, thus including the directional coupler by default. The LNA is programmed from the onboard FPGA in the same conditions used for Table 12.1, whereas the SiGe amplifier was set to three different $V_{DD,AMP}$ values to eliminate high-gain issues. A tone of 133.33 MHz was again generated and passed through the amplification chain, resulting in the frequency spectra as shown in Figure 12.6 with $V_{DD,AMP} = 2.5$ V. We compare the results obtained with a signal with input powers of -60 and -80 dBm, in between which the only difference would ideally be a decrease in signal strength of 20 dB (assuming the amplifiers are not clipping). Although the strength of the signal in (b) is indeed lower, the noise is substantially increased, and some spurious tones are even larger than the signal. These results indicate that the system became highly non-linear and behaves very unexpectedly for different input powers.

To quantify this statement, we show the performance metrics at 300 K under the three bias conditions in Figure 12.7 and compare them to the system without the LNA connected. At higher input powers, the system is clipping, as can be seen from the constant output power for various input powers in Figure 12.7(a). This agrees with the expected higher gain of the two amplifiers (expected 90 dB at 2.5 V), but according to (b) the extracted gain is never even close to 80 dB. Even more surprising, the gain of the




Figure 12.6: Frequency spectra of a sampled sinusoid with $f_c = 133.33$ MHz after being amplified by both amplifiers at 300 K. (a) is obtained with an input power of -60 dBm, whereas (b) is taken with -80 dBm. In contrast to an expected performance reduction of 20 dB, the signal is severely degraded at lower input powers, i.e. the system is highly non-linear.




Figure 12.7: Performance of the FPGA-based ADC combined with the second stage SiGe amplifier and the CMOS LNA on the QuRO board, operating at 1.2 GSa/s, sampling a 133.33 MHz sinusoid at 300 K. (a) Measured output power versus input power of the sine in dBm, comparing the ADC+AMP to the ADC+AMP+LNA, the supply voltage of the SiGe AMP is swept to find a stable operating point. (b) Calculated gain; the gain of the AMP+LNA is far from the measured value reported in Table 12.1, except for the 1.1 V case. (c, d) SFDR and SNR in a ± 10 MHz bandwidth around f_c versus input power for the ADC+AMP and the ADC+AMP+LNA.

1.1 V system is higher than the 1.8 and 2.5 V cases (for low input powers), but it is close to previous measurements around 70 dB. Even more worrisome are the SNR and SFDR in (c, d), the 1.1 V case is pretty clean and shows indeed a performance improvement over the system without LNA, albeit the 10 dB SNR point shifts by only 2 dB, from -83 to -85 dBm. In the two other configurations, the performance is deteriorated; note especially the large step at -75 dBm. Clipping is apparently preventing the system from oscillating, or it forces the system to oscillate at f_c . This causes the performance to appear reasonable at high input powers, whereas it is not really so.

These results imply a serious issue with the combination of the two amplifiers; only when the gain is intentionally lowered, the performance is acceptable compared to a system with a single amplifier. This is in line with our previous findings that the expected gain and measured gain of the two amplifiers line up only for a low $V_{DD,AMP}$, as detailed in Table 12.1. Since there is no real benefit in using both amplifiers, for simplicity, power reduction and stability, we recommend the use of only a single amplifier, i.e. the SiGe

AMP would be preferred in that case. The main problem is the lowest measurable input power, which is not low enough to detect the signal from the qubits.

12.2 TOWARDS A DILUTION FRIDGE

Parallel to the measurements on the full chain presented in the previous section, we had to find a way to operate our QuRO board in a dilution fridge in order to allow for potential measurements with a quantum device, which needs to be operated at mK temperatures. Since the board is designed in a compact fashion for operation in a dip-stick, it is not directly suitable to be placed in the vacuum environment due to missing thermalization options. Copper heat-sink structures were designed and connected to the board with silver paste, as shown in Figure 12.8, covering both the FPGA and the LNA, i.e. the main sources of heat. A copper block on the bottom of the PCB enables the thermal link to the 4 K plate in the fridge.




Figure 12.8: (a) QuRO and (b) its adaptation for measurements in a dilution fridge, i.e. in vacuum instead of liquid helium. A copper heat sink is connected with silver paste on top of both the FPGA and the LNA to improve the thermal link to the cold (4 K) plate in the fridge, thermalization structures designed by A. Corna.




Figure 12.9: (a) FPGA power consumption and resulting die temperature at 4 K, measured in liquid helium (from Figure 5.10) and in a vacuum dilution fridge environment. (b) Total FPGA power consumption versus the number of activated oscillators running in the reconfigurable fabric. (c) Relative oscillation frequency of the main oscillator, i.e. the first oscillator to be turned on.

With the temperature diode in the FPGA, we characterized the self-heating and the corresponding quality of the thermal link to the cold plate by slowly increasing the FPGA power consumption with an oscillator farm. These results are compared to those reported for liquid helium in Subsection 5.2.6 and are presented in Figure 12.9(a). In a vacuum environment, the temperature rises much faster compared to a liquid cooling solution; when burning 400 mW, the temperature in the fridge is over $2\times$ compared to that in liquid helium.

The temperature's settling time is close to 20 s irrespective of the power consumed by the FPGA, as shown in Figure 12.10 for two realistic power consumption scenarios. In the first case, the FPGA is configured to program the LNA on the board. This implementation consumes little resources and thus burns little excess power compared to the FPGA being idle. The ADC configuration on the contrary burns a substantial amount of energy of > 800 mW, resulting in an FPGA temperature in the fridge close to 35 K.

In principle, these higher temperatures would not be a concern for the performance of the FPGA, which should actually improve. However, there are two related issues: stability of the system and the noise level of the amplifiers. Regarding system stability, the clock crystal is heated up and enters a regime where the output amplitude is severely decreased, see the 25 to 150 K temperature range in Figure 4.6(b). Unfortunately, this lower amplitude is not enough for the IO of the FPGA to switch, leaving the FPGA without a running clock. To solve this issue, we reverted to provide a room temperature clock through two RF lines (LVDS).

We also observed an instability in the LDO-FPGA combination, when the board was installed in the dilution fridge. While increasing the number of running oscillators in the farm, we expect the frequency of the oscillators to decrease slightly, due to increasing voltage drop and heat. Indeed, the frequency of the main oscillator in the farm (Figure 12.9(c)) goes down, but, in contrast to liquid helium measurements, the frequency decrease is more rapid, especially beyond 250 running oscillators. A similar characteristic was present at 300 K in our opened fridge, which should thus ideally be similar to 300 K dip-stick measurements. Therefore, we cannot attribute the effects to vacuum or




Figure 12.10: (a) Temperature measured from the internal FPGA diode under various conditions of operation with the corresponding power consumption indicated. (b, c) The settling time of the temperature (to 95% of the stable value) is close to 20 s.

thermalization issues, but they turned out to be caused by the impedance of our DC-lines ($3\ \Omega$) resulting in a larger voltage drop on the line than the LDO can compensate for. Operating at a higher input voltage alleviated most of these issues.

The second issue is the heat distribution from the various components, mainly from the FPGA and LDOs, over the PCB. The amplifiers will thus be operating at elevated temperatures, increasing their noise figure. We recharacterized the gain and noise figure from the CMOS LNA in Figure 12.11. The gain is slightly affected, but we have to note that the device under test is not the same in liquid and vacuum measurements, hence some variation is expected. This does, however, not explain the difference in shape at especially low and high frequencies. It seems that the positions of the poles and zeros have shifted, something we cannot easily attribute to the operation in the fridge. The noise figure is also worsened with respect to liquid helium measurements, and increases by roughly 0.15 dB (21 to 32 K). To observe the influence of the FPGA's heat on the LNA's noise performance, we measured the noise figure in the two extreme power consumption cases from Figure 12.10. Although there is a slight increase in noise figure when the ADC is running, the impact is minimal, indicating that the thermal link between LNA and FPGA is fairly weak. Also, the noise generated by the FPGA is not affecting the noise figure obtained from the LNA.

12.3 SUMMARY

In this final chapter, we have combined the various parts of the read-out chain and assessed its performance in a dilution refrigerator. Although the components tend not to operate well together, we can learn quite a few important lessons of operation of electronics under various conditions. In this case, we observed an unexpected limiting factor on the gain and corresponding performance of the two amplifiers when combined. Most




Figure 12.11: Comparison of the gain and noise figure of the CMOS LNA at 4 K as measured in liquid helium or the dilution fridge. The gain (a) is slightly lower when placed in the fridge, whereas the noise figure (b) is elevated. Both effects can be attributed to higher temperatures observed when operating in a vacuum environment. (b) The noise figure has been measured under the FPGA working conditions from Figure 12.10, i.e. with the FPGA used to only program the LNA, or used to also run the ADC resulting in higher temperatures.

likely a parasitic coupling is limiting the gain to around 90 dB, instead of the predicted 120 dB at 4 K. Although by itself this would not be a problem, a further combination with the ADC showed that the system is unstable for lower input powers and appears to be oscillating. Also the ADC is not perfect for the given task, since its performance is limited at higher input frequencies due to the interpolating RC -ramps. We can thus see many parts that need to be improved, and together with the poor thermalization of this first design, we can expect to do better in the future. A better optimized system, with proper thermalization structures, better shielding between amplification stages and an improved ADC for high frequencies, would allow to demonstrate the read-out of real spin qubits in a proper way.

