Master Thesis
An On-Chip Noise Thermometer

Microelectronics, Electronic Instrumentation group

J.P.G. van Dijk

Delft University of Technology

TU Delft Delft University of Technology

Challenge the future

An On-Chip Noise Thermometer

by

J.P.G. van Dijk

to obtain the degree of Master of Science at the Delft University of Technology, to be defended publicly on Monday January 25, 2016 at 02:30 PM.

Student number:	4065603	
Project duration:	September 1, 2014 – January 25, 2016	
Thesis committee:	Prof. Dr. K. A. A. Makinwa,	TU Delft, responsible full professor
	Prof. Dr. E. Charbon,	TU Delft
	Dr. F. Sebastiano,	TU Delft, daily supervisor

This thesis is confidential and cannot be made public until December 31, 2020.

TU Delft Delft University of Technology

Abstract

Temperature sensors are widely integrated in high performance systems, e.g. microprocessors. If the die temperature becomes too high, the processor must throttle its clock speed to prevent reliability problems. This application requires temperature sensors that can be integrated in the same advanced technologies as the microprocessor, and that can operate from the same low supply voltage. Since many temperature sensors are required, it is important that they only require a minimal amount of trimming.

Thermal noise can be exploited as a temperature sensing principle in any technology with resistive elements. Since thermal noise is a fundamental physical phenomenon, it is inherently accurate and linear. The noise power only depends on the measurement bandwidth, which can be easily calibrated.

Previous noise-based temperature sensors required very long measurement times to achieve decent resolution. This is because only small measurement bandwidths could be achieved with discrete measurement setups. By integrating the noise-thermo-meter on chip, and using the increasing speed of more advanced technologies to our advantage, resolution can greatly improve.

This thesis presents the first on-chip noise-based temperature sensor. A prototype chip has been fabricated in a standard 160 nm CMOS process and it achieves a resolution of 0.93 °C in 1 s. The sensor can achieve an inaccuracy of 6.8 °C (3σ) from -70 °C to 50 °C with a single point trim and an inaccuracy of 6.1 °C (3σ) from -70 °C to 95 °C with a two-point trim. The sensor occupies a die area of 0.15 mm^2 and consumes 1.9 mW from a 1.8 V supply. The design demonstrates that a pure electrical calibration should be possible, and both the resolution and accuracy are expected to improve in more advanced CMOS processes.

Contents

List of Figures

List of Tables

Acknowledgement

My master thesis work and personal development during the last one and a half year would not have been possible without the help of many people.

Firstly, I am very thankful to Prof. Dr. Kofi Makinwa for his support of my master thesis. Most importantly, he has shown me how to properly write a thesis and give a good technical presentation.

I would also like to express my gratitude to my daily supervisor, Dr. Fabio Sebastiano, for his continuous help and guidance throughout my project. We have had many interesting discussions and fruitful coffee-corner chats. This project would not have been a success without his patience and experience. He has often encouraged me, and provided bright ideas, when new problems were encountered.

Moreover, I would like to take this opportunity to also thank Dr. Robert van Veldhoven from NXP for his support of the project and Prof. Dr. Edoardo Charbon for taking the time to review my thesis and conduct my defense.

J.P.G. van Dijk
Delft, January 2016

1

Introduction

This chapter will start by addressing the intended application of the temperature sensor (Section 1.1). Next, a summary of previous research will be given in Section 1.2. Section 1.3 then discusses how a noise-based temperature sensor can improve on the current state-of-the-art.

1.1. Application

As predicted by Moore's law, the number of transistors in integrated circuits (ICs) has doubled approximately every two years. This development has been made possible by the continuing CMOS technology scaling. As a result, high performance microprocessors and systems-on-chip are available that contain billions of transistors operating with a clock speed exceeding 1 GHz.

Due to the increase in clock speed and the large number of integrated devices, the dynamic power consumption of ICs increases. Moreover, leakage currents are also larger for more advanced technologies and higher temperatures. This results in a dramatic increase of power density and die temperature. If no actions are taken, these self-reinforcing trends threaten chip reliability.

Therefore, state-of-the-art microprocessors embed many temperature sensors for thermal monitoring. When the on-chip temperature reaches a reliability limit, the microprocessor must slow down. Hence, the speed and accuracy of these on-chip temperature sensors directly affect the microprocessors performance. A measurement time of less than 10 ms is required, since the temperature can rise as fast as 0.5 °C/ms [3]. With a 5 °C error in the temperature reading, a typical microprocessor's power dissipation must be set 10 W below its actual thermal limit [4]. Considering that most microprocessors dissipate less than 100 W, this translates into a significant loss in computing power.

Since the temperature sensors need to be integrated in the microprocessors, they must be implemented in the same advanced technologies. In addition, due to the scaling of supply voltages in advanced technologies, the sensor must be able to operate from sub-1V supplies. Another requirement for the temperature sensor is that it needs a minimal amount of trimming in order to keep calibration costs down. This is very important, because modern microprocessors contain many temperature sensors [5].

In summary, thermal monitoring applications require accurate low-voltage temperature sensors that can be integrated in advanced CMOS technologies and require minimum calibration.

1.2. Existing approaches and their limitation

Different types of temperature sensors are available. A general temperature sensor survey is presented in [1] and shown in Figure 1.1. From this figure, a general trend of decreasing accuracy with smaller process nodes becomes clear. In this survey four main types of temperature sensors are distinguished: BJT-based, MOS-based, resistor-based and TD-based temperature sensors.

MOS-based temperature sensors

The previously mentioned survey shows that in general MOSFET-based temperature sensors have a low resolution, and are inaccurate even when using multiple trimming points. Therefore, this type of sensor is not further discussed.

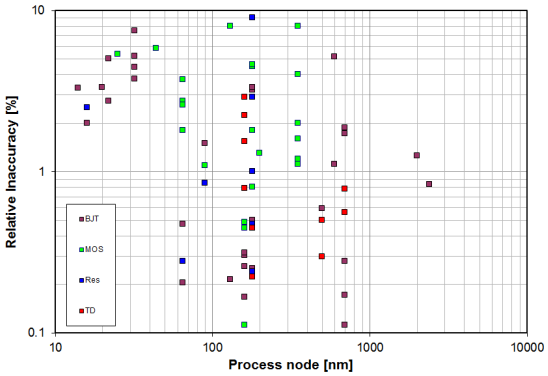

Figure 1.1: The relative inaccuracy versus the process node (Source: [1]).

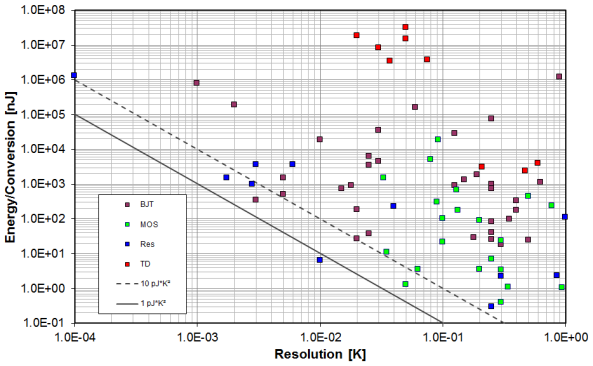

Figure 1.2: The energy per conversion versus the resolution (Source: [1]).

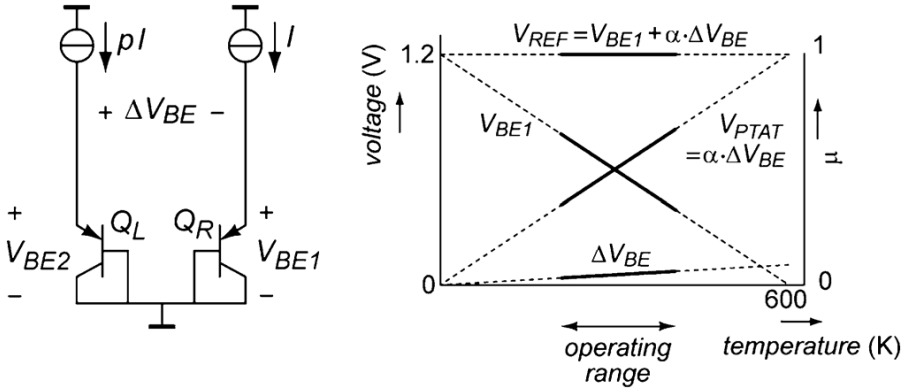

Figure 1.3: BJTs as temperature sensing elements (Source: [2]).

BJT-based temperature sensors

In a BJT-based temperature sensor, bipolar junction transistors (BJTs) are used to measure the temperature, see e.g. [2]. This is achieved by biasing two BJTs at different current densities with a fixed and known ratio. The resulting base-emitter voltage, V_{BE}, and the difference between the two base-emitter voltages, ΔV_{BE}, are used to measure the temperature. This is illustrated in Figure 1.3.

Paper [2] also demonstrates that voltage calibration can be used to trim the temperature sensor. Since this is based on electrical measurements, no temperature-stabilized environment is necessary.

In a CMOS technology, parasitic BJTs are used for the bipolar core of the sensor. While BJT-based temperature sensors can achieve an accuracy below 0.2 °C in mature CMOS technologies, their accuracy deteriorates when integrated in nanometer CMOS technologies. For example, in a 32 nm process an accuracy of only 4.5 °C (3σ) was achieved after a single point calibration and at least a 2-point calibration was necessary to achieve the required resolution [3]. This decrease in accuracy is due to the deterioration of BJT current-gain with scaling [6]. Paper [6] overcomes the beta limitation by using NPN transistors. However, these parasitic devices cannot easily be ported to more advanced technologies.

In summary, BJT-based temperature sensors do not scale well with technology because of the non-idealities of the parasitic devices. Furthermore, fitting the required base-emitter voltage in the decreasing supply voltage of scaled technologies becomes challenging, if not impossible.

Resistor-based temperature sensors

Resistor-based temperature sensors use the temperature dependence of a resistor to determine the temperature. As shown in Figure 1.2, these types of sensors can achieve a high resolution and are quite-energy efficient. However, since the resistors spread and their temperature dependence is non-linear, typically a multi-point thermal calibration is required. For example, in [7] a 3-point trim is required for a 0.15 °C inaccuracy (3σ).

Since many temperature sensors need to be integrated, it is clear that the extensive calibrations required for a resistor-based thermometer are a problem for the intended application.

TD-based temperature sensors

TD-based temperature sensors rely on the temperature-dependent thermal diffusivity (TD) of bulk silicon [8]. The thermal diffusivity of silicon is well-defined for the highly pure silicon used in ICs. This thermal diffusivity can be determined with the help of an electro-thermal filter (ETF). This consists of a heater and a relative temperature sensor located a certain distance from this heater. Its accuracy is then limited by variations in the spacing between these elements. Since lithography improves in newer technologies, this variation becomes smaller and the accuracy improves. However, their power consumption is relatively high, see Figure 1.2.

It can be concluded that TD-based temperature sensors are the first to show improving accuracy with more advanced technologies, but they are not energy-efficient.

1.3. Noise-based temperature sensor

As an alternative temperature sensing principle, thermal noise can be exploited to measure temperature. Any resistor generates thermal noise with a noise power spectral density independent of the resistor's composition or actual resistance given by:

$$P_n = 4kT \tag{1.1}$$

where k is Boltzmann's constant. Hence, any technology that offers a dissipating electrical element can be used to implement a thermal-noise-based temperature sensor. This allows for a pure CMOS implementation of the temperature sensor. In addition, even a resistor biased at zero volt generates noise, which allows for low-voltage operation. For these reasons, noise thermometry is ideal for advanced technologies.

Other advantages derive from the fact that thermal noise is a fundamental physical phenomenon. As a result, it is inherently accurate and linear. The proportionality factor only depends on a well known constant, see Equation 1.1. Furthermore, it has been shown that this type of thermometry can work over a very wide temperature range (2.1 K - 2473 K, see Section 2.1).

A common problem of the noise thermometers presented in literature is their very long measurement time. Because noise, a random signal, is being measured, more information leads to lower uncertainty. More information can in this case be obtained by measuring for a longer time, or over a wider bandwidth. As a result, the uncertainty is inversely proportional to $\sqrt{f_c t_m}$, where f_c is a measure for the bandwidth and t_m is the measurement time (see Section 2.2).

None of the published noise-based temperature sensors are integrated[1]. As a result, the measurement bandwidth has been limited to below 1 MHz due to parasitic capacitances associated with the resistors or their connections (see Section 2.1 for an overview). Furthermore, the bandwidth has often been limited because many of the setups had problems with EMI (electromagnetic interference), which is hard to distinguish from the noise to be measured.

1.3.1. Integrated noise-based temperature sensor

Integrating a noise thermometer on-chip has several advantages. First of all, many of the problems associated with long sensor leads will be removed. These problems include the additional resistance of these wires, signal leakage and EMI pick-up. Secondly, higher measurement bandwidths are possible. Analog circuits with bandwidths exceeding 100 MHz can be realized efficiently on-chip. With newer technologies, the maximum frequency of circuits is expected to increase, which helps in achieving even higher resolutions or lower measurement times. Therefore, this sensing principle is expected to scale well with technological advances. Finally, in ICs very low cost digital electronics for further signal processing are available.

Besides these advantages, it might be possible to integrate additional circuitry for trimming purposes. For the application, it is preferable if the calibration can be done electrically. This means that by providing additional electrical references, the actual temperature can be determined. By comparing this with the temperature given by the temperature sensor, a calibration can be performed. This obviates the need for a calibration using a temperature-stabilized environment, which is slow and therefore expensive.

As shown by Equation 1.1, the only unknown is the temperature to be measured. However, in order to measure the noise power spectral density as given by this equation, the noise power will be evaluated in a certain bandwidth. As a result, this bandwidth must be accurately known. However, many microprocessors already have an accurate temperature-independent clock signal available which can be used to measure or control this noise bandwidth. Besides the noise bandwidth, the gain of the readout circuit can also be calibrated electrically.

[1] An integrated noise-thermometer is proposed in [9], but only simulation results to verify the concept are presented.

There are, however, several challenges in implementing this temperature sensing principle on-chip. Measuring noise power requires both the measurement of noise current and noise voltage. Alternatively, only the noise voltage or noise current could be measured. But this requires knowledge of the resistance value over temperature, which is non-linear and subject to spread (see Section 1.2). On the other hand, measurement of a voltage (or current) implies the availability of a voltage (or current) reference. True current references do not exist and standard voltage references are based on the same principles as a BJT-based temperature sensor (see Section 1.2). As discussed, these devices are unsuitable for use in advanced technologies. These challenges need to be tackled in order to realize a true on-chip noise-based thermometer.

1.3.2. Applications

From the mentioned advantages, it is clear that on-chip noise-thermometry is a good candidate for thermal monitoring in advanced technologies. Only resistors are required, which are available in any technology, and there is no minimum voltage requirement. The sensing principle is accurate, and the resolution is expected to improve with more advanced technologies. It has also been shown that a fast and inexpensive electrical calibration of the sensor is possible.

Since noise-based temperature sensors have never been integrated on-chip, it cannot be anticipated what the required area and power consumption are. In case they are too large for the integration of many temperature sensors, as is required for thermal monitoring of modern microprocessors, a single noise-based temperature sensor could be integrated on the chip. This sensor could then be used for the background calibration of many less accurate conventional temperature sensors. In [10] a thermal monitoring solution is proposed that uses many relative temperature sensors and just a few absolute temperature sensors. Similarly, conventional temperature sensors could be used to provide relative temperature information, while the absolute temperature could be provided by a noise-based temperature sensor. Because of the possibility of an electrical calibration, the sensor could even be integrated solely for the purpose of performing a one-time calibration of other sensors and circuits.

As described before, noise-thermometry has a very wide sensing range. Practically, the operating range will only be limited by the readout electronics, but CMOS circuits operating in the range from 4 K up to 300 °C have been demonstrated ([11][12]). This opens the way to more applications. Such temperature sensors could for example be used in high temperature environments, as is the case for automotive applications and oil drilling [12]. In the low-temperature regime, such sensors could for example be used to monitor the supporting electronics for quantum computers, which operate near absolute zero [13].

1.4. Thesis outline and main objective

Since an integrated noise-thermometer has never been demonstrated before, it is hard to predict its performance. Therefore, the main goal of this thesis is to demonstrate the feasibility of performing on-chip noise thermometry.

Chapter 2 will provide an overview of prior (non-integrated) noise-thermometers. This chapter will conclude with the basic requirements for the noise-thermometer to be designed. Based on these requirements, a system level design is described (Chapter 3). Details on circuit implementation of this system are discussed in Chapter 4. Chapter 5 presents the measurement results of the fabricated chip. Conclusions and recommendations for future work follow in Chapter 6.

2

Noise Thermometry

This chapter will first discuss existing noise thermometers and related sensing techniques in Section 2.1. The measurement uncertainty of the different detectors are analysed in Section 2.2. Finally, Section 2.3 discusses how suitable those techniques are for on-chip integration, with specific attention to the intended application.

2.1. State-of-the-art

Noise-based temperature sensors have been used because of their independence of resistance material or composition. The only property of interest, the actual resistance value, can easily be measured with a high accuracy. For this reason, noise-thermometers have been used in industrial measurements were the sensor is subject to change, for instance in nuclear reactors. On the other hand, because of the fundamental nature, highly accurate temperature sensors could be designed. As a result, this sensing principle is particularly attractive for metrological applications, for instance for the measurements of fixed-point temperatures.

Different discrete noise thermometers have been proposed in the past, see Table 2.1. Besides measuring the absolute temperature, noise measurements have been used to determine Boltzmann's constant [14], impedance [15][16] and battery lifetimes [17]. From this table it can be seen that noise thermometers can function over a wide temperature range.

Since this work aims to realize a CMOS temperature sensor, techniques that require non-CMOS devices, such as super-conductors are not considered. Examples are shot noise thermometers requiring a tunnel junction [28], or noise thermometers employing Josephson junctions [26][14].

The block diagram of a general noise thermometer is shown in Figure 2.1. Theoretically, an amplifier is not required, but is often included because the noise signal is extremely small. The resistance generates a white noise

Table 2.1: Noise measurements found in literature, showing their range, measurement bandwidth, measurement time, and measurement uncertainty.

Paper	Temperature Range (K)		Noise BW (kHz)	Measurement time (s)	Relative uncertainty (RMS, %)
[18]	273 - 1300		41	10	0.13
[19]	90 - 100	(boiling point of oxygen)	100	25200	0.008
[20]	< 1770		500	1	0.3
[21]	2.1 - 4.3	(4He at various vapour pressures)	12	33000	0.01
[22]	2473	(reactor fuel rods)	-	-	-
[23]	77 - 1235	(liquid nitrogen, ice point, freezing point of tin, zinc, antimony and silver)	20...1000 (adj.)	240	0.3
[24]	273	(ice point)	40	3600	0.04
[24]	1233, 1357	(freezing point of copper)	95	3000	0.01
[25]	83	(liquid nitrogen)	20	7200	0.012
[26]	505, 693	(freezing point of zinc and tin)	200	> 16900	0.0039
[16]	273	(ice point)	20	14400	2.5
[27]	303	(freezing point of gallium)	40	684000	0.0031
[14]	273	(determining Boltzmann's constant)	638	68400	0.0017

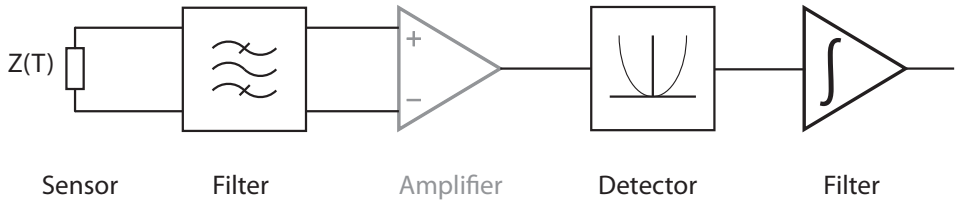

Figure 2.1: The block diagram of a general noise thermometer.

voltage with a power spectral density given by [29]:

$$S_v(f) = 4hf \cdot Re\{Z\} \cdot \left[\frac{1}{2} + \frac{1}{e^{\frac{hf}{kT}} - 1}\right] \approx 4kT \cdot Re\{Z\} = 4kTR \tag{2.1}$$

or equivalently for the noise current:

$$S_i(f) = 4hf \cdot Re\left\{\frac{1}{Z}\right\} \cdot \left[\frac{1}{2} + \frac{1}{e^{\frac{hf}{kT}} - 1}\right] \approx 4kT \cdot Re\left\{\frac{1}{Z}\right\} = 4kT/R \tag{2.2}$$

where h is Planck's constant, k is Boltzmann's constant and Z is the complex impedance. The approximation is only valid for low frequencies, such that $f \ll kT/h$, which is 6.25 GHz at 300 K [15]. The last simplification is only valid if the sensor is purely resistive, with a resistance R.

The resulting noise signal is subsequently amplified and filtered to obtain a noise signal whose power has a well-defined relation to temperature. This noise power is then measured using a detector. The last filter obtains the mean output of the detector and defines the measurement bandwidth. Measurement of the temperature then requires a known resistance value. Furthermore the bandwidth and gain of the various stages must be known.

2.1.1. Uncertainty and inaccuracy

For conventional temperature sensors, noise from the readout circuit is associated with the measurement uncertainty or resolution. However, since for a noise-based temperature sensor, noise is the signal to be measured, it is important to realize the effects of noise from the readout circuit.

Any additional noise present before the signal detection will appear as an error in the noise power measurement. Therefore, this type of readout noise results in inaccuracy of the temperature sensor. On the other hand, readout noise after the signal detection increases the uncertainty of the measurement.

Another important point is that the relation of the mean output of the detector to the variance of the input signal differs for different detectors. In case this relation is non-linear, the uncertainty and inaccuracy must be properly translated into the uncertainty and inaccuracy of the measured temperature. In Section 2.2, the uncertainty of various detectors is analyzed and these uncertainties are distinguished as 'detector/measurement uncertainty' and 'uncertainty in the calculated power' respectively.

2.1.2. Sensing techniques

This section describes the techniques that have been applied to calibrate the unknown gain factors and compensate for additional noise in the measurement setups.

The ratio-metric measurement and the switching technique

In order to deal with the large number of unknown gain factors (resistance, gain, effective noise bandwidth etc.), some of which are difficult to determine accurately, noise thermometers usually employ a reference resistor, R_0, at a reference temperature, T_0. Assuming both noise sources are measured in the same bandwidth, the temperature can be determined using:

$$\frac{T}{T_0} = \frac{v^2}{v_0^2}\frac{R_0(T_0)}{R(T)} \tag{2.3}$$

where v_0^2 is the integrated noise voltage from the reference resistor and v^2 the integrated noise voltage from the sensor resistor $R(T)$. This technique requires the measurement of the resistance ratio. Since the resistance is usually strongly temperature dependent, this measurement must be performed at every temperature measurement.

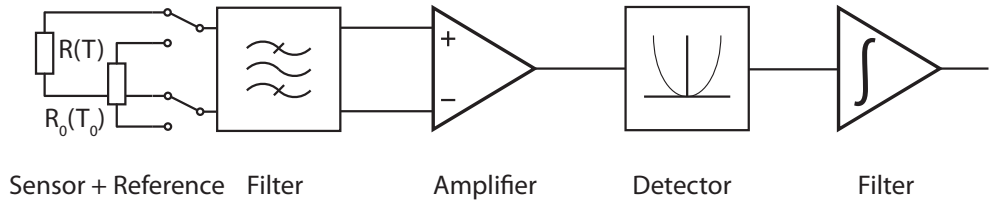

Figure 2.2: The block diagram of a general noise thermometer with the switching technique.

With the ratio-metric measurement, the reference resistor and reference temperature need to be determined only once. Better accuracy can be achieved by continuously switching between the sensor and the reference, see Figure 2.2. This switching technique can compensate for drift in the gain, and as a result the stability requirements are reduced. This technique has been used for almost all the noise thermometers reported in the literature.

An additional advantage can be obtained by balancing the noise power from the sensor with the noise power from the reference. Assume that an error (ξ) is made that depends on the measured noise power:

$$v_{out}^2 = v^2 \cdot (1 + \xi(v^2)) = v^2 \cdot \left(1 + \xi(v_0^2) + \frac{\partial \xi(v_0^2)}{\partial v^2}\frac{(v^2 - v_0^2)}{1!} + \frac{\partial^2 \xi(v_0^2)}{\partial {v^2}^2}\frac{(v^2 - v_0^2)^2}{2!} + ...\right) \tag{2.4}$$

where the Taylor series expansion has been applied. The measured temperature ratio is then (assuming $\xi(v_0^2) = 0$) [29]:

$$\frac{T}{T_0} \approx \frac{v^2}{v_0^2}\frac{R_0(T_0)}{R(T)} \cdot \left[1 + \frac{\partial \xi(v_0^2)}{\partial v^2}(v^2 - v_0^2)\right] \tag{2.5}$$

and it can be seen that the error term goes to zero as the noise power is matched. In case the noise powers are exactly matched in a full balancing method, almost any non-linearity in the readout circuit can be allowed. In addition, the readout noise that is common to the sensor and reference is removed. In case of incomplete balancing, additional measurements can be performed to compensate for the error [30].

Problems with the ratio-metric measurement and the switching technique

Equations 2.3 and 2.5 are only valid when both noise sources are measured in the same bandwidth. However, the noise bandwidth is affected by the pole caused by the sensing (or reference) resistance and the parasitic capacitance originating from the input capacitance of the amplifier and the capacitance of the wiring to the amplifier. Inaccuracies arise due to the fact that the pole, and hence the bandwidth, changes with the resistance variation over temperature [27].

The matching of the noise powers in the switching technique can be achieved by adjusting the reference resistance. This will further influence the pole locations. One can compensate for this problem by tuning the capacitance over the resistor, which requires a high and low frequency channel in the readout [19]. The low frequency channel is used to match the power spectral densities, at which point the high frequency channel can be used to compare the pole locations. Alternatively, the wiring could be carefully matched [14].

In addition, the noise from the readout circuit can depend on the source resistance, leading to errors. Solutions have been proposed that introduce a variable attenuator to match the noise powers [29]. While the trimming has usually been performed manually, also automatic control using feedback has been demonstrated [23].

Another problem with these techniques is that a reference noise source is required with the same spectrum as the sensor noise. Besides using a reference resistor at a reference temperature [19][24][26][27], other references have been used. For example, a programmable quantum voltage noise source (QVNS) can generate pseudo-random-noise waveforms with a white spectrum [26][14]. However, these devices use Josephson junctions which only operate at very low temperatures.

Shot noise has also been used as a reference [20][23]. The spectrum of shot noise is the same as for thermal noise, but the power spectral density is independent of temperature:

$$S_i(f) = 2qI \tag{2.6}$$

where q is the elementary charge, and I is the large signal bias current of a diode. In these papers it has been demonstrated that only a voltage reference is required for the temperature measurement. Other techniques exist that eliminate the need for a resistance measurement, i.e. by filtering the noise over a capacitor or by measuring the noise power.

Filtered noise

Filtering the resistance noise over a parallel capacitor [20] leads to an integrated noise level of kT/C, since the effective noise bandwidth is $\frac{1}{4RC}$:

$$ENBW = \frac{1}{|H(0)|^2} \cdot \int_0^\infty |H(f)|^2 df = \int_0^\infty \frac{1}{1+(2\pi fRC)^2} df = \frac{arctan(2\pi fRC)}{2\pi RC}\Big|_0^\infty = \frac{\pi/2}{2\pi RC} = \frac{1}{4RC} \quad (2.7)$$

which is only valid in case of sufficient settling [31]. Although now the capacitance value is unknown and needs to be measured, there is an advantage. Capacitors are generally far less temperature dependent and as a result this measurement needs to be performed only once. Even if the temperature dependence is too large, the capacitor can be kept at a constant temperature, away from the sensor.

For this type of readout, the parallel capacitor must be dominant in determining the noise bandwidth. This requires that the amplifier and detector have a higher bandwidth in order to perform an accurate measurement.

Noise power

Another noise measurement method that does not require a measurement of the resistance is obtained by measuring both the noise voltage and noise current associated with a resistance [18]. The product gives the total noise power, which is independent of the sensor resistance:

$$\begin{aligned} v_n &= \sqrt{4kTR \cdot B_v} \cdot K_v \\ i_n &= \sqrt{4kT/R \cdot B_i} \cdot K_i \\ P_n &= v_n \cdot i_n = 4kT \cdot K_v K_i \cdot \sqrt{B_v B_i} \end{aligned} \quad (2.8)$$

where K_v, K_i, B_v and B_i are the gain and bandwidth of the voltage and current readout. These constants have to be determined only once, in case they are temperature independent and have sufficiently low drift. The noise voltage and noise current can be determined in another bandwidth and over another resistor, as long as the ratio is stable and accurate.

The proposed method does however not work in case the sensor has a complex impedance, since $Re\left\{\frac{1}{Z}\right\} \neq \frac{1}{Re\{Z\}}$. A solution to this problem has been proposed in [16] that uses 4 channels to measure cross-spectra in order to determine the complex source impedance and absolute temperature.

Correlation

As explained in Section 2.1.2, the switching technique eliminates the effects of amplifier noise when the sensor and reference noise powers are exactly matched. It does, however, not compensate for differences in noise generated in the sensor leads that can be different from the noise generated in the reference leads.

A method that can compensate for the amplifier noise, and also for the noise of the leads, is by using a correlator. Here two amplifiers and a multiplier are used, see Figure 2.3. The multiplier, also called correlator, is the detector and removes the uncorrelated noise sources (this will be shown mathematically in Section 2.2.1). In order to also remove noise from the leads, a 4-wire connection must be used. With the 4-wire connection, the resistance can be measured more accurately and only the noise of the actual sensor is detected. This technique has been used in most of the noise thermometers, with the multiplier either in hardware [21][23][24][27] or software [25][17][26][14].

It is important to realize that only uncorrelated noise sources can be removed. In case for instance the amplifier injects a noisy current into the sensor resistance, this results in a noisy voltage over the sensor that is detected by both amplifiers. These noise voltages are correlated and hence not removed [32]. For this correlation technique to work, it is also important that the phase responses of the two channels are matched, otherwise also uncorrelated sources can contribute to the output signal [21].

2.1.3. Detectors

Different detectors can be used to measure the noise power. The input to the detector is a noise signal with zero mean and a certain variance. The output of the detector is a function of this noise variance. Therefore, detectors must inherently be non-linear devices. Several detectors are reviewed in the following.

ADC with DSP

The variance computation can easily be implemented in the digital domain. This requires that the noise signal is first digitized by an analog-to-digital converter. An additional advantage is that extremely sharp and accurate

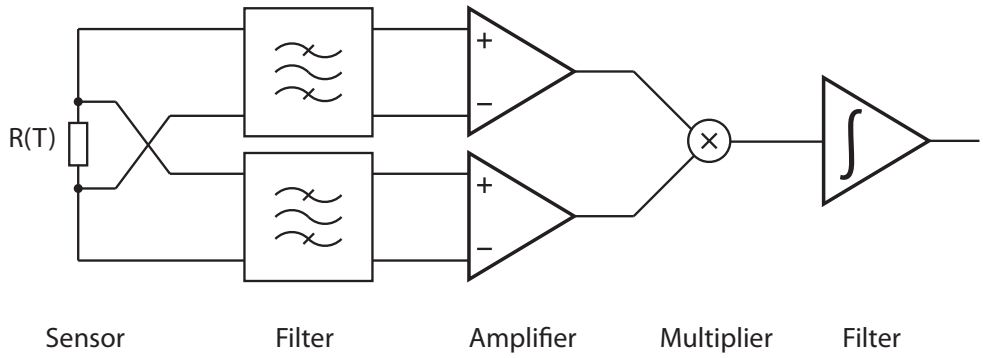

Figure 2.3: The block diagram of a general noise thermometer with a correlating detector.

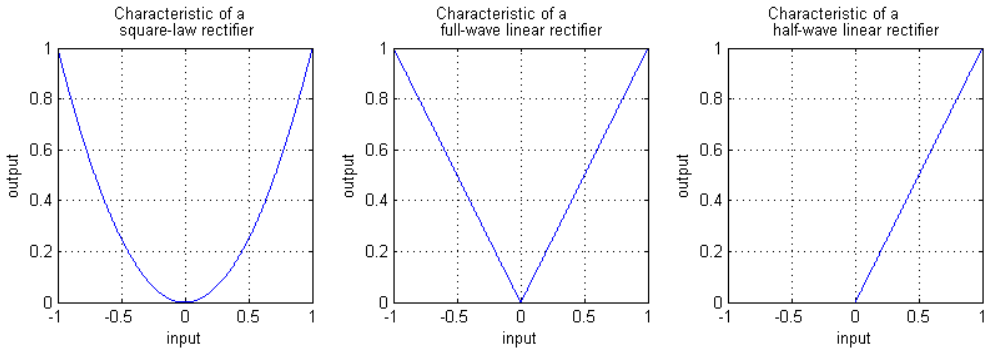

Figure 2.4: The input-output characteristic of the square-law rectifier and the linear rectifiers.

filters can be designed in the digital domain for the definition of the noise bandwidth, e.g. a 42^{nd} order filter is used in [25].

Existing ADC-based noise thermometers use high resolution converters, ranging from 14-bit [26][27] to even 20-bit [14]. However, it will be shown in Section 2.3.2 that the minimum number of bits required is much lower.

Square-law rectifier

In a square-law rectifier, the output signal is proportional to the square of the input signal, i.e.: $y \propto x^2$ (see Figure 2.4). Also True-RMS detectors fall in this category, since the input voltage is first squared. A special case is the correlating detector (see Section 2.1.2) where the noise signal is presented to two inputs of a general multiplier, also leading to square-law rectification. A square-law rectifier is the optimal detector for Gaussian signals, and therefore requires the shortest measurement time [29].

In noise thermometers, either digital voltmeters [18] or off-the-shelf integrated-circuits (AD425 [21], AD422 [23], AD534 [24]) have been used. Since a multiplication is performed, an additional reference is required. For circuit implementations based on translinear circuits this is a current [33][34]. Other implementations rely on transistor characteristics, so that the result of the conversion is proportional to device transconductance [35][36]. This cannot be accurately controlled, and also depends on the temperature (via the mobility).

Linear rectifier

Another detector used in noise thermometers is the linear rectifier or envelope detector. For a full-wave linear rectifier, the output signal is described by $y \propto |x|$, while for a half-wave linear rectifier only the positive signals are passed to the output (see Figure 2.4). After the detection, the output signal must be squared to obtain the power of the input signal.

For the presented noise thermometers, no circuit details have been specified, but the circuit likely simply consists of rectifying diodes in a half bridge configuration [19]. The diode threshold voltage can be overcome by using feedback [37][38]. The disadvantage of these circuits is that the amplifier requires a high unity gain frequency. In case 0.1% error can be tolerated, the gain of the amplifier should be around 1000x over the complete signal range. In case the noise bandwidth is 10 MHz, this means that the UGF must be around 10 GHz. Other implementations exist that do not use feedback, but these again rely on the characteristics of transistors [39][40].

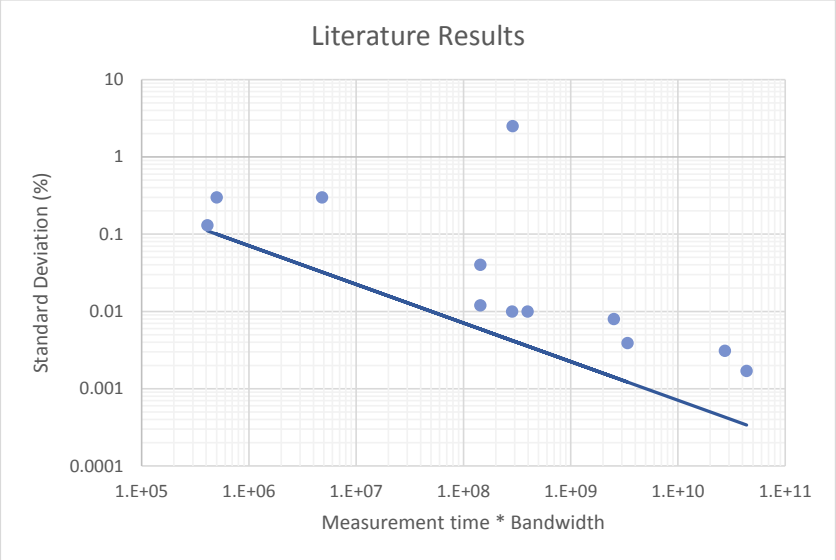

Figure 2.5: The measurement uncertainty of the noise measurements found in literature.

Threshold detector

In a threshold detector, the input signal is compared with a fixed threshold voltage. The number of times the input voltage exceeds this threshold in a certain amount of time is a measure of the noise power. An advantage of this method is that counting can be done with very high accuracy and a virtually unlimited length of scale. A disadvantage is that the output of the detector has a non-linear relation with the noise power and hence the absolute temperature.

The threshold detector was only used very early in the '60s [41][42][30], and no numbers on the accuracy or resolution were reported. Probably this type of detector has not been further investigated since it was mentioned that a 'special-purpose pulse height discriminator' was just as difficult to implement as an RMS-to-DC converter and this new method showed no improvement [15]. However, a threshold detector is simply a comparator, which is a very common circuit block these days. In later noise measurement systems, threshold detectors were used, not to measure the amount of noise, but to detect the presence of other types of noise or interference [24][17].

2.2. Uncertainty analysis

In Figure 2.5, the measurement uncertainty of the different noise measurements found in literature is shown. It can be seen that the uncertainty can only be decreased by increasing the product of the measurement time and measurement bandwidth. This section will show why that is the case, and that there are slight difference between the uncertainties that can be obtained with different detectors and measurement techniques.

2.2.1. Mathematical derivation

Throughout this section the following symbols and operators will used:

- $E[X(t)]$: the expectation of X;
- μ_X, σ_X^2: the mean and variance of X;
- $R_{XX}(\tau)$, $C_{XX}(\tau)$: the auto-correlation and auto-covariance of X;
- $S_X(f)$: the power spectral density (PSD) of X;
- $\mathcal{F}\{\}$, $\mathcal{F}^{-1}\{\}$: the Fourier and inverse Fourier transforms;
- $a \star b$: the convolution of a and b;
- $rect()$, $tri()$: the rectangular and triangular function (a convolution of two rectangular functions);
- α, α': the relative uncertainty in the measured power and in the detectors' output signal.

The relative uncertainty is defined as the standard deviation of the stochastic variable (a signal or noise power) normalized to the mean of the stochastic variable (a signal or noise power).

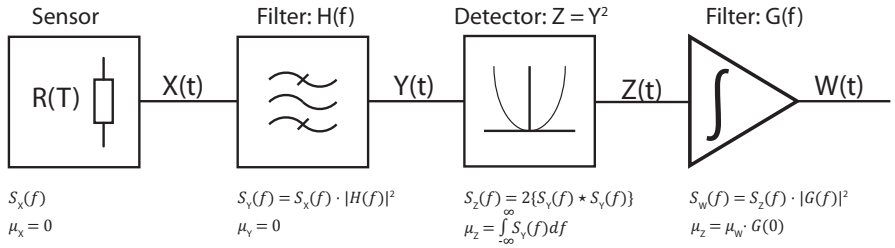

Figure 2.6: Measurement system using a square-law rectifier, with the signals shown.

Square-law rectifier

In this section an elaborate analysis of the measurement uncertainty of the square-law rectifier as detector in the measurement system will be performed, see Figure 2.6.

For the input signal a wide-sense stationary process[1] with zero mean is assumed:

$$\begin{aligned} E[X(t)] &= \mu_X = 0 \\ E[X(t)X(t+\tau)] &= \sigma_X^2 = R_{XX}(\tau) = C_{XX}(\tau) = \mathcal{F}^{-1}\{S_X(f)\} \end{aligned} \tag{2.9}$$

The noise signal is band-limited using a band-pass filter with a zero at DC with transfer function H(f):

$$S_Y(f) = S_X(f) \cdot |H(f)|^2 \tag{2.10}$$

The mean of the signal is multiplied by the DC-gain, and hence remains zero. This signal is then applied to the squarer which performs the operation $Z(t) = Y^2(t)$:

$$\mu_Z = E[Z(t)] = E[Y^2(t)] = \sigma_Y^2 = \int_{-\infty}^{\infty} S_Y(f)df \tag{2.11}$$

and it can be seen that the mean output of the detector is the power of the input signal. To determine the measurement uncertainty, the PSD of the output noise signal must be calculated. This is easiest via the time-domain using the auto-covariance:

$$\begin{aligned} C_{ZZ}(\tau) &= E[Z(t)Z(t+\tau)] - E[Z(t)]E[Z(t+\tau)] = E[Y^2(t)Y^2(t+\tau)] - E[Y^2(t)]^2 \\ &= 2E[Y(t)Y(t+\tau)]^2 = 2C_{YY}^2(\tau) \qquad \text{([43], Chapter 12)} \end{aligned} \tag{2.12}$$

Hence, the output PSD of the squarer is related to the PSD of the input signal as:

$$S_Z(f) = 2(S_Y(f) \star S_Y(f)) \tag{2.13}$$

To reduce the measurement uncertainty, this signal is filtered using a LPF with transfer function G(f):

$$S_W(f) = S_Z(f) \cdot |G(f)|^2 \tag{2.14}$$

and the mean of the signal will be multiplied by the DC-gain. A summary of the mean and variance of the various signals is given in Figure 2.6.

The relative uncertainty in the measured noise power, α, can be determined using these results:

$$\begin{aligned} \alpha^2 &= \frac{\sigma_W^2}{\mu_W^2} = \frac{\int_{-\infty}^{\infty} S_W(f)df}{\mu_W^2} \\ &= \frac{2\int_{-\infty}^{\infty} [S_Y(f) \star S_Y(f)] \cdot |G(f)|^2 df}{\left[\int_{-\infty}^{\infty} S_Y(f)df \cdot G(0)\right]^2} \\ &\approx 2 \cdot \frac{\int_0^{\infty} S_Y^2(f)df}{\left[\int_0^{\infty} S_Y(f)df\right]^2} \cdot \frac{\int_0^{\infty} |G(f)|^2 df}{G^2(0)} \end{aligned} \tag{2.15}$$

[1] The mean and auto-covariance are independent of time.

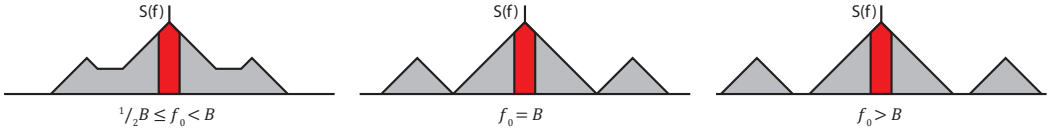

Figure 2.7: The possible power spectral densities described in Equation 2.20.

where the approximation $S_Y(f) \star S_Y(f) \approx \int_{-\infty}^{\infty} S_Y^2(f)df$ has been used [44].