13

CONCLUSION

Quantum computing promises to revolutionize current computing methods. The expected exponential speed-up of certain quantum algorithms over their classical counterparts is the main motivation that drives both academic and industrial research all over the world. To arrive at a viable quantum computer, however, a long road is still ahead of us. Many of the parts essential to such a system are in a premature state, from the quantum devices, the electronics, the cooling infrastructure, to the quantum algorithms. In this work, we explored a small part of this massively complex system, focussing on a potential solution for operating the electronics at low temperatures. Since the quantum bits require extremely low temperatures to exhibit quantum phenomena, a large-scale system will be limited by its surrounding infrastructure, such as the interconnects, cooling power and multiplexing. We see the co-integration of quantum and classical devices as the most promising way forward for large-scale systems, ideally in a tightly integrated (3D) solution. A few main steps are to be taken before such a system can ever be realized: operating electronics at low temperatures, handling qubits at higher temperatures, integrating all devices (quantum and classical) in one or two process technologies, and improving refrigeration techniques.

The electronics, playing a key role in quantum computers, is currently placed at room temperature, but for integration, it needs to be operated at cryogenic temperatures. The operation of electronics at cryogenic temperatures down to 4 K has been explored in the past for many applications, ranging from astronomy to space-exploration. For quantum computing, this field is relatively new and unexplored; hence in this work, we initiate this cryogenic exploration with the main contributions:

- One of the first studies to operate a complete read-out chain for spin qubits at cryogenic temperatures.
- The first systematic approach to build a complete digital front-end, comprising voltage regulators, clock generation and a reconfigurable digital core, that operates at deep-cryogenic temperatures.
- Broad overview of commercial elements that withstand the chills of cold, delivering a solid basis for future cryogenic electronic designers.

In this work, we have started a systematic exploration of the operation of cryogenic electronics from 4 K to 300 K. We studied the behaviour of a wide range of commercial components, required in any cryogenic system, useful and applicable not only in this work, but for any cryogenic system designer. Although the behaviour of several basic components, such as capacitors, resistors and transistors, deviates significantly over this extremely wide temperature range, and not one single commercial component is specified to operate at 4 K, several components do have a relatively stable behaviour. We advocate the use of NP0/C0G, acrylic and silicon capacitors, as well as metal film resistors, whenever possible, since they are most stable at low temperatures. Discrete MOS transistors exhibit a particular turn-on voltage on the source–drain junction, but further operate normally. Similarly to integrated MOS transistors, their threshold voltage and mobility increase, whereas the subthreshold slope gets steeper. Several crystal clock oscillators are perfectly functional, showing only a 0.1% frequency deviation over a 300 K temperature range. Also MOS-based operational amplifiers profit from functional CMOS transistors and allowed to create ad-hoc voltage regulators with a stable performance (within 0.1%) over temperature, compared to all commercial regulators that simply break below 60–90 K. Integrating all of these components led to one of the first demonstrations of field-programmable gate arrays to operate at 4 K. We characterized both Altera and Xilinx 28 nm FPGAs at low temperatures. Although the Altera device is not functional below 30 K, the Xilinx Artix 7 behaves fairly stable over an extremely wide temperature range, extending the standard temperature range by over 200 K. We believe that, especially in this premature research state, FPGAs play a vital role in the development of quantum systems, since they are easy to reprogram during system operation. For this very reason, we see FPGAs emerging in all quantum fields (although currently at room temperature) for a fast control of qubits.

Although discrete electronics will be sufficient during the coming years for small quantum systems, a scalable approach has to be found in CMOS technologies. Above all others, CMOS technology offers the best integration possibilities and the most promising implementation of a large electronic infrastructure consuming low power. Therefore, we explored the operation of CMOS transistors and (small) systems at deep-cryogenic temperatures. Indeed, the CMOS technology nodes tested by us were fully functional, even down to 100 mK, and our circuits, implemented in both 40 and 160 nm processes, were fully working with fairly stable performance over temperature. We designed both an analog voltage reference and some basic digital logic blocks and characterized them at 4 K. The voltage references employ MOS diodes, since bipolar transistors freeze out at 70–90 K, and are therefore fully operational with a temperature coefficient of 0.9 mV/K. Digital logic shows a limited speed-up, especially in newer CMOS technologies, and operates with constant speed over temperature when supplied with roughly 1 V.

In a co-approach of both discrete and integrated logic, we designed a read-out system for spin qubits, comprising a CMOS LNA, a SiGe amplifier and an FPGA with firmware-based ADC operating at 4 K. All elements of the system were functional, although a harmonious co-operation of the complete system turned out to be challenging, especially when operating in a dilution refrigerator. Nonetheless, this is one of the first systematic approaches of integrating a complete electronic front-end for the quantum–classical interface at deep-cryogenic temperatures. Both amplifiers show an increase in

gain and bandwidth, whereas noise temperature goes down. We were able to digitize a -85 dBm signal (at an SNR of 10 dB in a ± 10 MHz bandwidth) through the amplification chain in our FPGA-based ADC at 300 K. Cryogenic performance is slightly reduced and current performance figures are not yet sufficient to read qubits, although these results are promising for this initial version of the Quantum Read-Out system. The first tests in a dilution fridge of the current platform revealed a temperature deviation of over $2\times$ compared to liquid helium measurements, due to poor thermalization structures. Various improvements are proposed in the form of better shielding between amplification stages and a larger design that allows for a better thermal link to the 4 K stage.

We foresee a long road still ahead of us, not only in the development of cryogenic electronic systems, but later on in the integration with the quantum devices. For the near future, major steps can be taken along the road that we have paved so far, extending our read-out platform towards the complete system with also control and radio-frequency electronic components. Our discrete elements will have to be steadily replaced by integrated circuits, but leaving room for special technologies for the amplification (SiGe or BiCMOS) and the FPGA. We see FPGAs as a valuable component to implement and test the quantum error correction and quantum execution algorithms; a task that is not easily replaced by a dedicated circuit. For cryogenic electronics, the power consumption is the main challenge, especially when operating in dilution fridges with tight power budgets and complex thermalization constraints. We expect that for large-scale systems the electronics might need to be split over multiple temperature stages in order to distribute its power consumption. Thanks to the tough requirements for qubit read-out and control, the goal of a few mW of electronic power per qubit is still far from being reached.

In conclusion, cryogenic electronics is slowly maturing and has gained attention in the race for scalable quantum computers. Our explorations over the past four years open a new era of cryogenic systems, from small-scale over the past decades towards complete and integrated solutions for the read-out and control of quantum processors. Where in the past mainly noise sensitive devices were cooled down, we believe that viable quantum computers can only be scalable when the complete electronic stack is operated at low temperatures. The future of cryogenic electronics is ‘cold and dark’.

ABBREVIATIONS

2DEG	two-dimensional electron gas
AC	alternating current
ADC	analog-to-digital converter
ALM	adaptive logic module
AMP	amplifier
ASIC	application-specific integrated circuit
BGA	ball grid array
BJT	bipolar junction transistor
BRAM	block random-access memory
CMOS	complementary metal-oxide-semiconductor
CTAT	complementary-to-absolute-temperature
DAC	digital-to-analog converter
DC	direct current
DEM	dynamic element-matching
DNL	differential non-linearity
DRAM	dynamic random-access memory
DTMOS	dynamic threshold-voltage MOS
ENOB	effective number of bits
ENR	excess noise ratio
ERBW	effective resolution bandwidth
ESR	effective series resistance
FD-SOI	fully depleted silicon-on-insulator
FET	field-effect transistor
FFT	fast Fourier transform
FLL	frequency-locked loop
FPGA	field-programmable gate array
GaAs	gallium-arsenide
HBT	heterojunction bipolar transistor
HDL	hardware description language
HEMT	high-electron-mobility transistor
HPL	high-performance low-power
IC	integrated circuit
IIR	infinite impulse response
IM	intermodulation
INL	integral non-linearity
IO	input-output
JTAG	joint test action group
LAB	logic array block
LDO	low-dropout
LNA	low-noise amplifier
LP	low-power
LSB	least significant bit

LUT	look-up table
LVC MOS	low-voltage CMOS
LVDS	low-voltage differential signaling
MEMS	micro-electro-mechanical system
MMCM	mixed-mode clock manager
MOS	metal-oxide-semiconductor
NMOS	n-channel MOS
PC	personal computer
PCB	printed circuit board
PLL	phase-locked loop
PMOS	p-channel MOS
PSD	power spectral density
PSRR	power-supply rejection ratio
PTAT	proportional-to-absolute-temperature
QD	quantum dot
QEC	quantum error-correction
QEX	quantum execution
RF	radio frequency
RO	ring oscillator
RSFQ	rapid single flux quantum
SD	standard deviation
SET	single-electron transistor
SFDR	spurious-free dynamic range
SiGe	silicon-germanium
SMD	surface-mount device
SMU	source measure unit
SNDR	signal-to-noise-and-distortion ratio
SNR	signal-to-noise ratio
SoC	system on a chip
SOI	silicon-on-insulator
SPI	serial peripheral interface
SSR	single shot resolution
TDC	time-to-digital converter
THD	harmonic distortion
UART	universal asynchronous receiver-transmitter
USB	universal serial bus
VCO	voltage-controlled oscillator

BIBLIOGRAPHY

- [1] R. R. Schaller, *Moore's law: past, present and future*, IEEE Spectrum **34** (6), pp. 52–59 (1997).
- [2] C. A. Mack, *Fifty years of Moore's law*, IEEE Trans. Semicond. Manuf. **24** (2), pp. 202–207 (2011).
- [3] F. R. Braakman, P. Barthelemy, C. Reichl, W. Wegscheider and L. M. K. Vandersypen, *Long-distance coherent coupling in a quantum dot array*, Nat. Nanotechnol. **8** (6), pp. 432–437 (2013).
- [4] E. Kawakami, P. Scarlino, D. R. Ward, F. R. Braakman, D. E. Savage, M. G. Lagally *et al.*, *Electrical control of a long-lived spin qubit in a Si/SiGe quantum dot*, Nat. Nanotechnol. **9** (9), pp. 666–670 (2014).
- [5] D. Kim, Z. Shi, C. B. Simmons, D. R. Ward, J. R. Prance, T. S. Koh *et al.*, *Quantum control and process tomography of a semiconductor quantum dot hybrid qubit*, Nature **511** (7507), pp. 70–74 (2014).
- [6] J. I. Colless, A. C. Mahoney, J. M. Hornibrook, A. C. Doherty, H. Lu, A. C. Gossard *et al.*, *Dispersive readout of a few-electron double quantum dot with fast RF gate sensors*, Phys. Rev. Lett. **110** (4), p. 046805 (2013).
- [7] J. I. Colless and D. J. Reilly, *Cryogenic high-frequency readout and control platform for spin qubits*, Rev. Sci. Instrum. **83** (2), p. 023902 (2012).
- [8] R. Kalra, A. Laucht, C. D. Hill and A. Morello, *Robust two-qubit gates for donors in silicon controlled by hyperfine interactions*, Phys. Rev. X **4** (2), p. 021044 (2014).
- [9] J. T. Muhonen, J. P. Dehollain, A. Laucht, F. E. Hudson, R. Kalra, T. Sekiguchi *et al.*, *Storing quantum information for 30 seconds in a nanoelectronic device*, Nat. Nanotechnol. **9** (12), pp. 986–991 (2014).
- [10] J. T. Muhonen, A. Laucht, S. Simmons, J. P. Dehollain, R. Kalra, F. E. Hudson *et al.*, *Quantifying the quantum gate fidelity of single-atom spin qubits in silicon by randomized benchmarking*, J. Phys. Condens. Matter **27** (15), p. 154205 (2015).
- [11] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster *et al.*, *Demonstration of two-qubit algorithms with a superconducting quantum processor*, Nature **460** (7252), pp. 240–244 (2009).
- [12] D. Ristè, S. Poletto, M.-Z. Huang, A. Bruno, V. Vesterinen, O.-P. Saira *et al.*, *Detecting bit-flip errors in a logical qubit using stabilizer measurements*, Nat. Commun. **6**, p. 6983 (2015).
- [13] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey *et al.*, *Superconducting quantum circuits at the surface code threshold for fault tolerance*, Nature **508** (7497), pp. 500–503 (2014).
- [14] E. Jeffrey, D. Sank, J. Y. Mutus, T. C. White, J. Kelly, R. Barends *et al.*, *Fast accurate state measurement with superconducting qubits*, Phys. Rev. Lett. **112** (19), p. 190504 (2014).

- [15] K. Takeda, J. Kamioka, T. Otsuka, J. Yoneda, T. Nakajima, M. R. Delbecq *et al.*, *A fault-tolerant addressable spin qubit in a natural silicon quantum dot*, *Sci. Adv.* **2** (8), p. e1600694 (2016).
- [16] J. Yoneda, K. Takeda, T. Otsuka, T. Nakajima, M. Delbecq, G. Allison *et al.*, *A >99.9%-fidelity quantum-dot spin qubit with coherence limited by charge noise*, arXiv preprint [arXiv:1708.01454](https://arxiv.org/abs/1708.01454) (2017).
- [17] F. Luthi, T. Stavenga, O. Enzing, A. Bruno, C. Dickel, N. Langford *et al.*, *Evolution of nanowire transmons and their quantum coherence in magnetic field*, arXiv preprint [arXiv:1711.07961](https://arxiv.org/abs/1711.07961) (2017).
- [18] M. Veldhorst, C. Yang, J. Hwang, W. Huang, J. Dehollain, J. Muhonen *et al.*, *A two-qubit logic gate in silicon*, *Nature* **526** (7573), p. 410 (2015).
- [19] C. Yang, R. Leon, J. Hwang, A. Saraiva, T. Tanttu, W. Huang *et al.*, *Silicon quantum processor unit cell operation above one Kelvin*, arXiv preprint [arXiv:1902.09126](https://arxiv.org/abs/1902.09126) (2019).
- [20] J. M. Hornibrook, J. I. Colless, I. D. Conway Lamb, S. J. Pauka, H. Lu, A. C. Gossard *et al.*, *Cryogenic control architecture for large-scale quantum computing*, *Phys. Rev. Appl.* **3** (2), p. 024010 (2015).
- [21] C. Degenhardt, L. Geck, A. Kruth, P. Vliex and S. van Waasen, *CMOS based scalable cryogenic control electronics for qubits*, in ICRC, pp. 1–4, IEEE (2017).
- [22] X. Fu, L. Rieseboos, L. Lao, C. G. Almudever, F. Sebastiano, R. Versluis *et al.*, *A heterogeneous quantum computer architecture*, in *Computing Frontiers*, pp. 323–330, ACM (2016).
- [23] J. M. Chow, J. M. Gambetta, E. Magesan, D. W. Abraham, A. W. Cross, B. Johnson *et al.*, *Implementing a strand of a scalable fault-tolerant quantum computing fabric*, *Nat. Commun.* **5**, p. 4015 (2014).
- [24] H. Ying, B. R. Wier, J. Dark, N. E. Lourenco, L. Ge, A. P. Omprakash *et al.*, *Operation of SiGe HBTs down to 70 mK*, *IEEE Electron Device Lett.* **38** (1), pp. 12–15 (2017).
- [25] M. G. Castellano, L. Grönberg, P. Carelli, F. Chiarello, C. Cosmelli, R. Leoni *et al.*, *Characterization of a fabrication process for the integration of superconducting qubits and rapid-single-flux-quantum circuits*, *Supercond. Sci. Technol.* **19** (8), pp. 860–864 (2006).
- [26] O. A. Mukhanov, *Energy-efficient single flux quantum technology*, *IEEE Trans. Appl. Supercond.* **21** (3), pp. 760–769 (2011).
- [27] A. A. Prager, H. C. George, A. O. Orlov and G. L. Snider, *Experimental demonstration of hybrid CMOS-single electron transistor circuits*, *J. Vac. Sci. Technol. B Microelectron. Nanom. Struct.* **29** (4), p. 041004 (2011).
- [28] H. Al-Taie, L. W. Smith, B. Xu, P. See, J. P. Griffiths, H. E. Beere *et al.*, *Cryogenic on-chip multiplexer for the study of quantum transport in 256 split-gate devices*, *Appl. Phys. Lett.* **102** (24), p. 243102 (2013).
- [29] D. R. Ward, D. E. Savage, M. G. Lagally, S. N. Coppersmith and M. A. Eriksson, *Integration of on-chip field-effect transistor switches with dopantless Si/SiGe quantum dots for high-throughput testing*, *Appl. Phys. Lett.* **102** (21), p. 213107 (2013).

- [30] R. K. Puddy, L. W. Smith, H. Al-Taie, C. H. Chong, I. Farrer, J. P. Griffiths *et al.*, *Multiplexed charge-locking device for large arrays of quantum devices*, Appl. Phys. Lett. **107** (14), p. 143501 (2015).
- [31] S. R. Ekanayake, T. Lehmann, A. S. Dzurak, R. G. Clark and A. Brawley, *Characterization of SOS-CMOS FETs at low temperatures for the design of integrated circuits for quantum bit control and readout*, IEEE Trans. Electron Devices **57** (2), pp. 539–547 (2010).
- [32] A. Rhouni, O. Gevin, X. De La Broïse, J.-L. Sauvageot, V. Revéret and L. Rodriguez, *First-ever test and characterization of the AMS standard bulk 0.35 μm CMOS technology at sub-kelvin temperatures*, in J. Physics: Conference Series, 834, p. 012005, IOP Publishing (2017).
- [33] R. L. Patterson, A. Hammond and S. S. Gerber, *Evaluation of capacitors at cryogenic temperatures for space applications*, in Sym. Electr. Insul., pp. 468–471, IEEE (1998).
- [34] A. Teverovsky, *Performance and reliability of solid tantalum capacitors at cryogenic conditions*, in STI, NASA (2006).
- [35] P. C. van Niekerk and C. J. Fourie, *Cryogenic CMOS-based control system*, in AFRICON, pp. 1–7, IEEE (2007).
- [36] F. Teyssandier and D. Prêle, *Commercially available capacitors at cryogenic temperatures*, in WOLTE9 (2010).
- [37] R. Kirschman, *Survey of Low-Temperature Electronics*, in WOLTE11 (2014).
- [38] T. Zednicek, M. Barta, J. Petrzilek and M. Biler, *Low ESR SMD tantalum capacitors for aerospace applications*, in 1st Space Passive Component Days, AVX (2013).
- [39] C. Rogers, *MOST's at cryogenic temperatures*, Solid-State Electron. **11** (11), pp. 1079–1091 (1968).
- [40] S. Wu and R. Anderson, *MOSFET's in the 0° K approximation*, Solid-State Electron. **17** (11), pp. 1125–1137 (1974).
- [41] M. J. Deen, *Cryogenic temperature dependence of the voltage transfer characteristics of CMOS inverters*, Solid-state Electron. **31** (8), pp. 1299–1308 (1988).
- [42] F. Gago, J. J. Diaz, F. Garzón and J. Patrón, *Characterization of the OPA350 operational amplifier at cryogenic temperatures*, in Scientific Detectors for Astronomy, pp. 645–650, Springer (2005).
- [43] Z. Svindrych, Z. Janu, F. Soukup and R. Tichý, *Operational amplifiers operating in temperature range from 300 to 4.2 K*, Cryogenics **48** (3), pp. 160–165 (2008).
- [44] K. Hayashi, K. Saitoh, Y. Shibayama and K. Shirahama, *A current to voltage converter for cryogenics using a CMOS operational amplifier*, in J. Physics: Conference Series, 150, p. 012016, IOP Publishing (2009).
- [45] E. D. Buchanan, D. J. Benford, J. B. Forcione, S. H. Moseley and E. J. Wollack, *Cryogenic applications of commercial electronic components*, Cryogenics **52** (10), pp. 550–556 (2012).
- [46] J. Proctor, A. Smith, T. Jung and S. Woods, *High-gain cryogenic amplifier assembly employing a commercial CMOS operational amplifier*, Rev. Sci. Instrum. **86** (7), p. 073102 (2015).