In this equation the first term is the inverse of the correlation bandwidth (Δf_c, [44]) of $\sqrt{S_Y(f)} = \sqrt{S_X(f)} \cdot |H(f)|$. The second factor is the effective noise bandwidth of G(f). Hence, the uncertainty of the measurement can be simplified to:

$$\alpha = \frac{\sigma_W}{\mu_W} \approx \sqrt{2\frac{ENBW_{LPF}}{\Delta f_c}} \tag{2.16}$$

This clearly shows that the uncertainty depends on the measurement bandwidth (related to Δf_c). The dependence on the measurement time is hidden in the effective noise bandwidth of the low pass filter. Consider for example an integrator for which $ENBW = \frac{1}{2t_i}$.

By considering the practical case of band-limited thermal noise as input signal, the validity of the approximations in the previous analysis can be evaluated. The PSD of the input noise voltage is given by:

$$S_X(f) = 4kTR \tag{2.17}$$

After filtering this using a brick-wall filter with bandwidth B centered around the center frequency f_0:

$$S_Y(f) = \frac{1}{2}S_{Y_B}(f \pm f_0) \text{ where } S_{Y_B}(f) = 4kTR \cdot rect\left(\frac{f}{B}\right) \tag{2.18}$$

Using the Fourier transform the auto-covariance follows as:

$$C_{YY}(\tau) = C_{Y_BY_B}(\tau) \cdot cos(2\pi f_0 \tau) \text{ where } C_{Y_BY_B}(\tau) = 4kTRB \cdot sinc(B\tau) \tag{2.19}$$

The mean output of the detector is described by this equation for $\tau = 0$ (see Equation 2.11) and shows that the correct variance, $4kTRB$, will be detected. Using Equation 2.12 and Equation 2.13, the auto-covariance and PSD after the square-law detector are found:

$$\begin{aligned} C_{ZZ}(\tau) &= 2C_{YY}^2(\tau) = 2 \cdot [4kTRB \cdot sinc(B\tau) \cdot cos(2\pi f_0 \tau)]^2 \\ S_Z(f) &= (4kTR)^2 \cdot B \cdot \left[tri\left(\frac{f}{B}\right) \star \left(\delta(f) + \frac{1}{2}\delta(f \pm 2f_0\right)\right] \end{aligned} \tag{2.20}$$

For the applied bandpass filter the correlation bandwidth is simply B. However, as explained, using the correlation bandwidth leads to an approximation of the uncertainty. Which approximations are necessary to obtain this result will be shown now.

With the help of Figure 2.7, which shows the possible power spectral densities described in Equation 2.20, it can be seen that the terms around $\pm 2f_0$ can be removed in case $f_0 \geq B$. In the other case, neglecting the terms around $\pm 2f_0$ is only valid if the low-pass filter has a small enough bandwidth. To get rid of the convolution in equation 2.15 the following approximation has been made:

$$(S_Y \star S_Y)(f) = \int_{-\infty}^{\infty} S_Y(v)S_Y(f-v)dv \approx \int_{-\infty}^{\infty} S_Y^2(f)df \tag{2.21}$$

which is only valid since this approximation will be evaluated around DC, which is again only a valid approximation if the low-pass filter has a small enough bandwidth. Using these approximations, the following simplification follows for the term that was approximated to the correlation bandwidth:

$$\frac{2\left[S_Y(f) \star S_Y(f)\right]}{\left[\int_{-\infty}^{\infty} S_Y(f)df\right]^2} \approx \frac{1}{B} \cdot tri\left(\frac{f}{B}\right) \tag{2.22}$$

This equation evaluates to $1/B$ around DC, which is the same as would have been found when directly applying the correlation bandwidth. From this discussion it can be seen that the only assumption necessary to make the correlation bandwidth approximation is that the low-pass filter following the detector has a small enough bandwidth. Small enough means in this case much lower than the noise bandwidth, which is clearly the case for any of the noise thermometers in Table 2.1.

Correlating detector

Using the correlation technique as described in Section 2.1.2 is similar to using the square-law detector. The measurement uncertainty will however be influenced by the amount of uncorrelated noise that needs to be removed.

Instead of multiplying the same noise signal, the multiplier will multiply the noise signal corrupted with uncorrelated noise: $Z(t) = (Y(t) + N_1(t)) \cdot (Y(t) + N_2(t))$. And the mean output of the correlator follows as:

$$\begin{gathered}\mu_Z = E[Z(t)] = E[Y^2(t) + Y(t)N_1(t) + Y(t)N_2(t) + N_1(t)N_2(t)] \\ = E[Y^2(t)] + E[Y(t)](E[N_1(t)] + E[N_2(t)]) + E[N_1(t)]E[N_2(t)] = E[Y^2(t)]\end{gathered} \tag{2.23}$$

Which shows that only the power of the desired noise is measured. For the auto-correlation of the output signal (only the non-zero terms):

$$\begin{gathered}R_{ZZ}(\tau) = E[Z(t)Z(t+\tau)] \\ = E[Y^2(t)Y^2(t+\tau)] + E[Y(t)Y(t+\tau)](E[N_1(t)N_1(t+\tau)] \\ +E[N_2(t)N_2(t+\tau)]) + E[N_1(t)N_1(t+\tau)]E[N_2(t+\tau)N_2(t)]\end{gathered} \tag{2.24}$$

And then for the auto-covariance and PSD:

$$\begin{gathered}C_{ZZ}(\tau) = R_{ZZ}(\tau) - \mu_Z^2 = R_{YY}^2(\tau) + (R_{YY}(\tau) + R_{N1N1}(\tau))(R_{YY}(\tau) + R_{N2N2}(\tau)) \\ S_Z(f) = [S_Y(f) \star S_Y(f)] + [(S_Y(f) + S_{N1}(f)) \star (S_Y(f) + S_{N2}(f))]\end{gathered} \tag{2.25}$$

When solving this for band-limited white noise as input signal, exactly the same PSD as in Equation 2.20 and visualized in Figure 2.7 are obtained, but with a different power:

$$S_Z(f) = \frac{1}{2}\left[(4kTR)^2 + (4kTR + 4kTR_{eq,N1})(4kTR + 4kTR_{eq,N2})\right] \cdot B \cdot \left[tri\left(\frac{f}{B}\right) \star \left(\delta(f) + \frac{1}{2}\delta(f \pm 2f_0\right)\right] \tag{2.26}$$

where $R_{eq,N1}$ and $R_{eq,N2}$ are the equivalent resistors generating the noise signals $N_1(t)$ and $N_2(t)$ respectively. After passing this through the low-pass filter, the amount of power is higher by a factor (for $R_{eq} = R_{eq,N1} = R_{eq,N2}$):

$$\frac{\sigma_{correlator}^2}{\sigma_{square}^2} = 1 + \frac{R_{eq}}{R} + \frac{1}{2}\left(\frac{R_{eq}}{R}\right)^2 \tag{2.27}$$

The measurement uncertainty is then higher by the square root of this number, which is the same as reported in [29]:

$$\alpha = \sqrt{2\frac{ENBW_{LPF}}{\Delta f_c}} \cdot \sqrt{1 + \frac{R_{eq}}{R} + \frac{1}{2}\left(\frac{R_{eq}}{R}\right)^2} \tag{2.28}$$

The increase in measurement uncertainty can be understood by remembering that there is additional noise that needs to be filtered out. However, the additional noise does not lead to an error in the measurement. In conclusion, inaccuracy has been traded for uncertainty.

Switching technique

The previous section analysed the increased measurement uncertainty when considering the correlation technique. Similarly, the switching technique has a higher measurement uncertainty if there is additional noise that must be removed. Assuming again a square-law detector, the following two signals are obtained after detection in the two separate measurements:

$$\begin{aligned} Z_1(t) &= (Y_1(t) + N_1(t))^2 \\ Z_2(t) &= (Y_2(t) + N_2(t))^2 \end{aligned} \tag{2.29}$$

for which the resulting auto-covariance and PSD for measurement i is given by:

$$\begin{aligned} C_{ZiZi}(\tau) &= 2\left(C_{YiYi}(\tau) + C_{NiNi}(\tau)\right)^2 \\ S_{Zi}(f) &= 2\left((S_{Yi}(f) + S_{Ni}(f)) \star (S_{Yi}(f) + S_{Ni}(f))\right)^2 \end{aligned} \tag{2.30}$$

Which has again the same shape for band-limited white noise as input:

$$S_Z(f) = \left[4kTR + 4kTR_{eq}\right]^2 \cdot B \cdot \left[tri\left(\frac{f}{B}\right) \star \left(\delta(f) + \frac{1}{2}\delta(f \pm 2f_0\right)\right] \tag{2.31}$$

This is higher than the power spectral density given in Equation 2.20 by a factor:

$$\frac{\sigma^2_{switching}}{\sigma^2_{square}} = \left(1 + \frac{R_{eq}}{R}\right)^2 \tag{2.32}$$

However, the final quantity of interest, the signal power, is obtained by subtracting the two measurements. This subtraction increases the measurement uncertainty by a factor $\sqrt{2}$. In case the two measurements cannot be performed at the same time, the measurement uncertainty is increased by another factor $\sqrt{2}$ if the same total measurement time is used.

Full-wave linear rectifier

In [43] (Chapter 13) it is shown that the relative uncertainty of the signal out of a full-wave linear rectifier is two times lower than the relative uncertainty found for the square-law rectifier:

$$\alpha' = \frac{\sigma_W}{\mu_W} \approx \frac{1}{2}\sqrt{2\frac{ENBW_{LPF}}{\Delta f_c}} \tag{2.33}$$

However, the mean output of the detector is in this case for Gaussian signals:

$$\mu_Z = E[Z(t)] = 2\int_0^\infty \frac{z}{\sigma_x\sqrt{2\pi}} e^{-\frac{z^2}{2\sigma_x^2}} dz = -\sigma_x\sqrt{\frac{2}{\pi}} e^{-\frac{z^2}{2\sigma_x^2}}\Big|_0^\infty = \sigma_x\sqrt{\frac{2}{\pi}} \tag{2.34}$$

So, in order to obtain the actual noise power, σ_x^2, the output of the detector must be squared. The uncertainty in the noise measurement is then:

$$\alpha = \frac{\sigma_{\sigma_x^2}}{\mu_{\sigma_x^2}} = 2\frac{\sigma_W}{\mu_W} \approx \sqrt{2\frac{ENBW_{LPF}}{\Delta f_c}} \tag{2.35}$$

which is exactly the same as for the square-law detector.

Half-wave linear rectifier

Equivalently, for the half-wave linear rectifier the mean output of the detector for Gaussian signals is given by:

$$\mu_Z = E[Z(t)] = \int_0^\infty \frac{z}{\sigma_x\sqrt{2\pi}} e^{-\frac{z^2}{2\sigma_x^2}} dz = -\frac{\sigma_x}{\sqrt{2\pi}} e^{-\frac{z^2}{2\sigma_x^2}}\Big|_0^\infty = \frac{\sigma_x}{\sqrt{2\pi}} \tag{2.36}$$

The relative uncertainty of the signal out of the half-wave rectifier is $2\sqrt{2}$ times higher than that out of a full-wave rectifier:

$$\alpha' = \frac{\sigma_W}{\mu_W} \approx \sqrt{2}\sqrt{2\frac{ENBW_{LPF}}{\Delta f_c}} \tag{2.37}$$

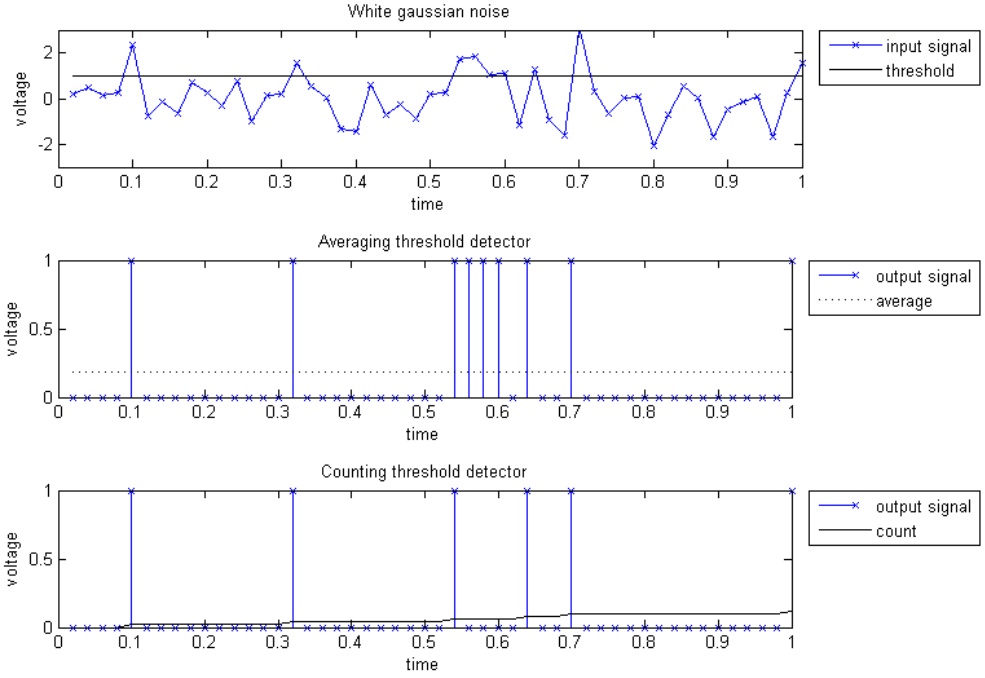

Figure 2.8: An illustration of the difference of the averaging and counting threshold detectors.

This is because compared to the full-wave rectifier, effectively half of the time the signal can be measured (for zero-mean noise). Also, the mean of the signal out of the detector is twice as low (compare Equation 2.34 and Equation 2.36). The resulting uncertainty in the calculated noise power is:

$$\alpha = \frac{\sigma_{\sigma_x^2}}{\mu_{\sigma_x^2}} = 2\frac{\sigma_W}{\mu_W} \approx 2\sqrt{2}\sqrt{2\frac{ENBW_{LPF}}{\Delta f_c}} \tag{2.38}$$

And it can be seen that the uncertainty is a factor $2\sqrt{2}$ higher, which can only be reduced to the same level as for the full-wave linear rectifier or square-law rectifier by measuring 8 times longer, or over an 8 times higher bandwidth.

Threshold detector

The threshold detector proposed in the literature (see Section 2.1.3) is an asynchronous continuous time system where the number of times the threshold is exceeded is counted, and not the relative amount of time the threshold is exceeded. For these two cases the output of the detector has a different relation to the variance of the input signal. This is illustrated in Figure 2.8. This figure clearly shows that if the input signal exceeds the threshold for a longer time, it is only counted once, while this is not the case when determining the average time the threshold is exceeded.

In [45] it is shown that for a Gaussian noise signal with spectrum $S_x(f)$ the expected number of passes per second through a threshold at V_t with a positive slope is given by:

$$E[N] = 2\sqrt{\frac{\int_0^\infty f^2 S_x(f)df}{\int_0^\infty S_x(f)df}} \cdot \frac{1}{2} e^{-\frac{V_t^2}{2\sigma_x^2}} \tag{2.39}$$

where the first term describes the expected number of zeros per second and the second term corrects for the threshold voltage. This threshold detector will be referred to as the counting threshold detector. A possible implementation is by using a comparator followed by an edge-sensitive counter.

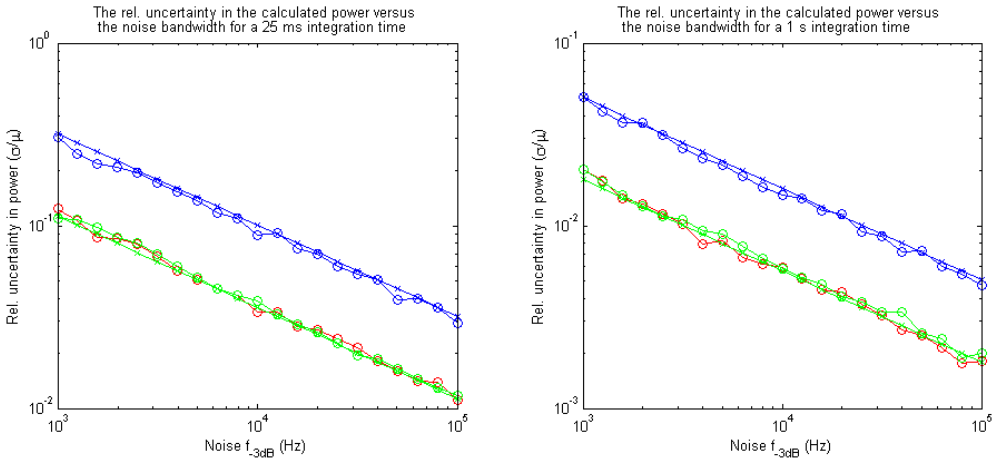

Figure 2.9: Simulation results (circles) for the square-law rectifier (red), half-wave linear rectifier (blue) and the full-wave linear rectifier (green), with the theoretical lines (crosses), for two different integration times.

If, on the other hand, the relative time is determined that the noise is above the threshold, the following average is expected:

$$E[N] = P(X > V_t) = \frac{1}{2} - \frac{1}{2} erf\left(\frac{V_t}{\sqrt{2}\sigma_x}\right) \tag{2.40}$$

This threshold detector will be referred to as the averaging threshold detector. A possible implementation is by using a comparator followed by a low-pass filter, either analog or digital.

The equations reported in this section can be used to translate the detector's uncertainty to the uncertainty in the noise power, σ_x, if the uncertainty in the output of the detector is known.

2.2.2. Simulations

The previous section analysed the uncertainty of the measurement of noise power using various detectors. The uncertainty of the threshold detector was not discussed, since this cannot easily be calculated and no clear derivations were found in the literature. To evaluate its uncertainty, a simulator was designed in MATLAB that could simulate the various detectors.

The system as depicted in Figure 2.6 is simulated. White Gaussian noise is generated over a bandwidth 200 times wider than the -3 dB frequency of the band-limited noise that is obtained from this signal. The band-limited noise is applied to the various detector and the average output is recorded. The simulations have been repeated 200 times for different noise bandwidths and measurement times to estimate both the average output and the measurement uncertainty.

Verification of the theory

For the square-law rectifier (Section 2.2.1), the half-wave linear rectifier (Section 2.2.1) and the full-wave linear rectifier (Section 2.2.1) the expected measurement mean and standard deviation were found, and therefore also the expected uncertainty in the measured power was obtained, see Figure 2.9. In addition, the ideal square-law rectifier was compared with the correlating detector (Section 2.2.1) and a switching detector using a square-law rectifier (Section 2.2.1). Again, the expected uncertainty in the detected power was obtained, for different signal-to-noise ratios (SNR), see Figure 2.10. The SNR is defined as the amount of noise from the sensor compared to the amount of additional noise.

Threshold detectors

This section will first discuss the simulated average output of the detector and the uncertainty of this signal. Next, these simulation results are used to obtain the uncertainty in the noise power that can be calculated from the detector's output.

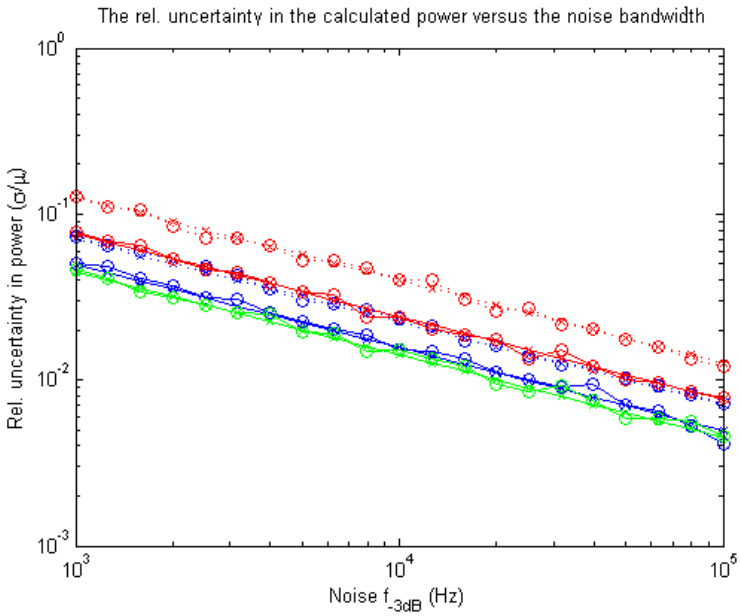

Figure 2.10: Simulation results (circles) for the square-law rectifier (green), switching detector (red) and the correlating detector (blue), with the theoretical lines (crosses). The dotted lines are for an SNR of 1, while the solid lines are for an SNR of 5.

Mean and uncertainty of the detector's output The simulations of the threshold detectors confirmed the theoretical formulas describing the mean output of the detector as presented in Section 2.2.1. However, the simulation of the counting threshold detector revealed a problem with this type of detector. For a simple first order filter as used in the simulator, the expected number of counts as given by Equation 2.39 is infinite. This is because even the highly attenuated noise at high frequencies passes through an ideal comparator, which has infinite gain. In the MATLAB simulation, the bandwidth is then limited by the sampling frequency of the discrete-time simulation.

The simulated uncertainty in the output of the threshold detectors is shown in Figure 2.11 and Figure 2.12 for the counting threshold detector and the averaging threshold detector respectively. The uncertainty has been determined for 4 different threshold voltages.

From the simulation results it can be seen that the threshold voltage has an influence on the uncertainty. The plots also show fitting lines that were obtained using the following reasoning. For the counting threshold detector, a counting event occurs at random if the band is not too small [45]. Therefore, this random process can be modelled with a Poisson distribution with mean and variance N, where N is the expected count in the total measurement time [30]. As a result, the uncertainty in the detector's output is expected to follow:

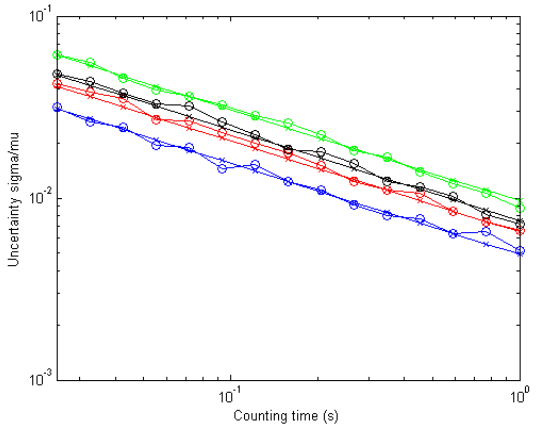

Figure 2.11: Simulation results (circles) for 4 different threshold for the counting threshold detector, with fitting lines (crosses).

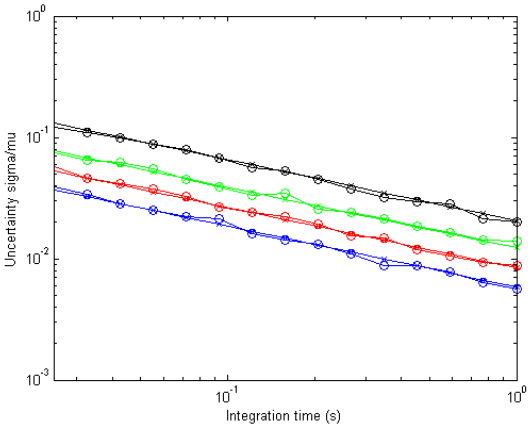

Figure 2.12: Simulation results (circles) for 4 different threshold for the averaging threshold detector, with fitting lines (crosses).

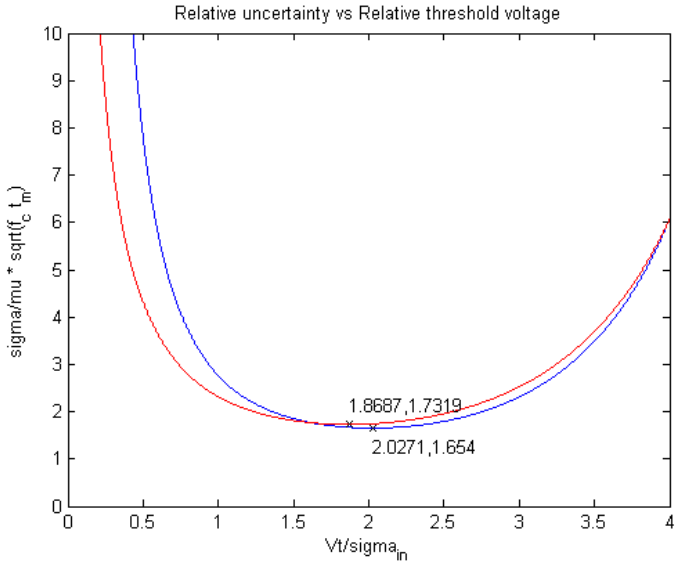

Figure 2.13: Estimates of the uncertainty in the calculated power for the threshold detectors (red: averaging threshold detector; blue: counting threshold detector).

$$\alpha' = \frac{\sigma}{\mu} = \frac{1}{\sqrt{N}} \tag{2.41}$$

For the averaging threshold detector, the outcome of each comparison is either a zero or a one, where the probability of a one, P, is given by Equation 2.40. The sampling of this random process can be modelled with a Bernoulli process with mean P. As a result, the the uncertainty in the detector's output after N samples is expected to follow:

$$\alpha' = \frac{\sigma}{\mu} = \sqrt{\frac{1-P}{P}} \cdot \frac{1}{\sqrt{N}} \tag{2.42}$$

By introducing scaling factors in P and N, the simulation results for the two threshold detectors could be fitted fairly well to Equation 2.42. This shows that the counting threshold detector can also be better modelled by a Bernoulli process than by a Poisson process, which was not expected.

Uncertainty in the noise power Using the formulas that fitted the uncertainty in the detector's output and the formulas for the mean output of the detectors, the uncertainty in the calculated power can be estimated. Since the noise power is proportional to the absolute temperature, this is the uncertainty in which we are ultimately interested. The resulting uncertainty depends on the threshold voltage as shown in Figure 2.13. It can be seen that the uncertainty is about the same for the two types of threshold detectors. The best uncertainty is obtained around $V_t \approx 2\sigma_x$.

Sampled threshold detectors

In the simulations performed so far, the threshold detector was assumed 'infinitely' fast, as was suggested in the literature (see Section 2.2.1). However, in that case there is a high correlation among the noise samples due to the limited noise bandwidth and it might be possible to reduce the speed of the threshold detector without any cost in the uncertainty. In the simulator the continuous-time comparator has been replaced by a clocked comparator with a finite sampling frequency.

It was found that when reducing the sampling frequency, the mean output of the counting threshold detector no longer follows the formula in Equation 2.39, but reducing the sampling speed does work for the averaging threshold detector.

The plot of the simulated uncertainty in the power versus the sampling frequency of the averaging threshold detector is shown in Figure 2.14. For this simulation the estimated optimum threshold voltage is applied and it can be seen that for higher sampling frequencies the uncertainty asymptotically reaches the value that was found

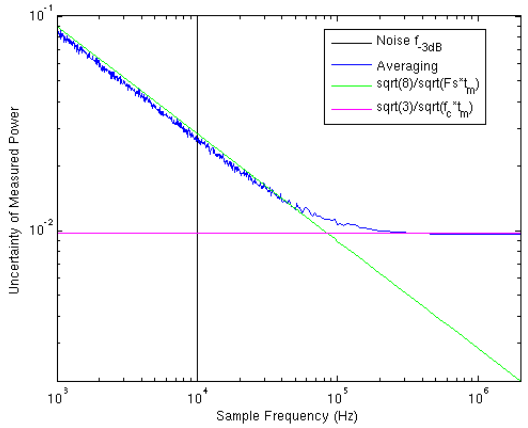

Figure 2.14: The simulated uncertainty in the power versus the sampling frequency of the averaging threshold detector.

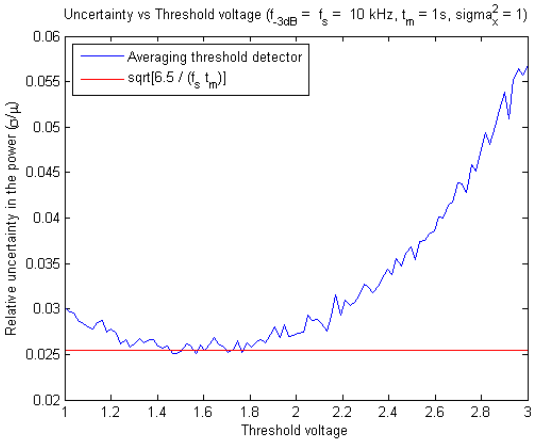

Figure 2.15: The simulated uncertainty in power for different threshold voltages for the averaging threshold detector.

Table 2.2: A summary of the uncertainty of the various detectors. The reported uncertainties assume either an integration or counting time of t_m. In the table N_0 is used for the counting threshold detector to denote the number of zero crossings per second.

Detector	Measurement		Calculated power
	Mean (μ)	Rel. Uncertainty (σ/μ)	Rel. Uncertainty (σ/μ)
Square-law rectifier	σ_x^2	$\sqrt{\frac{1}{f_c t_m}}$	$\sqrt{\frac{1}{f_c t_m}}$
Correlating detector	σ_x^2	$\sqrt{\frac{1}{f_c t_m}} \cdot \sqrt{1 + \frac{R_{eq}}{R} + \frac{1}{2}\left(\frac{R_{eq}}{R}\right)^2}$	$\sqrt{\frac{1}{f_c t_m}} \cdot \sqrt{1 + \frac{R_{eq}}{R} + \frac{1}{2}\left(\frac{R_{eq}}{R}\right)^2}$
Switching detector	σ_x^2	$\sqrt{\frac{1}{f_c t_m}} \cdot \sqrt{2\left(1 + \frac{R_{eq}}{R}\right)^2}$	$\sqrt{\frac{1}{f_c t_m}} \cdot \sqrt{2\left(1 + \frac{R_{eq}}{R}\right)^2}$
Full-wave linear rectifier	$\sigma_x\sqrt{\frac{2}{\pi}}$	$\sqrt{\frac{1}{f_c t_m}} \cdot \sqrt{\frac{1}{4}}$	$\sqrt{\frac{1}{f_c t_m}}$
Half-wave linear rectifier	$\sigma_x\sqrt{\frac{1}{2\pi}}$	$\sqrt{\frac{1}{f_c t_m}} \cdot \sqrt{2}$	$\sqrt{\frac{1}{f_c t_m}} \cdot \sqrt{8}$
Counting threshold detector	$\frac{1}{2}N_0 e^{-\frac{V_t^2}{2\sigma_x^2}}$		$\approx \sqrt{\frac{1}{f_c t_m}} \cdot \sqrt{2.7}$ at $V_t \approx 2.03\sigma_x$
Averaging threshold detector ($f_s \gg f_c$)	$\frac{1}{2} - \frac{1}{2}erf\left(\frac{V_t}{\sqrt{2}\sigma_x}\right)$		$\approx \sqrt{\frac{1}{f_c t_m}} \cdot \sqrt{3}$ at $V_t \approx 1.87\sigma_x$
Averaging threshold detector ($f_s < f_c$)	$\frac{1}{2} - \frac{1}{2}erf\left(\frac{V_t}{\sqrt{2}\sigma_x}\right)$		$\approx \sqrt{\frac{1}{f_s t_m}} \cdot \sqrt{6.5}$ at $V_t \approx 1.6\sigma_x$

in Figure 2.13. For lower sampling frequencies, the uncertainty starts to increase around $10f_{-3dB}$ and follows the asymptote set by the sampling frequency instead of the correlation bandwidth.

From this plot it can be concluded that for a certain sampling frequency, the best inaccuracy can be achieved by using very wide-band noise. Consider for instance the sampling frequency of 1 MHz. For the simulated noise bandwidth, the best possible accuracy is set by the purple asymptote, while with a higher noise bandwidth the uncertainty set by the green asymptote could be achieved for the same sampling frequency.

A simulation of the relative uncertainty in power for different threshold voltages for the averaging threshold detector is shown in Figure 2.15. This is for the case of very wide-band noise and a simulation time of 1 s. This shows an optimum threshold voltage around 1.6σ of approximately $\frac{\sqrt{6.5}}{\sqrt{f_s t_m}}$.

Summary

A summary of the mean output of each detector is given in Table 2.2. This table also shows the uncertainty in the measured output of the detector. Finally, using the equation describing the mean output of the detector, this uncertainty can be translated to the uncertainty in the calculated noise power, which is also included in the table.

For the averaging threshold detector two rows are added, one for when the sampling frequency is much higher than the noise bandwidth, and one for when the sampling frequency is lower than the noise bandwidth.

2.3. Application considerations

For the targeted application, i.e. thermal monitoring (see Section 1.1), the following specifications are targeted. The noise-thermometer should achieve a resolution of 0.2 °C$_{rms}$ in a 0.1 s measurement time. The targeted accuracy is 0.5 °C using only an electrical calibration.

Based on these specifications, this section will discuss what kind of calibration should be used for the sensor. Additionally, the type of detector used for measuring the noise power is chosen. These choices are the basis for the system level architecture as described in the next chapter.

2.3.1. Choice for calibration

During the thermal monitoring, the references that must be supplied externally must be kept to a minimum. In order to achieve this, factory trimming might be needed. Only capacitors are suitable for trimming purposes, because of their low temperature coefficient and slow aging. The sensor should be calibrated electrically to prevent an expensive trimming using a temperature stabilized environment. For this calibration, any reference can be supplied externally, but the goal is to keep the amount of pads to a minimum to reduce the costs.

To understand which references are minimally needed during operation or calibration, different methods of noise measurement are analysed:

- *Measure the noise voltage*, $v_n = \sqrt{4kTBR}$: the voltage should be measured and the factor BR should be calibrated. The factor to be calibrated is a capacitor ($[VA^{-1}s^{-1}] = [F^{-1}]$).
- *Measure the noise current*, $i_n = \sqrt{4kTB/R}$: the current should be measured and the factor B/R should be calibrated. The factor to be calibrated is an inductor ($[AV^{-1}s^{-1}] = [H^{-1}]$). These are not practical components in integrated circuits.
- *Measure the noise power*, $p_n = 4kTB$ (see also Section 2.1.2): measuring the power requires a measurement of the voltage and current. The unknown factor to be calibrated is the bandwidth.

In the first method, considering that a voltage reference is needed for the voltage measurement, actually only a charge needs to be calibrated ($[As] = [C]$). For this charge an accurately known charge can be used, i.e. the elementary charge, which is found in the equations describing shot noise. In the literature a noise-thermometer has been presented that only requires a voltage reference (see Section 2.1.2). Since the addition of shot noise will increase the complexity of the circuit, for this design the options without shot noise are considered. A similar trick cannot be applied in case of the current measurement, since then a flux reference is required. For this Josephson junctions can be used, but these require super conduction.

Direct implementations of the proposed measurement methods require a reference voltage to be supplied together with a capacitor or, in the case of the noise power measurement, a reference voltage, current and time. These references can also be obtained indirectly using other references, but not all references are equally desirable.

During operation of the sensor, it is undesirable to supply a reference current, since true current references do not exist. A voltage reference is also not optimal, since this is implemented using a bandgap reference which in essence is already a temperature sensor. The requirement for a stable and very accurate external reference resistor or capacitor is also not desirable because of the increased component count. Moreover, the former requires a Kelvin connection for an accurate measurement, which requires additional pads, and the latter cannot be implemented accurately due to parasitic capacitances.

It can be concluded that during operation of the sensor an external reference voltage is most desirable, and as a result, factory trimming of some internal capacitor is required. This trimming can be performed most accurately, and with the least amount of additional pads, by supplying a reference clock and reference current.

A note on the requirement of a reference voltage: the reference voltage does not have to be independent of temperature. For example, a well-defined PTAT voltage is also usable, since the RMS noise voltage scales with the square root of the temperature, resulting in a temperature dependent ratio of the RMS voltage to reference voltage:

$$\left.\begin{array}{l} \sigma \propto \sqrt{T} \\ V_{ref} \propto T \end{array}\right\} \Rightarrow \frac{\sigma}{V_{ref}} \propto \frac{1}{\sqrt{T}} \tag{2.43}$$

Furthermore, when shot noise is used in combination with thermal noise, it should be possible to replace this reference voltage by a reference timebase. In that case, an internal capacitor needs to be trimmed in factory, and during operation only an accurate external clock must be supplied. Although this is ultimately the best, possible implementations will not be discussed in this thesis.

2.3.2. Choice of detector

Different detectors and sensing techniques have been discussed and analysed in this chapter. The ratio measurement or switching technique as discussed in the literature study are not directly applicable on-chip, since no reference temperature is available and shot noise will not be used. Therefore, it is important that the transfer function of the chosen detector is accurate. The most flexible solution is to use an ADC such that any kind of detector or filter can be realized in the digital domain. However, this solution is quite energy-inefficient, as will be shown in the following.

Assuming thermal noise with a variance σ_{th}^2 is applied to the input of an ADC, the required input range of the ADC is at least $V_{full} = 6\sigma_{th}$. The resulting quantization noise for an N-bit ADC is given by [46]:

$$\sigma_q^2 = \frac{q^2}{12} = \frac{1}{12}\left(\frac{V_{full}}{2^N}\right)^2 = \frac{6^2}{12 \cdot 2^{2N}} \cdot \sigma_{th}^2 \tag{2.44}$$

And the resulting temperature error due to quantization noise is:

$$\frac{\Delta T}{T} = \frac{\sigma_q^2}{\sigma_{th}^2} = \frac{3}{4^N} \tag{2.45}$$

This shows that a 7-bit ADC is sufficient to measure up to about 250 °C with a 0.1 °C inaccuracy due to quantization noise. To achieve a resolution of about 0.1 °C in 0.1 s the sample rate should be in the order of 100 MS/s (using the formulas in Table 2.2).

Another possibility is implementing the detector in the analog domain and using an ADC to digitize only the detector output. The required conversion rate of this ADC is then only 10 S/s for a measurement every 0.1 s. For again a temperature error of 0.1 °C in a 200 °C temperature range, the required number of bits is only 10 (resulting in an ADC step size of approx. 0.2 °C), assuming that the detector outputs a signal proportional to the noise power.

Considering the same energy per conversion step, an ADC after a detector will consume about 1 million times less power than using an ADC as a detector. Using the ADC as a detector is therefore very likely inferior in terms of power consumption.

Comparing the analog implementations, the square-law rectifier and the full-wave linear rectifier give the smallest statistical uncertainty (see Table 2.2). The half-wave linear rectifier is eight times worse in this respect, and since the implementation complexity is almost the same, the half-wave linear rectifier is discarded. It has been discussed that for both the square-law rectifier (or correlating detector) and the linear-rectifier the transfer is determined by the (temperature-dependent) circuit characteristics (see Section 2.1.3). Any deviation from the expected transfer results in an error. It will therefore be difficult to get an accurately defined transfer over a wide bandwidth using this approach. These detectors also all require an additional low-pass filter and ADC after the detector. This increases the complexity of the design, since also these blocks should have an accurate transfer function. The same is the case for a continuous time comparator.

This is not the case for both the counting and averaging threshold detectors when implemented with a clocked comparator. These comparators can directly be interfaced with the digital logic. Furthermore, by using positive feedback in the decision making process, it is easier to reach a higher speed. These detectors however have a slightly lower statistical uncertainty than the square-law and full-wave rectifier circuits. Another advantage of using a comparator as a detector is that there are only two possible output values. Therefore, there is no need to accurately know the gain, and there is no non-linearity involved. This also holds for any pre-amplifier that might be necessary for e.g. offset compensation [47].

In conclusion, the only disadvantage of the threshold detector is the slightly higher measurement uncertainty. However, the other detectors are difficult to implement while ensuring accuracy over a wide band, and therefore likely a higher overall measurement uncertainty can be achieved.

As shown in Section 2.2.2, it is advantageous to choose the sample frequency lower than the noise bandwidth to reduce the correlation between the noise samples. It was also shown that the counting threshold detector no longer produces the expected noise power in this case. Although one could determine the new transfer characteristic, the averaging threshold detector will simply be used instead.

3
System Level Design

This chapter describes the system level design of a noise thermometer based on a threshold detector.

3.1. The general architecture

3.1.1. Architectures

Different architectures are possible for threshold detectors, see Figure 3.1.

First of all, a fixed threshold voltage could be used and the temperature could be derived from the average output of the detector. In this case, the threshold voltage is only optimal for a single temperature. Alternatively, it is possible to use feedback in order to adapt the threshold voltage to keep it near the optimal threshold at all temperatures. Since the feedback architecture can be derived from the open-loop architecture with some modifications, it was chosen to first implement the open-loop architecture to evaluate feasibility. In addition, this would not strongly affect the performance since the uncertainty has a relatively flat optimum, as can be seen in Figure 2.15.

The threshold detector requires the RMS value of the noise to be about the same as the threshold voltage. The noise voltage will be low, even when measuring the noise of a large resistance over a wide bandwidth. For example, a 100 kΩ resistor generates about 1 mV_{rms} noise over a 1 GHz effective noise bandwidth, which requires the parasitic capacitance to be smaller than 2.5 fF. Since the noise voltage will be low, either the noise must be amplified or the threshold voltage must be very small. More generally, a combination of noise amplification and threshold division can be used. This general architecture is shown in Figure 3.2.

3.1.2. Architecture requirements

As mentioned in Section 2.3, an inaccuracy of 0.5 °C and a resolution of 0.2 $°C_{rms}$ in 0.1 s are targeted. This section derives the requirements on the general architecture to achieve these specifications.

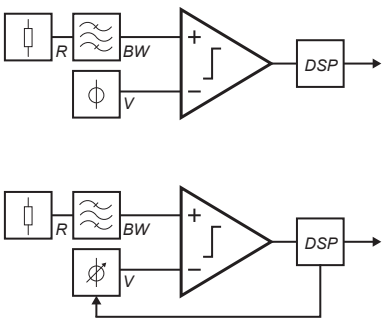

Figure 3.1: Two architectures based on the threshold detector.

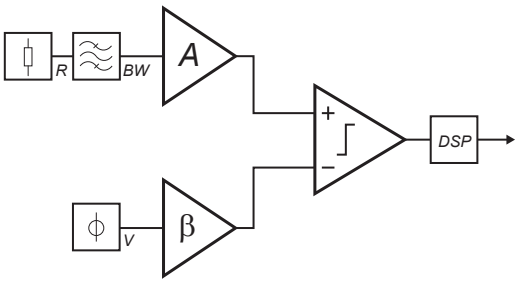

Figure 3.2: The general architecture of the noise thermometer, using either an amplifier, a divider or a combination.