- [47] W. A. Basit, S. El-Ghanam, A. A. El-Maksood and F. Soliman, *Performance of shunt voltage regulators based on Zener diodes at cryogenic temperatures*, Phys. Sci. Res. International **1** (2), pp. 15–24 (2013).
- [48] R. Patterson and A. Hammoud, *Performance of a 3.3 V SOI voltage regulator under extreme temperatures*, in STI, NASA (2006).
- [49] V. Giordano, S. Grop, B. Dubois, P.-Y. Bourgeois, Y. Kersalé, G. Haye *et al.*, *New-generation of cryogenic sapphire microwave oscillators for space, metrology, and scientific applications*, Rev. Sci. Instrum. **83** (8), p. 085113 (2012).
- [50] P. Clapera, S. Ray, X. Jehl, M. Sanquer, A. Valentian and S. Barraud, *Design and cryogenic operation of a hybrid quantum-CMOS circuit*, Phys. Rev. Appl. **4** (4), p. 044009 (2015).
- [51] A. Hollmann, D. Jirovec, M. Kucharski, D. Kissinger, G. Fischer and L. R. Schreiber, *30 GHz-voltage controlled oscillator operating at 4 K*, arXiv preprint [arXiv:1804.09522](https://arxiv.org/abs/1804.09522) (2018).
- [52] D. Sheldon, G. Burke, A. Argueta, A. Bakshi, N. Aranki and M. Sadigursky, *Cryogenic operation of Field Programmable Gate Arrays*, in JPL TRS 1992+, NASA (2011).
- [53] J. Chang and C. Viswanathan, *Low frequency noise in CMOSTs at cryogenic temperatures*, in Low Temperature Semiconductor Electronics, pp. 133–136, IEEE (1989).
- [54] A. Bakhshi, *FPGA-based instrumentation withstands the chill of deep space*, Xcell J. **80**, pp. 24–27 (2012).
- [55] T. Liu, D. Gong, S. Hou, C. Liu, D. S. Su, P. K. Teng *et al.*, *Cryogenic digital data links for the liquid argon time projection chamber*, J. Instrum. **7** (1), p. C01091 (2012).
- [56] I. D. Conway Lamb, J. I. Colless, J. M. Hornibrook, S. J. Pauka, S. J. Waddy, M. K. Frechtling *et al.*, *An FPGA-based instrumentation platform for use at deep cryogenic temperatures*, Rev. Sci. Instrum. **87** (1), p. 014701 (2016).
- [57] P. Campagne-Ibarcq, E. Flurin, N. Roch, D. Darson, P. Morfin, M. Mirrahimi *et al.*, *Persistent control of a superconducting qubit by stroboscopic measurement feedback*, Phys. Rev. X **3** (2), p. 021008 (2013).
- [58] D. Riste, M. Dukalski, C. Watson, G. De Lange, M. Tiggelman, Y. M. Blanter *et al.*, *Deterministic entanglement of superconducting qubits by parity measurement and feedback*, Nature **502** (7471), pp. 350–354 (2013).
- [59] Y. Liu, N. Ofek, K. Geerlings, M. Hatridge, R. Schoelkopf and M. Devoret, *Real-time digital processing of qubit readout and feedback control*, in APS Meeting Abstracts, p. F25.010 (2013).
- [60] E. Mount, D. Gaultney, G. Vrijsen, M. Adams, S.-Y. Baek, K. Hudek *et al.*, *Scalable digital hardware for a trapped ion quantum computer*, Quantum Inf. Process. **15** (12), pp. 5281–5298 (2016).
- [61] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis *et al.*, *Extending the lifetime of a quantum bit with error correction in superconducting circuits*, Nature **536** (7617), pp. 441–445 (2016).
- [62] C. Favi and E. Charbon, *A 17 ps time-to-digital converter implemented in 65 nm FPGA technology*, in FPGA, pp. 113–120, ACM (2009).

- [63] J. Hamlet, K. Eng, T. Gurrieri, J. Levy and M. Carroll, *Modeling of circuits with strongly temperature dependent thermal conductivities for cryogenic CMOS*, Microelectron. J. **42** (7), pp. 936–941 (2011).
- [64] B. Farley, J. McGrath and C. Erdmann, *An all-programmable 16 nm RFSoc for digital-RF communications*, IEEE Micro **38** (2), pp. 61–71 (2018).
- [65] D. Ristè, C. C. Bultink, K. W. Lehnert and L. DiCarlo, *Feedback control of a solid-state qubit using high-fidelity projective measurement*, Phys. Rev. Lett. **109** (24), p. 240502 (2012).
- [66] P. Campagne-Ibarcq, *Measurement back action and feedback in superconducting circuits*, Ph.D. thesis, Ecole Normale Supérieure (ENS) (2015).
- [67] D. T. Sank, *Fast, accurate state measurement in superconducting qubits*, Ph.D. thesis, University of California (2014).
- [68] B. Okcan, G. Gielen and C. van Hoof, *A third-order complementary metal–oxide–semiconductor sigma-delta modulator operating between 4.2 K and 300 K*, Rev. Sci. Instrum. **83** (2), p. 024708 (2012).
- [69] H. Zhao and X. Liu, *A low-power cryogenic analog to digital converter in standard CMOS technology*, Cryogenics **55–56**, pp. 79–83 (2013).
- [70] Y. Creten, P. Merken, W. Sansen, R. Mertens and C. van Hoof, *A cryogenic ADC operating down to 4.2 K*, in ISSCC, pp. 468–469, IEEE (2007).
- [71] Y. Yao, *Design and implementation of high-speed low-power analog-to-digital and digital-to-analog converters*, Ph.D. thesis, Auburn University (2010).
- [72] 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 J. Solid-State Circuits **44** (7), pp. 2019–2025 (2009).
- [73] 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*, Rev. Sci. Instrum. **81** (2), p. 024702 (2010).
- [74] H. Zhao, Y. Zhao and Z. Zhang, *A cryogenic SAR ADC for infra-red readout circuits*, J. Semicond. **32** (11), p. 115015 (2011).
- [75] 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*, J. Shanghai Jiaotong University (Science) **18** (5), pp. 520–525 (2013).
- [76] Lattice Semiconductor, *Leveraging FPGA and CPLD digital logic to implement analog to digital converters* (2010).
- [77] Missing Link Electronics, *Integrated ADC for Altera Cyclone-IV devices* (2011).
- [78] J. Wu, S. Hansen and Z. Shi, *ADC and TDC implemented using FPGA*, in NSS/MIC, pp. 281–286, IEEE (2007).
- [79] M. Pałka, P. Moskal, T. Bednarski, P. Białas, E. Czerwiński, Ł. Kapłon *et al.*, *A novel method based solely on field programmable gate array (FPGA) units enabling measurement of time and charge of analog signals in positron emission tomography (PET)*, Bio-Algorithms and Med-Systems **10** (1), pp. 41–45 (2014).
- [80] Xilinx, *7 series FPGAs and Zynq-7000 all programmable SoC XADC* (2017).
- [81] B. Murmann, *ADC performance survey 1997-2018* (2018).

- [82] W. He, M. Stottinger, E. de la Torre and V. Diaz, *A self-tuned thermal compensation system for reducing process variation influence in side-channel attack resistant dual-rail logic*, in DCIS, pp. 1–6, IEEE (2015).
- [83] F. Balestra, L. Audaire and C. Lucas, *Influence of substrate freeze-out on the characteristics of MOS transistors at very low temperatures*, Solid-state Electron. **30** (3), pp. 321–327 (1987).
- [84] E. Simoen, B. Dierickx and C. L. Claeys, *Low-frequency noise behavior of Si NMOSTs stressed at 4.2 K*, IEEE Trans. Electron Devices **40** (7), pp. 1296–1299 (1993).
- [85] S. Broadbent, *CMOS operation below freezeout*, in Low Temperature Semiconductor Electronics, pp. 43–47, IEEE (1989).
- [86] J. T. Watt, B. J. Fishbein and J. D. Plummer, *A low-temperature NMOS technology with cesium-implanted load devices*, IEEE Trans. Electron Devices **34** (1), pp. 28–38 (1987).
- [87] H. Hanamura, M. Aoki, T. Masuhara, O. Minato, Y. Sakai and T. Hayashida, *Operation of bulk CMOS devices at very low temperatures*, IEEE J. Solid-State Circuits **21** (3), pp. 484–490 (1986).
- [88] A. Kamgar, *Subthreshold behavior of silicon MOSFETs at 4.2 K*, Solid-State Electron. **25** (7), pp. 537–539 (1982).
- [89] L. Deferm, E. Simoen and C. Claeys, *The importance of the internal bulk-source potential on the low temperature kink in NMOSTs*, IEEE Trans. Electron Devices **38** (6), pp. 1459–1466 (1991).
- [90] M. Aoki, K. Yano, T. Masuhara and K. Shimohigashi, *Fully symmetric cooled CMOS on (110) plane*, IEEE Trans. Electron Devices **36** (8), pp. 1429–1433 (1989).
- [91] E. Simoen, A. Mercha, Y. Creten, P. Merken, J. Putzeys, P. de Moor *et al.*, *Impact of irradiations performed at liquid helium temperatures on the operation of 0.7 μm CMOS devices and read-out circuits*, in RADECS, pp. 369–375, IEEE (2003).
- [92] E. Simoen and C. Claeys, *Impact of CMOS processing steps on the drain current kink of NMOSFETs at liquid helium temperature*, IEEE Trans. Electron Devices **48** (6), pp. 1207–1215 (2001).
- [93] W.-Y. Liu, Q. Feng and R.-J. Ding, *Impact of kink effect on CMOS readout circuits for cryogenic operation*, Laser & Infrared **37**, pp. 990–992 (2007).
- [94] L. Varizat, G. Sou, M. Mansour, D. Alison and A. Rhouni, *A low temperature 0.35 μm CMOS technology BSIM3.3 model for space instrumentation: Application to a voltage reference design*, in MetroAeroSpace, pp. 74–78, IEEE (2017).
- [95] G. S. Fonseca, L. B. de Sá and A. C. Mesquita, *Extraction of static parameters to extend the EKV model to cryogenic temperatures*, in SPIE Defense + Security, p. 98192B, SPIE (2016).
- [96] H. Zhao and X. Liu, *Modeling of a standard 0.35 μm CMOS technology operating from 77 K to 300 K*, Cryogenics **59**, pp. 49–59 (2014).
- [97] N. C. Dao, A. El Kass, M. R. Azghadi, C. T. Jin, J. Scott and P. H. Leong, *An enhanced MOSFET threshold voltage model for the 6–300 K temperature range*, Microelectron. Reliab. **69**, pp. 36–39 (2017).

- [98] N. Yoshikawa, T. Tomida, M. Tokuda, Q. Liu, X. Meng, S. Whiteley *et al.*, *Characterization of 4 K CMOS devices and circuits for hybrid Josephson–CMOS systems*, IEEE Trans. Appl. Supercond. **15** (2), pp. 267–271 (2005).
- [99] M. Kakumu, D. Peters, H. Liu and K. Chiu, *Design optimization for deep-submicron CMOS device at low temperature operation*, in IEDM, pp. 211–214, IEEE (1990).
- [100] Y. Feng, P. Zhou, H. Liu, J. Sun and T. Jiang, *Characterization and modelling of MOSFET operating at cryogenic temperature for hybrid superconductor-CMOS circuits*, Semicond. Sci. Technol. **19** (12), p. 1381 (2004).
- [101] Z. Zhu, A. Kathuria, S. G. Krishna, M. Mojarradi, B. Jalali-Farahani, H. Barnaby *et al.*, *Design applications of compact MOSFET model for extended temperature range (60–400 K)*, Electron. Lett. **47** (2), pp. 141–142 (2011).
- [102] P. Martin, A. Royet, F. Guellec and G. Ghibaudo, *MOSFET modeling for design of ultra-high performance infrared CMOS imagers working at cryogenic temperatures: Case of an analog/digital 0.18 μm CMOS process*, Solid-State Electron. **62** (1), pp. 115–122 (2011).
- [103] G. De Geronimo, A. D’Andragora, S. Li, N. Nambiar, S. Rescia, E. Vernon *et al.*, *Front-end ASIC for a liquid argon TPC*, in NSS/MIC, pp. 1658–1666, IEEE (2010).
- [104] A. Akturk, M. Holloway, S. Potbhare, D. Gundlach, B. Li, N. Goldsman *et al.*, *Compact and distributed modeling of cryogenic bulk MOSFET operation*, IEEE Trans. Electron Devices **57** (6), pp. 1334–1342 (2010).
- [105] S. H. Hong, G. B. Choi, R. H. Baek, H. S. Kang, S. W. Jung and Y. H. Jeong, *Low-temperature performance of nanoscale MOSFET for deep-space RF applications*, IEEE Electron Device Lett. **29** (7), pp. 775–777 (2008).
- [106] A. Siligaris, G. Pailloncy, S. Delcourt, R. Valentin, S. Lepilliet, F. Danneville *et al.*, *High-frequency and noise performances of 65 nm MOSFET at liquid nitrogen temperature*, IEEE Trans. Electron Devices **53** (8), pp. 1902–1908 (2006).
- [107] Z. Chen, H. Wong, Y. Han, S. Dong and B. Yang, *Temperature dependences of threshold voltage and drain-induced barrier lowering in 60 nm gate length MOS transistors*, Microelectron. Reliab. **54** (6), pp. 1109–1114 (2014).
- [108] A. H. Coskun and J. C. Bardin, *Cryogenic small-signal and noise performance of 32 nm SOI CMOS*, in IMS, pp. 1–4, IEEE (2014).
- [109] H. Achour, B. Cretu, J.-M. Routoure, R. Carin, R. Talmat, A. Benfdila *et al.*, *In depth static and low-frequency noise characterization of n-channel FinFETs on SOI substrates at cryogenic temperature*, Solid-State Electron. **98**, pp. 12–19 (2014).
- [110] A. Beckers, F. Jazaeri, A. Ruffino, C. Bruschini, A. Baschiroto and C. Enz, *Cryogenic characterization of 28 nm bulk CMOS technology for quantum computing*, in ESSDERC, pp. 62–65, IEEE (2017).
- [111] M. Shin, M. Shi, M. Mouis, A. Cros, E. Josse, G.-T. Kim *et al.*, *Low temperature characterization of 14 nm FDSOI CMOS devices*, in WOLTE, pp. 29–32, IEEE (2014).
- [112] M. Shin, M. Shi, M. Mouis, A. Cros, E. Josse, G.-T. Kim *et al.*, *In depth characterization of electron transport in 14 nm FD-SOI CMOS devices*, Solid-State Electron. **112**, pp. 13–18 (2015).

- [113] J. D. Cressler and H. A. Mantooth, *Extreme environment electronics*, CRC Press (2012).
- [114] F. Balestra and G. Ghibaudo, *Brief review of the MOS device physics for low temperature electronics*, Solid-state electronics **37** (12), pp. 1967–1975 (1994).
- [115] K. Masaki, K. Taniguchi, C. Hamaguchi and M. Iwase, *Temperature dependence of electron mobility in Si inversion layers*, Jpn. J. Appl. Phys. **30** (11R), p. 2734 (1991).
- [116] 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*, J. Appl. Phys. **67** (4), pp. 1950–1952 (1990).
- [117] M. A. Pertijs and J. Huijsing, *Precision temperature sensors in CMOS technology*, Springer (2006).
- [118] K. Ueno, T. Hirose, T. Asai and Y. Amemiya, *A 300 nW, 15 ppm/°C, 20 ppm/V CMOS voltage reference circuit consisting of subthreshold MOSFETs*, IEEE J. Solid-State Circuits **44** (7), pp. 2047–2054 (2009).
- [119] A.-J. Annema, *Low-power bandgap references featuring DTMOSTs*, IEEE J. Solid-State Circuits **34** (7), pp. 949–955 (1999).
- [120] K. Souri, Y. Chae and K. A. Makinwa, *A CMOS temperature sensor with a voltage-calibrated inaccuracy of $\pm 0.15^\circ\text{C}$ (3σ) from -55°C to 125°C* , IEEE J. Solid-State Circuits **48** (1), pp. 292–301 (2013).
- [121] K. Souri, Y. Chae, F. Thus and K. Makinwa, *A 0.85 V 600 nW all-CMOS temperature sensor with an inaccuracy of $\pm 0.4^\circ\text{C}$ (3σ) from -40 to 125°C* , in ISSCC, pp. 222–223, IEEE (2014).
- [122] F. Assaderaghi, D. Sinitsky, S. A. Parke, J. Bokor, P. K. Ko and C. Hu, *Dynamic threshold-voltage MOSFET (DTMOS) for ultra-low voltage VLSI*, IEEE Trans. Electron Devices **44** (3), pp. 414–422 (1997).
- [123] J.-K. Lee, N.-J. Choi, C.-G. Yu, J.-P. Colinge and J.-T. Park, *Temperature dependence of DTMOS transistor characteristics*, Solid-State Electron. **48** (1), pp. 183–187 (2004).
- [124] W. Liao, C. Shih and P. Li, *Thermal stability of SiGe dynamic threshold pMOSFET, in IPFA*, pp. 137–140, IEEE (2005).
- [125] M. Terauchi, *Temperature dependence of the subthreshold characteristics of dynamic threshold metal–oxide–semiconductor field-effect transistors and its application to an absolute-temperature sensing scheme for low-voltage operation*, Jpn. J. Appl. Phys. **46** (7R), p. 4102 (2007).
- [126] B. Yousefzadeh, S. H. Shalmany and K. A. Makinwa, *A BJT-based temperature-to-digital converter with $\pm 60\text{ mK}$ (3σ) inaccuracy from -55°C to $+125^\circ\text{C}$ in $0.16\text{ }\mu\text{m}$ CMOS*, IEEE J. Solid-State Circuits **52** (4), pp. 1044–1052 (2017).
- [127] G. Wang and G. C. Meijer, *Temperature characteristics of bipolar transistors fabricated in CMOS technology*, Sens. Actuators, A **87** (1), pp. 81–89 (2000).
- [128] F. Sebastiano, L. J. Breems, K. A. Makinwa, S. Drago, D. M. Leenaerts and B. Nauta, *A 1.2 V 10 μW NPN-based temperature sensor in 65 nm CMOS with an inaccuracy of 0.2°C (3σ) from -70°C to 125°C* , IEEE J. Solid-State Circuits **45** (12), pp. 2591–2601 (2010).