Resolution

Even though it is shown in Figure 2.15 that the optimum threshold is around $V_t = 1.6\sigma_x$, the slightly less optimum $V_t = 2\sigma_x$ is assumed during the system and circuit design (Chapter 4). The slightly less optimal value was estimated in Section 2.2.2, and only after the measurement results of the chip were obtained, the optimum threshold voltage was again simulated and found to be closer to $V_t = 1.6\sigma_x$. For $V_t = 2\sigma_x$ the relative uncertainty in the noise power is approximately:

$$\frac{\sigma}{\mu} = \sqrt{\frac{8}{f_s t_m}} \tag{3.1}$$

The results obtained when using this value are still valid. However, the circuit can be designed for a better resolution when using a slightly lower threshold voltage, but then the maximum allowed offset is also lower.

From this equation, it follows that in order to achieve the targeted 0.2 $°C_{rms}$ resolution in 0.1 s, the readout circuit must process 180 MS/s.

Accuracy

Any additional noise (σ_{add}^2) directly adds to the noise power to be measured (σ_x^2). Hence, the error in the measured temperature, which is directly proportional to the noise power, is:

$$\frac{\Delta T}{T} = \frac{\sigma_{add}^2}{\sigma_x^2} \tag{3.2}$$

Since the ratio of the input noise to threshold voltage is determined by the threshold detector (see Equation 2.40), an offset voltage (V_{os}) can be translated into a change in the amount of noise power (σ_{os}^2):

$$\frac{V_t + V_{os}}{\sqrt{2\left(\sigma_x^2 + \sigma_{os}^2\right)}} = \frac{V_t}{\sqrt{2\sigma_x^2}} \tag{3.3}$$

And the error in the measured temperature follows as:

$$\frac{\Delta T}{T} = \frac{\sigma_{os}^2}{\sigma_x^2} = \left(\frac{V_{os}}{V_t}\right)^2 + 2\left(\frac{V_{os}}{V_t}\right) \approx 2 \cdot \frac{V_{os}}{V_t} \tag{3.4}$$

Since a total inaccuracy of 0.5 °C is targeted, the individual error sources are designed to be lower than 0.1 °C at 300 K. Using these equations, the maximum noise and offset of the various blocks in the general architecture can be determined. Comparator metastability and gain errors are modelled as an offset. The resulting requirements are summarized in Table 3.1. This table assumes a practical value of 1 V for the external reference voltage. The generated noise level is assumed to be 200 μV_{rms}. These values are close to the values that will be used.

There are additional requirements on the amplifier. The linearity of the amplifier is important since any nonlinearity will alter the variance of the signal. In addition, in a continuous time system, the amplifier bandwidth (and all parasitic poles) must be well-known since they alter the noise bandwidth and hence the variance of the signal. For example, an additional pole with a frequency uncertainty of 10% should be at a frequency at least about 300 times higher than the first pole in order not to change the temperature reading by more than 0.1 K.

Another important note on the amplification is that the linearity and actual gain in the chain after the comparator do not affect the reading, because the results only depend on the sign of the difference between the noise and threshold. This enables the use of a comparator consisting of multiple open-loop stages.

Description	Formula	Size
Error at room temperature (T = 300 K)	ΔT	0.1 K
Relative error	$\alpha = \Delta T/T$	1 / 3000
Applied reference	V_{ref}	1 V
Generated noise	v_n	200 μV_{rms}
Optimal threshold for generated noise	$V_t = 2v_n$	400 µV
Amplifier gain/divider division	$\gamma = V_{ref}/V_t$	2500
Maximum offset (referred to the noise source)	$V_{os,vt} = V_t \cdot \alpha/2$	67 nV
Maximum offset (referred to the voltage reference)	$V_{os,vref} = V_t \cdot \alpha/2 \cdot \gamma$	167 µV
Maximum readout noise (referred to noise source)	$v_{n,vt} = v_n \cdot \sqrt{\alpha}$	3.7 μV_{rms}
Maximum readout noise (referred to the voltage reference)	$v_{n,vref} = v_n \cdot \sqrt{\alpha} \cdot \gamma$	9.1 mV_{rms}
Minimum gain (metastability)	$\approx V_{out}/V_t$	144 dB for a 1V output
Amplifier/divider ratio accuracy	$\alpha/2$	167 ppm
Resolution at room temperature (T = 300 K)	T_{rms}	0.2 K_{rms}
Sample frequency	f_s	180 MS/s

Table 3.1: Specifications for the general architecture.

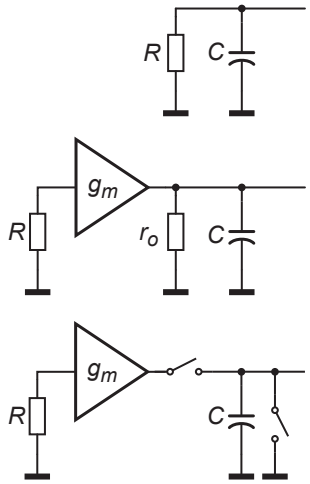

Figure 3.3: Different input structures that can be used to generate a well-defined amount of noise. From top to bottom: a parallel capacitor, a GmC-filter and a boxcar filter.

3.1.3. Noise generation & calibration

The generated noise variance depends on the resistor value and the noise bandwidth, which must be calibrated at least at one point. Higher-order filters are not suitable, since their higher component count makes calibration more difficult. Opamp-RC filters are not suitable since a high gain is required in the complete passband for accuracy. A switched capacitor filter is impractical, since this requires an additional anti-aliasing filter and oversampling of the desired 180 MS/s. Possible input structures are shown in Figure 3.3 and a summary of their properties is given in Table 3.2.

From the table it can be seen that the GmC-filter and the boxcar filter can generate a higher output noise level, because of the active component. These amplifiers will however introduce parasitic poles and additional noise. The resulting accuracy is lower, and the calibration is more involved. Therefore, a simple parallel capacitor is used as filter.

Filter	ENBW	Output noise level	To be calibrated
Parallel capacitor	$ENBW = 1/(4RC)$	$v_n^2 = kT/C$	C
GmC-filter	$ENBW = 1/(4r_oC)$	$v_n^2 = kT/C \cdot gain^2 \cdot R/r_o$	$g_m^2 \cdot r_o \cdot R/C$
Boxcar-filter [48] ($r_o := \tau/C$)	$ENBW = 1/(\lambda\tau)$	$v_n^2 = kT/C \cdot gain^2 \cdot R/r_o \cdot 4/\lambda$	$g_m^2 \cdot \tau/C \cdot R/C$

Table 3.2: Properties of the different filters.

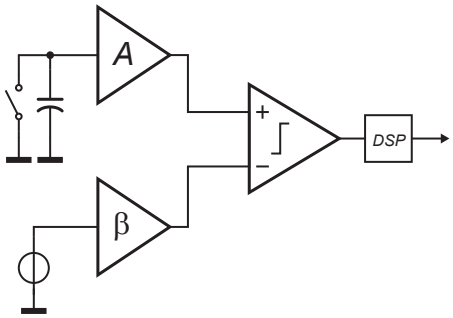

Figure 3.4: The noise generation circuit and the readout circuit.

The problem that the additional poles of the readout circuit changes the amount of noise, can easily be solved for this circuit by sampling the noise on the capacitor. The transfer function of the system that processes these samples then no longer influences the noise bandwidth of the samples. The resulting circuit does not even require a resistor, since the switch has a finite resistance and the actual resistor value is not important.

The resulting noise generation circuit, including the rest of the blocks for this architecture, is shown in Figure 3.4. This circuit only requires a one-time calibration of the capacitor value. This can be performed using a voltage, current and time as required. During operation only a reference voltage is required.

3.2. Achieving the accuracy requirements

To obtain an accurate design where only the capacitor value must be calibrated, the accuracy requirements of Section 3.1.2 must be achieved. From that section three main challenges become clear:

- The amplifier/divider must have an accurate ratio, i.e. a maximum deviation of 167 ppm is targeted. Section 3.2.1 discusses how this requirement can be met. Based on the results, this section will also make the choice for either amplification or division.
- The noise readout circuit must have a low offset, i.e. the maximum input referred offset is 67 nV. Section 3.2.2 discusses how this requirement can be met.
- The noise readout circuit must be low noise, i.e. the maximum input referred noise is 3.7 μV_{rms}. Section 3.2.3 discusses how this requirement can be met.

3.2.1. Accurate ratio

The general architecture of Figure 3.4 requires either an accurate amplifier, an accurate divider, or a combination of the two. An accurate amplification factor requires an active circuit with feedback. The accuracy will not only be limited by the accuracy of the feedback network, but also by the non-idealities of the amplifier. The feedback network of this amplifier should be an accurate divider. Therefore, it is chosen to implement only the divider for the threshold voltage. The disadvantage is that the comparator should provide more gain, but as explained before (Section 3.1.2) the requirements on the amplification are more relaxed in the comparator.

An accurate divider can be obtained by using PWM (pulse width modulation) [49], compensation circuitry [50], or dynamic element matching (DEM) [51][52]. With the required division ratio of more than 2000, more than 2000 elements are required, which makes this approach impractical. A better solution would be to cascade multiple division stages, for example cascading multiple DEM stages each dividing by 2. Different switched-capacitor circuits can be used to create an exact ratio of 2. By using a switched-capacitor circuit even more hardware can be saved, since a single stage can be reused multiple times in a cyclic fashion [53][54]. A problem with all these methods is that the division circuit, or the required averaging filter, adds noise.

However, the chosen noise source generates kT/C noise by sampling noise on a capacitor (see Section 3.1.3). Exactly the same is true for a passive switched-capacitor voltage divider, which is shown in Figure 3.5. This circuit contains two equally sized capacitors. Initially, one of these capacitors is charged with the reference voltage, while the other capacitor is discharged. Subsequently, the capacitors are shorted together and the charges are redistributed. Half the reference voltage is now stored on each capacitor. Then one capacitor is discharged, and the capacitors are again shorted. At that point only a quarter of the reference voltage is stored on the capacitors. These cycles of discharging and shorting the capacitors repeat until the desired voltage on the capacitors is obtained.

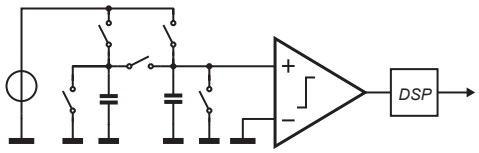

Figure 3.5: A cyclic switched-capacitor divider to generate the threshold voltage and the noise.

Because of the repeated redistribution of charge, this circuit will be referred to as the redistribution circuit in the rest of the thesis.

Each time a switch in this circuit opens, noise is sampled on the capacitors. The amount of noise generated by this circuit is predictable, and only depends on the capacitor size. Therefore, this circuit can be used to combine the division of the reference voltage with the generation of noise. The output of this circuit contains the divided reference voltage with the noise superimposed on it. Now the comparator should compare the output of this circuit with 0 V, which is also illustrated in Figure 3.5. This is equivalent to comparing zero-mean noise with a threshold voltage.

3.2.2. Low offset

For the obtained architecture, as shown in Figure 3.5, the offset of the readout circuit must be low, i.e. the maximum input referred offset is 67 nV. This section will discuss how that can be achieved.

The comparator will have offset and flicker noise. Their effect can be reduced using circuit techniques based on sampling (auto-zeroing, correlated double sampling) or modulation (chopping). These are however only applicable for linear systems. A solution is to use a pre-amplifier that uses these techniques. The input referred error due to the offset and flicker noise are reduced by the gain of this amplifier. In addition to the comparator, also the switched-capacitor noise generation circuit will have an offset, for example due to charge injection mismatch.

Since the additional pre-amplifier stage has to read out an unbuffered charge, chopping of this amplifier is not directly applicable as this would destroy the charge. Besides, since the system is a sampled system running at a high sample rate, the circuit techniques based on sampling are the most logical choice. However, it will be shown that both input offset sampling (IOS) and output offset sampling (OOS) cannot easily be implemented. Another option would be to simply use a high-pass filter, but also the implementation of this option has some challenges. Finally, it will be shown that system-level chopping with CDS demodulation can be used to remove the offset and flicker noise elegantly in this system.

Output offset sampling

Two possible implementations of OOS are shown in Figure 3.6. The circuit on the top has a single capacitor which is charged to the amplified offset voltage in the auto-zero phase ($\phi_{az} = 1$). In normal operation ($\phi_{az} = 0$) this voltage will be subtracted from the output, since the capacitor is in series with the amplifier output. In the circuit on the bottom, the offset voltage is stored on C1 in the auto-zero phase ($\phi_1 = 1$). In normal operation the amplifier output voltage (with an offset) is sampled on C2, and the voltages are subtracted using a difference amplifier. For the circuit on the top, the amplifier has to settle to the required accuracy in both the auto-zero and operational phases. For the circuit on the bottom, incomplete settling can be allowed, since the load capacitance can be the same in both phases.

The sampling of the offset on the capacitor results in noise sampling. For a low enough input referred noise, a large gain or a large load capacitance is required. In combination with the required sample rate of 180 MS/s, this results in tough requirements for the pre-amplifier. These requirements are a bit relaxed if incomplete settling can be allowed, as suggested by the other OOS implementation, but then the timing jitter must be very low.

Input offset samping

The amplifier specifications can be relaxed by using feedback, as is the case for IOS, see Figure 3.7. When $\phi_{az} = 1$ the amplifier offset is sampled on C_{az} which will be subtracted from the input voltage in normal operation. This circuit also allows for the amplifier to operate in open-loop configuration when the input sample is amplified, providing as much gain as possible.

This circuit however has several disadvantages. The amplifier's finite gain and input capacitance will limit the accuracy of the offset compensation. Moreover, the residual charge injection mismatch offset due to the switches is directly at the input. Minimum size switches already give more offset than can be allowed, and an extra auto-zero stage following this stage is required to compensate for this offset.

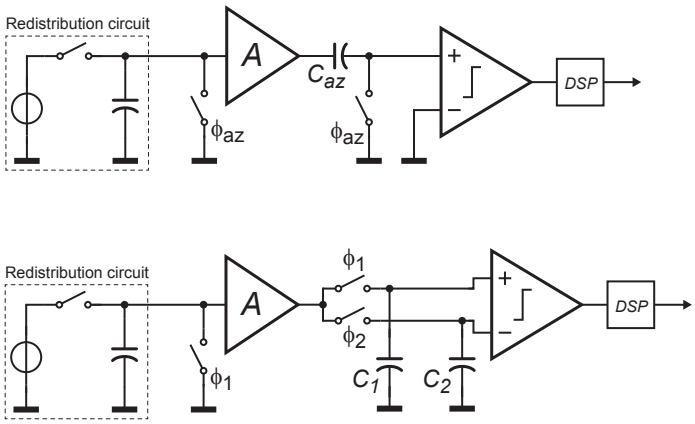

Figure 3.6: Two circuits for output offset sampling.

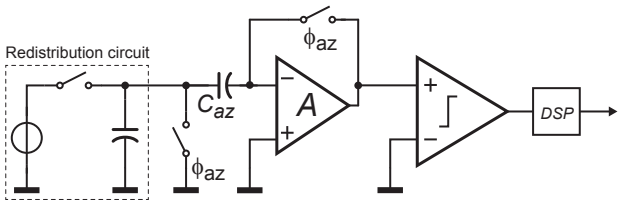

Figure 3.7: Circuit for input offset sampling.

Another problem is that an impractically large capacitor (around 500 pF) is required to make the sampled kT/C noise negligible, since it is directly at the input. Alternatively, the additional sampled noise could be considered extra generated source noise. This requires that the offset is sampled each cycle and that the noise generated by the pre-amplifier is well known, which is difficult to achieve.

High-pass filtering

Alternatively, to get rid of the offset and flicker noise of the amplifier, a high-pass filter can be used. Since the sample containing the DC threshold voltage is applied only when the sample is to be processed, this DC value is not filtered out. The circuit is shown in Figure 3.8, which also shows the waveforms at different points in the circuit. The pulse train shows the moments at which the comparator should take a sample. The high-pass filter functions as a differentiator and can therefore be used to determine if a rising or falling edge is applied at its input. The polarity of the sample then follows from the polarity at the output of the filter, which is sampled by the comparator.

The implementation of this technique has several challenges. In order to limit the amount of noise from the pre-amplifier, its output must be low-pass filtered. To prevent inter-symbol-interference, the different filters must

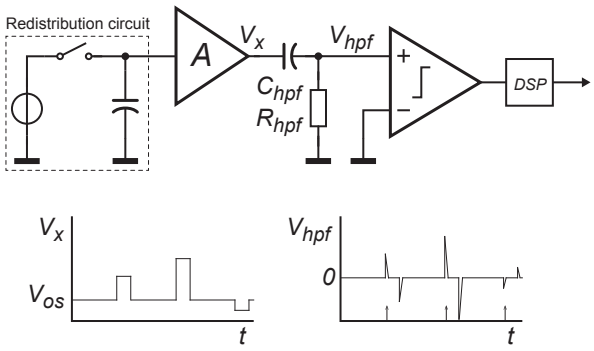

Figure 3.8: Circuit using a high-pass filter to remove offset and flicker noise.

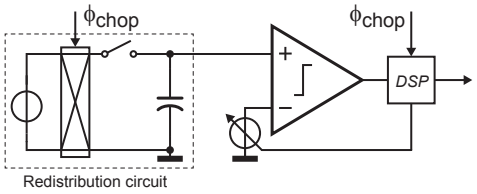

Figure 3.9: Circuit using low frequency chopping and an injected DC voltage to remove the circuit offset.

be reset or settle to sufficient accuracy. The amplifier, if not settled, must be reset in such a way that the offset is not periodically removed, since then it becomes a high frequency periodic pulse that passes through the filter similar to the input samples.

System-level chopping

The proposed solution, system-level chopping, will now be explained by first demonstrating how low-frequency system-level chopping can be used to remove the offset. Next, it will be demonstrated how this circuit can be adapted to also remove the high-frequency flicker noise.

Low-frequency chopping Theoretically, it is possible to get rid of any offset using two different threshold voltages and a correction in the calculation of the temperature. Half of the measurement time the noise (σ_n) is compared with threshold voltage 1 (V_{t1}) which has an offset (V_{os}). Equivalently, the second half of the measurement time a different threshold voltage (V_{t2}) is used, but the offset and thermal noise remain the same. From the output of the comparator the following ratios are determined:

$$\begin{aligned} \mu_1 &= \frac{V_{t1} + V_{os}}{\sigma_n} \\ \mu_2 &= \frac{V_{t2} + V_{os}}{\sigma_n} \end{aligned} \tag{3.5}$$

The temperature (σ_n) can be determined by subtracting these measurement values (μ_1 and μ_2), which removes the unknown offset from the equation:

$$\mu_1 - \mu_2 = \frac{(V_{t1} + V_{os}) - (V_{t2} + V_{os})}{\sigma_n} = \frac{V_{t1} - V_{t2}}{\sigma_n} \tag{3.6}$$

For this measurement to work, it is required that the voltage with which the noise is compared ($V_{t1} + V_{os}$ and $V_{t2} + V_{os}$) is not too high or too low. Otherwise the threshold detector would be working very far from the optimum point, resulting in a very high measurement uncertainty.

An additional compensating DC voltage can be injected into the circuit, such that the circuit keeps operating in the right region for both threshold voltages. The actual value of the injected voltage does not need to be known, since it will be cancelled like the circuit offset. The two different required threshold voltages can easily be obtained using an input chopper ($V_{t1} = +V_t$ and $V_{t2} = -V_t$). The resulting circuit is shown in Figure 3.9.

A disadvantage of this circuit is that higher frequency flicker noise will not be compensated. However, the main advantage of this circuit is that it removes the offset from the entire readout circuit, including the redistribution circuit.

It is important to realize that this is not simply possible for the offset sampling circuits presented before. Auto-zeroing of the redistribution circuit's offset requires redistribution of a 0 V input. This generates the offset at the output of the redistribution circuit, which can then be stored on an auto-zero capacitor for compensation. However, this generated offset will contain sampled kT/C noise with the same variance as the actual signal, since it is generated by the same circuit. The auto-zeroing will not work properly, since the offset cannot be obtained separately from the signal.

High-frequency chopping with CDS demodulation As explained, system-level chopping can easily remove the offset of the entire system, whereas this is not simply the case for auto-zeroing circuits. The described problem that the offset cannot be obtained separately from the signal, is a problem for regular signals. For the noise-thermometer the signal is however noise. In the auto-zero phase and the measurement phase the following samples

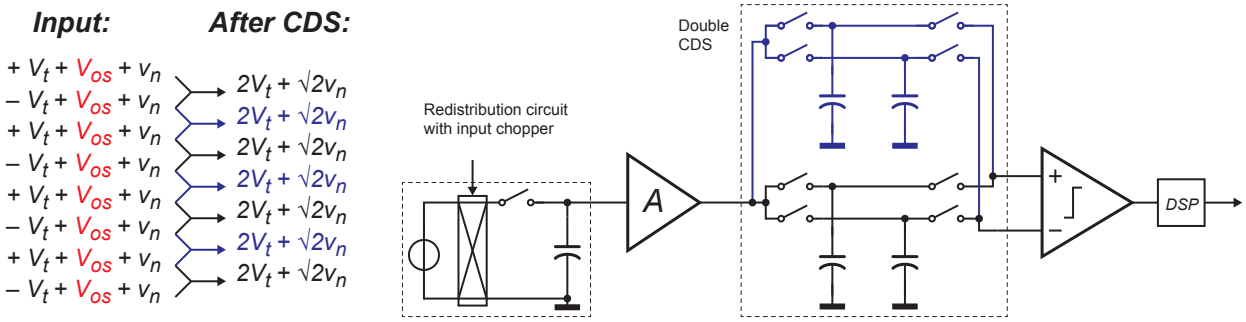

Figure 3.10: Circuit using an input chopper and two correlated double samplers to remove the offset and flicker noise.

and resulting output are obtained:

$$\begin{aligned} V_1 &= V_t + V_{os} + \sqrt{\sigma_{n1}^2} \\ V_2 &= 0 + V_{os} + \sqrt{\sigma_{n2}^2} \end{aligned} \tag{3.7}$$

$$V_{out} = V_1 - V_2 = V_t + \sqrt{\sigma_{n1}^2 + \sigma_{n2}^2 - 2\sigma_{n1,n2}}$$

where $\sigma_{n1,n2}$ is the co-variance of the noise samples. In case the samples are uncorrelated ($\sigma_{n1,n2} = 0$) and they have the same variance (σ_n), this reduces to:

$$V_{out} = V_t + \sqrt{2\sigma_n^2} \tag{3.8}$$

As a result, a sample with twice the noise variance is obtained, superimposed on the threshold voltage. This can be used as input to the detector, if the auto-zeroing is performed each cycle and the noise samples are uncorrelated. But, as already discussed, it is challenging to implement either IOS or OOS.

Alternatively, system-level chopping could again be applied, but at the same frequency as the sample frequency. For uncorrelated samples with the same variance this leads to:

$$\begin{aligned} V_1 &= +V_t + V_{os} + \sqrt{\sigma_n^2} \\ V_2 &= -V_t + V_{os} + \sqrt{\sigma_n^2} \end{aligned} \tag{3.9}$$

$$V_{out} = V_1 - V_2 = 2V_t + \sqrt{2\sigma_n^2}$$

From the equation describing the output voltage, it is clear that a correlated double sampler (CDS) is required for the demodulation.

This circuit does not suffer the same disadvantage as the low-frequency chopper, since this system will also attenuate high-frequency flicker noise. Moreover, no DC voltage injection is required, since the offset is removed before the detection.

From Equation 3.9 it seems like two samples from the redistribution circuit are required, to generate a sample that can be detected, which would reduce the resolution of the sensor. However, since each sample ($V_t + \sqrt{\sigma_n^2}$ and $-V_t + \sqrt{\sigma_n^2}$) is a new realization of the random noise process, each generated sample could be used as an input to the detector. Therefore, the measurement resolution is not reduced if the difference between every pair of samples is processed by the comparator. This statement was validated using a MATLAB simulation.

The proposed circuit with an input chopper and two correlated double samplers is shown in Figure 3.10. On the left the figure shows how the offset is removed and how the double CDS results in the same number of output samples as generated samples.

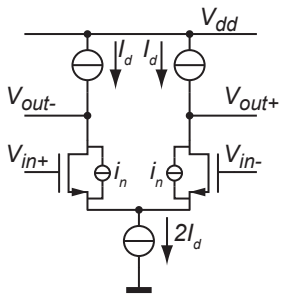

Figure 3.11: A simple differential pair as amplifier, with its noise sources.

Conclusion

In this section it has been shown that input and output offset sampling circuits cannot easily be implemented, especially if also the offset of the redistribution circuit must be removed. It has been demonstrated that system-level chopping can remove all offsets. Due to the specific properties of the system and the signal, i.e. a threshold detector and uncorrelated noise samples, the system-level chopping can run at the sampling frequency to also attenuate flicker noise. Moreover, this solution does not come with a penalty on the resolution, and has therefore been chosen for this system.

3.2.3. Low noise

For the obtained architecture, as shown in Figure 3.5, the noise of the readout circuit must be low, i.e. the maximum amount of input referred noise is about 3.7 μV_{rms}. This section will discuss how that can be achieved.

It will be shown that the required noise level cannot even be obtained with a simple and idealized differential pair, as the power consumption would become too high. The situation is even worse, since a charge must be read-out. It will be shown that in that case the input capacitance as determined by the technology is limiting. The most suitable charge readout amplifier will be shown.

Finally, as the amount noise from the readout circuit will be too high for the targeted accuracy, a digital noise compensation scheme is proposed.

Analysis of a differential pair

For a simple differential pair (see Figure 3.11) the input noise PSD can be approximated by $v_n^2 = 2 \cdot \frac{2}{3} \cdot \frac{4kT}{g_m}$. If the required settling time is given by $\tau = \frac{1}{\alpha \cdot f_s}$ (f_s is the frequency of the samples and α is the required settling, e.g.: $\alpha = 3$ for 3τ settling), the ENBW of the amplifier is $ENBW = \frac{\alpha \cdot f_s}{4}$. The total input referred noise is then:

$$P_{n,in} = 2\frac{2}{3}\frac{4kT}{g_m} \cdot \frac{\alpha f_s}{4} = \frac{4\alpha kT f_s}{3\left(\frac{g_m}{I_d}\right) I_d} \tag{3.10}$$

Whereas the maximum allowed input referred noise for a certain error $\Delta T/T$ and noise generating capacitor C is given by:

$$P_{n,in} = \frac{kT}{C} \cdot \left(\frac{\Delta T}{T}\right) \tag{3.11}$$

Equating these two equations and isolating the transistor drain current gives:

$$I_d = \frac{4\alpha}{3} \frac{C \cdot f_s}{\left(\frac{g_m}{I_d}\right) \cdot \left(\frac{\Delta T}{T}\right)} \tag{3.12}$$

For α = 3, C = 100 fF, f_s = 180 MS/s, g_m/I_d = 20 and $\Delta T/T$ = 0.1/300 this gives I_d = 11 mA which could easily result in self-heating of about 4 °C (assuming 100 °C/W at 1.8 V).

The resolution of the sensor can be approximated by (see Section 3.1):

$$\left(\frac{t_{rms}}{T}\right) = \frac{\sqrt{8}}{\sqrt{f_s \cdot t_m}} \tag{3.13}$$

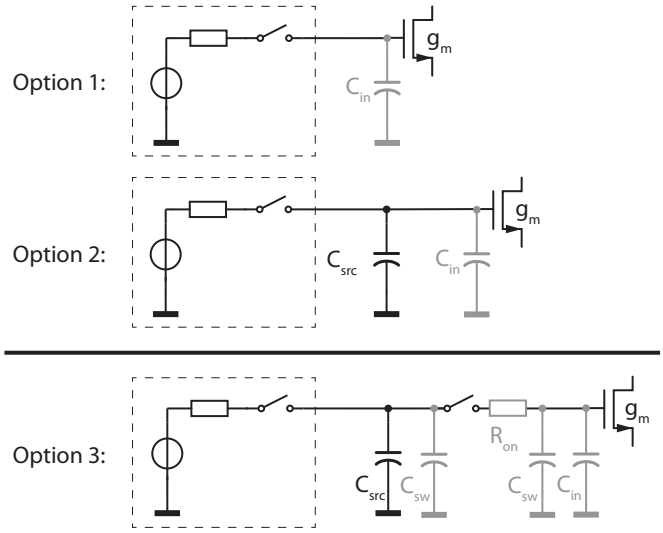

Figure 3.12: Different possible readout connections. The redistribution circuit is modelled as shown inside the dotted box, and the parasitics are shown in gray.

where t_m is the measurement time and hence the energy required from the supply is $E = 2I_d \cdot V_{dd} \cdot t_m$. Assuming the following relation between the temperature error due to noise and the temperature resolution, $\left(\frac{\Delta T}{T}\right) = \left(\frac{t_{rms}}{T}\right) \cdot \beta$ (the targeted resolution is slightly lower than the targeted accuracy, $\beta > 1$), the energy spent in the differential pair for a single temperature measurement can be simplified to:

$$E = \frac{64\alpha\beta^2}{3} \frac{V_{dd}}{\left(\frac{g_m}{I_d}\right)} \frac{C}{\left(\frac{\Delta T}{T}\right)^3} \tag{3.14}$$

This equation shows that it is best to use the smallest possible capacitor and lowest supply voltage. In addition it is important to note that the energy increases with the cube of the required accuracy or resolution. It can also be seen that the energy consumption is independent of the sample frequency.

Amplifier input capacitance

The main problem with the simplified analysis in the previous section is the neglected input capacitance. Since a large current is required, the transistor should be wide and has a large input capacitance. This is a problem since the source signal is a charge on a small capacitor and this capacitor must be accurately known. The input capacitance of the amplifier is in parallel with the source capacitance.

Different readout strategies are possible, see Figure 3.12. The first option is to use the input capacitance of the amplifier as the source capacitance. However, this capacitance is highly voltage and temperature dependent. This can be improved by adding an extra capacitor, which is the second option. In these situations the source noise is given by (C_{src} is the optionally added capacitor):

$$v_{n,src}^2 = \frac{kT}{C_{src} + C_{in}} \tag{3.15}$$

Another option is to connect the amplifier later using a switch (option 3 in Figure 3.12). The switch non-idealities like C_{sw} and R_{on} will be neglected for now. In that case the source capacitance is not changed by the input capacitance of the amplifier. However, when the amplifier is connected the source signal is attenuated due to charge redistribution. The source noise, at the input of the amplifier is given by:

$$v_{n,src}^2 = \frac{kT}{C_{src}} \cdot \left(\frac{C_{src}}{C_{in} + C_{src}}\right)^2 = \frac{kTC_{src}}{(C_{in} + C_{src})^2} \tag{3.16}$$

The noise of the readout circuit is ultimately limited by the noise of the input transistor:

$$v_{n,gm}^2 = \frac{8}{3} \frac{kT}{g_m} \cdot ENBW \tag{3.17}$$

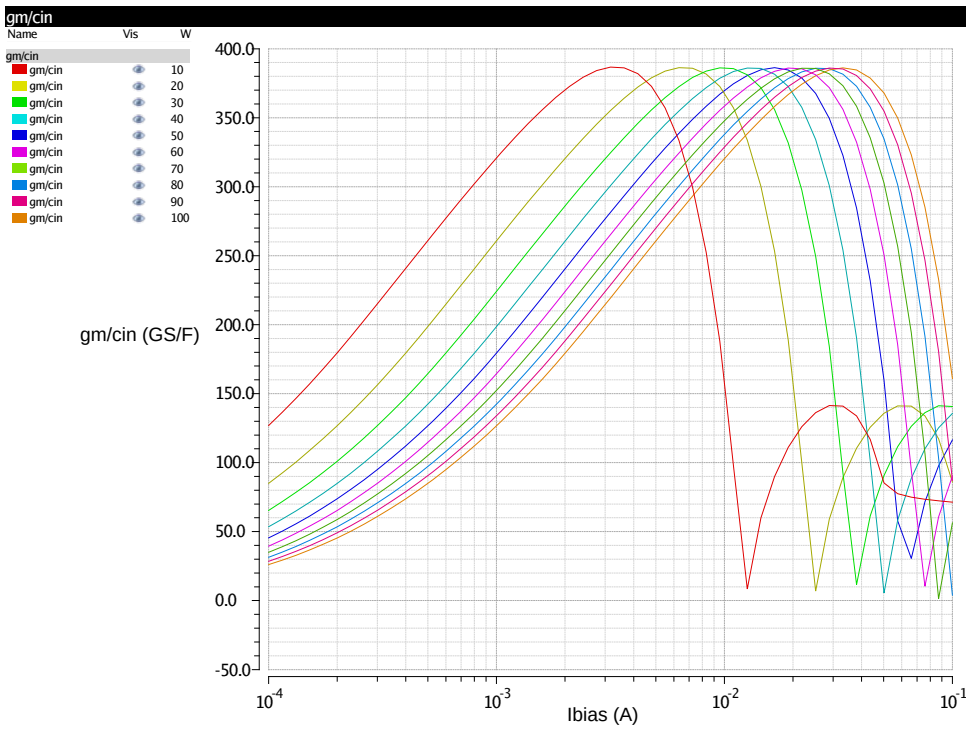

Figure 3.13: The input capacitance and transconductance of an NMOS transistor in the CMOS14 process for different bias currents and transistor widths.

The resulting temperature accuracy is given by:

$$\begin{aligned}
\text{Option 1/2: } \left(\frac{\Delta T}{T}\right) &= \frac{1}{SNR} = \frac{v_{n,gm}^2}{v_{n,src}^2} = \frac{8}{3}\frac{C_{in}}{g_m}\frac{C_{in}+C_{src}}{C_{in}} \cdot ENBW \\
&\longrightarrow \frac{8}{3}\frac{C_{in}}{g_m} \cdot ENBW \qquad \text{(Optimum: } C_{src} = 0\text{)} \\
\text{Option 3: } \left(\frac{\Delta T}{T}\right) &= \frac{1}{SNR} = \frac{v_{n,gm}^2}{v_{n,src}^2} = \frac{8}{3}\frac{C_{in}}{g_m} \cdot \frac{(C_{in}+C_{src})^2}{C_{in}C_{src}} \cdot ENBW \\
&\longrightarrow \frac{8}{3}\frac{C_{in}}{g_m} \cdot 4ENBW \qquad \text{(Optimum: } C_{src} = C_{in}\text{)}
\end{aligned} \tag{3.18}$$

In this equation C_{in}/g_m is ultimately limited by the technology:

$$\begin{aligned}
\frac{g_m}{W} &= \frac{\mu C_{ox} V_{gt}}{L} \\
\frac{C_{in}}{W} &= C_{ox} L \cdot \zeta \\
max\left(\frac{g_m}{C_{in}}\right) &= \frac{\mu V_{gt}}{L_{min}^2 \cdot \zeta}
\end{aligned} \tag{3.19}$$

These equations are for a transistor in strong-inversion, where ζ is factor that represents a correction for overlap capacitances etc. It can be seen that this will improve for newer technologies with a lower minimum channel length. A simulation of the input capacitance and transconductance of an NMOS transistor in the CMOS14 process for different bias currents and transistor widths shows that $max\left(\frac{g_m}{C_{in}}\right) = 386$ µS/fF (see Figure 3.13).

When the input capacitance of the amplifier can only be less than a quarter of the total source capacitance for option 2, the circuit of option 3 has a better SNR (see Equation 3.18). The resulting maximum ENBW is then 12 MHz for 0.1 K inaccuracy.

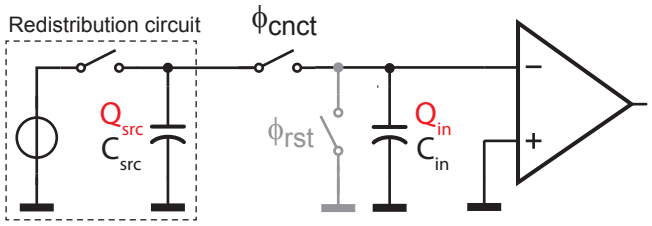

Figure 3.14: Residual charge on the parasitic input capacitance influences the charge sample.

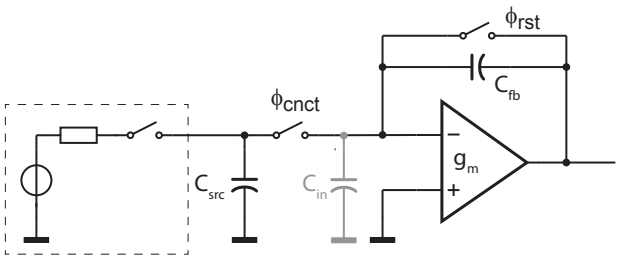

Figure 3.15: Readout using negative feedback.

A more elaborate model that includes the effects of the parasitics shown in gray in Figure 3.12 is discussed in Appendix A. The modelled effects include the noise from the connect switch and the inaccuracy in the redistribution due to the voltage-dependence of the parasitics. Simulations on that model confirm that the achievable accuracy of both architectures is about the same.

That model also shows that if a switch is used to connect the pre-amplifier to the redistribution circuit, this switch must be implemented using an NMOS with the largest possible overdrive voltage. The common-mode input voltage should then be kept low to keep the design simple (no clock boosting etc.). The analysis also reveals that operating at the maximum $\frac{g_m}{C_{in}}$ complicates the design and is less energy-efficient.

Since both architectures can achieve the same accuracy, it is chosen to connect the amplifier later using a switch, as that decouples the design of the pre-amplifier from the design of the redistribution circuit. As a result, there is more freedom in the design of the redistribution circuit. For instance, it is then easier to readout both capacitors of the redistribution circuit.

The proposed readout circuit

Connecting the pre-amplifier using a switch results in an issue which is illustrated in Figure 3.14. After the sample has been read out, residual charge is left on the parasitic input capacitance of the pre-amplifier. When the new sample is connected, a redistribution of charge takes place, which effectively results in a weighted average of multiple samples. As a result the noise samples will be correlated, which is not allowed.

This can be prevented by discharging the parasitic input capacitance before connecting the sample, as is illustrated by the switch in gray. However, kT/C noise will be sampled on the parasitic input capacitance when the switch opens. When the sample is now connected, again redistribution of charge takes place. The samples are no longer correlated, but accuracy is lost by the additional noise.

This problem can be solved by compensating for the charge that is present on the parasitic input capacitance. In case of the open-loop architecture discussed in the previous section, the voltage due to this reset charge can easily be measured, but in order to compensate for its effect the ratio of the source capacitance to parasitic input capacitance must be accurately known, which is not the case.

This can be solved by using a closed-loop architecture where both the reset charge and the charge of the input sample are moved into the same feedback capacitor, allowing for a compensation of the reset charge. This closed-loop circuit is shown in Figure 3.15. The details on how exactly the reset charge will be compensated, will be provided in Section 3.3.2.

The circuit works as follows. While the readout amplifier is disconnected ($\Phi_{cnct} = 0$) the feedback capacitor is reset ($\Phi_{rst} = 1$) and the redistribution circuit divides the input voltage. Next, in the amplification phase ($\Phi_{cnct} = 1$ and $\Phi_{rst} = 0$) the charge on the source capacitor is moved to the feedback capacitor.

The noise at the output of the circuit due to the applied input noise sample is given by:

$$v_{n,src}^2 = \frac{kT}{C_{src}} \cdot \left(\frac{C_{src}}{C_{fb}}\right)^2 \tag{3.20}$$

Whereas the noise at the output due to the amplifier's input transistor is given by:

$$v_{n,gm}^2 = \frac{8}{3}\frac{kT}{g_m} \cdot ENBW \cdot \left(\frac{C_{fb} + C_{in} + C_{src}}{C_{fb}}\right)^2 \tag{3.21}$$

The resulting temperature accuracy is:

$$\left(\frac{\Delta T}{T}\right) = \frac{1}{SNR} = \frac{v_{n,gm}^2}{v_{n,src}^2} = \frac{8}{3}\frac{C_{src}}{g_m}\left(\frac{C_{fb} + C_{in} + C_{src}}{C_{src}}\right)^2 \cdot ENBW \tag{3.22}$$

As can be seen from the equation, this circuit does not seem to be directly limited by the technology's C_{in}/g_m. A more elaborate comparison involving this circuit (closed-loop) and circuit option 3 of Figure 3.12 (open-loop) is also shown in Appendix . That model includes the flicker noise, as that imposes a limit on the lowest ENBW that should be used to filter the noise. From the resulting bandwidth, the expected resolution is obtained. Besides the measurement error due to noise from the readout circuit, also self-heating is modelled.

The results show that both architectures have similar characteristics for the different sources of error and uncertainty, i.e. this implementation is similarly limited by the technologies C_{in}/g_m. At the optimum found, the self-heating results in an error around 0.02 °C, the resolution is about 0.4 °C$_{rms}$ in 1 s, and the noise of the readout circuit results in an error of about 4.5 °C. Some resolution has been traded for a reduced noise error.

Achieving the noise requirement

It has been shown that it is not possible to achieve the noise requirement in an amplifier. However, for the noise-thermometer, noise from the readout circuit behaves as an offset that can be corrected for. In order to be able to correct later for this error in the digital domain, additional circuitry must be added such that the amount of readout noise can estimated.

An auto-zeroing like scheme would require the measurement of this readout noise in absence of the signal (the noise generated in the redistribution circuit). Measurement of this noise could be performed using the threshold detector that is already present. It would however be very difficult to generate a threshold voltage with a negligible amount of noise for the auto-zeroing action.

Instead, one could generate the threshold voltage with different amounts of noise using capacitors of a different size (different by a factor α). This allows for the removal of a fixed readout noise level, using the following scheme:

$$\begin{aligned} &\text{Measurement 1:} && v_{n,meas1}^2 = v_{n,readout}^2 + kT/C \\ &\text{Measurement 2:} && v_{n,meas2}^2 = v_{n,readout}^2 + kT/C \cdot \alpha \\ &\text{Calculate:} && v_{n,meas2}^2 - v_{n,meas1}^2 = kT/C \cdot (\alpha - 1) \end{aligned} \tag{3.23}$$

from which the temperature can be determined. In case the readout noise is temperature dependent, this calibration must be performed for each temperature measurement.

Conclusion

In this section it has been shown that the required noise level cannot be obtained with a simple and idealized differential pair. For a charge readout circuit, the input capacitance as determined by the technology is even more limiting.

A charge readout circuit that can be connected to the redistribution circuit using a switch has been chosen, since this decouples the design and there is no penalty in the achievable accuracy. It has been shown that the readout circuit must then employ feedback to compensate for reset errors, but this again does not have a penalty on the achievable accuracy.

For the resulting readout circuit, the temperature error due to noise from this circuit is at best about 4.5 °C, and this was only achieved by reducing the resolution. To compensate for the noise error, a digital noise compensation scheme was proposed.

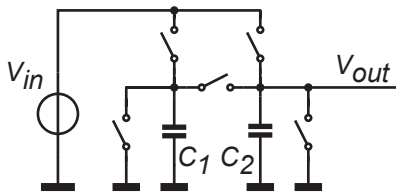

Figure 3.16: The redistribution circuit to generate the threshold voltage and noise.