- [129] J. D. Cressler, J. Comfort, E. F. Crabbe, G. L. Patton, J. M. Stork, J.-C. Sun *et al.*, *On the profile design and optimization of epitaxial Si-and SiGe-base bipolar technology for 77 K applications. I. Transistor DC design considerations*, IEEE Trans. Electron Devices **40** (3), pp. 525–541 (1993).
- [130] F. Akyurek and B. Bayram, *Digital output ROIC with single slope ADC for cooled infrared applications*, Microsyst. Technol. **23** (4), pp. 1091–1099 (2017).
- [131] M. Rahman and T. Lehmann, *A cryogenic DAC operating down to 4.2 K*, Cryogenics **75**, pp. 47–55 (2016).
- [132] M. T. Rahman and T. Lehmann, *A self-calibrated cryogenic current cell for 4.2 K current steering D/A converters*, IEEE Trans. Circuits Syst. II Express Briefs **64** (10), pp. 1152–1156 (2017).
- [133] N. Sahin and M. B. Yelten, *A 0.18 μm CMOS X-band low noise amplifier for space applications*, in NGCAS, pp. 205–208, IEEE (2017).
- [134] L. Najafizadeh, C. Zhu, R. Krithivasan, J. D. Cressler, Y. Cui, G. Niu *et al.*, *SiGe BiCMOS precision voltage references for extreme temperature range electronics*, in Bipolar/BiCMOS Circuits and Technology Meeting, pp. 1–4, IEEE (2006).
- [135] Y. Yao, F. Dai, R. C. Jaeger and J. D. Cressler, *A 12-bit cryogenic and radiation-tolerant digital-to-analog converter for aerospace extreme environment applications*, IEEE Trans. Ind. Electron. **55** (7), pp. 2810–2819 (2008).
- [136] L. Varizat, G. Sou, M. Mansour and D. Alison, *A cryogenic 0.35 μm CMOS technology BSIM3.3 model for space instrumentation: Application to a bandgap design*, IEEE Aerosp. Electron. Syst. Mag. **33** (8), pp. 24–29 (2018).
- [137] R. Patterson, *Evaluation of voltage reference circuits and N-channel field effect transistors at low temperatures*, in STI, NASA (2002).
- [138] L. Najafizadeh, J. S. Adams, S. D. Phillips, K. A. Moen, J. D. Cressler, S. Aslam *et al.*, *Sub-1 K operation of SiGe transistors and circuits*, IEEE Electron Device Lett. **30** (5), pp. 508–510 (2009).
- [139] J. Schrankler, J. Huang, R. Lutze, H. Vyas and G. Kirchner, *Cryogenic behavior of scaled CMOS devices*, in IEDM, pp. 598–600, IEEE (1984).
- [140] I. Kato, H. Oka, S. Hijiya and T. Nakamura, *1.5 μm gate CMOS operated at 77 K*, in IEDM, pp. 601–604, IEEE (1984).
- [141] J. Laramée, M. Aubin and J. Cheeke, *Behavior of CMOS inverters at cryogenic temperatures*, Solid-state Electron. **28** (5), pp. 453–456 (1985).
- [142] M. J. Deen, *Digital characteristics of CMOS devices at cryogenic temperatures*, IEEE J. Solid-State Circuits **24** (1), pp. 158–164 (1989).
- [143] E. A. Gutierrez-D, L. Deferm and G. Declerck, *DC characteristics of submicrometer CMOS inverters operating over the whole temperature range of 4.2–300 K*, IEEE Trans. Electron Devices **39** (9), pp. 2182–2184 (1992).
- [144] Y. Mii, S. Rishton, Y. Taur, D. Kern, T. Lii, K. Lee *et al.*, *Experimental high performance sub-0.1 μm channel nMOSFET's*, IEEE Electron Device Lett. **15** (1), pp. 28–30 (1994).
- [145] E. Simoen and C. Claeys, *The cryogenic operation of partially depleted silicon-on-insulator inverters*, IEEE Trans. Electron Devices **42** (6), pp. 1100–1105 (1995).

- [146] J. B. Burr, *Cryogenic ultra low power CMOS*, in *Low Power Electronics*, pp. 82–83, IEEE (1995).
- [147] F. Assaderaghi, W. Rausch, A. Ajmera, E. Leobandung, D. Schepis, L. Wagner *et al.*, *A 7.9/5.5 ps room/low temperature SOI CMOS*, in *IEDM*, pp. 415–418, IEEE (1997).
- [148] K. J. Hass and H. C. Shaw, *Cryogenic operation of ultra low power CMOS*, in *Symp. VLSI Design*, 4, pp. 1–7, NASA (2000).
- [149] D. Foty, *An evaluation of deep-submicron CMOS design optimized for operation at 77 K*, in *ISIE*, 3, pp. 1139–1144, IEEE (2005).
- [150] I. V. Vernik, T. A. Ohki, M. B. Ketchen and M. Bhushan, *Performance characterization of PD-SOI ring oscillators at cryogenic temperatures*, in *SOI Conference*, pp. 1–2, IEEE (2010).
- [151] D. Loss and D. P. DiVincenzo, *Quantum computation with quantum dots*, *Phys. Rev. A* **57** (1), p. 120 (1998).
- [152] J. J. Pla, K. Y. Tan, J. P. Dehollain, W. H. Lim, J. J. Morton, D. N. Jamieson *et al.*, *A single-atom electron spin qubit in silicon*, *Nature* **489** (7417), p. 541 (2012).
- [153] J. Elzerman, R. Hanson, L. W. Van Beveren, S. Tarucha, L. Vandersypen and L. Kouwenhoven, *Semiconductor few-electron quantum dots as spin qubits*, in *Quantum Dots: a Doorway to Nanoscale Physics*, pp. 25–95, Springer (2005).
- [154] S. Weinreb, J. Bardin, H. Mani and G. Jones, *Matched wideband low-noise amplifiers for radio astronomy*, *Rev. Sci. Instrum.* **80** (4), p. 044702 (2009).
- [155] G. Cortes-Medellin, A. Vishwas, S. Parshley and D. B. Campbell, *Fully cryogenic phased array prototype camera for the arecibo radio telescope*, in *Ground-based and Airborne Instrumentation for Astronomy V*, p. 91479Q, SPIE (2014).
- [156] F. Bruccoleri, E. A. Klumperink and B. Nauta, *Wide-band CMOS low-noise amplifier exploiting thermal noise canceling*, *IEEE J. Solid-State Circuits* **39** (2), pp. 275–282 (2004).
- [157] Agilent Technologies, *Noise figure measurement accuracy – The Y-factor method* (2004).

PUBLICATIONS

JOURNALS

- [J1] M. Turchetti, **H. Homulle**, F. Sebastiano, G. Ferrari, E. Charbon and E. Prati, *Tunable single hole regime of a silicon field effect transistor in standard CMOS technology*, Applied Physics Express **9** (1), p. 014001 (2015).
- [J2] **H. Homulle**, F. Powolny, P. Stegehuis, J. Dijkstra, D.-U. Li, K. Homicsko, D. Rimoldi, K. Muehlethaler, J. Prior, R. Sinisi, E. Dubikovskaya, E. Charbon and C. Bruschini, *Compact solid-state CMOS single-photon detector array for in vivo NIR fluorescence lifetime oncology measurements*, Biomedical Optics Express **7** (5), pp. 1797–1814 (2016).
- [J3] **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 **63** (11), pp. 1854–1865 (2016).
- [J4] **H. Homulle**, S. Visser, B. Patra, G. Ferrari, E. Prati, F. Sebastiano and E. Charbon, *A reconfigurable cryogenic platform for the classical control of quantum processors*, Review of Scientific Instruments **88** (4), pp. 045103-1–9 (2017).
- **H. Homulle**, S. Visser, B. Patra, G. Ferrari, E. Prati, F. Sebastiano and E. Charbon, *A reconfigurable cryogenic platform for the classical control of scalable quantum computers*, arXiv preprint **arXiv:1602.05786** (2016).
- [J5] B. Patra, R. Incandela, J. van Dijk, **H. Homulle**, L. Song, M. Shahmohammadi, R. Staszewski, A. Vladimirescu, M. Babaie, F. Sebastiano and E. Charbon, *Cryo-CMOS circuits and systems for quantum computing applications*, IEEE Journal of Solid-State Circuits **53** (1), pp. 309–321 (2018).
- [J6] **H. Homulle**, S. Visser, B. Patra and E. Charbon, *Design techniques for a stable operation of cryogenic field-programmable gate arrays*, Review of Scientific Instruments **89** (1), pp. 014703-1–8 (2018).
- **H. Homulle**, S. Visser, B. Patra and E. Charbon, *FPGA design techniques for stable cryogenic operation*, arXiv preprint **arXiv:1709.04190** (2017).
- [J7] **H. Homulle**, L. Song, E. Charbon and F. Sebastiano, *The cryogenic temperature behavior of Bipolar, MOS and DT MOS transistors in standard CMOS*, IEEE Journal of the Electron Devices Society **6**, pp. 263–270 (2018).