3.3. The final architecture

The accurate division of the threshold voltage will be performed with a redistribution circuit (Section 3.2.1). This circuit also generates the noise samples and can be calibrated using a capacitor measurement. This redistribution circuit is shown in Figure 3.16 and its details will be discussed in Section 3.3.1. Also the digital noise calibration will be discussed in that section.

For the offset and flicker noise compensation, an input chopper in combination with a CDS will be used (Section 3.2.2). The only requirement is that the noise samples generated by the redistribution circuit are uncorrelated. It will be shown in Section 3.3.1 that this will be true. The CDS circuit will be further discussed in Section 3.3.3.

The low-noise pre-amplifier stage for the charge readout (Section 3.2.3) will be discussed in Section 3.3.2, and finally an overview of the entire system will be provided in Section 3.3.4.

3.3.1. Redistribution circuit

In this section the redistribution circuit will first be analysed for the situations of mismatch and finite speed. After this, the simulations are discussed that find the inaccuracy after applying DEM, and the uncertainty after applying digital noise calibration.

Circuit analysis

The redistribution circuit will be analyzed for the case of 2 different capacitors, C1 and C2 (see Figure 3.16). In the following analysis, the cycle number will be shown between brackets in superscript. The initial charging cycle is defined as cycle $i = 0$. Then, for each iteration $i > 0$, the even cycles $(2i - 2)$ describe the state after discharging one capacitor, and the odd cycles $(2i - 1)$ describe the state after redistributing the charges.

In the first cycle, one of the capacitors is charged (C1 in this case) while the other capacitor is discharged:

$$\begin{aligned} Q_{C1}^{(0)} &= C1 \cdot V_{in} \\ Q_{C2}^{(0)} &= 0 \end{aligned} \tag{3.24}$$

The noise sampled when the switches are opened is:

$$\begin{aligned} q_{n,C1}^{2^{(0)}} &= kT \cdot C1 \\ q_{n,C2}^{2^{(0)}} &= kT \cdot C2 \end{aligned} \tag{3.25}$$

Next, the capacitors are shorted together and the charge is redistributed to:

$$\begin{aligned} Q_{C1}^{(2i-1)} &= (Q_{C1}^{(2i-2)} + Q_{C2}^{(2i-2)}) \cdot \frac{C1}{C1 + C2} \\ Q_{C2}^{(2i-1)} &= (Q_{C1}^{(2i-2)} + Q_{C2}^{(2i-2)}) \cdot \frac{C2}{C1 + C2} \end{aligned} \tag{3.26}$$

The noise after opening the redistribution switch is given by (redistributed + new noise sampling):

$$\begin{aligned} q_{n,C1}^{2^{(2i-1)}} &= (q_{n,C1}^{2^{(2i-2)}} + q_{n,C2}^{2^{(2i-2)}}) \cdot \left(\frac{C1}{C1 + C2}\right)^2 + kT \cdot \frac{C1 \cdot C2}{C1 + C2} \\ q_{n,C2}^{2^{(2i-1)}} &= (q_{n,C1}^{2^{(2i-2)}} + q_{n,C2}^{2^{(2i-2)}}) \cdot \left(\frac{C2}{C1 + C2}\right)^2 + kT \cdot \frac{C1 \cdot C2}{C1 + C2} \end{aligned} \tag{3.27}$$

After the redistribution, one of the capacitors is discharged, and only the charge on this capacitor is changed (C2 in this case):

$$\begin{aligned} Q_{C1}^{(2i)} &= Q_{C1}^{(2i-1)} \\ Q_{C2}^{(2i)} &= 0 \end{aligned} \tag{3.28}$$

The noise after opening the switches is given by (discharging + new noise sampling):

$$\begin{aligned} q_{n,C1}^{2^{(2i)}} &= q_{n,C1}^{2^{(2i-1)}} \\ q_{n,C2}^{2^{(2i)}} &= kT \cdot C2 \end{aligned} \tag{3.29}$$

After this, the charge will again be redistributed and discharged repeatedly until the required output voltage is obtained.

When assuming nominally identical capacitors (C1 = C2 = C), the expected output voltage, after N cycles, is given by:

$$V_{out} = V_{in} \cdot \frac{1}{2^N} \tag{3.30}$$

When analysing the noise, it can be seen that both capacitors start with a random charge (see Equation 3.25). Then, after the first redistribution, this is (filling in Equation 3.27):

$$\begin{aligned} q_{n,C1}^{2^{(1)}} &= (q_{n,C1}^{2^{(0)}} + q_{n,C2}^{2^{(0)}}) \cdot \left(\frac{C1}{C1+C2}\right)^2 + kT \cdot \frac{C1 \cdot C2}{C1+C2} = kT \cdot \frac{C1^2}{C1+C2} + kT \cdot \frac{C1 \cdot C2}{C1+C2} = kT \cdot C1 \\ q_{n,C2}^{2^{(1)}} &= (q_{n,C1}^{2^{(0)}} + q_{n,C2}^{2^{(0)}}) \cdot \left(\frac{C2}{C1+C2}\right)^2 + kT \cdot \frac{C1 \cdot C2}{C1+C2} = kT \cdot \frac{C2^2}{C1+C2} + kT \cdot \frac{C1 \cdot C2}{C1+C2} = kT \cdot C2 \end{aligned} \tag{3.31}$$

And it can be seen that both capacitors end up with the same noise variance as before the redistribution. Next, during the discharging phase, the charge on one of the capacitors is removed. But when the discharge switch opens, noise is sampled, and the same noise variance is again available on the capacitor. As a result the amount of noise does not change with the number of redistribution cycles.

Speed requirement

The voltage as given in Equation 3.30 is only obtained in case the circuit is infinitely fast. With a finite settling speed, this voltage is only obtained with a certain accuracy. In case of N redistribution cycles, the final output voltage is given by Equation 3.32, assuming the charging, discharging and redistribution have the same time constant τ. This equation also assumes that incomplete settling always results in an output voltage that is too low, which is not always the case.

$$V_{out} \approx V_{in} \cdot \frac{1}{2^N} \cdot \left(1 - e^{-\frac{t}{\tau}}\right)^{2N} \tag{3.32}$$

As an example, for only 10 redistribution cycles and an accuracy of 1% already 7.6 tau settling is required.

Another speed requirement follows from the requirement that the noise samples must be uncorrelated. In case the switching speed is too high compared to the settling speed, the resulting noise spectrum is not white. However, simulations have shown that already at a switching frequency of less than 0.8 times the -3 dB frequency, the temperature error will be less than 0.06 °C due to correlation between the samples (see Figure 3.17). This requirement corresponds to a 7.9 tau settling speed.

Analysis of DEM schemes

Because of mismatch between the capacitors, $C1 \neq C2$, the charge will not be equally distributed when shorting the capacitors together. As a result, when always discharging the same capacitor, the obtained threshold voltage is either too low or too high. This can be solved by applying a form of dynamic element matching (DEM).

For this circuit this means that not always the same capacitor will be discharged. The capacitor that is discharged in a certain cycle is determined by the redistribution or switching scheme. By applying multiple different redistribution schemes in different DEM cycles more accurate division ratios can be obtained [54].

This section analyses the effect of DEM on the accuracy of the redistribution circuit.

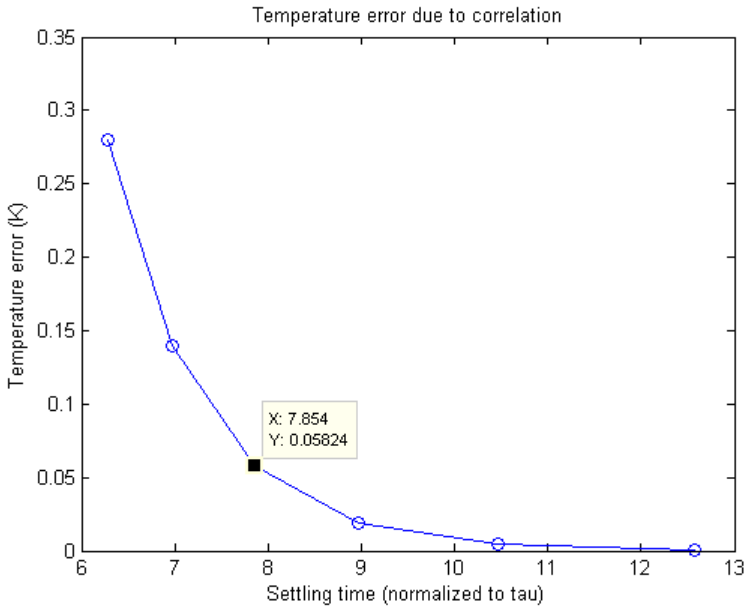

Figure 3.17: A simulation of the temperature error due to correlation for different settling speeds.

A MATLAB model was designed using the equations of Section 3.3.1. The mismatch is included by defining $C1 = C + \Delta C$ and $C2 = C - \Delta C$, where C is the nominal capacitor value and ΔC is the difference in capacitance. The relative error is defined as $\alpha = \Delta C/C$.

For this application not only the expected output voltage is of interest, but also the amount of generated noise is of importance. This depends on how the redistribution stops, there are two possibilities, both of which are simulated:

- First of all, one of the capacitors could be read out before discharging one of the capacitors, but after the redistribution switch has opened. At this point the large signal voltage on both capacitors is the same, while the noise is different.
- Secondly, both capacitors could be read out at the end of the redistribution, before the redistribution switch is opened. Now the capacitors are still shorted together and a single voltage is available.

For the simplest hardware implementation always the same capacitor should be discharged. In this case the inaccuracy in the threshold voltage after N redistribution cycles is given by:

$$\frac{\Delta V}{V} \approx N \cdot \frac{\Delta C}{C} \tag{3.33}$$

and the noise is either $\frac{kT}{2C}$ (both capacitors are read), $\frac{kT}{C+\Delta C}$ (C1 is read) or $\frac{kT}{C-\Delta C}$ (C2 is read), which follows from Equation 3.31.

A simulation has shown that by choosing wisely which capacitor should be discharged in which cycle, this can be improved to:

$$\frac{\Delta V}{V} \approx \alpha_1 \cdot \frac{\Delta C}{C} \tag{3.34}$$

where α_1 depends on the actual number of cycles, but is always smaller than 1 for up to at least 12 redistribution cycles.

This can be improved by using two different DEM schemes to generate the threshold voltage. By choosing the switching schemes wisely such that in case of mismatch the output voltage is too low for one of the schemes, and too high for the other, the average of the threshold voltages can get closer to the ideal threshold voltage.

Simulations have been performed for the usage of two different DEM schemes in case the redistribution switch is not opened and in case it is opened. In case the redistribution switch is opened, the switching schemes are chosen such that in one case one of the capacitors is read out and in the other case the other capacitor, in order to get the noise level as close as possible to the ideal noise level.

However, better results were obtained in case the redistribution switch is not opened. The following relation for the error in the threshold voltage was found:

$$\frac{\Delta V}{V} \approx \alpha_2 \cdot \left(\frac{\Delta C}{C}\right)^2 \tag{3.35}$$

where α_2 depends on the actual number of cycles, but is always smaller than 3 for up to at least 12 redistribution cycles. The generated noise is independent of any mismatch between the capacitors. This option was implemented in the final design.

Digital noise calibration

The digital noise calibration as described in Section 3.2.3 comes with a penalty on the achievable resolution. This section will analyse that penalty.

During the first measurement, a certain amount of noise is generated and is read out with a certain SNR[1]. During the second (calibration) measurement, less noise is generated, but the readout circuit adds the same amount of noise. The amount of generated noise is different by a factor α, which is obtained by increasing the size of the capacitors in the redistribution circuit by a factor α.

After these two measurements, the amount of noise found in each of the measurements is (normalized to the amount of noise generated in the redistribution circuit):

$$\begin{aligned} M_1 &= 1 + \frac{1}{SNR} \\ M_2 &= \frac{1}{\alpha} + \frac{1}{SNR} \end{aligned} \tag{3.36}$$

And the temperature is found from:

$$M = M_1 - M_2 = 1 - \frac{1}{\alpha} \tag{3.37}$$

In case for the first and second measurement respectively N_1 and N_2 noise samples are used, the uncertainty in these measurements can be determined. Recalling that the resolution of the measurements is inversely proportional to the square root of the number of samples (see Equation 3.1):

$$\begin{aligned} \frac{\sigma_{M1}}{M_1} &= \frac{1}{\sqrt{N_1}} \\ \frac{\sigma_{M2}}{M_2} &= \frac{1}{\sqrt{N_2}} \end{aligned} \tag{3.38}$$

The uncertainty in M is found to be:

$$\frac{\sigma_M}{M} = \frac{\sqrt{\sigma_{M1}^2 + \sigma_{M2}^2}}{M} = \frac{\sqrt{\frac{M_1^2}{N_1} + \frac{M_2^2}{N_2}}}{M_1 - M_2} \tag{3.39}$$

Whereas if no digital noise calibration was required (high enough SNR), this would have been $\frac{\sigma_M}{M} = \frac{1}{\sqrt{N_1+N_2}}$ for the same total measurement time.

From Equation 3.39 it can be seen that the resolution can be optimized by changing the time allocated to each of the two measurements (N_1 and N_2), and by correctly choosing the amount of generated noise in the second measurement (α). A MATLAB simulation was performed to find the effect of these parameters on the resolution of the measurements. The results are shown in Figure 3.18, this plot was generated using an SNR of 100.

From the plot it can be seen that for very large α there will be hardly any penalty on the resolution due to the calibration, if this calibration measurement is allocated only a very short time. However, a very large capacitor would be required, which is impractical. In case $\alpha = 2$, which would only require a capacitor of double the size, it is found that it is optimal to allocate approximately one third of the time to the calibration measurement. The resolution is in that case 3 times lower and as a result a 9 times longer measurement is required for the same resolution. In case only a two times longer measurement time is permitted to achieve the same resolution with digital noise calibration, α should be at least equal to 6. In the final circuit α can be selected as 2, 3, 4, 5 or 6.

[1]The SNR, or signal-to-noise ratio, is in this thesis defined as the amount of noise generated in the redistribution circuit ('the signal'), compared to the amount of noise from the readout circuit.

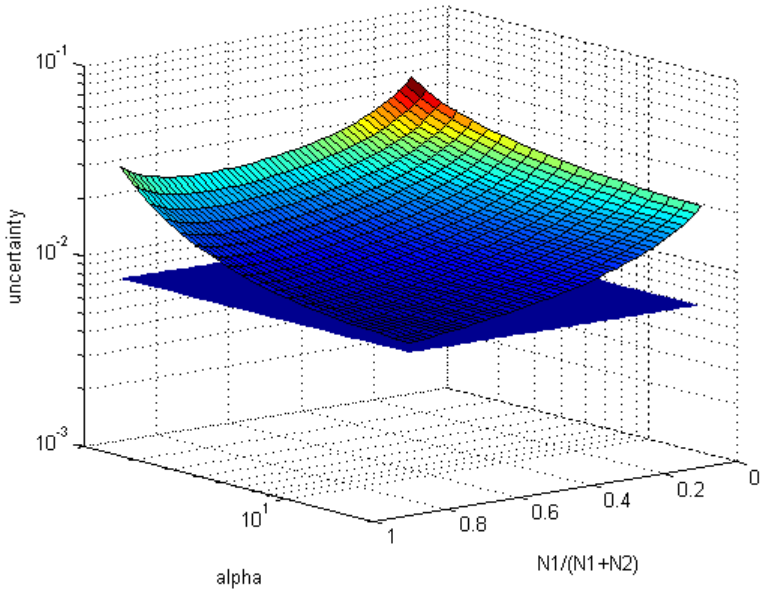

Figure 3.18: The plot shows the uncertainty in the measurement after digital noise calibration. The blue plateau shows the uncertainty in case no digital noise calibration is required.

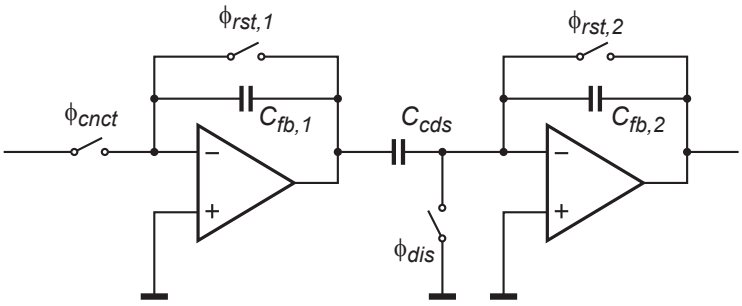

Figure 3.19: The complete pre-amplifier circuit.

3.3.2. Pre-amplifier

This section discusses the architecture of the pre-amplifier that is required to correctly read out the redistribution circuit. The pre-amplifier core has already been introduced in Section 3.2.3. It was mentioned that a reset charge compensated is possible for this circuit, but not how this could be implemented.

The complete circuit containing the previously shown pre-amplifier and the additional amplifier required for this reset charge compensation is shown in Figure 3.19. It works as follows. At first the redistribution circuit is disconnected from the pre-amplifier and a sample is generated. In phase 1 the first stage is reset and all capacitors are discharged. After opening the reset switch, a residual charge will be present on the feedback capacitor and the parasitic input capacitance. During the second phase, time is given to the pre-amplifier to move all charge into the feedback capacitor $C_{fb,1}$. This charge results in an output voltage of the first stage that appears over C_{cds}. At the end of the second phase, this voltage is sampled on C_{cds} when the reset switch of the second stage opens. In the third phase, the input charge is connected and moves to $C_{fb,1}$. This extra charge results in an output voltage on top of the voltage due to the reset charge. Only the change in output voltage of the first stage has to be compensated by the second stage to move its virtual ground back. Hence, the output voltage of the second stage only depends on the charge that was stored on the redistribution capacitors.

This solution is similar to the active noise cancellation circuit presented in [55] where also the sampled noise is stored on another capacitor and later removed in a CDS kind of fashion.

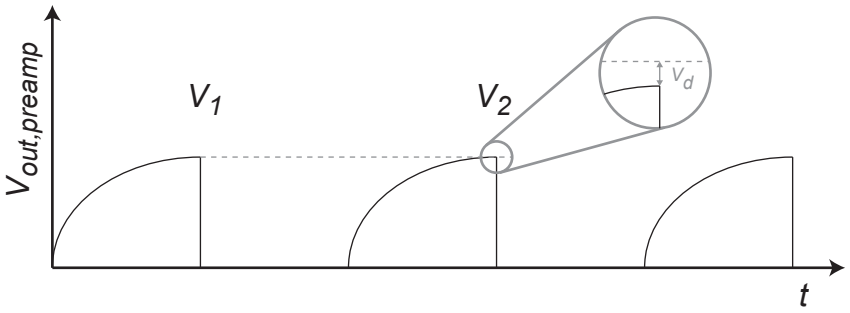

Figure 3.20: A sketch of the output voltage of the pre-amplifier. The CDS must resolve $V_d = V_1 - V_2$.

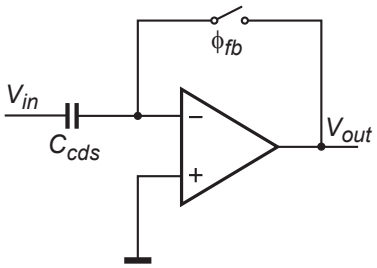

Figure 3.21: A simple CDS circuit.

3.3.3. Correlated double sampler

The correlated double sampler is responsible for the offset compensation of all the previous stages. Since the previous stages can have an input referred offset that is much larger than can be tolerated, a small voltage difference must be detected between the samples that can have a relatively large common voltage. Furthermore the output of the pre-amplifier as discussed in Section 3.3.2 settles slowly. This output voltage and the voltage difference that must be resolved are sketched in Figure 3.20. This exponential settling can be described by:

$$V_{in}(s) = \frac{V_{final}}{s\,(1 + s\tau)} \tag{3.40}$$

The problem with the simple CDS as shown in Figure 3.21 is the difference in effective settling speed in case the feedback of the CDS is closed compared to when it is open. When closed in feedback, the amplifier can be approximated by a resistance with value $1/g_m$. The transfer from input voltage to capacitor voltage is given by:

$$\frac{V_C(s)}{V_{in}(s)} = \frac{1}{1 + sC/g_m} \tag{3.41}$$

For the slowly settling input voltage, the voltage over the capacitor can be described by:

$$V_C(s) = \frac{V_{final}}{s\,(1 + s\tau)} \frac{1}{1 + s\frac{C}{g_m}} \tag{3.42}$$

When subsequently the same input voltage is applied with the amplifier in open-loop, the voltage at the amplifier input is given by:

$$V(s) = V_{in}(s) - V_C(s) = V_{final} \cdot \frac{\frac{1}{\tau}}{\frac{1}{\tau} + s} \frac{1}{\frac{g_m}{C} + s} = V_{final} \cdot \left[\frac{\frac{1}{1-\tau\frac{g_m}{C}}}{\frac{g_m}{C} + s} - \frac{\frac{1}{1-\tau\frac{g_m}{C}}}{\frac{1}{\tau} + s} \right] \tag{3.43}$$

Which in time domain is:

$$V(t) = \frac{V_{final}}{1 - \tau\frac{g_m}{C}} \cdot \left[e^{-\frac{g_m}{C}t} - e^{-\frac{t}{\tau}} \right] \cdot u(t) \tag{3.44}$$

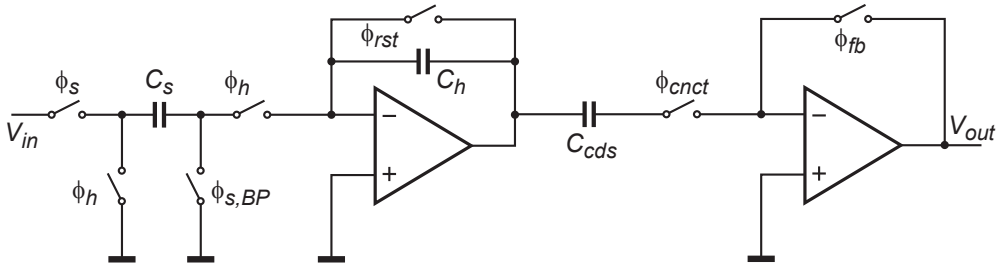

Figure 3.22: A S&H-circuit in front of the CDS circuit.

If, for example, the input voltage settles to only t = 3τ with τ = 100 ns and the final error (given by V(t)) can only be 10 ppm of V_{final} (these are realistic values), then the unity gain frequency (g_m/C) should be about 8 GHz. This is however not possible given that the sampling capacitor should be big enough such that the charge injection mismatch is small enough (in the order of pF).

This problem can be solved by inserting a S&H-circuit in front of the CDS (see Figure 3.22). In that case both samples from the pre-amplifier are sampled on the same capacitor of the S&H-circuit and moved to the same hold capacitor, so no errors will occur at this point.

After the samples are moved to the hold capacitor, the CDS-circuit is connected to the output of the S&H-circuit. When the CDS is in feedback, the sample will settle with $\tau_c = (1/g_{m,sh} + 1/g_{m,cds}) \cdot C_{cds}$. When the CDS is in open-loop, the sample will settle with $\tau_o = (1/g_{m,sh}) \cdot C_{par} \ll \tau_c$. Therefore, the final settling accuracy can be approximated by:

$$V(t) = -V_{final} \cdot e^{-t/\tau_c} \tag{3.45}$$

which is a single-pole system. Furthermore, there is quite some time to settle since the sample is being held until the next sample arrives.

3.3.4. Overview

The total system, showing all circuits described in the previous sections (3.3.1, 3.3.2 and 3.3.3), is shown in Figure 3.23. The implementation details will be discussed in the next chapter.

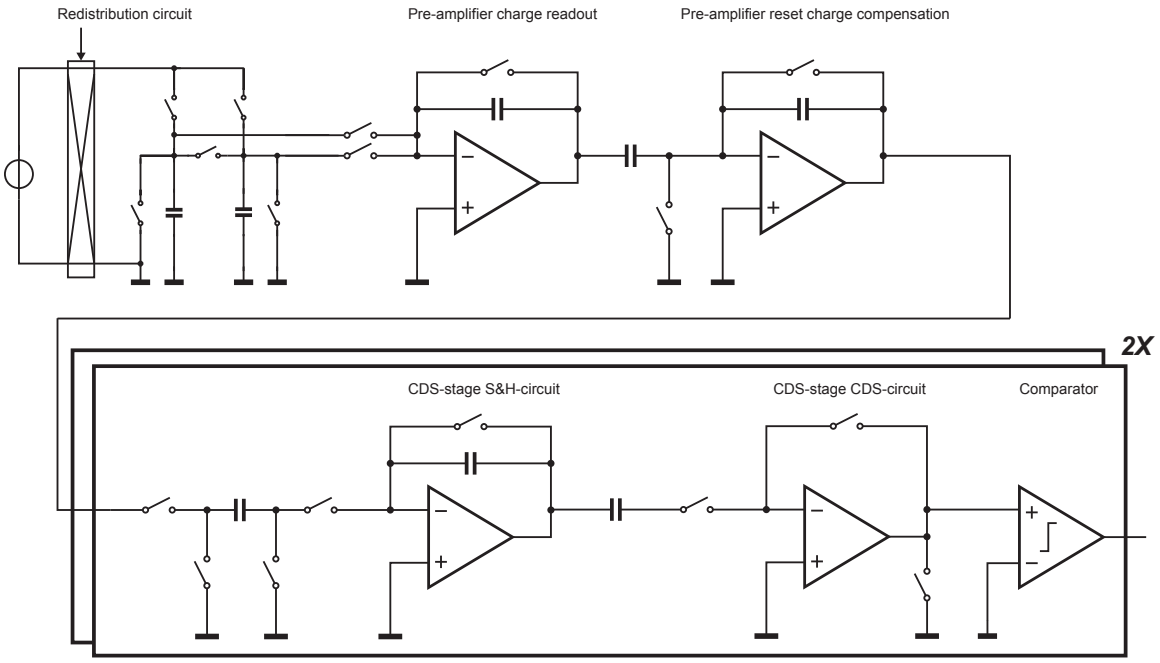

Figure 3.23: The total system.

4
Circuit Implementation

This chapter discusses the circuit implementation of the system level blocks discussed in the previous chapter. The first four sections (4.1, 4.2, 4.3 and 4.4) discuss the general circuit. The last two sections (4.5 and 4.6) discuss the supporting circuitry that make the circuit work and allow it to be debugged. The last section (4.7) discusses the general chip layout.

4.1. Redistribution Circuit

As discussed in the previous chapter, the redistribution circuit is responsible for generating an accurate threshold voltage with a well-defined amount of noise. Firstly, in order to achieve this accuracy, the redistribution circuit must support different switching schemes and both capacitors must be read out. Secondly, in order to provide offset compensation, an input chopper must be implemented. At last, in order to provide means for digital noise calibration, the capacitor size must be variable. The differential circuit that has all the switches to support these requirements is shown in Figure 4.1.

4.1.1. Capacitor size

The threshold voltage must be accurate to 167 ppm. For the highest SNR it is best to use the smallest possible capacitor, but the matching properties get worse with smaller capacitors. MATLAB simulations have shown that with two different switching schemes the threshold voltage inaccuracy due to capacitor mismatch is given by:

$$\frac{\Delta V}{V} = \alpha \cdot \left(\frac{\Delta C}{C}\right)^2 \tag{4.1}$$

where α is a factor that depends on the actual number of redistribution cycles, and is always less than 3 for up to at least 12 redistribution cycles. Metal fringe capacitors will be used because of their linearity, temperature stabiliy and matching. In the CMOS14 process the capacitor mismatch (3σ) is given by:

$$\frac{\Delta C}{C} = \frac{3.5 \cdot 10^{-2}\sqrt{fF}}{\sqrt{C}} + 9 \cdot 10^{-4} \tag{4.2}$$

In the worst case ($\alpha = 3$) this requires a capacitor of at least 29 fF, or 40 fF differentially. To provide some headroom, for the effect of parasitics, charge injection mismatch etc., a unit capacitor of 100 fF is used in the differential circuit which could theoretically give an 82 ppm accuracy.

The differential RMS output noise voltage of the redistribution circuit with unit capacitors of size C is $\sqrt{kT/C}$. Due to the CDS, the RMS voltage of the samples that will actually be detected is $v_n = \sqrt{2kT/C}$ (subtraction of two uncorrelated noise samples). The optimal threshold voltage is given by $V_{t,opt} \approx 2v_n = 2\sqrt{2kT/C}$. Due to the CDS, the required threshold voltage to generate is then $V_t = \sqrt{2kT/C}$ (doubling of the threshold voltage). Hence, for optimal detection at room temperature and C = 100 fF, the required threshold voltage is: $Vt \approx 288$ µV.

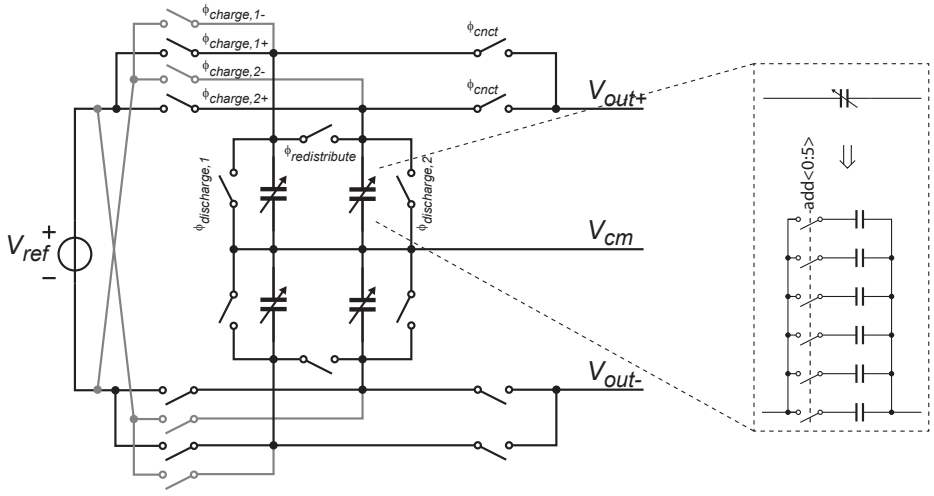

Figure 4.1: The redistribution circuit, showing all required switches.

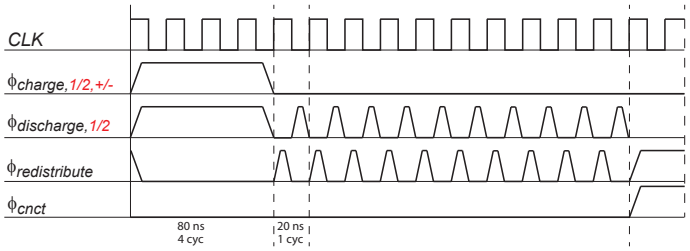

Figure 4.2: The timing diagram.

4.1.2. Timing requirements

In case N redistribution cycles are used, the voltage is divided by a factor 2^N. To allow for a reasonably large reference voltage, 11 redistribution cycles are implemented. In that case the required input reference voltage is: 288 µV · 2048 = 590 mV.

As shown in Section 3.3.1, the accuracy of the output voltage depends on the settling speed as shown in Equation 4.3. With 12 tau settling, the final output voltage is accurate to at least 135 ppm. It is also shown in that section that the correlation between the samples is then small enough.

$$V_{out} = V_{in} \cdot \frac{1}{2^N} \cdot \left(1 - e^{-\frac{t}{\tau}}\right)^{2N} \tag{4.3}$$

The timing requirements are derived from the design of the pre-amplifier (see Section 4.2.3). With this pre-amplifier 280 ns are available to generate the threshold voltage. With a clock frequency of 50 MHz, exactly 14 clock cycles are available for the redistribution, where only 11 clock cycles are required. The timing diagram is shown in Figure 4.2. Of each clock cycle, 10 ns is used for the (dis)charging, and the other 10 ns is allocated to the redistribution. Keeping in mind that these phases must be non-overlapping, the requirement for sufficient settling is: $\tau < 0.8$ ns.

4.1.3. Switch sizing

The charging switches implement the input chopper. As a result, their charge injection mismatch will not be compensated, however the offset voltage will be divided by 2048. To keep the charge injection mismatch to a minimum, minimum size NMOS-only switches are used.

The on-resistance of a minimum size switch is not sufficiently low over all corners to achieve sufficient settling in the available time. Therefore, the 3 unused clock cycles are added to the charging phase (see also Figure 4.2). However, with the variable capacitor at its maximum value, the settling speed is still too low. Therefore also the charging switch is variable together with the capacitor, which results in a constant settling speed independent of

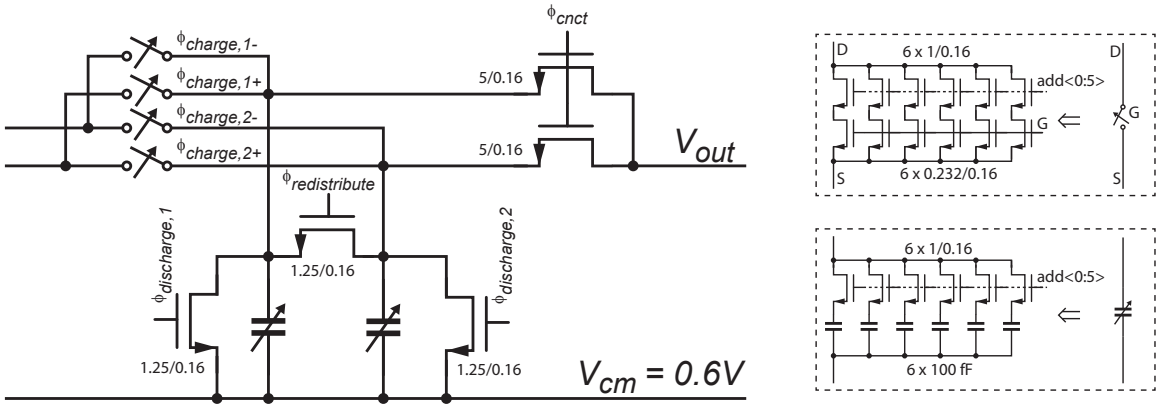

Figure 4.3: The schematic of the redistribution circuit, single-ended equivalent.

the capacitor size. The larger switches can have more charge injection mismatch, but this will then be injected over a larger capacitor.

The redistribution and discharge switches were chosen to provide sufficient settling with all capacitors enabled ($\tau < 0.8$ ns).

For the readout switches (controlled by Φ_{cnct}), it is best to use the widest possible switch, since their on-resistance will contribute to the unwanted readout noise. However, to keep the voltage dependent part of their parasitic below the maximum ΔC that the DEM scheme can handle, a width of 5 µm is used. To lower the on-resistance, a low common-mode voltage of 0.6 V is used, resulting in an on-resistance of 188 Ω nominally.

The circuit diagram showing all transistor sizes (W/L, both in µm) is shown in Figure 4.3. This schematic only shows half of the differential circuit.

4.1.4. Layout considerations

In the layout special care was taken to ensure matching of the capacitors and the charging switches. In both cases a unit cell was created which is placed in a two-dimensional array with dummies on the sides.

The unit cell of the capacitor array is shown in Figure 4.4, which also shows part of the neighbouring capacitors. The unit cell contains the 100 fF fringe capacitor and the switch to configure which capacitors are used. In the vertical direction are the unit capacitors that make one variable capacitor, whereas the different variable capacitors are laid out in the horizontal direction.

The bottom-plate of one variable capacitor is next to the top-plate of another variable capacitor such that the parasitics in the horizontal direction mainly result in an extra capacitance to the common-mode (bottom-plate). An even number of fingers is used for the capacitors such that the parasitics in the vertical direction are also mainly to the common-mode.

Furthermore, a lot of attention was paid to minimize the parasitics (cross-)coupled over the redistribution switch. These capacitors are illustrated in Figure 4.5. Their effect will now be explained.

After the redistribution, there is no voltage over the parasitic, $V_{Cpar,1} = 0$, and hence there is no charge on this parasitic. When discharging one of the capacitors, a voltage will appear over the parasitic. Charge is taken from the capacitor that is not being discharged to provide the charge required for $C_{par,1}$. This charge is destroyed when the redistribution switch closes again. Hence, as a result of this parasitic, charge 'leaks' away and the final output voltage will be too low. Something similar happens for $C_{par,2}$. Before discharging one of the capacitors, the voltage over the parasitic is $V_{Cpar,2} = 2 \cdot V_C$, while after discharging the capacitor $V_{Cpar,2} = V_C$ (V_C is the voltage over a single capacitor). The excess charge from the parasitic is dumped in the capacitor that is not being discharged. As a result, there is too much charge and the final output voltage will be too high.

Theoretically, the effects of these parasitics could perfectly cancel each other. But due to process variations, the size of the parasitic is quite unknown and it is chosen to simply minimize these parasitics. This is achieved by moving the drain and source contacts further away from the gate of the redistribution switch, by applying shielding, and by careful routing.

In order to minimize clock feed-through, the digital clock signals are shielded from the analog signals. This was especially important in the array of charge switches, since the clock signals are routed along the analog signals. In addition, a large M1-M5 shield is placed around the capacitor array, since all digital clock signals are routed

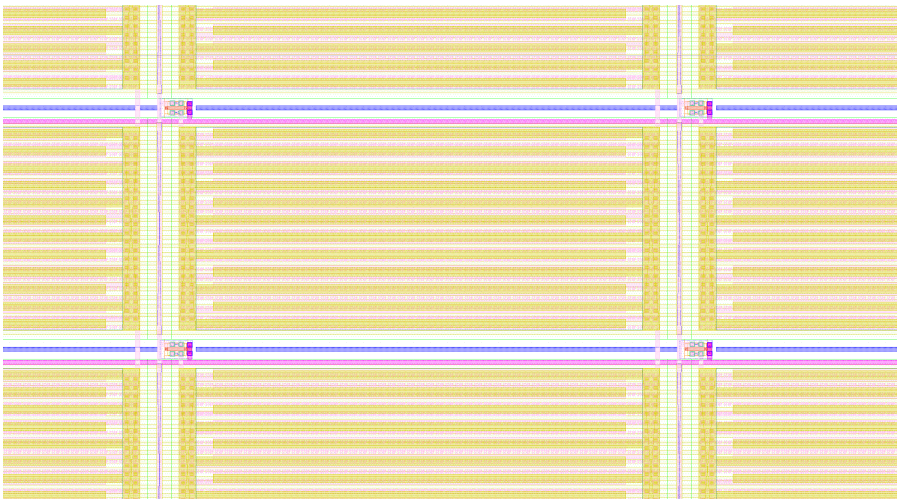

Figure 4.4: The unit cell of the capacitor array.

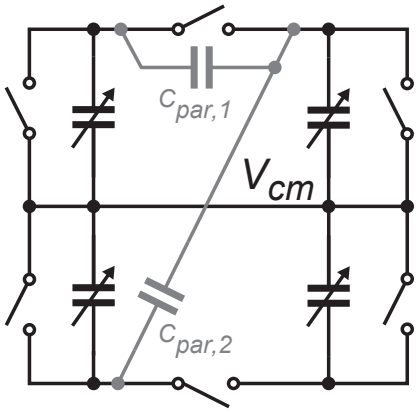

Figure 4.5: Illustration of the parasitic over the redistribution switch ($C_{par,1}$), and the parasitic cross-coupled over the redistribution switch ($C_{par,2}$).

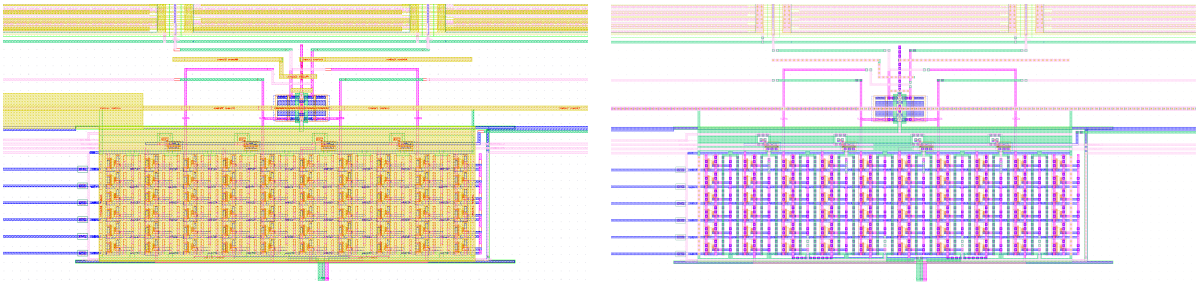

Figure 4.6: The layout of the switches for the redistribution circuit. The picture on the right has metal 5 hidden.

around this block.

Figure 4.6 shows most of the routing and shielding of this redistribution circuit. An improved layout can be found in Appendix B. That layout has even less parasitics between the capacitor nodes and better matches the parasitics of the different capacitor nodes. However, because of the additional shielding, there is a higher parasitic to ground. This layout was not used, since the clock generator could not drive the increased clock load.

4.1.5. Simulation results

A Monte-Carlo analysis was performed on the extracted layout. The simulation uses two special switching schemes and the input chopper to determine the final output voltage. These switching schemes are optimized for the case of capacitor mismatch, but ignore the effects of charge injection. The simulation results are shown in Table 4.1 for an input voltage of 512 mV. Simulations for the nominal (at 27 °C) and slow corner (at 125 °C), and for the minimum and maximum value of the variable capacitor were performed. V1 and V2 show the voltages for respectively switching scheme 1 and 2. The a and b suffix are for the chopper polarity, and V1/V2 without suffix show the output voltage when combining these chopped values using an ideal CDS.

It can be seen that with this layout the chopped values are close to $\pm$250 µV, with a maximum fixed offset of about 35 µV. By performing the chopping, the standard deviation of the output voltage already reduces to the sub-µV level. This is because the chopping deals with the charge injection mismatch offset.

The voltages for the two switching schemes are averaged to determine the final threshold voltage out of the redistribution circuit, Vt. It can be seen that in the nominal corner the standard deviation of the output voltage decreases even more. This shows that the switching schemes do compensate for the capacitor mismatch. Although the different switching schemes also help in the slow corner, the results are less impressive. The last row in the table shows the temperature error based on 3σ of the threshold voltage. This is below 0.2 K in the nominal corner and around 1.0 K in the slow corner.

	Single capacitor				**All capacitors**			
	Nominal		**Slow, 125 °C**		**Nominal**		**Slow, 125 °C**	
	Mean	**Std.dev.**	**Mean**	**Std.dev.**	**Mean**	**Std.dev.**	**Mean**	**Std.dev.**
V1a	259.4 µV	46.04 µV	257 µV	36.12 µV	253.3 µV	11.43 µV	255.4 µV	9.517 µV
V1b	-239 µV	45.84 µV	-241.8 µV	35.99 µV	-244.5 µV	11.37 µV	-248.3 µV	9.529 µV
V1	498.3 µV	456.9 nV	498.8 µV	933.8 nV	497.8 µV	284.1 nV	503.6 µV	408.8 nV
V2a	283.8 µV	40.33 µV	280.4 µV	30.19 µV	260.5 µV	9.661 µV	261.4 µV	7.938 µV
V2b	-215.4 µV	40.53 µV	-219.1 µV	30.33 µV	-239.1 µV	9.736 µV	-243.8 µV	8.105 µV
V2	499.2 µV	452.9 nV	499.5 µV	472.7 nV	499.6 µV	284.8 nV	505.2 µV	479.2 nV
Vt	498.8 µV	53.62 nV	499.1 µV	326.9 nV	498.7 µV	3.459 nV	504.4 µV	325.1 nV
Terror		0.17 K		1.04 K		0.01 K		1.04 K

Table 4.1: Simulation results of the Monte-Carlo analysis on the extracted layout.

4.2. Pre-Amplifier

The pre-amplifier amplifies the input sample, which is the charge stored on the redistribution capacitors. As explained in Section 3.3.2 the pre-amplifier consists of 2 stages. The first stage uses feedback to read out the charge

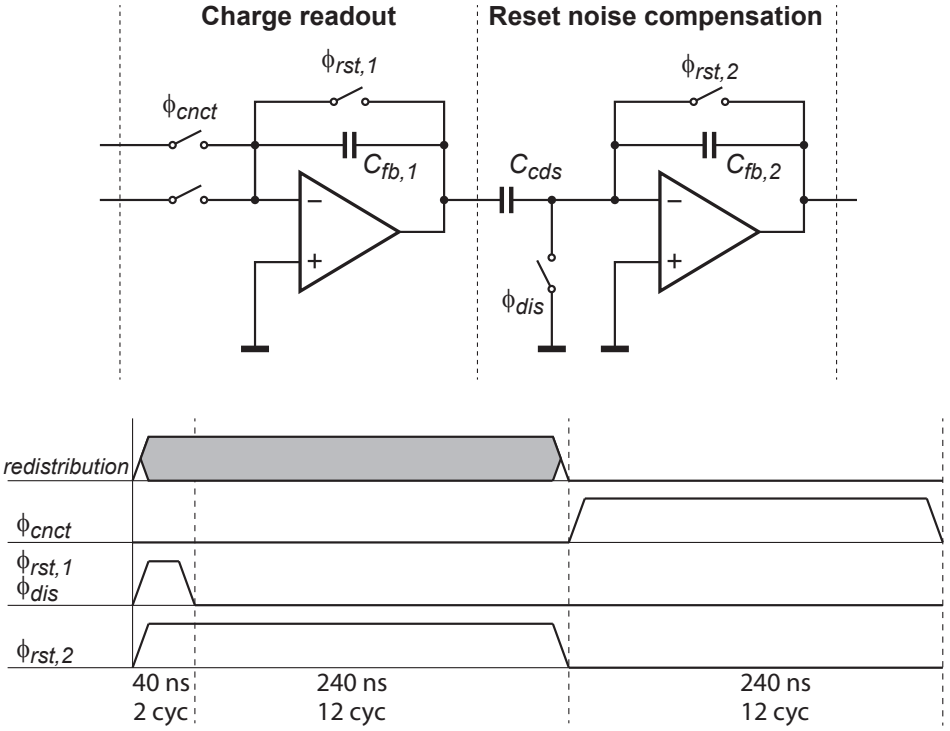

Figure 4.7: A single-ended equivalent of the pre-amplifier with the timing diagram.

and provide some amplification. The second stage compensates for the reset charge in the feedback capacitor of the first stage.

A single-ended equivalent of the circuit and the timing diagram are shown in Figure 4.7. This circuit is designed for low-noise, since any noise from the readout circuit results in a temperature offset error. Furthermore, this block is designed for maximum gain without clipping the output, such that the noise contribution of the next block is negligible.

4.2.1. First stage of the pre-amplifier

To relax the noise requirements on the second stage, the gain of the first stage is set to 10x. This requires a feedback capacitor of 20 fF. With this gain, the output swing of the first stage is limited to less than 100 mV and a telescopic amplifier can be used. This is advantageous for the noise, since there'll be no additional noise from the folding transistors.

Furthermore, to get low noise the input pair uses NMOS transistors, since they can provide a higher transconductance for the same input capacitance. The transistors have a length of 0.24 µm, which is slightly more than the minimum length. As a result, for this process, the flicker noise corner is at a much lower frequency (50% lower) while the transconductance is almost the same as for a minimum length transistor (6% lower). As discussed in Section 4.1.3, the common-mode voltage of the redistribution circuit is at 0.6 V. Still the highest possible V_{gt} is used, since this results in the highest g_m/C_{in}.

Next, the transistor width, and hence the required bias current, were chosen such that the SNR (based on thermal noise) is highest for the 100 fF input capacitance. Figure 4.8 shows a plot of the maximum ENBW that can be allowed to achieve a 0.1 K temperature error due to readout noise for different sizes of the input capacitor. The blue lines show the range of values of the input capacitor that can be used for normal operation and for the digital noise calibration. It can be seen that for the chosen transistor sizing, the 100 fF input capacitance is slightly past the optimum. This is better for the digital noise calibration, since that assumes that the readout noise contribution remains the same for different values of the input capacitor (see Section 3.3.1). That condition is best met past the optimum (see Section 3.2.3). The resulting width and current (80 µm, 100 µA) are close to the optimum as predicted by the model in Appendix A.

A plot of the output noise PSD of the telescopic amplifier is shown in Figure 4.9. This plot was created with the load capacitor disconnected in order to clearly show the thermal and flicker noise. The plot also shows the

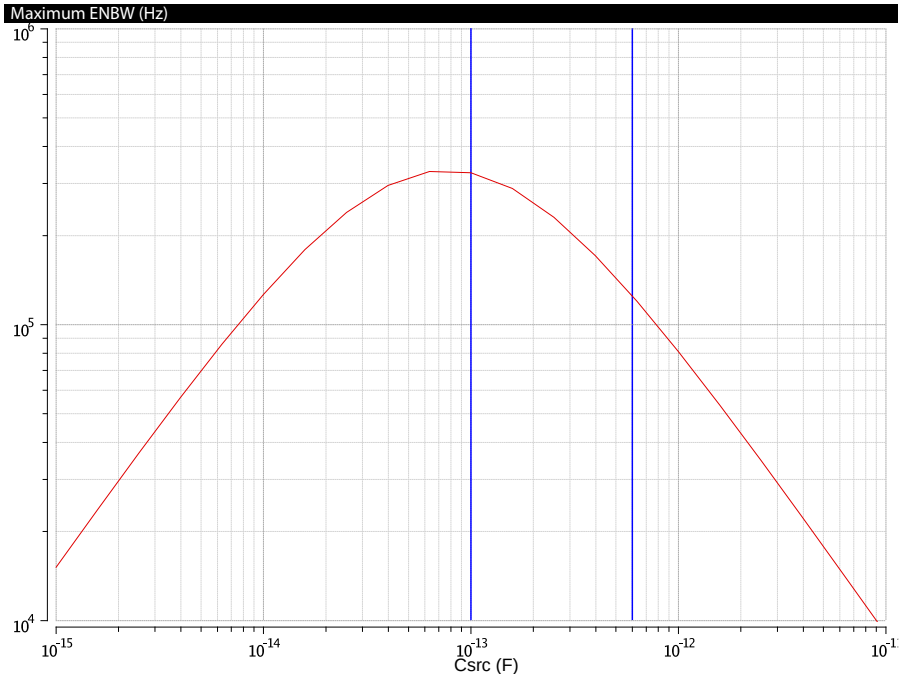

Figure 4.8: The maximum ENBW that can be allowed to achieve a 0.1 K temperature error due to readout noise for different sizes of the input capacitor.

contribution of the on-resistance of the input switches. It can be seen that the noise contribution of the other sources is lower than the contribution of the input pair. The flicker noise corner is around 2 MHz. The load capacitor (10 pF) is chosen such that the -3 dB frequency is at this point.

The NMOS cascode transistors have a small voltage headroom since the input and output common-mode voltage are the same. To ensure that this headroom does not get too small, extra gain is required in the common-mode loop. A simple amplifier is used to control the output common-mode voltage. The circuit diagram of the first stage of the pre-amplifier, including its bias circuit, CMFB circuit and feedback circuit, is shown in Figure 4.10. All transistor sizes are shown as W/L, both in µm.

4.2.2. Second stage of the pre-amplifier

To provide the largest possible output swing, a folded-cascode amplifier is used for this stage. Furthermore, the output common-mode voltage is chosen half-way the supply-voltage, at 0.9 V. In this case the maximum differential output voltage is around $\pm$ 1.2 V. The maximum gain of this stage was then derived from the worst-case maximum input voltage. This is mainly determined by the maximum input signal ($V_t + 3 \cdot v_n$) and the maximum offset of the first stage, which are 0.9 mV and 4.3 mV respectively. This allows for a total gain of about 230x for both stages together. To leave some margin, the second stage provides an additional gain of 20x. Using a Monte-Carlo analysis it was verified that the output voltage does not exceed 1.2 V.

The folded-cascode amplifier was designed for low-noise. Therefore, PMOS transistors are used for the input pair and more current is spent in these transistors compared to the output branches. The resulting output noise PSD, including the transistor contributions, is shown in Figure 4.11. It can be seen that the noise is again dominated by the input pair, by design. The total contribution of this amplifier to the readout noise is about the same as the contribution of the input switches of the first stage of the pre-amplifier.

With the required transconductance to meet the noise requirement, the unity gain frequency is already such that the amplifier can settle to more than 12τ accuracy. The circuit diagram with the transistor sizes and branch currents is shown in Figure 4.12. Again all transistor sizes are shown as W/L, both in µm. To maximize the output swing, a switched-capacitor circuit with transmission gates is used for the common-mode feedback. Its circuit diagram is shown in Figure 4.13.

4.2.3. Timing requirements

Simulations have shown that 3τ settling of the first stage (with the input capacitor connected) is sufficient for the reset noise compensation. Hence, with f_{-3dB} = 2 MHz, 240 ns is required for settling. For the reset phase 40 ns

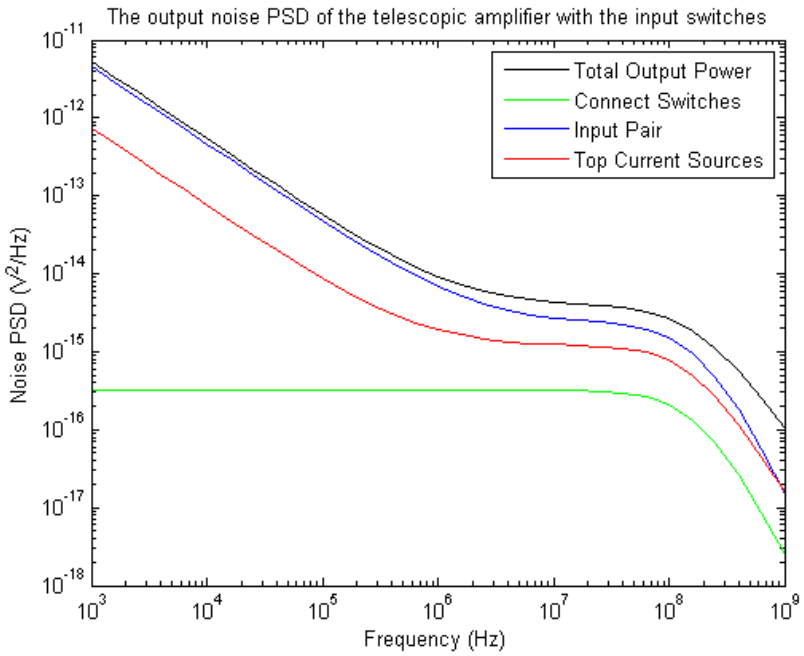

Figure 4.9: The output noise PSD of the telescopic amplifier with the input switches.

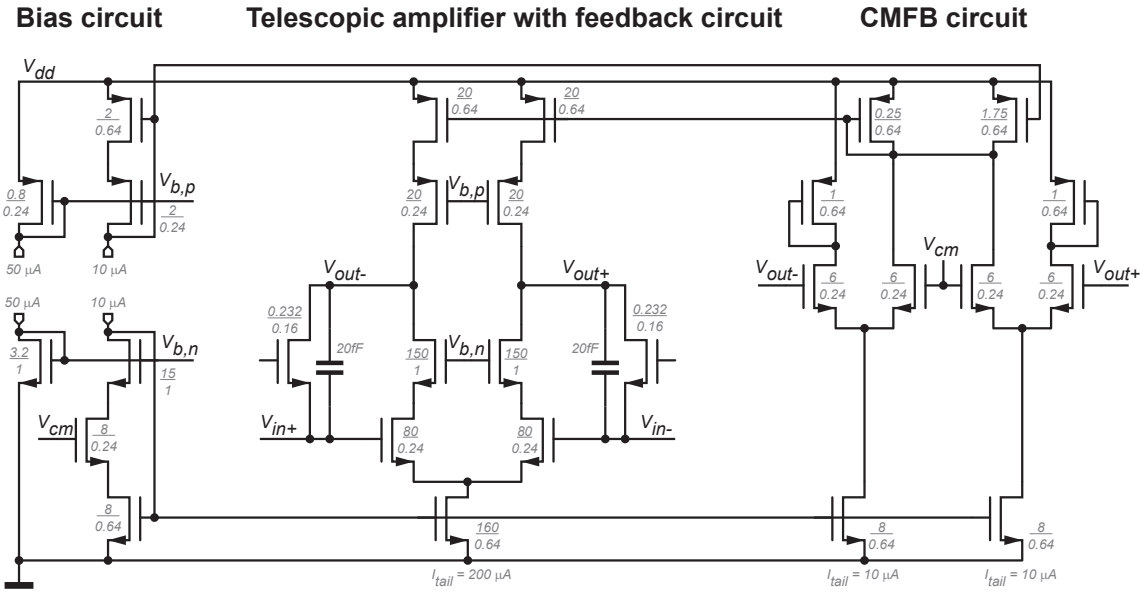

Figure 4.10: The first stage of the pre-amplifier, including its bias circuit, CMFB circuit and feedback circuit.

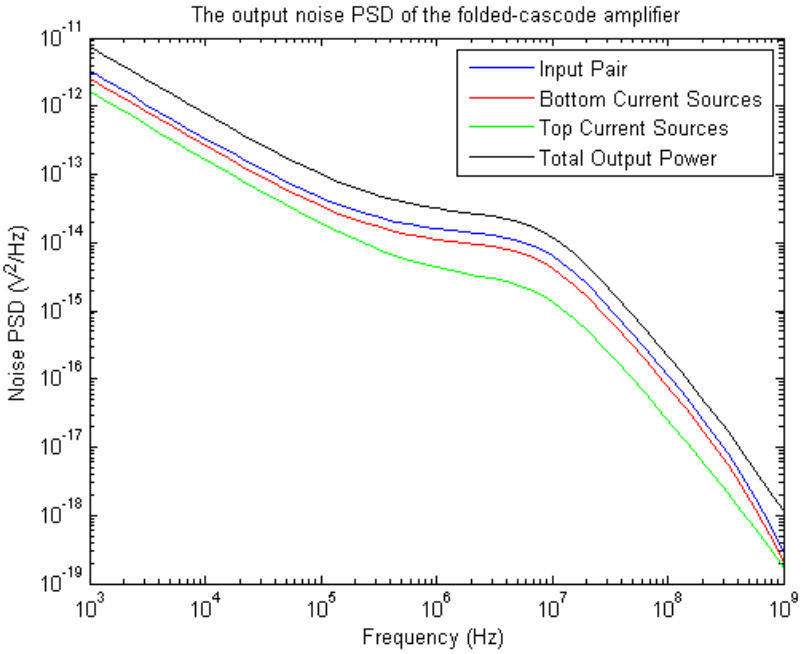

Figure 4.11: The output noise PSD of the folded-cascode amplifier.

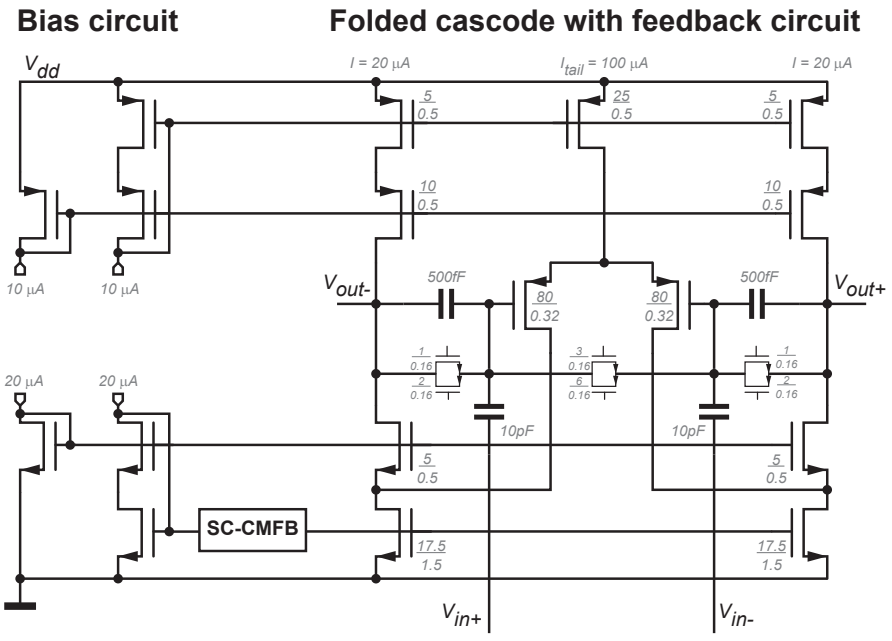

Figure 4.12: The second stage of the pre-amplifier, including its bias circuit and feedback circuit.

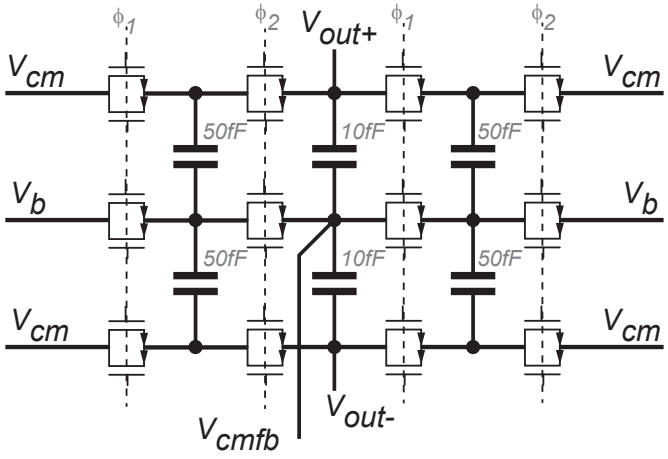

Figure 4.13: The switched-capacitor common-mode feedback circuit, all switches are minimum size.

Device	Param	Noise Contribution
Charging switches (noise source)	Sth	99.10 %
Pre-amplifier	Sth	0.502 %
Pre-amplifier	Sfl	0.344 %
Input switches	Sth	0.034 %
Second amplifier	Sth + Sfl	<0.01 %
Pre-amplifier reset switches	Sth	<0.005 %

Table 4.2: The individual noise contributions of the total pre-amplifier circuit.

is used. Therefore, the time required for each sample is 40 ns + 2 · 240 ns = 520 ns, and 40 ns + 240 ns = 280 ns is available for the redistribution (see also Figure 4.7).

The required non-overlap conditions are as follows:

- The input charge must be connected after the reset switch of the second stage has opened.
- The reset switch of the second stage must close after the next block has sampled the output voltage.
- The clock phases of the switched-capacitor common-mode feedback circuit must be non-overlapping.

The common-mode feedback clock phases switch at the same time as the clock phase for the input connect switch. In this case the clock phases are the same for every input sample that is processed. If this were not the case, an error is made when later the samples are subtracted by the correlated double sampler.

4.2.4. Layout considerations

The layout of these stages has no special features: dummies are used to minimize offset, and all clock signals are shielded. Special care was taken to keep the input signals low-ohmic, since the resistance results in additional noise.

4.2.5. Simulation results

A timedomain PNOISE simulation was used to determine the individual noise contributions and to see the effect of the reset noise compensation. A summary of the results is shown in Table 4.2.

It can be seen that 0.9% of the noise is from the readout circuit. This corresponds to a temperature offset error of 2.7 K. It can also be seen that with the chosen timing, less than half of the noise is originating from flicker noise (the simulation uses an ideal CDS circuit). The noise contribution of the second stage of the pre-amplifier is negligible.

The last row in the table shows the contribution of the reset switches of the first stage of the pre-amplifier. Without the reset noise compensation, the contribution of these switches to the output noise was about the same as the contribution of the charging switches. Therefore, it can be concluded that the reset compensation is effective.

	SNSP	SNFP	NOM	FNSP	FNFP
−55 °C	3.00 K	3.18 K	2.52 K	4.08 K	3.96 K
25 °C	2.64 K	2.70 K	2.76 K	4.56 K	4.68 K
125 °C	2.88 K	2.94 K	3.06 K	5.88 K	6.00 K

Table 4.3: A schematic simulation of the total noise error over corners and temperatures.

The results of a simulation of the total readout noise over different corners and temperatures is shown in Table 4.3. It can be seen that the noise performance is worse in the fast NMOS corners, but by increasing the bias current the noise can be further reduced. These simulations are based on the schematic circuit. The extracted simulation is slightly worse and has a 3.0 K temperature error in the nominal case at room temperature.

4.3. CDS Stage

As explained in Section 3.3.3, also the CDS stage consists internally of 2 stages; a sample & hold stage, followed by the actual correlated double sampler. To double the sample rate, the difference between every pair of samples is taken and this requires double the circuit. A single-ended equivalent of the circuit and the timing diagram are shown in Figure 4.14.

The input voltage is first sampled on both sampling capacitors ($C_{s,a/b}$) and is then transferred to the hold capacitor ($C_{h,a/b}$), which was reset just before. After the charge has been transferred and this stage is in the hold-phase, the CDS stage is connected. The two CDS circuits have different clock phases in order to take the difference between different pairs of samples.

In order to perform the subtraction of two samples, the CDS stage is connected with the feedback switch closed for the first sample. This sample is then stored on $C_{cds,a/b}$. For the second sample the feedback switch is open, and the voltage difference of the two samples is available at the virtual ground. The amplifier then amplifies this voltage difference in open-loop configuration. Due to the finite time, the amplifier (an OTA) acts like an integrator, and hence a reset switch at the output is included to ensure that the output starts integrating from 0 V. Even though this switch injects charge, the offset improves.

4.3.1. Offset considerations

This stage is designed such that the total readout circuit achieves a low offset. An input referred offset of 94 nV can be allowed for the readout circuit for a temperature error of 0.1 K. The pre-amplifier provides a gain of about 190x, hence the residual offset referred to the input of this stage can be 18 µV.

The final offset is determined by two factors. First of all, by the quality of the offset compensation. This stage must compensate for the offset of the redistribution circuit, pre-amplifier and its own stages. Secondly, the final offset is determined by the part of the offset that is not compensated for. The comparator following this circuit has an offset that will not be compensated later, and hence this offset must be made negligible when referring it to the input of the readout circuit.

The quality of the offset compensation is limited by different factors. Firstly, any charge injection mismatch on the $C_{cds,a/b}$ capacitors leads to an offset. Secondly, finite settling of the voltage on the $C_{cds,a/b}$ capacitors results in an offset, since the settling speed is different when the second stage is in open-loop or in closed-loop configuration. Finally, a finite gain of the second stage amplifier results in an imperfect cancellation of its offset. This is because not the full offset appears at the amplifier's virtual ground when it is placed in unity gain feedback.

In order to deal with the charge injection mismatch problem, a capacitor of at least 3 pF must be used for $C_{cds,a/b}$ in case a minimum size transmission gate is responsible for the charge injection ($3\sigma_v$ = 17 µV).

The required settling can be approximated by assuming that the settling is infinitely fast in case the second stage is in open-loop configuration. This is a valid approximation, since the capacitor that has to be charged is the input parasitic of the second stage amplifier, which can be assumed to be much smaller than 3 pF. The settling error is then simply the error caused by incomplete settling in one phase, given by $V_{error} = V_{in} \cdot e^{-\frac{t}{\tau}}$.

Since the input voltage can be as large as $\pm$ 1.2 V (the previous stage was designed for maximum gain without clipping the output), at least 11.1τ settling is required for an error of 18 µV in the worst-case.

The settling speed is limited by the settling speed of the S&H amplifier, the CDS amplifier and the resistance of the connect switch. For $C_{cds,a/b}$ = 3 pF the maximum equivalent charging resistance is given by:

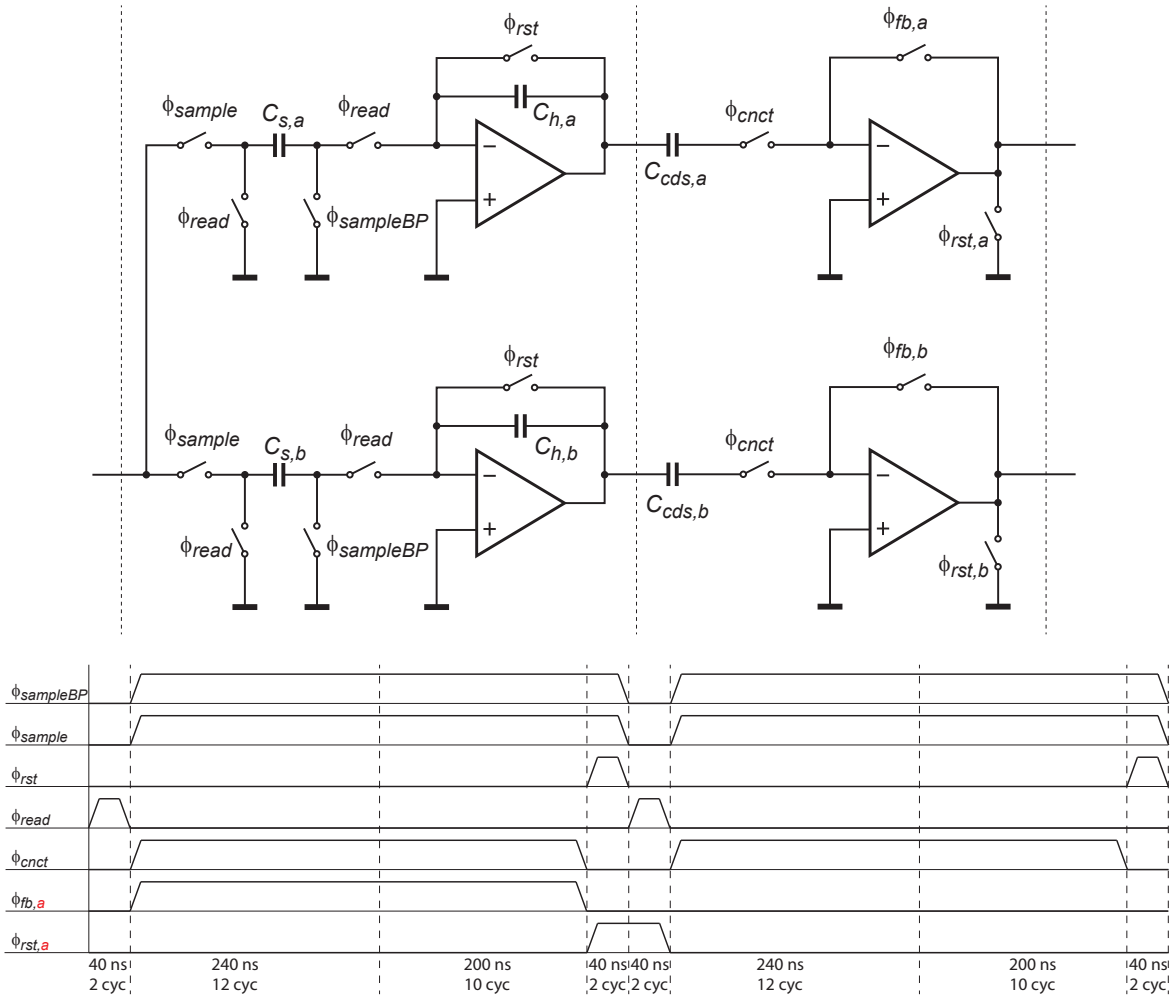

Figure 4.14: A single-ended equivalent of the CDS-stage with the timing diagram.

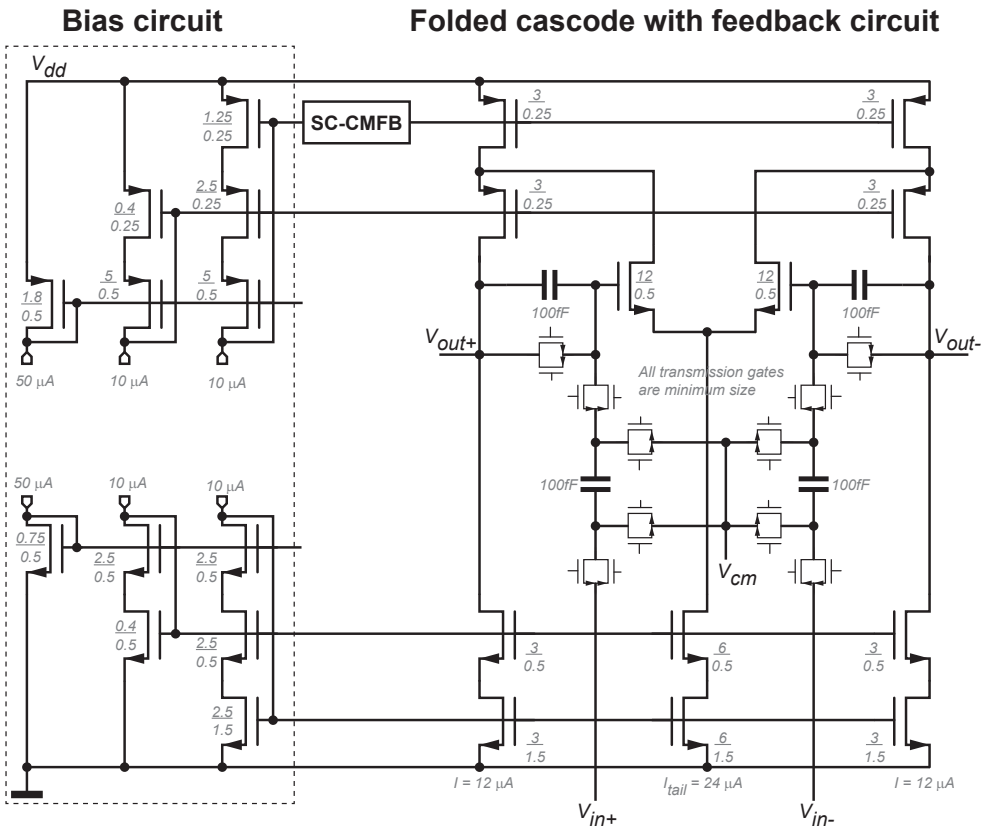

Figure 4.15: The S&H circuit, including its bias circuit and feedback circuit.

$$R_{max} = \frac{\tau}{C}$$
$$and \tag{4.4}$$
$$R_{max} = \frac{1}{g_{m,sh}} + \frac{1}{g_{m,cds}} + R_{on,cnct}$$

The maximum charging resistance is simply the sum of the resistances and inverses of transconductances in the charging path, since the amplifiers are in unity gain feedback. The on-resistance of the connect switch will be ignored for now. The time available for charging is maximized to 440 ns by optimizing the timing (see Section 4.3.4). Assuming the same transconductance for both amplifiers, the transconductance should be higher than 0.15 mS, or stated differently UGF > 8 MHz.

4.3.2. The S&H-stage

The S&H-stage cannot provide additional gain, since the input is already rail-to-rail and there is no additional offset compensation in this circuit. Again a folded-cascode amplifier is used with the common mode voltage at 0.9 V in order to maximize the output swing. The main requirement for the amplifier is the fast settling as discussed in the previous section. Simulations of the extracted circuit over process corners and temperature show that in the worst case UGF = 8.6 MHz.

The circuit diagram of this stage, including its bias circuit, and feedback circuit, is shown in Figure 4.15. All transistor sizes are shown as W/L, both in µm. The same SC-CMFB circuit is used as for the second stage of the pre-amplifier as shown in Figure 4.13. The switches in the sampling circuit are tied to the common-mode voltage, since otherwise there are problems with the settling of the input common-mode voltage.

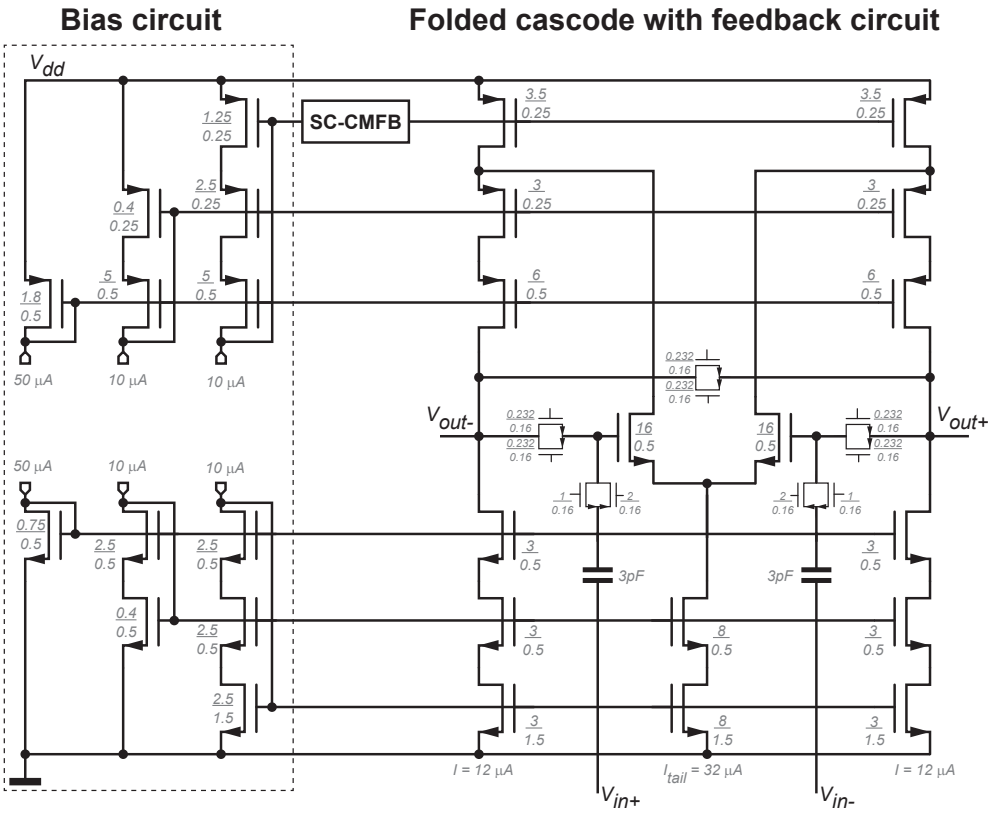

Figure 4.16: The CDS circuit, including its bias circuit and feedback circuit.

4.3.3. The CDS-stage

A 3 pF capacitor is used for $C_{cds,a/b}$. This requires that the charge injection mismatch is determined by a minimum size transmission gate. The connect switch cannot be minimum size, as then the settling requirement cannot be met. Therefore, the minimum size feedback switch is opened before the connect switch and will be the critical switch for the charge injection.

The residual input-referred offset due to the finite gain of this stage is: $V_{os,res} = V_{os}/(A+1)$. Therefore, a high DC-gain is important. Since the output of this stage is allowed to clip, as the offset is already compensated for, an extra cascode is used to boost the DC-gain. The resulting worst-case DC-gain of 75 dB allows for about 0.1 V offset. This is far more than required, and therefore the offset contribution of this amplifier stage can be neglected.

The UGF required for this stage depends not only on the settling requirement, but also on the required gain in open-loop configuration (the amplifier functions as an integrator). Therefore, this stage is made slightly faster than the previous stage (worst-case UGF is 13.1 MHz), such that the comparator offset requirements are relaxed. The resulting increased transconductance is the reason why the on-resistance of the connect switch was neglected when determining the requirements for the S&H-stage.

The circuit diagram of this stage, including its bias circuit and feedback circuit, is shown in Figure 4.16. All transistor sizes are shown as W/L, both in µm. Again the same SC-CMFB circuit is used as for the second stage of the pre-amplifier as shown in Figure 4.13. The bias circuit has a dashed line around it, since the same circuit is used for the two S&H amplifiers and the two CDS amplifiers.

4.3.4. Timing requirements

A sample must be processed every 520 ns. In order to maximize the hold time, only 40 ns is reserved for re-setting and another 40 ns for transferring the sampled charge to the hold capacitor. This was only possible by disconnecting the 3 pF load capacitor during this charge transfer.

The required non-overlap conditions are as follows:

- The bottom-plate sampling switch must close after the sample has been read and the read switch has opened.
- The sample switch must open after the bottom-plate sampling switch has opened.

- The sample must be read after the sampling and resetting have completed.
- The S&H-circuit can only be reset after the CDS circuit has been disconnected.
- The CDS feedback switch must be closed after the CDS reset switch and the S&H read switches have opened.
- The CDS reset switch must close after the CDS feedback switch has opened.
- The CDS connect switch must open after the next block has sampled the output voltage and after the feedback switch has opened.
- The clock phases of the switched-capacitor common-mode feedback circuit must be non-overlapping.

The common-mode feedback clock phases of the S&H-circuit switch at the same time as the S&H-circuit reset clock. In this way the clock phases are the same for each input sample that is processed. The common-mode feedback clock phases of the CDS circuit on the other hand change polarity each time the CDS circuit is disconnected. In this case the clock polarity switches when the CDS configuration changes from closed-loop to open-loop and vice versa.

4.3.5. Layout considerations

The layout of these stages has again no special features: dummies are used to minimize offset, and all clock signals are shielded. Also, no matching is required between any components of the two CDS channels.

4.3.6. Simulation results

In order to accurately determine the offset performance, the comparator circuit was connected to the output of the CDS stage and clocked. This is important, since the input capacitance of the comparator determines the gain in open-loop configuration and the comparator could kickback charge.

The resulting total gain (including the pre-amplifier) over corners and temperatures is shown in Table 4.4. In the worst-case, a comparator offset of 10 mV would translate into a temperature error of 0.07 K.

	SNSP	**SNFP**	**NOM**	**FNSP**	**FNFP**
-55°C	264k	246k	246k	187k	156k
25°C	223k	209k	211k	194k	169k
125°C	147k	168k	166k	157k	165k

Table 4.4: The total gain of the readout circuit from the pre-amplifier input to the comparator input.

A Monte-Carlo analysis on the total readout circuit (including the pre-amplifier, but not the redistribution circuit) over corners and temperatures was performed. A large signal was applied as input before applying a 0 V input signal in order to include the effect of previous samples on the current sample. The simulation results are shown in Table 4.5, which shows the worst-case sample of the Monte-Carlo analysis. It can be seen that the total error is always below 0.3 K and is below 0.1 K in the nominal corner.

	SLOW	**NOM**	**FAST**
-55°C	0.084 K		0.290 K
25°C		0.050 K	
125°C	0.227 K		0.253 K

Table 4.5: The offset performance of the readout circuit from the pre-amplifier to the CDS stage.

4.4. Comparator

For the metastability, the minimum gain of the total readout circuit should be 144 dB. This requires only an additional 100x gain from this comparator stage. The simple StrongARM latch as shown in Figure 4.17 is used. An elaborate noise analysis of this circuit is presented in [56]. The circuit works as follows. When CLK is low, the comparator is reset. The outputs are pulled high and the internal nodes X and X' charge to $V_{dd} - V_{t,n}$. When the CLK goes high, the input differential pair starts discharging nodes X and X' with a different rate depending on the input voltage. Then first only the cross-coupled NMOS pair start discharging the output nodes, and when

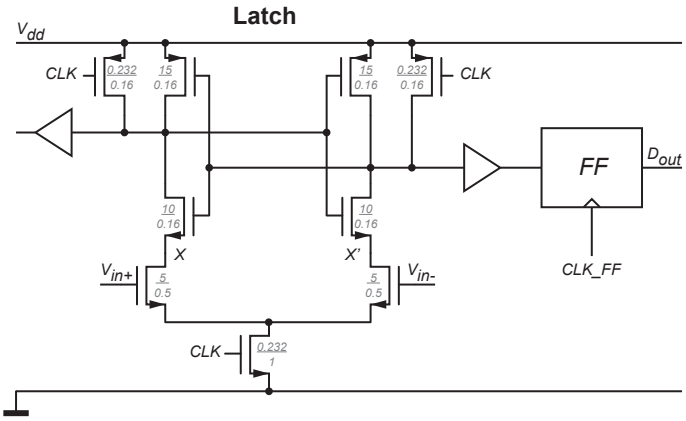

Figure 4.17: The latch, all transistor sizes are shown as W/L, both in μm.

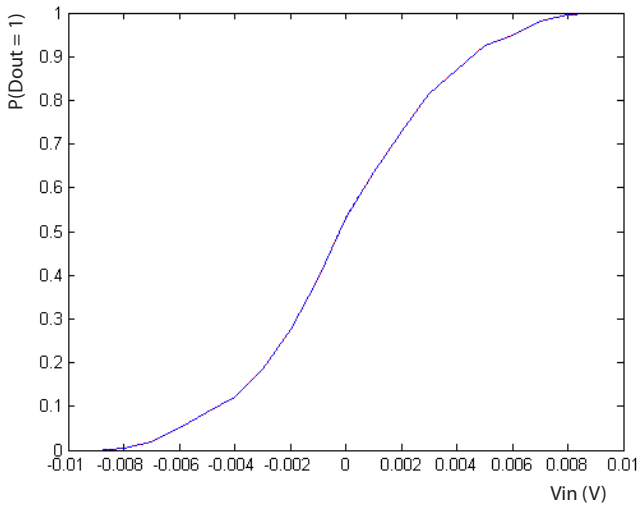

Figure 4.18: The probability that the output is high for a certain input voltage.

the output voltage has decreased enough also the PMOS transistors turn on. The cross-coupled inverter pair then regenerates the output voltage difference to digital levels.

The latch uses relatively large transistors to lower the input referred offset. The output of this latch is stored in a flipflop. For symmetry reasons, the latch output is first buffered and a dummy buffer is added to the unused output.

A Monte-Carlo analysis was performed to determine the offset performance using the methodology of [57]. A staircase input was applied to the comparator and the probability that the output was high for a certain input voltage was determined. This simulation was performed for an input voltage sweep from -10 mV to +10 mV and in the other direction in order to determine if there is a hysteresis. The resulting plot is shown in Figure 4.18, which shows no hysteresis. It can be seen that the offset is about 10 mV, leading to a worst-case temperature offset of 0.07 K as discussed in Section 4.3.6.

In the layout special care was taken to ensure matching of the transistors and parasitic capacitances. A post-layout simulation shows a fixed 0.3 mV offset without component variations.