- [J8] R. 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 **6**, pp. 996–1006 (2018).
- [J9] **H. Homulle** and E. Charbon, *Cryogenic low-dropout voltage regulators for stable low-temperature electronics*, Cryogenics **95**, pp. 11–17 (2018).
- [J10] **H. Homulle**, F. Sebastiano and E. Charbon, *Deep-cryogenic voltage references in 40-nm CMOS*, IEEE Solid-State Circuits Letters **1** (5), pp. 110–113 (2018).

CONFERENCE PROCEEDINGS

- [C1] F. Powolny, K. Homicsko, R. Sinisi, C. Bruschini, E. Grigoriev, **H. Homulle**, J. Prior, D. Hanahan, E. Dubikovskaya and E. Charbon, *Time-resolved imaging system for fluorescence-guided surgery with lifetime imaging capability*, SPIE. 9129, Biophotonics: Photonic Solutions for Better Health Care IV, p. 912938 (2014).
- [C2] P. Stegehuis, M. Boonstra, K. de Rooij, F. Powolny, R. Sinisi, **H. Homulle**, C. Bruschini, E. Charbon, C. van de Velde, B. Lelieveldt, A. Vahrmeijer, J. Dijkstra and M. van de Giessen, *Fluorescence lifetime imaging to differentiate bound from unbound ICG-cRGD both in vitro and in vivo*, SPIE. 9313, Advanced Biomedical and Clinical Diagnostic and Surgical Guidance Systems XIII, p. 931300 (2015).
- [C3] **H. Homulle**, F. Regazzoni and E. Charbon, *200 MS/s ADC implemented in a FPGA employing TDCs*, ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA), pp. 228–235 (2015).
- [C4] S. Burri, **H. Homulle**, C. Bruschini and E. Charbon, *LinoSPAD: a time-resolved 256×1 CMOS SPAD line sensor system featuring 64 FPGA-based TDC channels running at up to 8.5 giga-events per second*, SPIE. 9899, Optical Sensing and Detection IV, p. 98990D (2016).
- [C5] **H. Homulle**, S. Visser, B. Patra, G. Ferrari, E. Prati, C. Almudéver, K. Bertels, F. Sebastiano and E. Charbon, *CryoCMOS hardware technology a classical infrastructure for a scalable quantum computer*, ACM International Conference on Computing Frontiers, pp. 282–287 (2016).
- [C6] L. Song, **H. Homulle**, E. Charbon and F. Sebastiano, *Characterization of bipolar transistors for cryogenic temperature sensors in standard CMOS*, IEEE Sensors, pp. 1–3 (2016).
- [C7] E. Charbon, F. Sebastiano, A. Vladimirescu, **H. Homulle**, S. Visser, L. Song and R. M. Incandela, *Cryo-CMOS for quantum computing*, IEEE Electron Devices Meeting (IEDM), pp. 13–5 (2016).


- [C8] C. Bruschini, **H. Homulle** and E. Charbon, *Ten years of biophotonics single-photon SPAD imager applications: retrospective and outlook*, SPIE 10069, Multiphoton Microscopy in the Biomedical Sciences XVII, pp. 100691S (2017).
- [C9] E. Charbon, F. Sebastiano, M. Babaie, A. Vladimirescu, M. Shahmohammadi, R. Staszewski, **H. Homulle**, B. Patra, J. van Dijk, R. Incandela, L. Song and B. Valizadehpasha, *Cryo-CMOS circuits and systems for scalable quantum computing*, IEEE Solid-State Circuits Conference (ISSCC), pp. 264–265 (2017).
- [C10] F. Sebastiano, **H. Homulle**, B. Patra, R. Incandela, J. van Dijk, L. Song, M. Babaie, A. Vladimirescu and E. Charbon, *Cryo-CMOS Electronic Control for Scalable Quantum Computing*, Annual Design Automation Conference (DAC), pp. 13.1–13.6 (2017).
- [C11] F. Sebastiano, **H. Homulle**, J. van Dijk, R. Incandela, B. Patra, M. Mehrpoo, M. Babaie, A. Vladimirescu and E. Charbon, *Cryogenic CMOS interfaces for quantum devices*, IEEE International Workshop on Advances in Sensors and Interfaces (IWASI), pp. 59–62 (2017).
- [C12] R. Incandela, L. Song, **H. Homulle**, F. Sebastiano, E. Charbon and A. Vladimirescu, *Nanometer CMOS characterization and compact modeling at deep-cryogenic temperatures*, IEEE European Solid-State Device Research Conference (ESSDERC), pp. 58–61 (2017).
- [C13] **H. Homulle** and E. Charbon, *Performance characterization of Altera and Xilinx 28 nm FPGAs at cryogenic temperatures*, IEEE International Conference on Field-Programmable Technology (FPT), pp. 25–31 (2017).
- [C14] **H. Homulle** and E. Charbon, *Commercial crystal and MEMS oscillators characterized at deep-cryogenic temperatures (+poster)*, IEEE International Frequency Control Symposium (IFCS), pp. 133–137 (2018).

CONFERENCE TALKS / POSTERS (WITHOUT PAPER)


- [O1] **H. Homulle**, S. Visser, B. Patra, R. Incandela, B. Valizadeh, M. Babaie, F. Sebastiano and E. Charbon, *Paving the road towards a cryogenic qubit control platform (poster)*, ScaleQIT (2016).
- [O2] S. Visser, **H. Homulle** and E. Charbon, *A 1 GSa/s, reconfigurable soft-core FPGA ADC (poster)*, ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (2016).
- [O3] **H. Homulle**, *An FPGA-based cryogenic platform for the classical control of Quantum Computers (talk)*, Silicon Quantum Workshop (2016).

- [O4] **H. Homulle**, *Cryogenic reconfigurable electronics: a classical control interface for Quantum Computers (talk)*,
Workshop on Low Temperature Electronics (2016).
- [O5] C. Bruschini, **H. Homulle** and E. Charbon, *A decade of single-photon SPAD imagers in the biomedical Sciences (talk)*,
New Developments in Photodetection (2017).
- [O6] C. Bruschini, **H. Homulle** and E. Charbon, *All-digital, quantum biomedical imaging (talk)*,
Emerging Technologies – Communications, Microsystems, Optoelectronics, Sensors (2017).
- [O7] **H. Homulle**, *QuRO: The first entirely cryogenic interface between silicon spin qubits and a programmable classical system (talk)*,
Workshop on Low Temperature Electronics (2018).
- [O8] **H. Homulle** and E. Charbon, *QuRO: A compact read-out interface for spin qubits operating at deep-cryogenic temperatures (talk)*,
Silicon Quantum Workshop (2018).


CRYOCMOS GALLERY



(a) NXP 0.16 μm




(b) ST 40 nm




(c) TSMC 40 nm

Figure A.1: Chips designed for CMOS characterization: (a) NXP/SSMC 0.16 μm co-designed with F. Sebastiano, (b) ST 40 nm and (c) TSMC 40 nm co-designed with J. van Dijk and B. Patra.


CRYOPCB GALLERY




(a) CryoFPGA




(b) CryoChips




(c) CryoCapacitors




(d) CryoLDOs & CryoOSC




(e) CryoClone



(f) QuRO



(g) CryoFPGA MB



(h) CryoZynq

Figure A.2: PCBs designed for cryogenic testing: (a) the first cryogenic FPGA test board with Xilinx Artix 7 co-designed with B. Patra, (b) test board for CMOS characterization, (c) capacitor test board, (d) LDO and oscillator board, (e) Altera Cyclone V test board and (f) QuRO or Quantum Read-Out featuring the Artix 7, a clock oscillator, LDOs, a SiGe amplifier and a CMOS LNA. (g) CryoFPGA MotherBoard with the Artix 7 and interface connections for a DaughterBoard. (h) CryoZynq with the Zynq SoC (FPGA + dual-core ARM Cortex-A9), DDR3 memory, QSPI and SDIO.

CRYOPROBE

Connector box

Sample stage

Figure A.3: For the majority of experiments conducted in this thesis, several dip-sticks (CryoProbes) were designed. The second and third iterations house 30 coaxial lines (SMA) to a dedicated connector board (QTE) in the sample stage. For high-frequency, four SMP coaxial lines were installed.



Figure A.4: Our cryogenic test setup, the CryoProbe (Figure A.3) is dipped into the liquid helium or the helium vapours inside a dewar to reach a temperature between 300 K and 4 K. Instrumentation is connected to the room temperature connector box on the top.

SUMMARY

Quantum computing promises an exponential speed-up of computation compared to what is nowadays achievable with classical computers. In this way, it enables the evaluation of more complex models and the breaching of current security algorithms. For the operation of a quantum system, many questions remain to be answered. Currently, there are several quantum technologies that promise to be both reliable and scalable, two features required for large scale quantum operations. Common to all technologies is the operating temperature that needs to be close to absolute zero, i.e. below 100 mK, to suppress environment noise and allow the quantum properties to become 'visible'.

In order for any quantum processor to be operated, a so-called quantum-classical interface is required for the quantum bit (qubit) read-out and control. This interface consists of various electronic blocks, such as analog-to-digital converters, digital-to-analog converters, mixers, amplifiers and a digital controller. Especially the analog blocks require effort to meet the noise and stability constraints as not to disturb the very sensitive qubits and allow reading of the tiny signals. As the qubits live in extremely deep-cryogenic temperatures, long wires interface the cryogenic with the room temperature environment, where most of the electronics is situated. However, for a scalable system, heat injection becomes a serious problem, with many wires between 300 K and sub-Kelvin. Furthermore, such amount of interconnects is challenging to mechanically place in a dilution refrigerator.