4.5. Test circuits

4.5.1. Redistribution circuit

To be able to fully test the redistribution circuit, the size of the variable capacitor and the switching scheme used to redistribute the charge are programmable. In addition, to be able to test only the readout circuit, a copy of the circuit is placed on-chip that does not have the redistribution circuit. However, since the readout circuit expects a

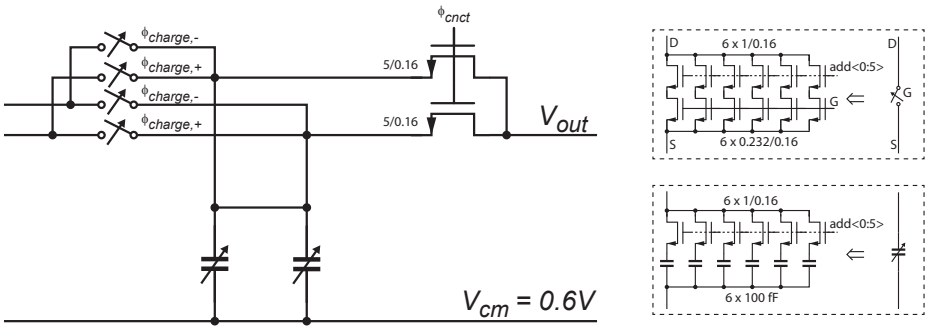

Figure 4.19: The stripped redistribution circuit that can not redistribute the charge.

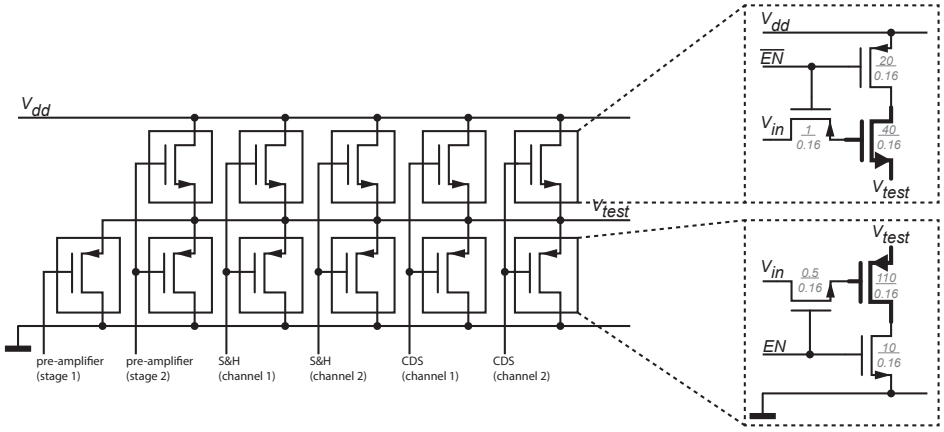

Figure 4.20: The source followers available to output internal voltages (single-ended equivalent).

charge on a variable size capacitor as input, the circuit as shown in Figure 4.19 is used (compare with Figure 4.3).

4.5.2. Readout circuit

The bias currents of the different blocks are all derived by mirroring an external 100 μA current. This current should have a PTAT characteristic such that the unity gain frequency of the different amplifiers remains almost constant over temperature. This current is generated off-chip, since this allows the current to be changed in case the readout circuit does not work properly.

Furthermore, the readout circuit can be debugged by outputting the output voltage of the different stages. This is done using source followers that should be loaded off-chip. Since the output voltage for all stages except the first pre-amplifier stage could be anywhere between the supply rails, both NMOS and PMOS source followers are placed, see Figure 4.20.

The source followers are designed for a load upto 20 pF (this requires an external load of 500 Ω). To supply the required current, relatively wide transistors with a large input capacitance are used, see Figure 4.20. In order not to disturb the readout circuit in normal operation, an additional switch is placed to disconnect the gate of the source follower from the output of the stage.

The source followers, the size of the variable capacitor and the switching scheme are programmed using a shift register.

4.6. Digital blocks

A clock generator was designed that generates the signals as shown in the timing diagrams of Figure 4.2 (redistribution circuit), Figure 4.7 (pre-amplifier) and Figure 4.14 (CDS stage). The clock generator also generates the phases for the switched capacitor common-mode feedback circuits and the comparator. The generated clock signals should adhere to the timing requirements as discussed in Section 4.1.2 (redistribution circuit), Section 4.2.3 (pre-amplifier) and Section 4.3.4 (CDS stage).

These signals and their non-overlap requirements are summarized in Figure 4.21. A red arrow means that the

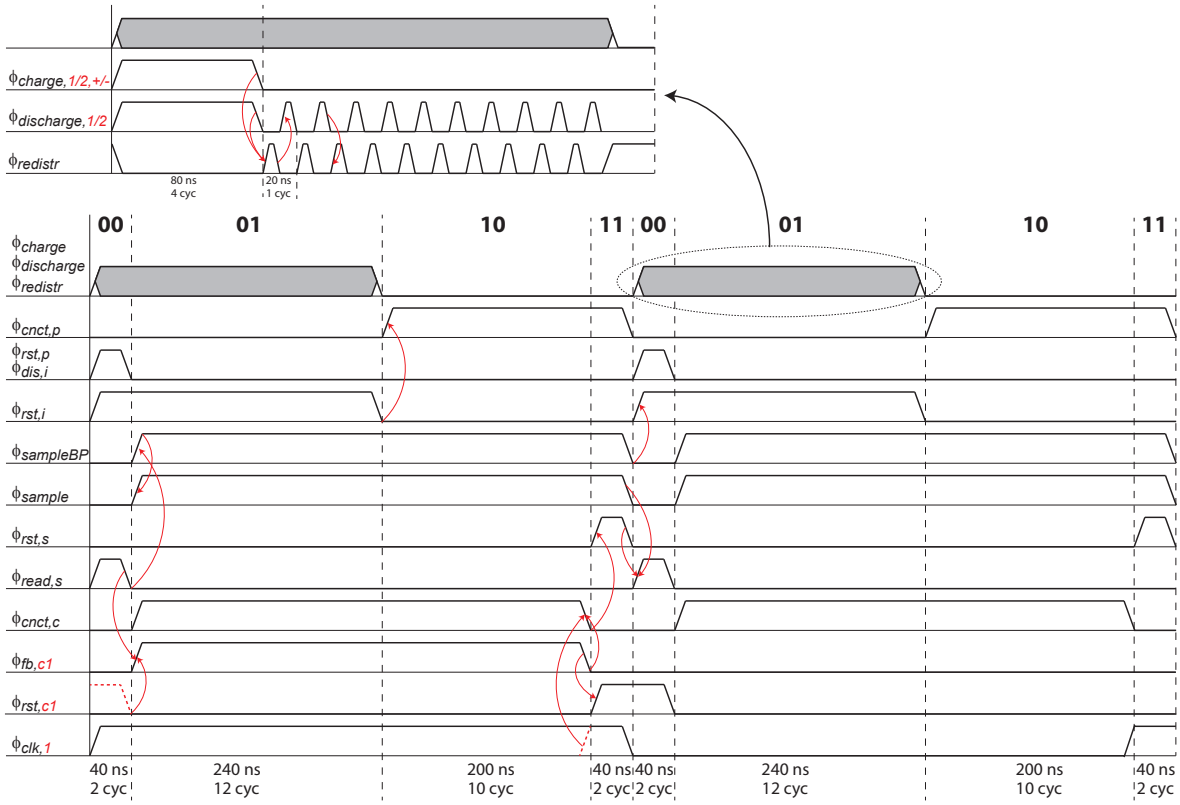

Figure 4.21: All required clocks phases and their non-overlap requirements.

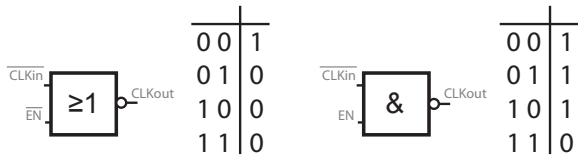

Figure 4.22: NAND and NOR truth tables.

clock edge pointed to should happen after the clock edge from which the arrow originates. To keep this diagram somewhat compact, the inverse clock signals and the switched capacitor common-mode feedback signals are not shown.

From this timing diagram different states can be distinguished, which are shown with their binary code (00, 01, 10 and 11) in the diagram. These states are derived from the 50 MHz external clock using some counters and state logic. The different clock phases are then derived from these state signals, but the resulting clock signals are still overlapping.

The non-overlap conditions are later added using NAND and NOR gates where the clock is applied to one input and the non-overlap condition to the other input, see the truth tables in Figure 4.22. E.g.: if a rising edge occurs on the input of a NAND gate (input clock), the output (output clock) will only go low when the other input is high (a certain condition is met).

The conditions are derived from other generated clock phases. Since also inverted clock signals must be generated for most phases, these inverted clock signals are used to set the conditions that must be met, see Figure 4.23 for an example. In this case it is ensured that the non-overlap condition holds for both the inverted and non-inverted clock phases.

For the redistribution circuit, the clock signals should not only be non-overlapping, but also programmable. Depending on the switching scheme either one capacitor should be discharged, or the other. Since there are 11 redistribution cycles, 11 bits are used to program the switching scheme. Before each complete redistribution, these 11 bits are loaded in parallel into a register. Then, for the generation of each discharging clock phase, this register

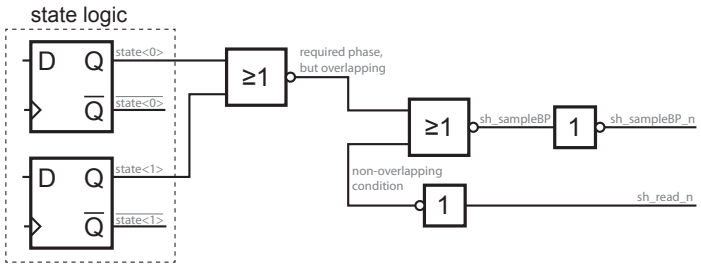

Figure 4.23: Example of the generation of the sampleBP phase and its inverse, where the non-overlap condition is derived from the S&H read phase.

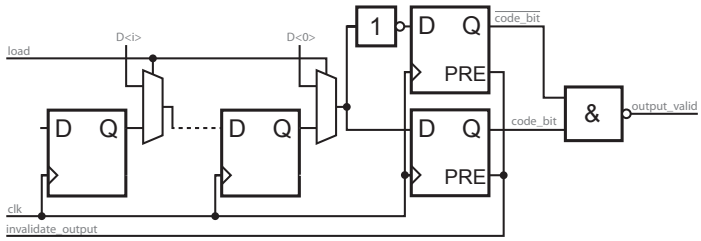

Figure 4.24: The parallel-in serial-out shift register, expanded with logic to determine whether the output is valid.

shifts out a single bit serially that determines which capacitor should be discharged.

To ensure that the correct bit is loaded, another output is added to the register that signals when the output is valid (see Figure 4.24). This signal is created using two flipflops with asynchronous presets. The output valid signal is then used as another condition that must be met in the generation of the non-overlapping signals. If this were not the case, the circuit could start discharging one capacitor and when then the output of the register changes continue with discharging another capacitor. This would corrupt the final output voltage.

4.7. Chip layout

The layout of the chip is shown in Figure 4.25. The padring and on-chip decoupling capacitors are not shown. A description of the numbered blocks is provided in table 4.6.

Number	Block description
1	Redistribution circuit
2	First stage of the pre-amplifier
3	2 x 10pF pre-amplifier load/CDS capacitor
4	Second stage of the pre-amplifier
5	S&H stage (channel 1)
6	S&H stage (channel 2)
7	2 x 3pF CDS capacitor (channel 1)
8	2 x 3pF CDS capacitor (channel 2)
9	CDS stage (channel 1)
10	CDS stage (channel 2)
11	Comparator (channel 1)
12	Comparator (channel 2)
13	Output driver
14	Clock generator
15	S&H/CDS bias voltage generator
16	Bias current mirrors
17	Shift register
18	Circuit copy, but without redistribution

Table 4.6: Description of the numbered blocks in Figure 4.25

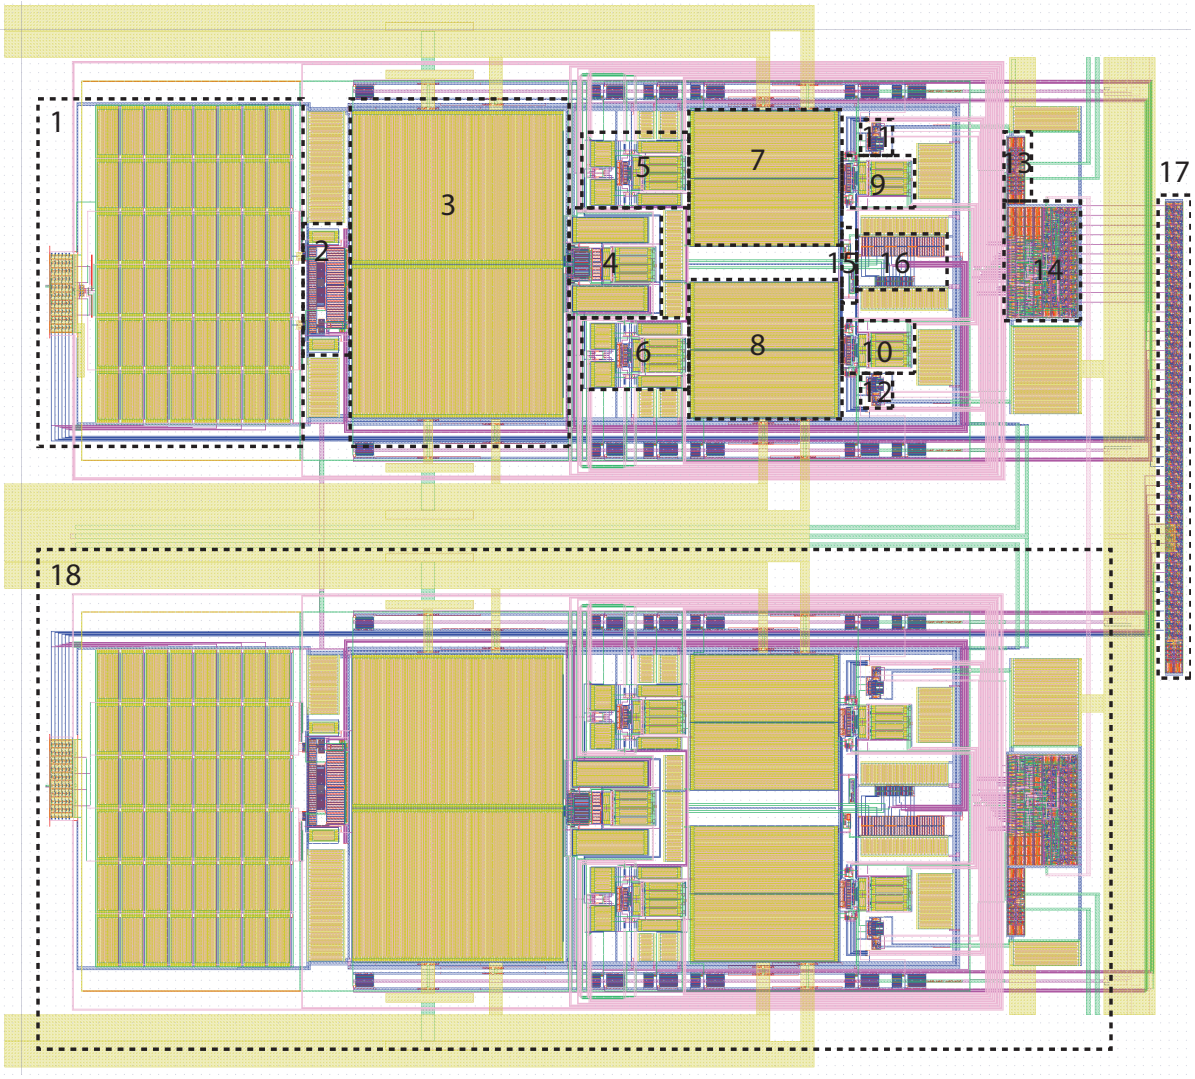

Figure 4.25: The layout of the chip.

5
Measurement Setup and Results

This chapter discusses the measurement setup and the measurement results.

5.1. Measurement Setup

The measurement setup was designed to measure 4 chips at a time. Each chip has 4 digital outputs and a programming interface. In order to process all the signals, an FPGA was placed on the PCB that communicates with a PC. To minimize the interference from this FPGA, its ground and power are isolated from the test chips. All digital lines from and to the FPGA pass through digital isolators. Furthermore, the test chip has a sensitive analog front-end, and high-speed digital electronics (50 MHz). Therefore, separate analog and digital power domains are used. The layout of the PCB, showing the main power domains, is shown in Figure 5.1.

5.1.1. Component selection

The digital supplies for the FPGA and the chip are generated using standard low-dropout voltage regulators. However, the analog supply of the test chip should have a low noise level. The chip also requires the common-mode voltages and the bias current to be generated externally. In addition, the test chip needs an accurate reference voltage which is used as the input to the redistribution circuit. The required voltages and currents with their specifications are summarized in Table 5.1. These specifications are derived from the simulated CMRR and PSRR and result in a temperature error of 0.01 K each. The noise level is over a bandwidth from about 1 Hz to 2 MHz.

For the analog supply voltage, an ADM1754-1.8 low-dropout voltage regulator is used. Using a bypass capacitor of 100 µF, the integrated noise from 1 Hz to 2 MHz is approximately 2.1 μV_{rms} including flicker noise.

The 0.9 V common mode voltage is derived from the analog supply voltage using a resistive divider. A buffer opamp would add noise and its bandwidth can not easily be limited by loading it with a large capacitor. Hence, no buffer is used, but the divider is designed to have less than 100 Ω output impedance. The noise of the resistors is

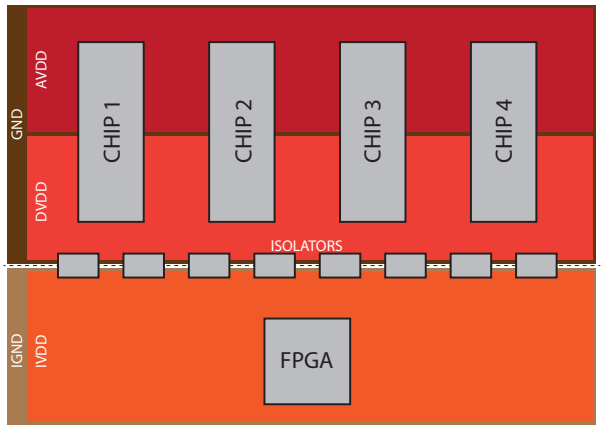

Figure 5.1: The layout of the test PCB.

	Nominal value	Noise level	Accuracy
Analog supply voltage	1.8 V	40 μV_{rms}	
Common mode voltage 1	0.6 V	8.0 μV_{rms}	
Common mode voltage 2	0.9 V	0.11 V_{rms}	
Reference voltage (w/ redistribution)	579 mV	2.4 mV_{rms}	9.7 µV
Reference voltage (w/o redistribution)	283 µV	1.2 μV_{rms}	4.7 nV
PTAT bias current	100 µA at 27 °C	9.7 nA_{rms}	

Table 5.1: Summary of the required voltages and currents with their specifications.

filtered by adding a 100 nF capacitor over the output. The resistor noise is then 0.2 μV_{rms}. Due to this capacitor also the noise contribution from the ADM1754-1.8 becomes band-limited to about 18 kHz.

The 0.6 V common mode voltage is derived from the LTC6655-3.0 precision reference using a resistive divider. Again the divider has an output impedance of less than 100 Ω. The resulting current is close to the maximum the LTC6655-3.0 can deliver. A 100 nF capacitor is used to filter the noise. The resistor noise is again 0.2 μV_{rms} and the noise contribution from the LTC6655-3.0 becomes band-limited to about 17 kHz. Also the noise is divided by the resistive divider, resulting in 2.0 μV_{rms} noise.

The reference voltage will be supplied by an external bench-top voltage supply. Since a high accuracy is required, a Kelvin connection is used. In case there is no redistribution, an even higher accuracy is required that the external supply cannot deliver. In that case the voltage is divided on the PCB by a MAX5492 precision-matched resistor-divider (1:11 divider). Its output voltage can be measured externally by the Agilent 34420A nanovolt meter (which has a 35 nV accuracy, or 0.075 K).

In case of no redistribution, the required noise level is again quite low. The resistor noise is limited to 0.2 μV_{rms} using a 100 nF capacitor. The output impedance of the resistor divider is about 0.8 $k\Omega$, resulting in an ENBW of 3.0 kHz. Because of the division, at the input a noise PSD of 0.24 $\mu V_{rms}/\sqrt{Hz}$ can be allowed. The Keithley 2400 sourcemeter specifies an output noise of 5 μV_{pp} in 0.1 Hz to 10 Hz, which is approximately 0.26 $\mu V_{rms}/\sqrt{Hz}$.

For all these voltages, the actual noise level will be even lower, since an additional 100 nF capacitor is placed in close proximity to all the chips for all these voltages. This ensures that high frequency current spikes are properly filtered and do not cause a voltage drop.

The PTAT bias current is generated using the LM134 temperature sensor with current output. The noise bandwidth is limited using a 100 nF capacitor to ground. Together with the input impedance of the chip of 1.54 $k\Omega$ (650 µS), the ENBW becomes 1.6 kHz. This results in about 1.2 nA_{rms} noise. This bias current source is shared between all 4 chips using an analog multiplexer (CD74HC4052). In this way it is guaranteed that all 4 test chips have the same bias current.

The test chip requires an input clock frequency of 50 MHz. This is generated using the ASEMPC-50MHz MEMS clock generator. A 50 Ω transmission line with series and parallel termination is used to avoid reflections and to lower the clock voltage from 3.3 V to about 1.8 V. The generated clock is also routed to the FPGA (via a digital isolator) and the FPGA can be used to generate a clock signal for the chips in case a different clock frequency is required for testing.

These 50 MHz signals pass through a high speed digital isolator (ISO7241M, max 150 Mbps). The other digital I/O from and to the test chip are at a frequency lower than 2 MHz and use a lower speed digital isolator (ISO7240C and ISO7241C, max 25 Mbps). These isolators require a supply voltage of at least 3.3 V and hence level shifters (SN74LV4T125) are used to communicate with the test chip at 1.8 V.

A schematic of the PCB can be found in Appendix D.

5.1.2. Overview

An overview of the measurement setup is shown in Figure 5.2. The PCB is connected to a PC running LabVIEW using a USB-UART-TTL cable. The LabVIEW program also controls the oven (using serial communication) and the instruments (via a USB-GPIB cable).

The analog and digital power are supplied by an Agilent E3631A. The reference voltage is supplied by a Yokogawa GS200. Furthermore, the on-board generated PTAT current is not used, and instead another Yokogawa GS200 is used to provide the bias current of the chips. This allows for a temperature-independent current to be used for the bias current.

The oven is used to sweep the temperature over a certain range. Since the temperature inside the oven will fluctuate, a large aluminium block is used to attenuate these fluctuations. A good thermal contact between the

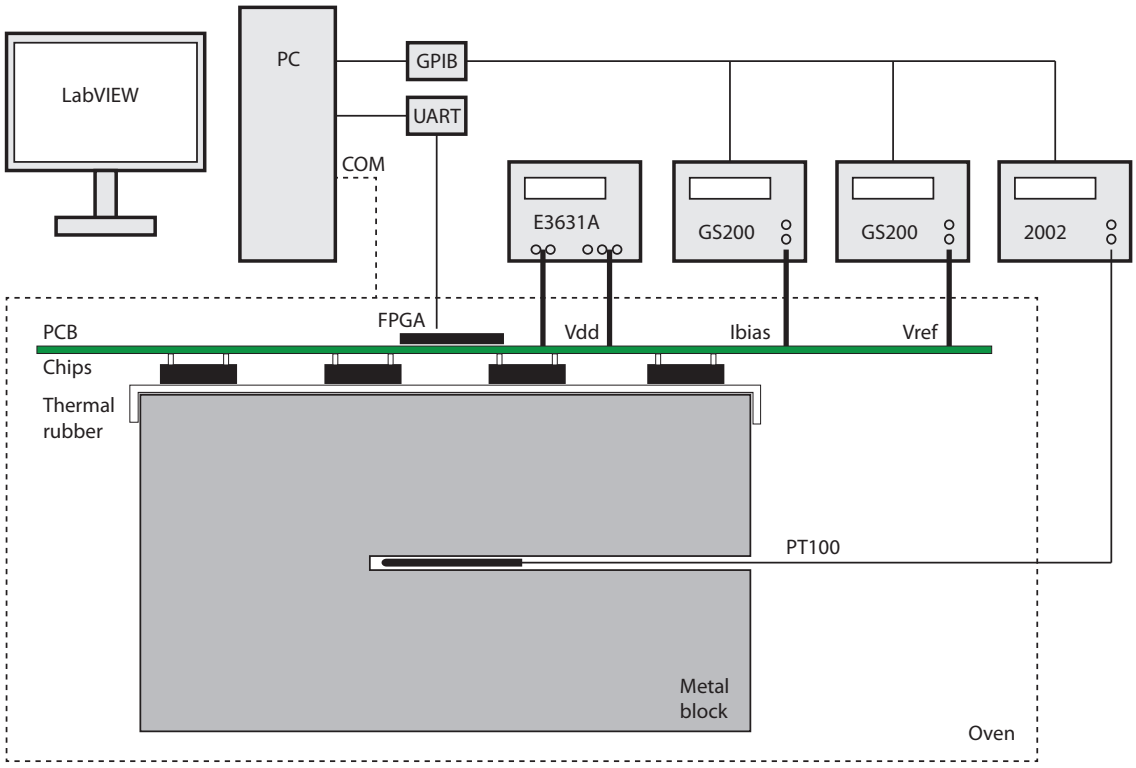

Figure 5.2: The measurement setup, the oven boundary is marked with the dotted line.

chips and the aluminium block is obtained using thermal rubber. A PT100 temperature sensor (connected to a Keithley 2002) inside this metal block is used as an accurate temperature reference.

5.2. Measurement Issues

During the measurements a few issues with the PCB and the chip were encountered. These issues will be discussed in this section.

5.2.1. Socket dependency

While testing the PCB, it was observed that the same chip would give different results depending on which socket was used to measure the chip. The observed change in bitstream average corresponds to a change in measured $\frac{V_t}{\sigma}$ of as large as 3% in some cases. This would correspond to a temperature error of about 20 °C.

The source of the socket dependency was not found, although it did seem to get worse with higher clock frequencies. The noise on the digital supply could be the cause of the socket dependency, see Section 5.2.3. In order to allow for a fair comparison of the chips, only a single socket was used to measure all the chips.

5.2.2. Power consumption discrepancy

According to the simulations, a single circuit inside the chip should draw about 940 µA from the analog supply for a bias current of 100 µA. A Monte-Carlo analysis with process and mismatch variations shows that the lowest expected current consumption is about 820 µA (see Figure 5.3). However, it was found that all chips draw even less current (see Figure 5.4).

Leakage of the external bias current through the ESD protection could not explain such a large difference. The cause of this discrepancy has not been found, so the chips have been tested at different bias current levels.

5.2.3. On-chip digital circuitry

During the measurements a design flaw in the on-chip clock generation circuit was found. The logic inverters/buffers driving the analog switches are powered from the digital supply. As a result, this noise is injected into the circuit. This problem was reduced by adding extra decoupling capacitors on the PCB close to the digital supply

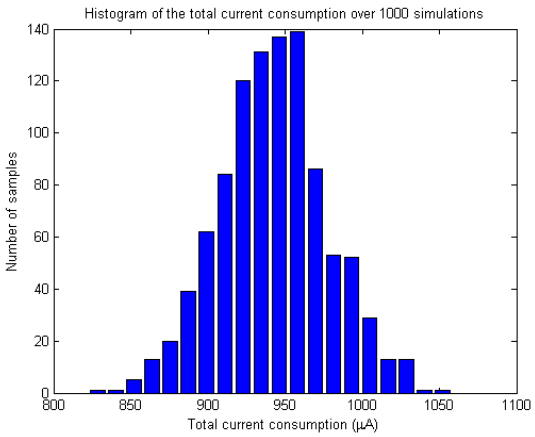

Figure 5.3: Simulated current consumption.

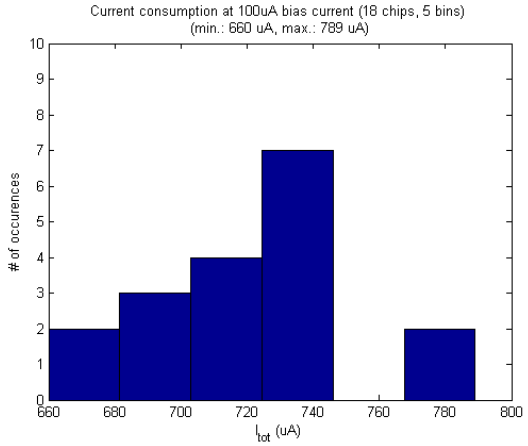

Figure 5.4: Measured current consumption.

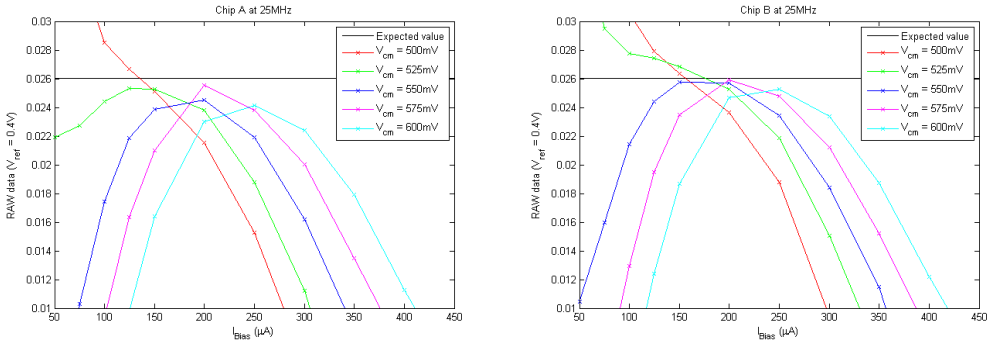

Figure 5.5: Bias current and common-mode voltage dependency.

pins.

Another design flaw is the absence of a reset signal for the on-chip clock generation circuit. As a result, the starting phase is undefined. If there would only be a single circuit on the chip, this would not have been a problem. However, two circuits, each with their own clock generator, have been integrated on the chip. Therefore, both circuit can start up in a different phase which could be different each time the chip is powered. Since there is some interference between the two circuits, it was observed that the output could change after reconnecting the digital supply. This problem has been solved by powering only one of the two on-chip circuits during measurements.

The design flaws described in this section are fixed in a new tape-out, see Appendix C for the details.

5.2.4. Bias dependency

During the measurements, a strong dependency on the bias current and the 0.6 V common-mode voltage were observed, see Figure 5.5. From these plots it can be seen that an optimum bias current exists for which the sensor's output is close to the expected value. Applying either a lower or a higher bias current results in a significant change in the sensor's output. The optimum bias current in turn depends on the common-mode voltage, i.e. for a lower common-mode voltage a lower bias current is optimal.

Measurements of the output voltage of the pre-amplifier with an oscilloscope have shown that the amplifier settling speed depends on this common-mode voltage. This can also be found in the simulations. A simulation of the phase margin frequency for different bias currents and common-mode voltages is shown in Figure 5.6. The decrease in phase margin frequency with higher common-mode voltages is due to the fact that the input transistor is no longer properly biased in these points.

The biasing problem can be explained with the help of Figure 5.7, which depicts part of the pre-amplifier. With an increase in common-mode voltage, the voltage at the source of the input transistor increases, while the drain voltage remains unchanged, since it is set by the fixed cascode voltage. As a result, the input transistor moves into the linear region and has a lower transconductance.

This is, however, not the only problem with the current biasing scheme. If on the other hand the common-

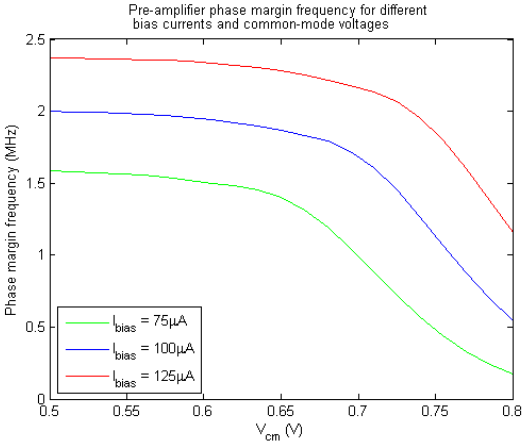

Figure 5.6: A simulation of the phase margin frequency for different bias currents and common-mode voltages.

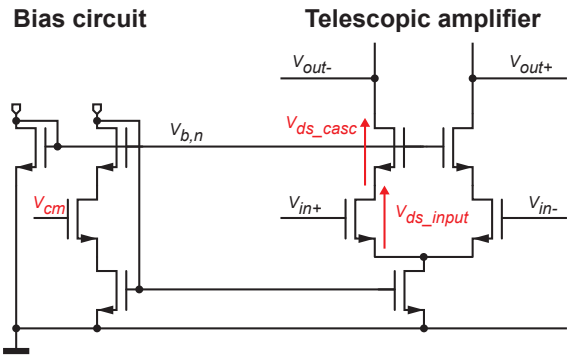

Figure 5.7: Part of the pre-amplifier.

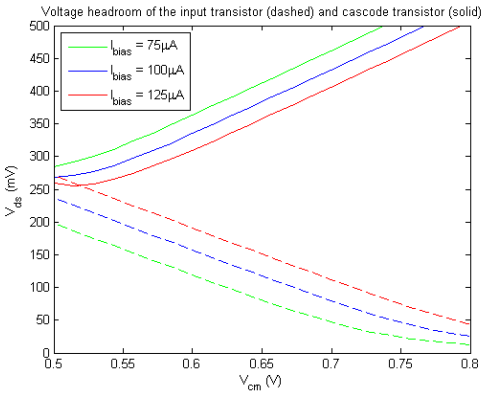

Figure 5.8: The simulated voltage headrooms.

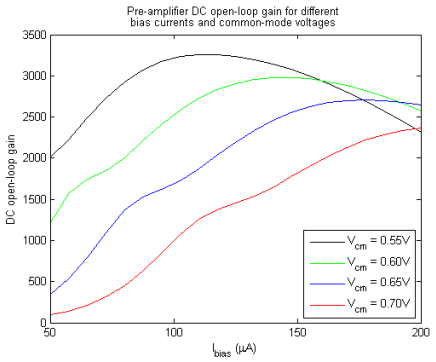

Figure 5.9: A simulation of the open-loop DC gain for different bias currents and common-mode voltages.

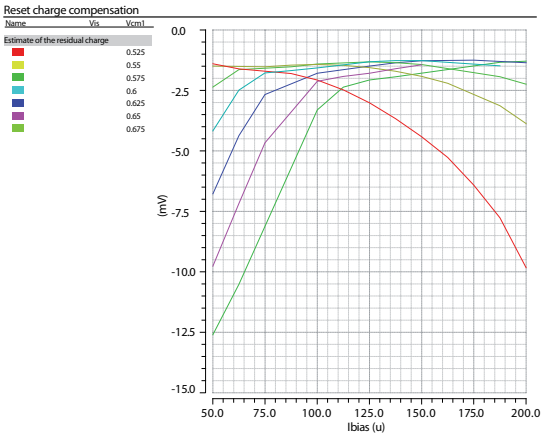

Figure 5.10: A simulation of the reset charge compensation for different bias currents and common-mode voltages.

mode voltage decreases, the drain-source voltage of the cascode transistor decreases and the DC-gain drops. Equivalently, the bias current influences the voltage headroom of the input transistor and the cascode transistor (see Figure 5.8). A simulation of the open-loop DC gain for different bias currents and common-mode voltages is shown in Figure 5.9. As a result, the reset charge compensation fails as is shown in Figure 5.10. This plot agrees with the measured bias current and common-mode voltage dependence, i.e. it shows the same parabolic dependence on the current and the location of the peak moves down with the common-mode voltage.

The measurements also show that the optimum bias conditions are slightly different from what was expected from the simulations. This is due to the fact that the minimum redistribution capacitor has almost twice the expected capacitance due to parasitics. As a result, the feedback of the pre-amplifier is different. Therefore, the loopgain is reduced which is compensated by the slight decrease in common-mode voltage and increase in current. However, also the phase margin frequency is about half of the expected value. Therefore, the clock frequency was halved to 25 MHz.

The biasing problems described in this section are fixed in the new tape-out, see Appendix C for the details.

Behaviour over temperature

For the oven measurements, the common-mode voltage was decreased from 0.60 V to 0.55 V, as this would require a lower bias current closer to the designed value, see Figure 5.5. For the bias current a PTAT characteristic should be used in order to keep the settling speed of the different stages more or less constant over temperature (see Section 4.5.2). With the oven measurements it was however found that the chip again experiences biasing problems (see Figure 5.11). From the measurements with different bias currents it was found that it is best to use a temperature independent bias current of around 100 µA or 125 µA (see Figure 5.11 and 5.12), since they show the most linear behaviour for the two tested chips.

Due to the biasing problem, the plot of the measured temperature versus actual temperature shows a curvature. This curvature can be removed using a master curve as a characteristic to translate the measured temperature back to the actual temperature. However, due to this biasing problem, the spread from chip to chip will be larger and as a result the accuracy will be worse than expected. The impact of the biasing problem on the accuracy will now be estimated.

From the measurements of Figure 5.12, the sensitivity to the bias current over temperature was estimated by determining the change in measured temperature when the bias current was increased from 100 µA to 125 µA. The results are plotted in Figure 5.13 and it can be seen that the sensitivity to the bias current is in the order of 1 °C/µA.

According to a Monte-Carlo simulation (1000 runs with process and mismatch variations), the current consumption of the pre-amplifier can spread by ±14.7% from chip to chip (±3σ). The resulting expected temperature spread from chip to chip due to the biasing problem is shown in Figure 5.14. Since anyway a one-point trim will be necessary, this will partially correct for the spread due to the biasing problem. The resulting spread after a one-point trim is also shown in the plot, and is about 8 °C (3σ).

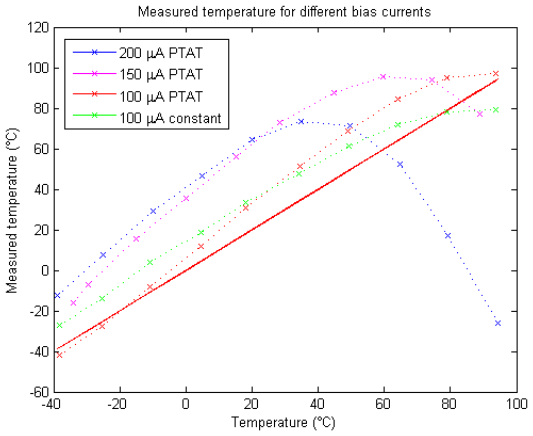

Figure 5.11: The use of a PTAT current versus a constant current. The red line is the ideal transfer.

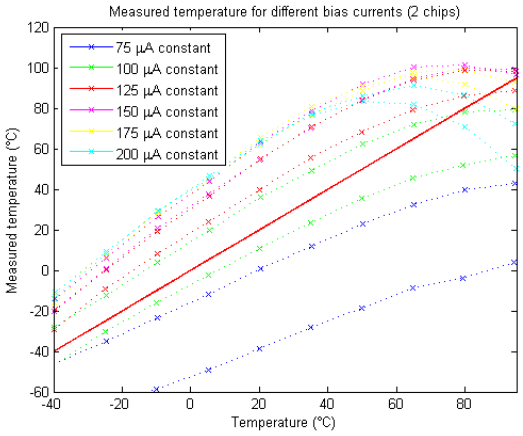

Figure 5.12: The use of different constant currents for 2 chips. The red line is the ideal transfer.

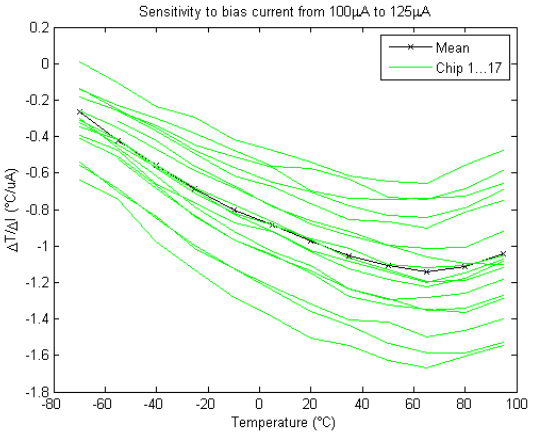

Figure 5.13: Estimate of the sensitivity to the bias current over temperature.

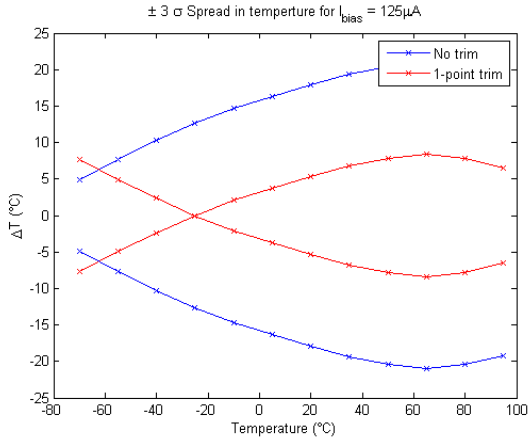

Figure 5.14: Estimate of the spread in temperature due to the biasing problem, with and without trim.