Therefore, in this work, we propose to implement the electronics not at room temperature, but at a temperature much closer to the qubits, for example at 4 K. This not only reduces significantly the wiring between room temperature and the qubits, but we can also benefit from lower electronic noise at such a temperature. The operation of conventional electronics almost 200°C below its normal temperature range is not trivial as device properties alter significantly and most circuits no longer operate as intended.

In CMOS processes, the main technology in the integrated electronics world, the behaviour of the transistors, required for the implementation of any circuit, deviates considerably at low temperatures. The transistor's threshold voltage goes up, the mobility increases and the subthreshold slope becomes steeper, to name just a few of these deviations. Although there are improvements in performance, there are also some counter effects, and characterization of the transistors is needed to observe the changes. Once devices are characterized at such low temperatures, new models can be built and circuits can be simulated and adapted to operate properly at cryogenic temperatures. Luckily, there are various commercially available devices that can already withstand the chills of cold. We demonstrated various commercially available devices, such as a field-programmable gate array (FPGA) implemented in a 28 nm CMOS process, to be operating without major concerns at 4 K. Its properties alter only slightly, within 5 to 10%, and all tested circuit implementations, working at 300 K, also worked at 4 K.

We combined both commercially available devices, that operate 200 K below their specified temperature range, with custom designed CMOS circuits to implement a cryogenic read-out platform for spin qubits. This system comprises amplifiers, an ADC, an FPGA, voltage regulators and a clock generator. It allows to amplify the tiny signal from

the qubits and digitize it directly in the FPGA in order to process the data locally at 4 K. This system is one of the first systematic attempts at operating a part of the quantum–classical interface at cryogenic temperatures and forms the basis for future systems comprising the complete electronic interface for qubits to operate at such low temperatures.

One of the main problems to tackle for cryogenic electronic systems is their power consumption. Power budgets are simply limited to roughly 1 or 2 Watts at 4 K and exponentially lower at deeper cryogenic temperatures, thus limiting the size of large-scale electronic systems. Our approach of combined commercial and custom circuits will have to be steadily replaced by a single (custom) technology that meets both power and scalability constraints. One of the best candidates is CMOS, a technology that the industry has relied upon for several decades and benefits from many optimizations thanks to Moore’s law.

SAMENVATTING

Kwantumcomputatie maakt het mogelijk exponentieel sneller te kunnen rekenen dan wat tegenwoordig met klassieke computers gebruikelijk is. Zodoende kunnen complexere wiskundige modellen worden geëvalueerd en beveiligingsalgoritmen worden gebroken. Om zo'n kwantumcomputer te laten werken moeten echter nog vele vragen worden beantwoord. Op dit moment zijn er verscheidene kwantumtechnologieën die beloven zowel betrouwbaar als schaalbaar te zijn. Beide factoren zijn cruciaal voor het maken van grootschalige kwantumoperaties. Al deze technologieën hebben één belangrijke gemeenschappelijke eigenschap, namelijk dat ze dichtbij het absolute nulpunt moeten opereren, dat wil zeggen bij minder dan 100 mK, om omgevingsruis te onderdrukken en de kwantumeigenschappen 'zichtbaar' te maken.

Om een kwantumprocessor werkzaam te krijgen, is een zogenoemde kwantum-klas-sieke verbinding nodig die beide werelden laat communiceren. Deze verbinding maakt het mogelijk om een kwantumbit uit te lezen en te manipuleren, en bestaat daartoe uit verscheidene elektronische circuits, zoals analoog-digitaalomzetters, digitaal-analoog-omzetters, mixers, versterkers, en een digitaal besturingsapparaat. Vooral de analoge onderdelen hebben veel aandacht nodig om de ruis- en stabiliteitseisen te behalen, om derhalve de zeer gevoelige kwantumbits niet te verstoren en het mogelijk te maken de kleine signalen van deze bits te lezen. Omdat de kwantumbits dus in een extreem koude omgeving leven, zijn lange kabels benodigd tussen de cryogene en kamertemperatuur omgevingen, daar naar waar de meeste elektronica zich bevindt. Echter voor een schaalbaar systeem wordt de warmte-injectie een serieus probleem, met zoveel kabels tussen 300 K en temperaturen onder 1 K. Daar komt het mechanische probleem om zoveel kabels in een mengkoelmachine te plaatsen nog bij.

Daarom stellen we in dit werk voor om de elektronica niet op kamertemperatuur te plaatsen, maar op een temperatuur die veel dichter bij die van de kwantumbits ligt, bijvoorbeeld op 4 K. Dit vermindert niet alleen significant het aantal kabels tussen kamertemperatuur en de kwantumbits, maar we kunnen ook profiteren van de lagere elektronische ruis bij zulke temperaturen. De werking van conventionele elektronica bijna 200°C onder zijn normale temperatuurbereik is niet triviaal aangezien het gedrag van schakelingen aanzienlijk verandert en de meeste niet meer werken zoals verwacht kan worden.

In CMOS processen, de belangrijkste technologie in de wereld van geïntegreerde schakelingen, verandert het gedrag van de transistoren, nodig voor de realisatie van elk denkbaar circuit, aanzienlijk bij lage temperaturen. De drempelspanning van de transistoren gaat omhoog, evenals zijn mobiliteit en de subdrempelcurve wordt steiler, om maar een aantal voorbeelden van deze veranderingen te noemen. Ook al zijn er een aantal verbeteringen in de prestaties, andere effecten doen die verbeteringen gedeeltelijk te niet, en derhalve is het nodig om de transistoren te karakteriseren om al deze veranderingen waar te nemen. Met deze karakterisaties bij lage temperaturen kunnen de transistoren gemodelleerd worden en circuits gesimuleerd. Dit maakt het mogelijk de circuits te optimaliseren zodat ze ook bij cryogene temperaturen goed blijven werken. Gelukkig zijn er ook een aantal commercieel beschikbare apparaten die de kou al kunnen weer-

staan, ook al zijn ze daar niet voor ontworpen. We hebben verscheidene commerciële elektronische apparaten aan koude temperaturen blootgesteld en laten zien dat ze zelfs dan nog goed functioneren, zoals een veldprogrammeerbare schakeling (FPGA) geïmplementeerd in een 28 nm CMOS proces. Deze werkt zonder grote problemen bij 4 K, met slechts kleine prestatie veranderingen die tussen de 5 en 10% liggen. Alle door ons geteste programmeerbare schakelingen die bij 300 K werkten, deden het eveneens bij 4 K.

We hebben commercieel beschikbare apparaten, die 200 K onder hun gespecificeerde temperatuurbereik functioneren, gecombineerd met zelf ontworpen CMOS circuits om zodoende een cryogeen uitleesplatform voor spin kwantumbits te bouwen. Dit systeem bestaat uit versterkers, een ADC, een FPGA, spanningsregelaars en een klokgenerator. Het maakt het mogelijk een klein signaal komend van de kwantumbits te versterken en direct daarna te digitaliseren in de FPGA om het daar bij 4 K te kunnen verwerken. Het ontwerp van het systeem is een van de eerste systematische pogingen om een deel van de kwantum–klassieke verbinding bij cryogene temperaturen te realiseren en vormt derhalve de basis voor een toekomstig systeem die de gehele elektronische verbinding van kwantumbits bij lage temperaturen doet functioneren.


Eén van de grootste problemen die moet worden aangepakt bij dit soort cryogene systemen is het energieverbruik. Vermogensbudgetten zijn simpelweg beperkt tot grofweg 1 a 2 Watt bij 4 K en exponentieel lager bij diepere cryogene temperaturen, wat direct een limiet legt op de grootte van het elektronische systeem. Onze aanpak, die commerciële met zelf ontworpen circuits combineert, moet in de toekomst gestaag vervangen worden met een enkele technologie die voldoet aan zowel beperkingen op het gebied van energie als schaalbaarheid. Eén van de beste kandidaten daarvoor is CMOS, een technologie waar de industrie de afgelopen decennia op is gebouwd en die veel heeft geprofiteerd van optimalisaties zoals gedicteerd door de wet van Moore.

ABOUT THE AUTHOR



Harald Homulle was born on the first of October 1990 in The Hague, The Netherlands. After obtaining his high-school diploma (gymnasium) from the Dalton The Hague, he decided to join the TU Delft in pursuit of a bachelor in Electrical Engineering. During the final bachelor assignment, he worked on a study for object recognition for unmanned aerial vehicles. Succeeding the finished bachelor in 2012, he continued with a master in the field of microelectronics. During the master's thesis project, he did an internship at EPFL, Lausanne, Switzerland, on the topic of fluorescence lifetime imaging and the corresponding development of FPGA-based time-to-digital converters. The master was concluded in 2014 with the thesis entitled *Development of a multichannel TCSPC system in a Spartan 6 FPGA*.

Staying in the group of Edoardo Charbon, we endeavoured in the topic of cryogenic electronics. Being the first student of the group to work in this topic, he had to start from scratch in all senses. Luckily, he learned to undertake cryogenic experiments from Enrico Prati, chip design for CMOS characterization from Fabio Sebastiano, and a lot of other relevant things from the students that joined the group along the way. Finally, all of the above can be found in this dissertation. Besides these academic interests, he enjoys to cycle, hike, play chess, listen to soundtracks and classical music, and explore the world.



Welcome
Quantum