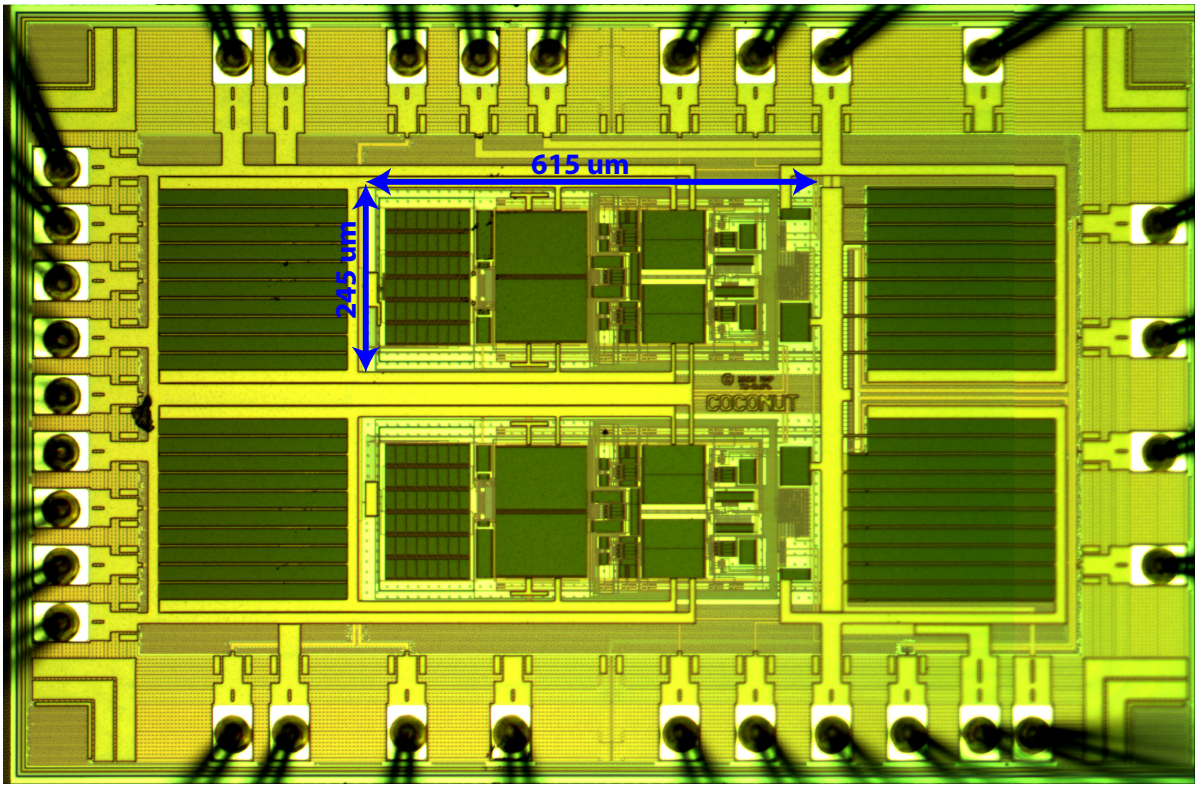

Figure 5.15: Chip micrograph with the dimensions of the sensor core.

5.3. Measurement Results

As discussed in the previous section, a few problems with the measurement setup and chips were encountered. To reduce the effect of these problems, the chips run at 25 MHz instead of 50 MHz with the common-mode voltage lowered from 0.60 V to 0.55 V. Also, both a constant bias current of 100 µA and 125 µA are used to test the chips, instead of a PTAT current with a nominal value of 100 µA. In total 17 chips were measured in the oven, and the results of these measurements are presented in this section.

The chips were fabricated in a standard 0.16 µm 1P5M CMOS process from SSMC. A micrograph of the chip is shown in Figure 5.15 with the dimensions of the temperature sensor core. It can be seen that the active area occupies 0.15 mm^2.

A measurement of the current consumption at a 1.8 V supply showed a 945 µA and 132 µA current from the analog and digital supply respectively. This was measured for 4 chips for a bias current of 125 µA while operating at 25 MHz. As a result the total power consumption of the chip is 1.9 mW.

5.3.1. Noise properties

The noise spectrum and the noise distribution of the system were analysed. For the noise spectrum, an FFT of the output bitstream was performed. The results are shown in Figure 5.16. This plot clearly shows that the noise is white and contains only a DC component, which is the threshold voltage.

To determine the noise distribution, the input reference voltage, and hence the internal threshold voltage, was swept and the average of the output bitstream was determined. The results are shown in Figure 5.17. This plot clearly shows that the noise has a normal distribution.

These two checks verify that white Gaussian noise is generated and measured using this system.

5.3.2. Resolution

The measurement resolution has been determined for different threshold voltages as an optimum is expected. As explained in Section 2.2.2 the resolution should follow:

$$\frac{\sigma_T}{T} = \frac{\sqrt{\alpha}}{f_s \cdot \tau_m} \tag{5.1}$$

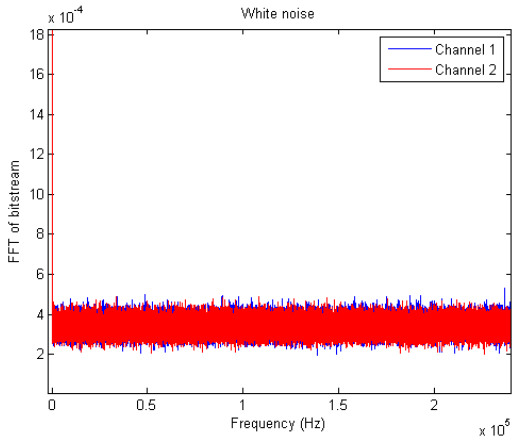

Figure 5.16: Measured noise spectrum.

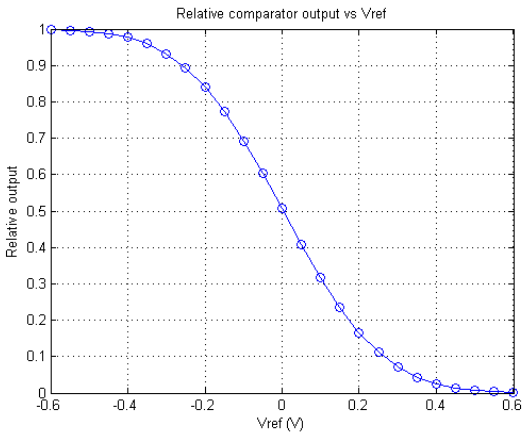

Figure 5.17: Measured noise cumulative distribution function.

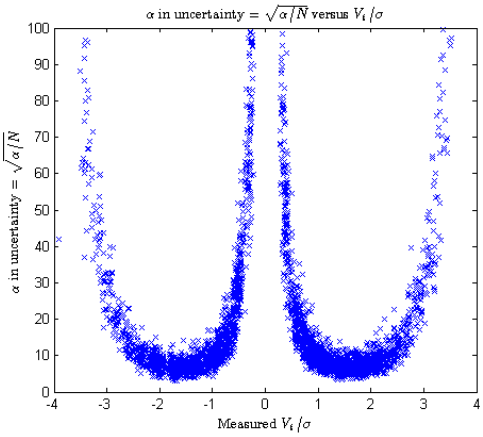

Figure 5.18: Resolution versus threshold voltage, many low resolution points.

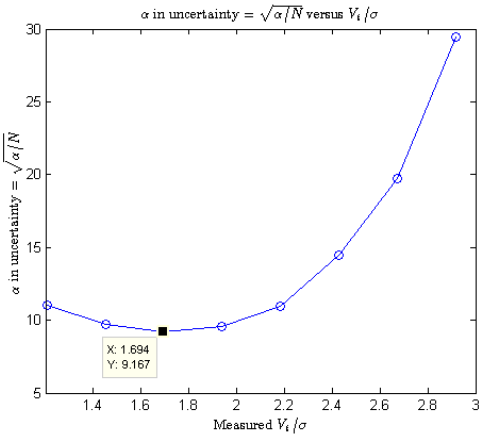

Figure 5.19: Resolution versus threshold voltage, a few high resolution points.

where α depends on the type of detector and the threshold voltage. The measured value of α has been plotted against the threshold voltage relative to the noise level in Figure 5.18 and Figure 5.19. As explained in Section 2.2.2, an optimum of $\alpha = 6.5$ is expected around $V_t/\sigma_n = 1.6$. The optimum found in the measurements is close to this optimum, but the resolution is a little bit worse.

For the optimum found in the measurements, the resolution is 0.926 K_{rms} at 300 K when measuring for 1 second at a clock frequency of 25 MHz (note that the internal sample rate is only 0.96 MHz).

It has also been verified that the resolution increases with the square root of the measurement time. These results are plotted in Figure 5.20 for up to almost 7 minutes. For this measurement not exactly the optimum threshold voltage has been used. From this plot it can be seen that there is no clear sign of flicker noise or drift, and as a result the resolution keeps improving even when measuring for several minutes.

5.3.3. Accuracy

This section will discuss the accuracy of the temperature sensor and the effect of the different compensation techniques on the sensor's accuracy.

Offset

The internal chopper and CDS cannot be disabled to see the effect of this offset compensation technique on the performance. However, the residual offset can be estimated using an off-chip system level chopper. This can easily be obtained with the current measurement setup by providing both a positive and negative reference voltage to the chip. The reported offset values are referred to the output of the redistribution circuit, or equivalently the input of the pre-amplifier.

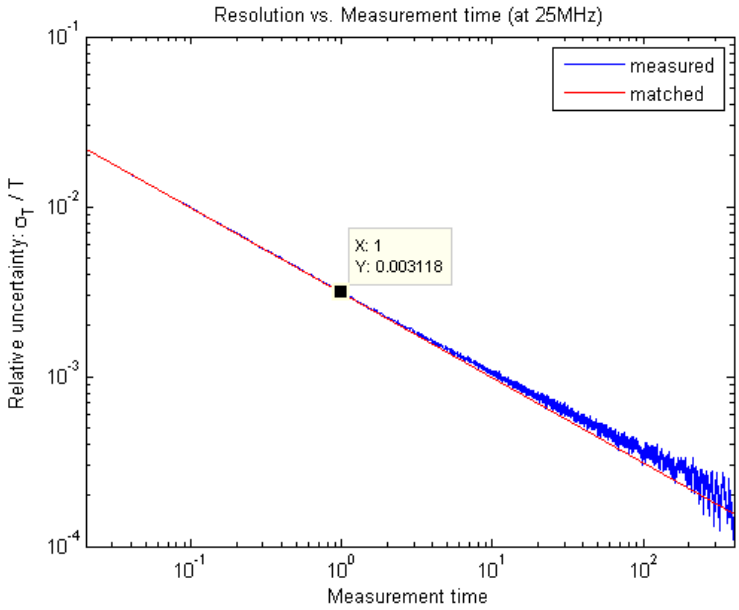

Figure 5.20: The temperature resolution versus measurement time.

Even in case the CDS is perfect, a residual offset is present due to the different chopping switches that are active in the different phases. Therefore, their residual charge injection mismatch is not compensated (see Section 4.1.3). Depending on the redistribution code, this charge can be injected on a different capacitor leading to a different offset. The capacitor that is initially charged is selected by the LSB (least significant bit) of the redistribution code and is therefore expected to be dominant in determining the offset. The residual offset over temperature for different redistribution codes is shown in Figure 5.21 for two different chips.

It can be seen that code 0 and 127 give the same offset, and code 1920 and 2047 give the same offset. The redistribution codes that give a similar offset have the same MSB, but a different LSB. So, either the MSB selects the capacitor that is initially charged, or somehow the capacitor that is discharged in the very last cycle influences the offset. Even though the order of bits in the labVIEW code, UART transmission, FPGA buffer and shift register programming has been checked carefully, it could be that the codes are somewhere reversed. In that case the MSB determines which capacitor is initially charged, and therefore has an influence on the residual offset.

From the plots it can be seen that this residual offset increases with temperature, which is in line with what is expected from charge injection mismatch. At higher temperatures, the threshold voltage decreases, V_{gt} is higher, and as a result more charge is injected. When selecting a larger capacitor for the redistribution, the capacitor size and the charging switch size are increased by the same factor. If the capacitor size would be doubled, the width of the charging switches is doubled, and the expected charge injection mismatch increases by a factor $\sqrt{2}$. The residual offset is therefore expected to be a factor $\sqrt{2}$ lower. In the measurements also a reduction of the residual offset is observed for a larger capacitor, see Figure 5.22.

Further analysis of the effect of different codes (see Section 5.3.3) reveals that it is unlikely that the code is reversed. A residual offset due to what capacitor is selected in the very last cycle is only possible if the offset is different in the two chopping phases such that the CDS cannot fully compensate the error. Since the voltage is different in the two chopping phases, a voltage dependent offset is likely the cause. This could again be caused by charge injection. Since the offset of previous cycles is divided, the last cycle is dominant in determining the offset.

Another possibility is that a large offset is introduced due to noise injection from the digital logic, see Section 5.2.3. The 'noise' on the digital supply is caused by, and hence synchronous with, the circuit. As a result, this noise could manifest as an offset. In case the noise on the digital supply is different during the two chopping phases, a residual offset will be present. The amount of clock feed-through depends on which capacitor is selected and the size of this capacitor. As this design flaw is fixed in the new tape-out (see Appendix C), it can be checked if this is the cause of the offset.

To verify the previous statements, the effects of the codes has been measured with a higher resolution and for multiple different codes for one of the chips. The codes measured can be found in Table 5.2 which contains both the binary and decimal representation. It can be seen that the LSB, MSB, MSB-1 and intermediate bits as a

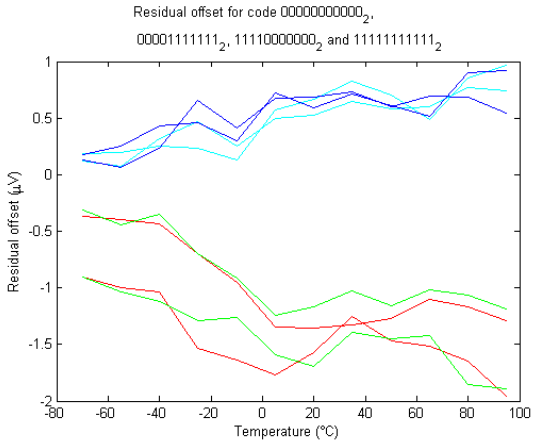

Figure 5.21: The residual offset for different redistribution codes for 2 chips: code 0 (red), code 127 (green), code 1920 (cyan), code 2047 (blue).

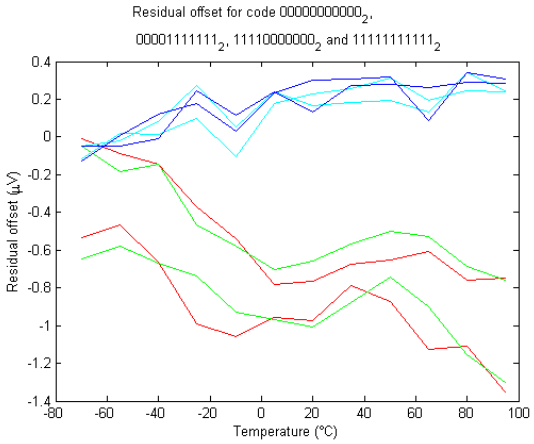

Figure 5.22: The residual offset for different redistribution codes for 2 chips and double the capacitor size: code 0 (red), code 127 (green), code 1920 (cyan), code 2047 (blue).

Table 5.2: The measured codes in decimal and binary representation.

Codes	
Binary	**Decimal**
0 0 00000000 0	0
0 0 00000000 1	1
0 0 11111111 0	510
0 0 11111111 1	511
0 1 00000000 0	512
0 1 00000000 1	513
0 1 11111111 0	1022
0 1 11111111 1	1023
1 0 00000000 0	1024
1 0 00000000 1	1025
1 0 11111111 0	1534
1 0 11111111 1	1535
1 1 00000000 0	1536
1 1 00000000 1	1537
1 1 11111111 0	2046
1 1 11111111 1	2047

Table 5.3: The effects of the LSB, MSB, MSB-1 and intermediate bits on the offset.

	single capacitor		**double capacitor**	
	mean(ΔV_{os})	**std(ΔV_{os})**	**mean(ΔV_{os})**	**std(ΔV_{os})**
Effect of LSB:	-147.7 nV	33.1 nV	-87.6 nV	41.0 nV
Effect of middle bits:	244.2 nV	36.2 nV	143.0 nV	44.8 nV
Effect of MSB-1:	351.2 nV	33.7 nV	164.4 nV	34.3 nV
Effect of MSB:	1501.6 nV	31.1 nV	872.4 nV	34.6 nV

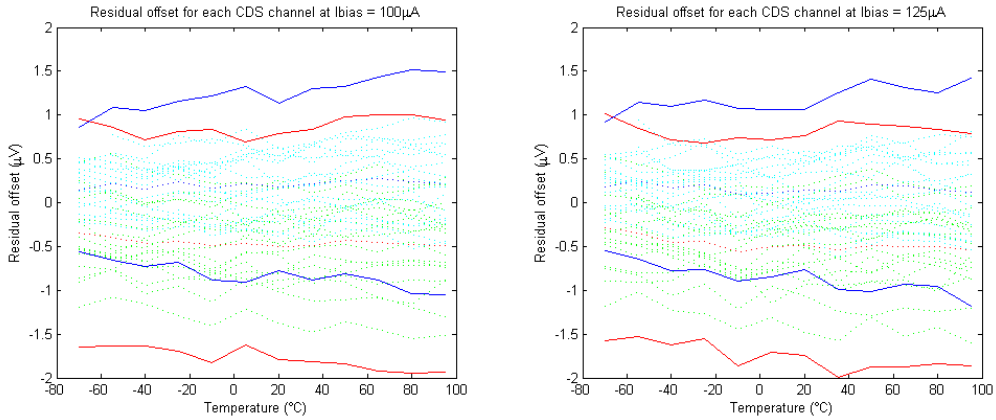

Figure 5.23: The residual offset for each CDS channel for all chips for two different bias currents. Channel 1 is shown in green with the mean (dotted) and 3σ boundaries (solid) in red. Channel 2 is shown in cyan with the mean (dotted) and 3σ boundaries (solid) in blue.

group have been varied. The effects on the offset are summarized in Table 5.3. From this table it can be seen that the LSB (selects the capacitor that is initially charged) has a small influence on the residual offset, which is due to the input chopper as explained before. The MSB (selects the capacitor that is discharged in the last cycle) on the other hand has the largest effect on the offset. Possible explanations for this effect have already been provided. It is however expected that the effect of the MSB-1 would be half of this value, since this offset is divided by 2 during the redistribution. From the table it can be seen that the effect is approximately 4 times lower. It can also be seen that when the capacitor size is doubled the sensitivities to the different bits is reduced.

Another source of offset is the imperfect offset compensation in the CDS. The residual offset for the two CDS channels has been determined over temperature. For this measurement the redistribution codes 127 and 1920 are used in order to minimize the effect of the specific redistribution codes on the offset. The results of the measurements are plotted in Figure 5.23 for two different bias currents.

As discussed in Section 4.3.1, a residual offset could be present due to insufficient settling. From the plots it can however be concluded that this source of error is not dominant, since an increase in bias current does not reduce the offset. It is however interesting to see that there is a fixed offset between the two channels. While the mean of the offset is positive for channel 1, it is negative for channel 2.

The fact that the mean of the offset is non-zero can have different causes related to layout differences for the positive and negative sides of the stages, which is then not perfectly differential. The different polarity for this fixed offset could be explained by considering that in the layout the second CDS channel was obtained by mirroring the first CDS channel instead of taking a direct copy.

In this section significant offsets were found (94 nV corresponds to 0.1 K error; Section 4.3.1), that depend on the selected redistribution code and channel. When considering the case that both channels are read and two redistribution codes are used that have a different MSB (e.g. 127 and 1920), the residual offset of all chips over temperature as shown in Figure 5.24 is obtained. It can be seen that the offset is around 1 µV (3σ), which corresponds to a temperature error of 1.1 K (3σ). The plot on the right is for a larger redistribution capacitor and the total offset is about the same. Therefore, the residual offset is dominated by the offset from the CDS which is more than 3 times larger than expected (see Section 4.3.6 and Section 4.4). However, as already pointed out the layout can probably be further improved to reduce this error.

For now anyway a 1-point trim will be required that also compensates for the temperature independent part of the offset. After a 1-point trim the temperature error due to the residual offset is 0.5 K (3σ).

Redistribution codes

The previous section already discussed the effects of the redistribution code on the offset, but the main reason for the use of different redistribution codes is to get the average threshold voltage as accurate as possible. A wrong threshold voltage will manifest as a gain error and the detected amount of noise changes. Since always both capacitors are read out, the generated amount of noise is fixed. Hence, any change in measured noise due to the redistribution code is actually a change in threshold voltage.

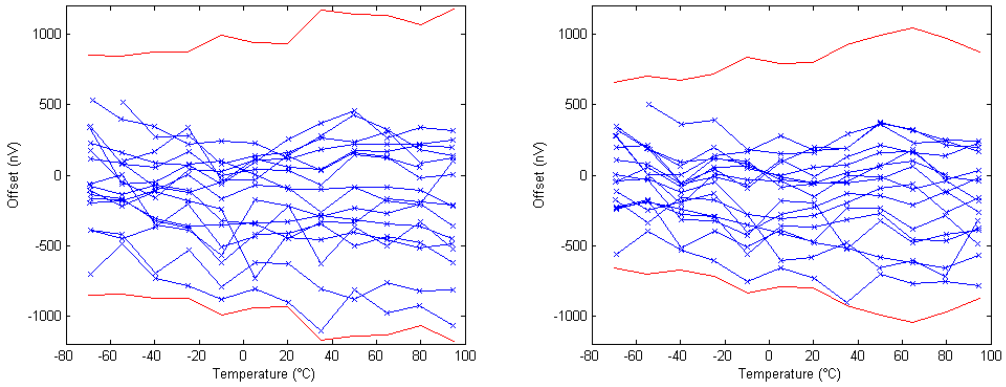

Figure 5.24: The residual offset of all chips over temperature with 3σ boundaries in red; left: single capacitor, right: double the capacitor size.

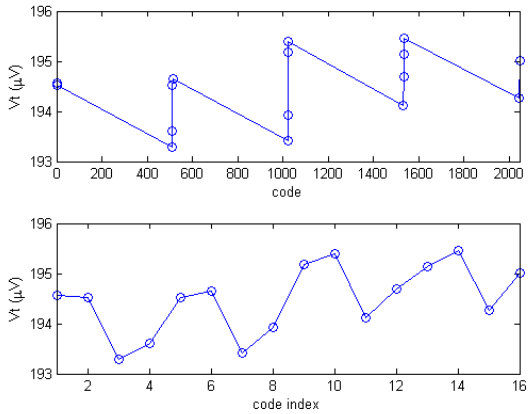

Figure 5.25: The measured threshold voltage for different redistribution codes for a single chip.

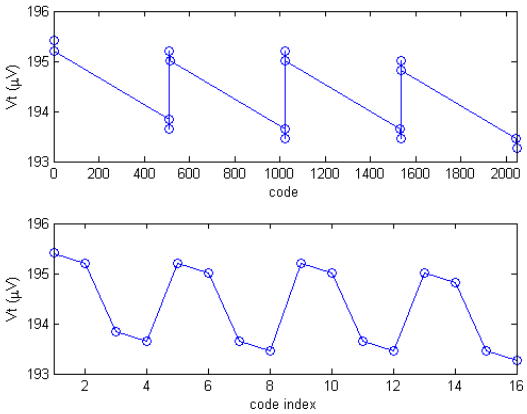

Figure 5.26: The simulated threshold voltage for different redistribution codes in case of 0.1% capacitor mismatch.

When again considering the chip for which all the codes of Table 5.2 are tested, the threshold voltage variation as depicted in Figure 5.25 is obtained. The results of the MATLAB simulation on the effect of these codes on the threshold voltage is shown in Figure 5.26. For this simulation, the mismatch was chosen such that the change in threshold voltage has about the same magnitude as was measured. The capacitor mismatch in that case is 0.1%, which is lower than the maximum expected mismatch of 0.6% for $\Delta C/C$ (3σ) (see Section 4.1.1).

The plots on the top show the threshold voltage versus the actual code, e.g. 512, while the plots on the bottom show the threshold voltage versus the measurement number (the points are equally spaced) in order to clearly see the differences. The features of the simulation and measurement results are similar. However, the simulation shows a generally decreasing trend in threshold voltage for higher codes, while the measurements show an increasing trend. As a result, codes 0 and 2047 that always discharge the same capacitor, C1 and C2 respectively, are not the worst as expected. One explanation is that the MATLAB simulation is very simple and only takes the capacitor mismatch into account while also other effects influence the threshold voltage, see for instance Section 4.1.4. Another explanation is that some residual offset could be present in these plots, even though these plots are obtained with the external system level chopping enabled.

For all the chips that were measured in the oven, only codes 127 and 1920 were used in the measurements. Since these codes are expected to give a different result in the presence of mismatch, the capacitor mismatch can be estimated. The resulting capacitor mismatch versus temperature is shown in Figure 5.27. Again on average there is a capacitor mismatch and hence the layout can be further optimized to remove this. Besides that, the observed mismatch is a little bit smaller than expected.

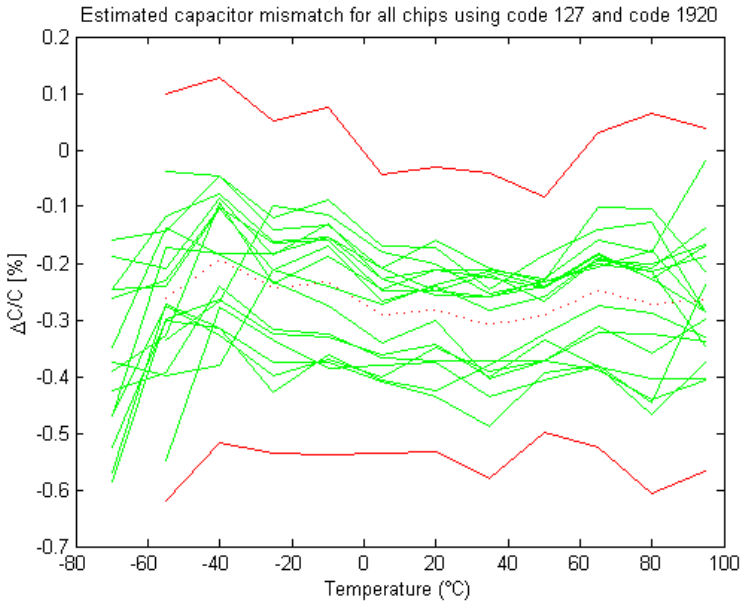

Figure 5.27: The measured capacitor mismatch versus temperature, with mean and 3σ lines in red.

Table 5.4: The size of the variable capacitor.

# Capacitors	1	2	3	4	5	6
Capacitance in the circuit (fF)	100	200	300	400	500	600
Capacitance extracted (fF)	231.2	338	444	552	658	764

Digital noise compensation

For the digital noise compensation, the capacitor size can be varied in order to change the amount of generated noise. The capacitor sizes are a bit larger than expected. The extracted capacitor sizes are summarized in Table 5.4 which shows increments of 106.6 fF instead of 100 fF and a larger starting capacitance, mostly due to wiring parasitics.

Furthermore, as discussed before (in Section 5.2.4), the quality of the reset charge compensation of the pre-amplifier will vary with the size of the source capacitance due to the change in settling speed. Due to this design flaw, the amount of readout noise will vary with the size of the source capacitance, and as a result the simple digital noise compensation as proposed in Section 3.2.3 is likely not possible.

The amount of noise measured for different capacitor sizes has been plotted in Figure 5.28 for all chips at three different temperatures, along with the expected amount of generated noise (kT/C). Since the amount of noise is plotted against 1/C, roughly a straight line is expected, which is the case when the extracted capacitances are used. Since in addition to the generated noise readout noise is present, the measured noise level should be higher than the generated noise as depicted by the solid black lines. This is clearly not always the case due to the biasing problem, and therefore digital noise calibration cannot be applied effectively.

As can be seen from the plot, some sort of digital noise calibration might be possible in case of the 125 µA bias current and the two smallest capacitors, since there the measured amount of noise is generally larger than the generated amount of noise. The expected amount of noise for the two smallest capacitors is:

$$\begin{aligned} v_{n,C1}^2 &= \frac{kT}{C1} + v_{n,readout}^2 \cdot c \\ v_{n,C2}^2 &= \frac{kT}{C2} + v_{n,readout}^2 \end{aligned} \tag{5.2}$$

where the factor c has been introduced, because the readout noise will depend on the size of the source capacitance. Ideally, $C = 1$ for the digital noise calibration, i.e. the readout noise is independent of the source capacitor. Mathematical solutions for this set of equations can be obtained when assuming known values for C1 (single capacitor) and C2 (double capacitor). When taking the extracted values for C1 and C2, the values for $v_{n,readout}^2$

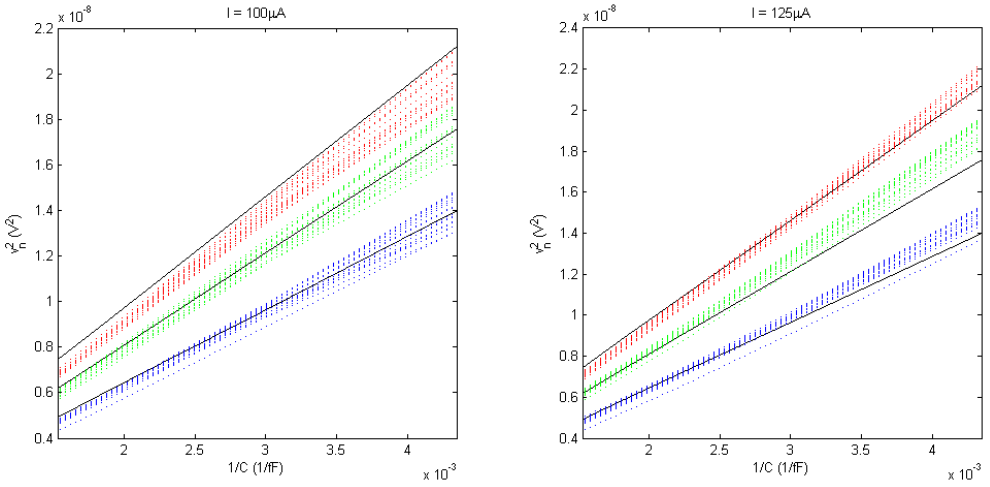

Figure 5.28: The amount of noise measured for all chips versus 1/C with the expected amount of generated noise in black. The blue curves are at -40 °C, the green curves are at 20 °C and the red curves are at 80 °C.

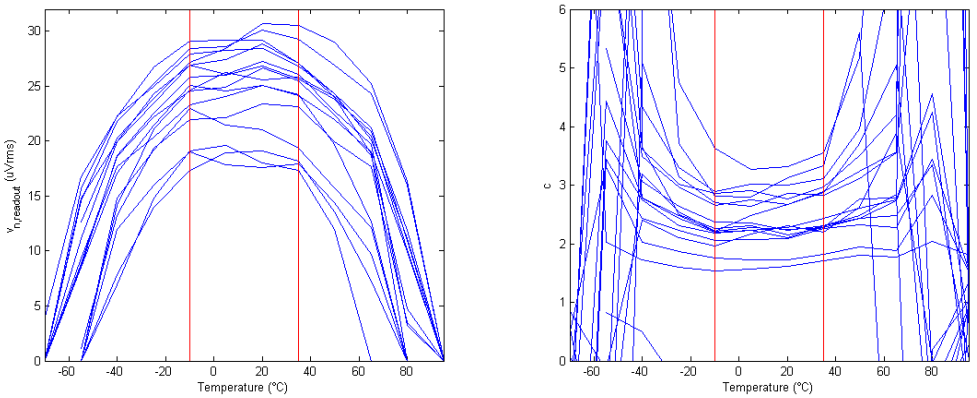

Figure 5.29: The estimates of the readout noise and its dependence on the source capacitance for all chips. Outside the boundaries in red, the amount of noise from the readout circuits starts to deviate a lot from the expected value, or even becomes negative.

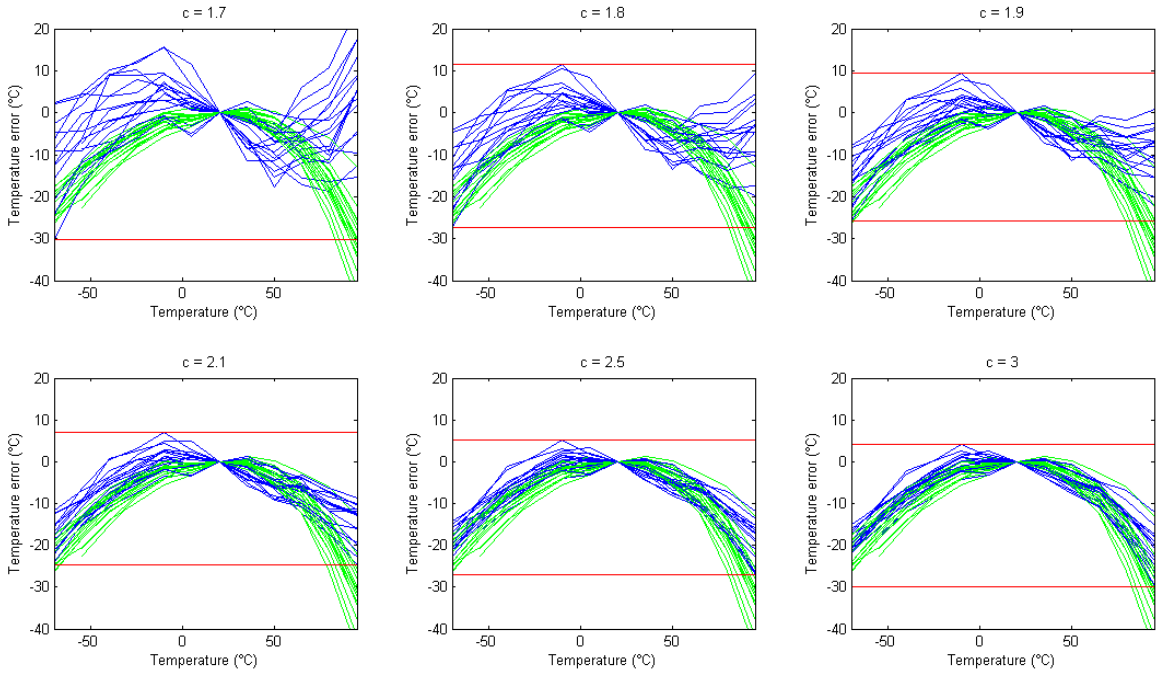

Figure 5.30: The temperature error after 1-point trim when performing digital noise calibration with a fixed value of the factor c. Green: without noise compensation, blue: with noise compensation, red: upper and lower limits of the temperature error.

and the factor c as shown in Figure 5.29 are obtained. It can be seen that this only works reasonably well for temperatures around room temperature, where c spreads from about 2 to 3, and the readout noise is close to the expected value. Also, in some cases the readout noise power is negative which is of course not possible. This is due to the fact that spread in the values of C1 and C2 are ignored. When on the other hand considering a fixed factor c, the temperature can be found from:

$$T = \frac{v_{n,C1}^2 - c \cdot v_{n,C2}^2}{k \cdot (\frac{1}{C1} - \frac{c}{C2})} \tag{5.3}$$

The factor c can then be optimized for the least amount of spread and curvature in the temperature error after a 1-point trim. Since the digital noise compensation might be able to compensate for the non-linearity, this was not removed using a master curve when optimizing for c. These results are shown in Figure 5.30 and it can be seen that $c = 2.1$ gives the smallest spread and curvature. This value is close to what was found using the previous approach.

Comparing the amount of readout noise estimated from the measurements with what was simulated, it can be seen that the measurement results are not far off. From the simulations a 3.0 K error due to readout noise was expected (see Section 4.2.5). This was for a source capacitance of 100 fF, and therefore corresponds to 20 μV_{rms} of readout noise. Considering that the chip has biasing problems and a different bias current is used, the estimates of Figure 5.29 can be reasonable.

5.3.4. Results

The previous sections discussed the effects of the CDS and redistribution codes on the offset and the amount of noise. Also the possibility of using digital noise calibration was shortly discussed. This section will discuss the actual temperature accuracy that can be obtained when considering these different techniques.

Due to the biasing problem, the plot of the measured temperature versus actual temperature shows a curvature. This curvature will be removed using a master curve as a characteristic to translate the measured temperature back to the actual temperature. This master curve will be generated by averaging the polynomial fits generated for each chip. Since the noise generating capacitor will spread from chip to chip and its value cannot be measured directly, a 1-point trim will be applied.

Table 5.5: The 3σ spread in temperature for different settings and temperature ranges.

CHOPPER	# CAPS	CODE	100 µA		125 µA	
			-70 °C...50 °C trim at -10 °C	-70 °C...95 °C trim at 20 °C	-70 °C...50 °C trim at 5 °C	-70 °C...95 °C trim at 65 °C
yes	1	127	12.59 °C	17.16 °C	7.48 °C	12.05 °C
		1920	11.85 °C	15.25 °C	7.06 °C	12.38 °C
		127 & 1920	12.16 °C	15.91 °C	7.04 °C	11.83 °C
	2	127	7.96 °C	11.59 °C	5.49 °C	9.59 °C
		1920	8.00 °C	10.83 °C	5.15 °C	9.48 °C
		127 & 1920	7.69 °C	11.03 °C	5.07 °C	9.29 °C
no	1	127	13.98 °C	18.64 °C	8.66 °C	12.66 °C
		1920	12.61 °C	15.86 °C	7.69 °C	14.31 °C
		127 & 1920	12.63 °C	16.06 °C	7.58 °C	12.06 °C
	2	127	10.21 °C	11.96 °C	7.74 °C	10.81 °C
		1920	10.31 °C	11.95 °C	6.92 °C	10.68 °C
		127 & 1920	8.35 °C	11.33 °C	6.75 °C	9.76 °C

Since many settings will be evaluated, not all master curves and plots of the temperature error versus temperature after a 1-point trim will be shown. Instead, the values for the 3σ spread in temperature are summarized in a table, see Table 5.5. For the master-curve, a 5-th order polynomial is used. Since the temperature sensor behaves especially bad at high temperatures, the performance is also evaluated over a smaller temperature range. The temperature at which the 1-point trim is performed differs for the different temperature ranges and bias currents and is also included in the table. The inaccuracies are quite high, but since already 8 °C spread was expected due to the biasing problem these large inaccuracies are not surprising.

The performance is evaluated with and without external system level chopper (denoted by 'CHOPPER' in the table) to see the effect of offset on the accuracy. It can be seen that without the chopper the performance improves quite a bit when both codes are used instead of only a single code. This is expected, since the code has a large influence on the offset. Enabling the external chopper when both codes are used does not improve the accuracy a lot. This is because the error due to offset is estimated to be around 0.5 K (see Section 5.3.3), while it can be seen from the table that the total accuracy is a lot worse.

From the table it can also be seen that the temperature sensor is more accurate for a 125 µA bias current than for a 100 µA bias current. The accuracy also improves when a larger source capacitance is used. With a larger source capacitance, the amount of generated noise decreases and less readout noise and offset can be tolerated. However, these errors are largely removed using the master curve and single point trim. With a larger source capacitance the spread is however smaller, since the parasitics are less dominant.

The temperature sensor performs best for a 125 µA bias current and with two unit capacitors used in the redistribution. For the final performance specification the external chopper cannot be used. The master curve and temperature error after no trim and 1-point trim are shown in Figure 5.31. The accuracy achieved using a 1-point trim over the reduced range can also be achieved over the full range from -70 °C to 95 °C if a 2-point trim is used, see Figure 5.32.

As discussed before, digital noise calibration can possibly be applied for this bias current by using one and two unit capacitors in the redistribution. Using Equation 5.3 with a fixed value for the factor c of 2.1, the master curve and temperature error after no trim and 1-point trim as shown in Figure 5.33 are obtained. It can be seen that this form of digital noise calibration does not improve the accuracy of the temperature sensor.

The performance of the temperature sensor is summarized in Table 5.6 which also contains a comparison of this work to other state-of-the-art temperature sensors and noise thermometers.

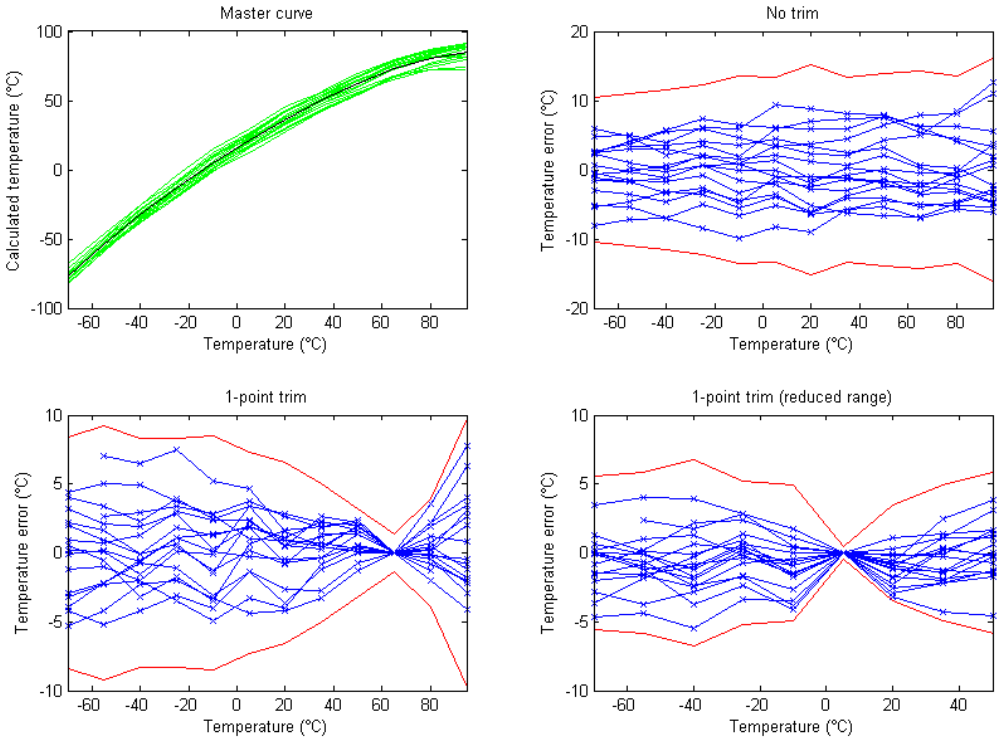

Figure 5.31: The master curve and temperature error after no trim and 1-point trim, for a 125 µA bias current and two capacitors.

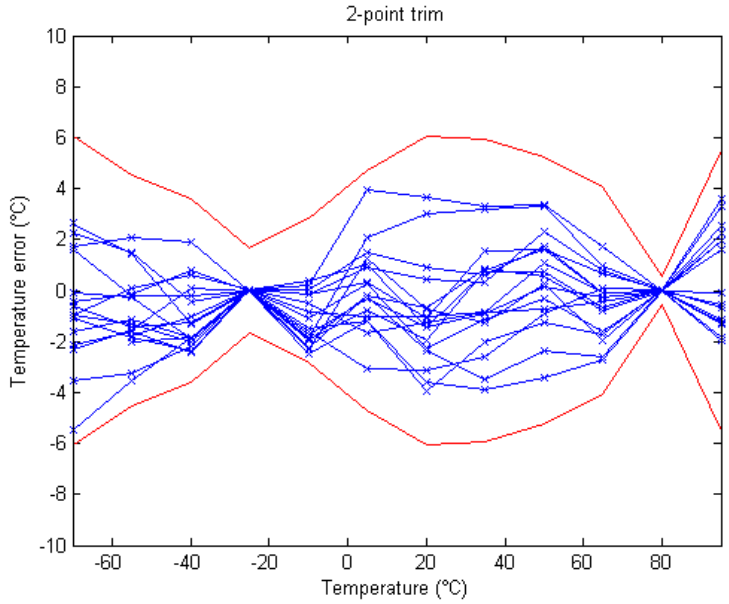

Figure 5.32: The temperature error after a 2-point trim, for a 125 µA bias current and two capacitors.

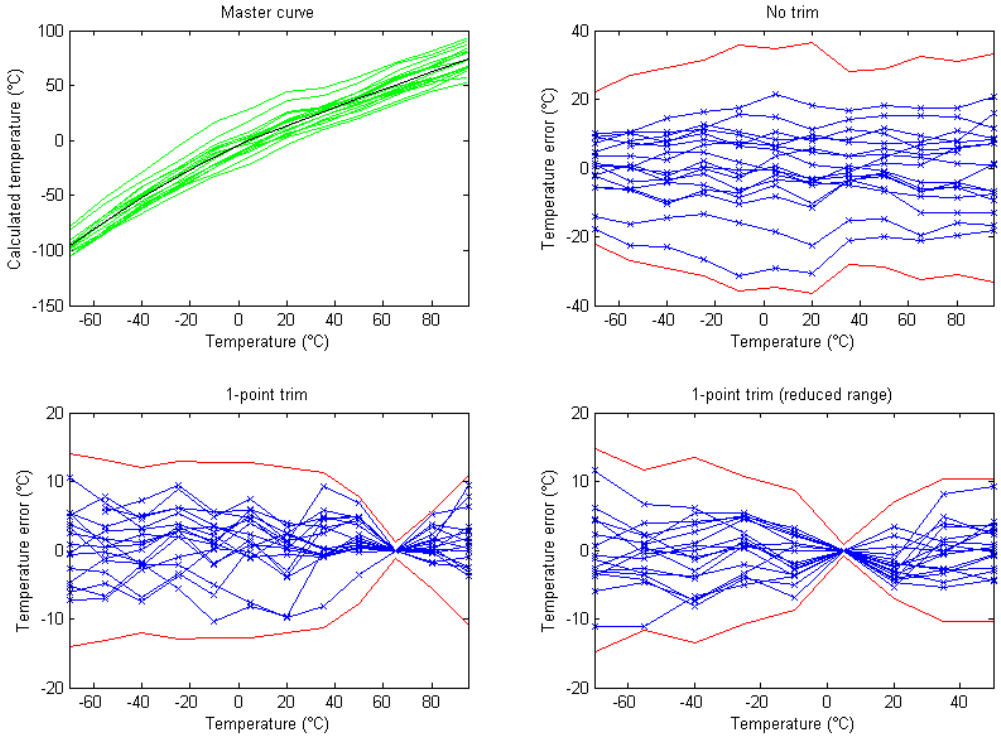

Figure 5.33: The master curve and temperature error after no trim and 1-point trim, for a 125 µA bias current after digital noise calibration.

Table 5.6: Performance comparison table: * Worst deviation from a straight line for a single sample, ** 3-point trim.

	This work		[18]	[2]	[7]	[8]
Technology	160 nm		*discrete*	160 nm	180 nm	160 nm
Sensor Type	Noise		Noise	BJT	RES	TD
Area (mm^2)	0.15		-	0.08	0.35	0.008
Inaccuracy (untrimmed; 3σ) (°C)	15.2	16.1	-	0.6	-	2.4
Inaccuracy (single temp. trim; 3σ) (°C)	6.8	9.8	≈2.5*	0.15	-	0.65
Inaccuracy (two temp. trim; 3σ) (°C)	-	6.1	-	-	0.15**	-
Temperature Range (°C)	-70...50	-70...95	0...1027	-55...125	-40...85	-40...125
Resolution (RMS) (°C)	0.93		0.4	0.02	0.006	0.21
Speed (S/s)	1		0.1	190	10	1000
Power (mW)	1.9		-	0.006	0.036	3.1

6

Conclusion and Future Work

6.1. Conclusion

This thesis describes the work performed to obtain the first fabricated on-chip noise-based temperature sensor.

First, a thorough analysis of different detectors has been presented. This analysis reveals that the temperature uncertainty that can be obtained with the different detectors is about the same. For the first time an averaging threshold detector has been used for the measurement of noise power in a noise-thermometer, since this detector is the easiest to implement and calibrate for a wide bandwidth.

Based on the threshold detector, a system level design has been obtained. It has been shown that it is possible to design the system such that it can be electrically calibrated and only requires a reference voltage during operating. Such a system has however three important requirements:

- *The external reference voltage must be divided with sufficient accuracy.* This thesis demonstrates the redistribution circuit which is capable of generating both the noise and the threshold voltage. As a result, the following comparator can have an unknown speed and gain, and almost any non-linearity.
- *The offset of the readout circuit must be sufficiently low.* This thesis uses a standard chopper with a CDS as demodulator. It has been demonstrated that in this application, due to the specific properties of noise and the quantity to be measured, it is possible to implement this without halving the sample rate.
- *The noise of the readout circuit must be sufficiently low.* This thesis demonstrates that the noise of various charge readout circuits is limited by the technology. However, this is expected to improve rapidly for more advanced technologies with shorter transistors. In addition, a simple digital noise compensation scheme has been proposed to compensate for the residual readout noise.

The circuit has been designed and although not implemented, the designed circuit allows for an electrical calibration of the chip by measurement of a single capacitor. The resulting chip has been fabricated, and the measurement results of the prototype chip have been presented.

The measurement results show that the inaccuracy of the sensor is worse than was expected from the simulations. The problem has been identified and should be fixed in the new tape-out. The measured temperature resolution agrees well with the simulated values. Although the measurement time is currently quite long, this will also improve for more advanced technologies that allow for higher measurement bandwidths.

6.2. Future work

There is still a lot of work required before the on-chip noise-thermometer can be considered a true competitor. The next steps include:

- *Measurement of the new chip.* During the measurements several problems in the chip were encountered. A new chip is being fabricated and the first step is to compare the performance of this chip with the simulation results. Because of the problems, the effectiveness of the digital noise calibration could not be evaluated yet. Therefore the achievable accuracy of this circuit still needs to be measured.

- *Improve the sample rate with die area.* It should be possible to reduce the area of the sensor by selectively employing denser capacitors. If the area is small enough, multiple circuits can be integrated to further improve the resolution by increasing the number of noise samples.
- *Break the g_m/C_{in} limit.* Positive feedback can be employed to reduce the input capacitance of an amplifier. It should be further investigated if this can be used to remove the g_m/C_{in} limitation of the proposed charge readout circuit.
- *Demonstrate electrical calibration.* The additional advantages of a noise-based temperature sensor have not yet been demonstrated. For an electrical calibration, some capacitor measurement circuit should be integrated on the chip. Besides, the requirement for an external reference voltage should be removed by using shot noise in conjunction with thermal noise.
- *Optimize for thermal calibration.* If the temperature sensor is to be used as a relative temperature sensor, or if it is allowed to have a one-point trim in a temperature stabilized environment, the design can be highly simplified. In that case the division ratio of the redistribution circuit does not need to be accurate and the specifications on the noise and offset of the readout circuit are relaxed. So, at the cost of electrical calibration, it should be possible to highly decrease the area and power consumption.
- *Port to more advanced technologies.* The expected improvement in accuracy and resolution with more advanced technologies should be confirmed by designing a noise-based temperature sensor in a more advanced technology. As the matching is expected to improve, more noise can be generated in the redistribution circuit which relaxes the requirements on the circuit. Besides, already in a 65 nm process the technological limit on the readout noise is expected to be about 10 times lower, which might remove the need for digital noise calibration.
- *Extend the temperature range.* It should be investigated how applicable this sensing principle is in other temperature ranges. Based on the results, a general design strategy can be proposed to adopt this sensing principle in different temperature ranges, for different applications.

A

Pre-amplifier models

Modelling the amplifier connect switch

This section describes the models that were used to compare the architecture where the amplifier is always connected to the source capacitor and the architecture where the amplifier is connected by a switch, see Figure A.1. In order to perform a fair comparison, the main parasitics that affect the redistribution accuracy (C_{sw} and C_{in}) are estimated and the amount of parasitics that can be tolerated is estimated (from the required accuracy of the redistribution circuit).

In case the amplifier is always connected to the source, the transistor gate capacitance (C_{in}) is the main parasitic, while in the architecture with switch it is the switch parasitics (C_{sw}). Transistor simulations have shown that in the selected operating point the switch parasitics are more voltage dependent than the gate capacitance. Therefore, a smaller switch parasitic can be allowed (about half).

As explained before, the maximum transistor gate capacitance limits its transconductance (g_m) and therefore its noise. In this model the technology limit of 386 µS/fF is used, but also an operating point with half this maximum value is evaluated. This is also of interest, since the required V_{gs} and I_d are much lower.

A similar relation holds for the connect switch where the switch cannot be too wide as the parasitics would become too big. This limits the on-resistance (R_{on}) which contributes to the noise of the readout circuit. For an NMOS switch, the on-resistance to parasitic ratio (around V_{gs} = 1.5 V) is approximately 1250 µS/fF, while for a PMOS this is only 250 µS/fF.

In the comparison, the energy consumption is fixed by using a fixed transconductance for the input transistor and a fixed noise bandwidth (an integration time of 500 ns). For the noise of the amplifier, γ (in $v_n^2 = 4\gamma kT/g_m$) is increased from the usual 2/3 to 2 to include the noise of the rest of the amplifier, which has been assumed ideal upto this point (see Figure 3.11).

In the simulation of this model, all possible switch on-resistances and source capacitor sizes are considered. Only the results that have small enough parasitics for the required redistribution accuracy are used. Among these, the results with the best accuracy are shown in the plot for each source capacitor size. In case the amplifier is always connected, the optimum is of course at the minimum required extra source capacitance, so only a single

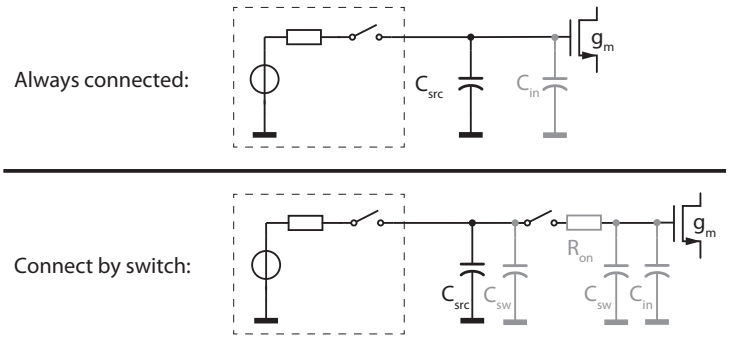

Figure A.1: Different possible readout connections. The redistribution circuit is modelled as shown inside the dotted box, and the parasitics are shown in gray.

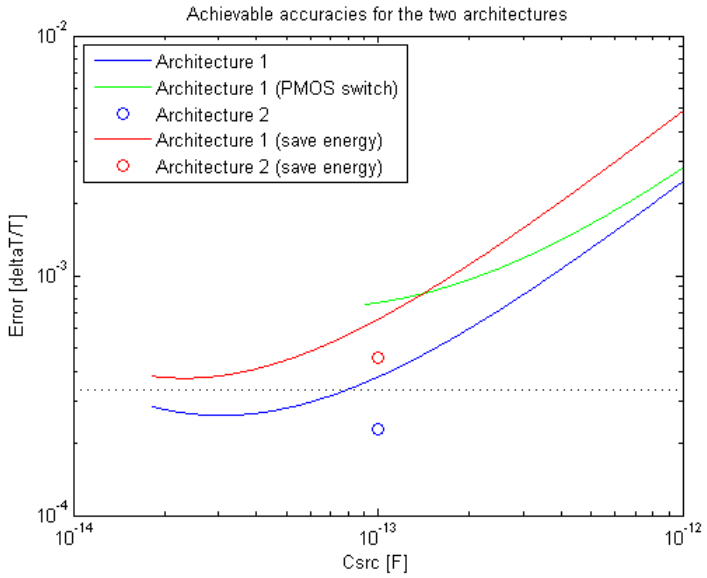

Figure A.2: The achievable accuracies when the amplifier is connected by a switch compared to when the amplifier is always connected. The red entries with the description 'save energy' denote the case where half of the maximum g_m/C_{in} is used.

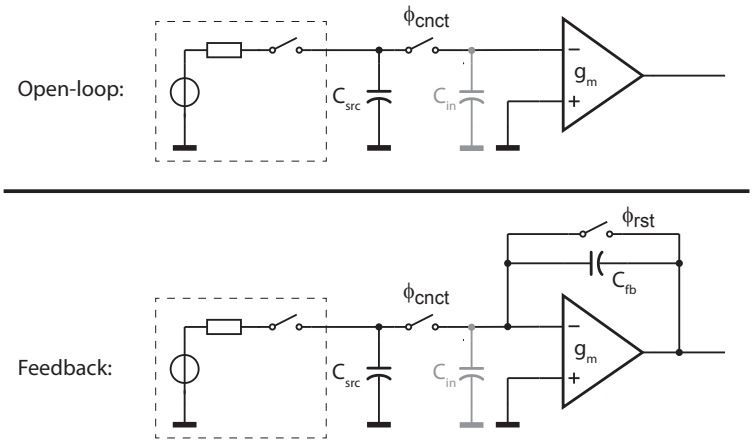

Figure A.3: Open or closed-loop readout. The redistribution circuit is modelled as shown inside the dotted box, and the parasitics are shown in gray.

point is calculated. The results are shown in Figure A.2.

It can be seen that a PMOS switch performs a lot worse than an NMOS switch. With an NMOS switch, the input voltage should be kept low to keep the design simple (no clock boosting etc.). The options denoted by 'save energy' are easier to implement, since they require a low V_{gs}. At this point the achievable accuracy of both architectures is about the same.

Modelling the amplifier feedback

In this section another model is presented that compares the open-loop architecture with the feedback architecture, see Figure A.3. Since both architectures require a switch to connect the amplifier to the redistribution circuit, the switch parasitics and switch noise are ignored in this model.

Both the self-heating and the flicker noise corner could present a serious issue. Hence, these effects are included in this model, using simulated values of g_m, V_{gs}, f_{knee} and C_{in} (see Figure A.4).

The flicker noise corner imposes a limit on the lowest ENBW that can be used to filter the noise. The ENBW in turn determines the speed of the circuit, and hence the maximum sample rate. This sample rate determines the best resolution that can be achieved in a limited amount of time. In this regard, the closed-loop architecture might have an advantage, as larger transistors can be used, since the effect of the input capacitance is reduced by the negative feedback.

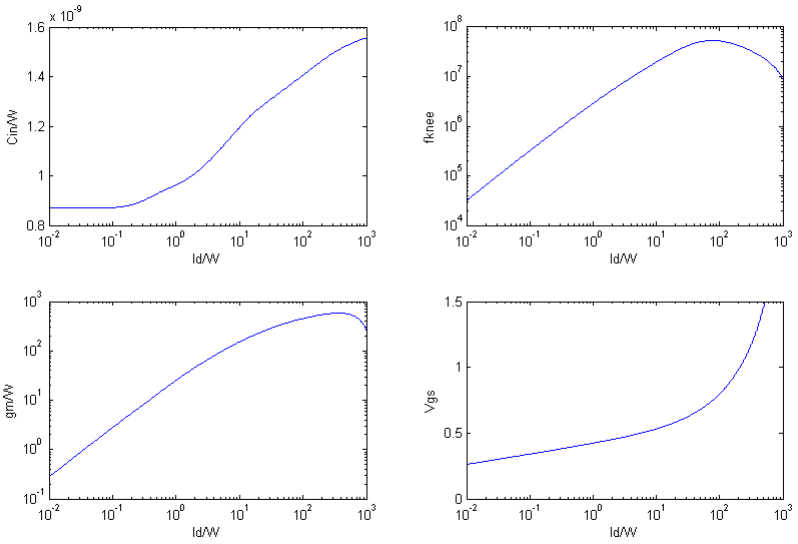

Figure A.4: The simulated behaviour of g_m, V_{gs}, f_{knee} and C_{in} that is used in the models.

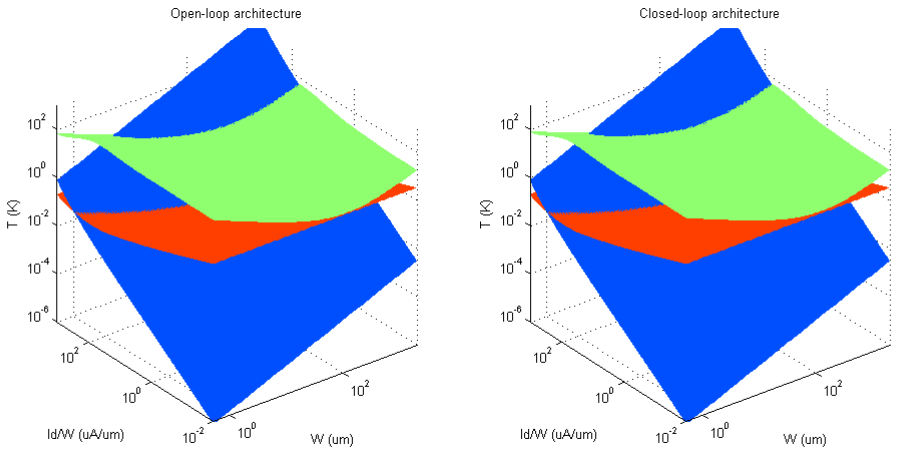

Figure A.5: The modelled error due to self-heating (blue) and noise (green), and the resolution (red) for the two architectures for different bias current densities and transistor widths.

The self-heating is modelled by estimating the power consumption using the required V_{gs} and I_d of the input pair of the amplifier. For a low V_{gs}, the power supply voltage can possibly be lower than 1.8 V, while for a high V_{gs} a folded cascode amplifier might be required which has an increased power consumption. By taking the self-heating into account, using the technology's optimum g_m/C_{in} will no longer be the best option.

To compare these errors for the two architectures, different bias current densities and transistor widths are evaluated to find the optimum. The results are shown in Figure A.5. It is interesting to see that both architectures show similar characteristics and have almost the same optimum for the total error. In both cases this optimum is around a width of 100 µm and a drain current of 80 µA. At this optimum, the self-heating results in an error around 0.02 °C, the resolution is about 0.4 °C$_{rms}$ in 1 s, and the noise of the readout circuit results in an error of about 4.5 °C.

B

Improved Redistribution Layout

In order to reduce any parasitics between the capacitor nodes, additional shielding has been applied. This shielding has been applied in a structural way in order to better match the parasitics of the different capacitor nodes. In the routing of the lines from the capacitors to the switches, metal 1, metal 3 and metal 5 have been exclusively used for shielding. Metal 2 is used for the routing in one direction, while metal 4 is used for the routing in the orthogonal direction. These routes are equally spaced and have additional shielding in between. All the routes also have equal length. This grid pattern of horizontal and vertical routes is interconnected via small holes in the metal 3 shield. The routing in this new layout can be seen in Figure B.1 (upto, but not including metal 3), Figure B.2 (upto, but not including metal 5) and Figure B.3 (all metal layers).

However, because of the additional shielding, there are more parasitic capacitances to ground. As a result, the minimum size capacitance that can be selected for the redistribution is larger and the amount of generated noise is lower. This can be solved by slightly reducing the size of the unit capacitors. Since the additional parasitic capacitances are mostly metal-to-metal capacitors, they will have similar properties (matching, temperature stability) as the metal fringe capacitors used as the main redistribution capacitors. Another option is to reduce the size of these parasitics by increasing the spacing of the routing and shielding.

With this new layout for the redistribution circuit, the clock loading has increased due to the extra shielding. As a result, the clock generator was no longer fast enough to ensure non-overlapping clock signals over all temperatures and corners. This can be solved by increasing the strength of the clock buffers in the clock generator.

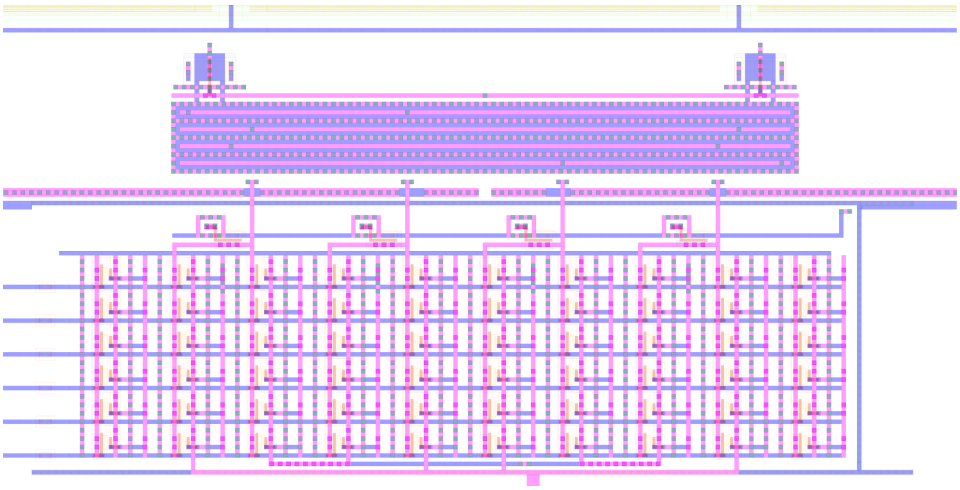

Figure B.1: The routing, shown upto, but not including metal 3.

Figure B.2: The routing, shown upto, but not including metal 5.

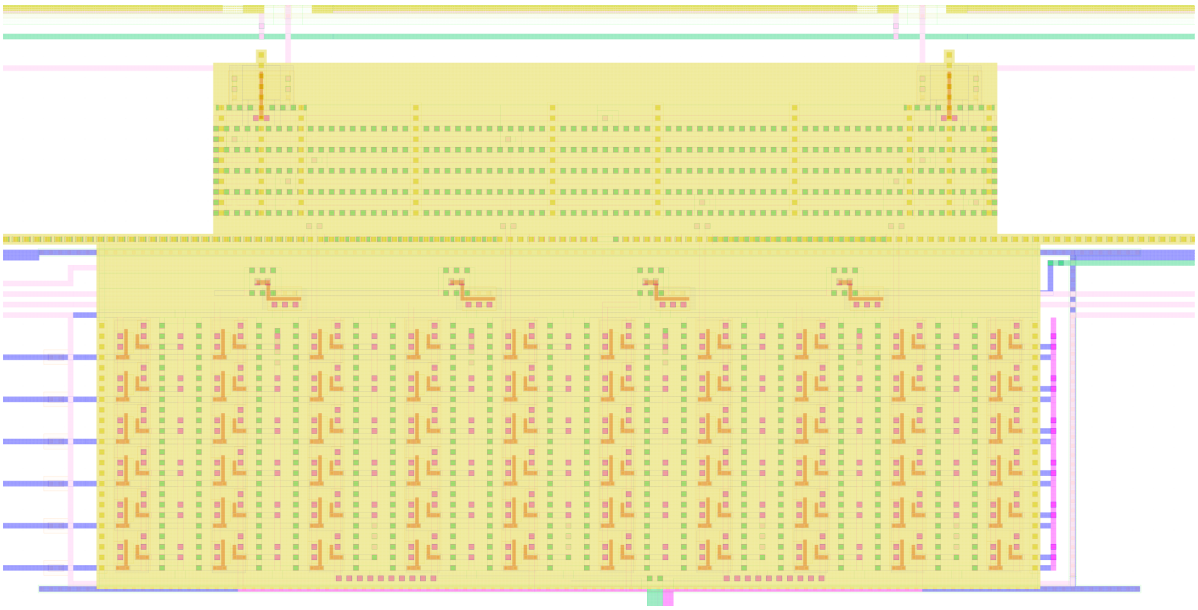

Figure B.3: The routing, shown with all metal layers visible.

C

Tape-out Rev. 2

This appendix discusses the updates in the second tape-out:

- Pre-amplifier updates (see Section C.1):
  - Updated the biasing circuit.
  - Introduced a variable feedback capacitor.
- Clock generator updates (see Section C.2):
  - The final clock buffers are now powered from the analog supply.
  - A reset signal for circuit synchronization has been added.
- Shift register updates:
  - Added bits to configure the pre-amplifier feedback and the clock generator reset.
  - Minor change in the layout and decoupling.
- Top level updates (see Section C.3):
  - Re-allocated part of the digital supply decoupling capacitors for analog supply decoupling
  - Added decoupling for the external reference voltage and 0.6V common-mode voltage.
  - The top-level now contains twice the same circuit, both with redistribution.

C.1. Pre-amplifier

First of all, the biasing circuit has been updated for a reduced bias current and common-mode voltage dependence, see Figure C.1. Now, the NMOS cascode biasing depends on the common-mode voltage and bias current, ensuring enough headroom for the pre-amplifier input transistor. Figure C.2 shows the simulated drain-source voltages of the input transistor and cascode transistor for the old and new tape-out and the resulting open-loop gain is shown in Figure C.3.

Secondly, to ensure a settling speed independent of the size of the source capacitor a variable feedback capacitor has been introduced in the pre-amplifier. The schematic of the capacitor DAC is shown in Figure C.4. Unit capacitors of about 13 fF are used that can be enabled in groups, the left-most capacitor has half the size of a unit-capacitor to allow for a fine-tuning if necessary. On both sides a dummy capacitor has been placed. Figure C.5 shows the simulated settling speed versus the number of unit capacitors that are active in the redistribution circuit.

The layout of the updated pre-amplifier with CDAC is shown in Figure C.6.

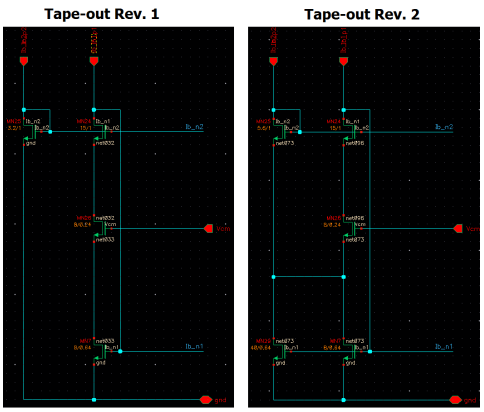

Figure C.1: The biasing circuit.

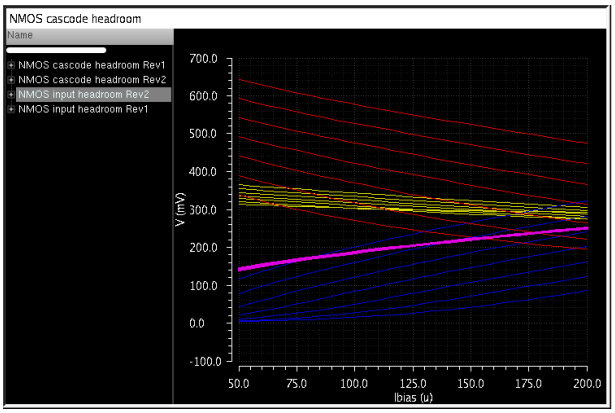

Figure C.2: The simulated drain-source voltages of the input transistor and cascode transistor.

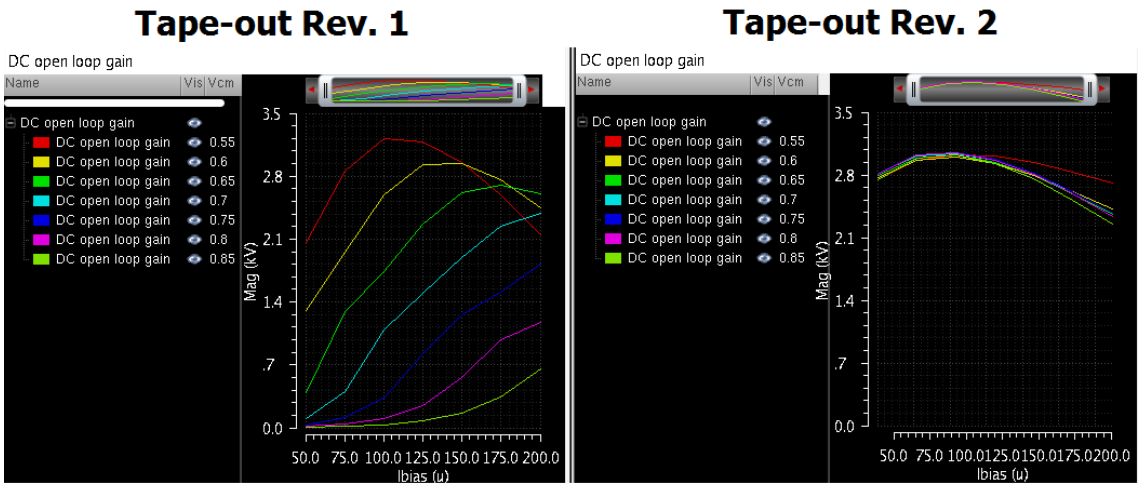

Figure C.3: The open-loop gain of the pre-amplifier.

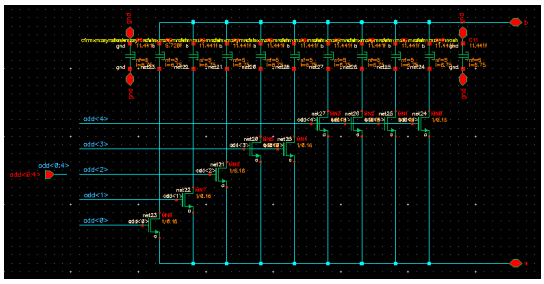

Figure C.4: The schematic of the capacitor DAC.

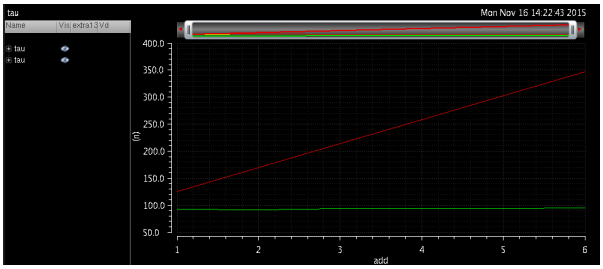

Figure C.5: The settling speed versus the number of unit capacitors that are active in the redistribution circuit (red: Rev. 1; green: Rev. 2).

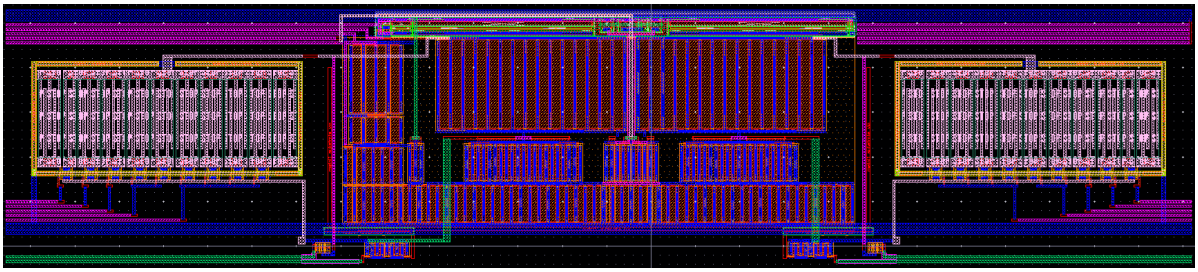

Figure C.6: The layout of the updated pre-amplifier with CDAC.

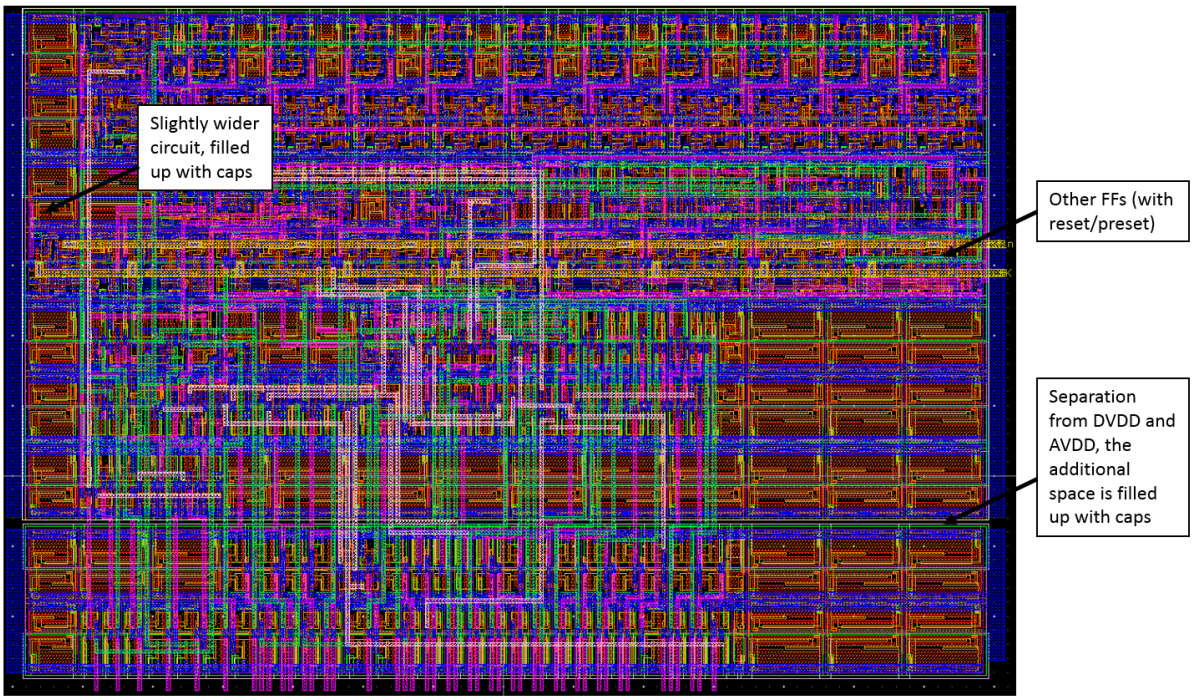

Figure C.7: The updates in the layout of the clock generator.

C.2. Clock generator

The final clock buffers and inverters that drive the switches in the analog circuit are now powered from the analog supply whereas before they were powered from the digital supply. Also a reset signal for circuit synchronization has been added by replacing all the state and counter flipflops by flipflops with an asynchronous reset/preset. The updates in the layout are highlighted in Figure C.7. Since some wires have an increased resistance/capacitance, the complete extracted circuit has been verified over corners/temperatures. Another simulation of repeatedly resetting the circuit shows that the circuit now always starts in the same state after a reset.

C.3. Top level

C.3.1. Decoupling

Since the clock buffers are now powered from the analog supply, part of the digital supply decoupling capacitors were re-allocated for analog supply decoupling. In addition, decoupling for the external reference voltage and 0.6 V common-mode voltage has been added. The change in decoupling has been summarized in Table C.1. The capacitance shown is for each circuit inside the chip, except in case of the common-mode voltage which is shared for the two circuits.

Table C.1: A summary of the change in on-chip decoupling capacitance.

	Tape-out Rev. 1	Tape-out Rev. 2
Analog supply	56.1 pF	67.3 pF
Digital supply	56.1 pF	22.4 pF
Reference voltage	0 pF	15.1 pF
Common-mode voltage	0 pF	11.2 pF

The layout of the new chip, including padring, is shown in Figure C.8. Comparing with the die photo in Figure 5.15, it is clear that the on-chip decoupling has changed.

C.3.2. Circuit

In the first tape-out (Rev. 1), two sensor circuits were incorporated in a single chip, with the only difference that the redistribution circuit would not redistribute the voltage for the second circuit. This circuit change was added,

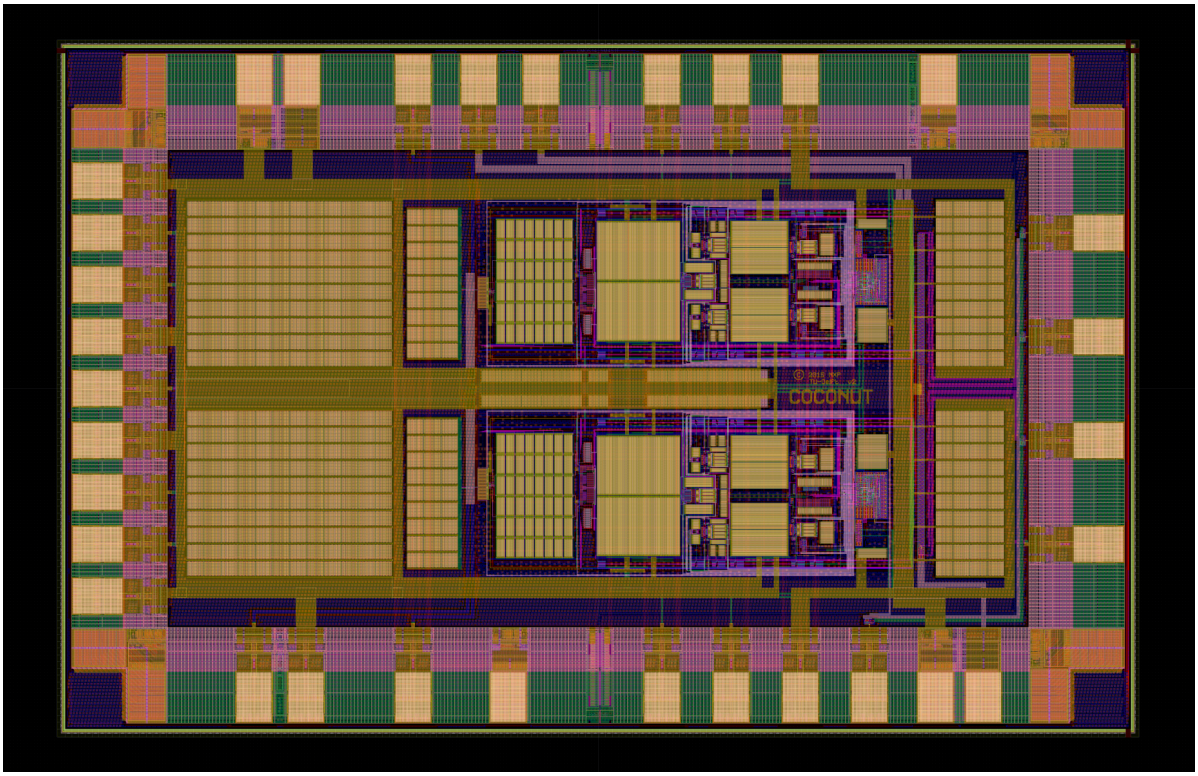

Figure C.8: The layout of the new chip, including padring.

because it was not certain that the redistribution circuit would function properly. In case it would not work, at least the readout circuit could be tested. However, as seen from the measurements the redistribution circuit functions. Therefore, for the second tape-out (Rev. 2), the full circuit with redistribution is placed twice. Alternatively, only a single circuit could be placed, but this was not done since anyway the same padring and PCB for testing would be used, which allows for two circuits. As a result, more samples will be available for testing.

As described in Section 5.2.3, the two circuits could interfere with each other. Therefore, it is important to be able to shut down one circuit. During the measurements, it was observed that the digital circuit cannot simply be shut off by disconnecting the power, it needs to be shorted to ground. Alternatively, in the new tape-out the reset signal of the clock generator could remain asserted.

Equivalently, the analog supply can be grounded. However, at that point the PMOS test followers of the circuit that is still powered are no longer usable. This is because the PMOS follower that is no longer powered has a diode (the source-bulk diode) from the test-pad to ground. As this only removes this test functionality, this is not a problem during normal operation. Besides, even if the analog part remains powered, no interference is expected when the circuit is not clocked.

D

Schematic of the PCB

The schematics of the PCB are shown in Figure D.1, D.2, D.3 and D.4.

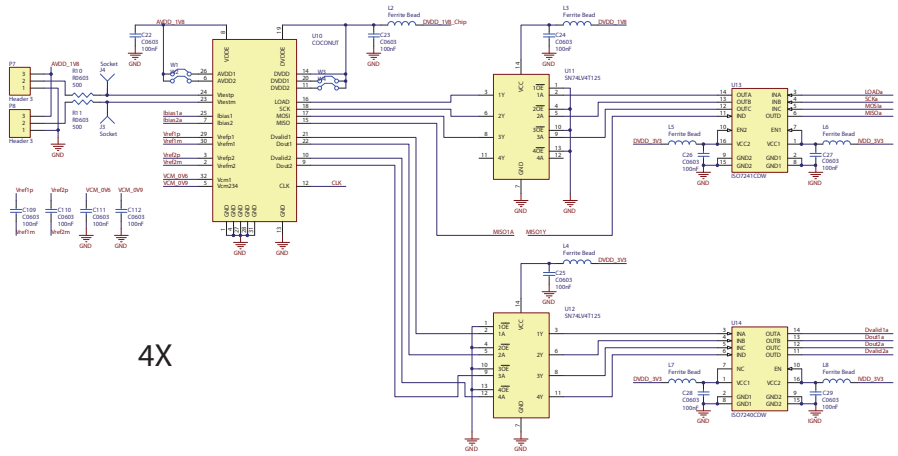

Figure D.1: The socket for the chip with the level shifters and isolators.

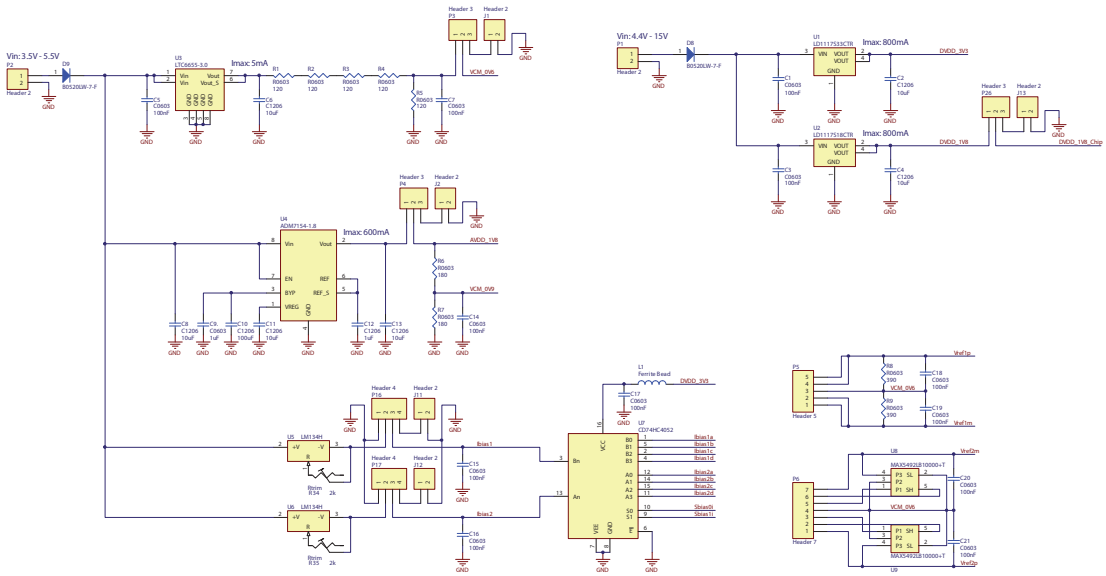

Figure D.2: Analog and digital supplies, and the common-mode voltage, reference voltage and bias current circuits.

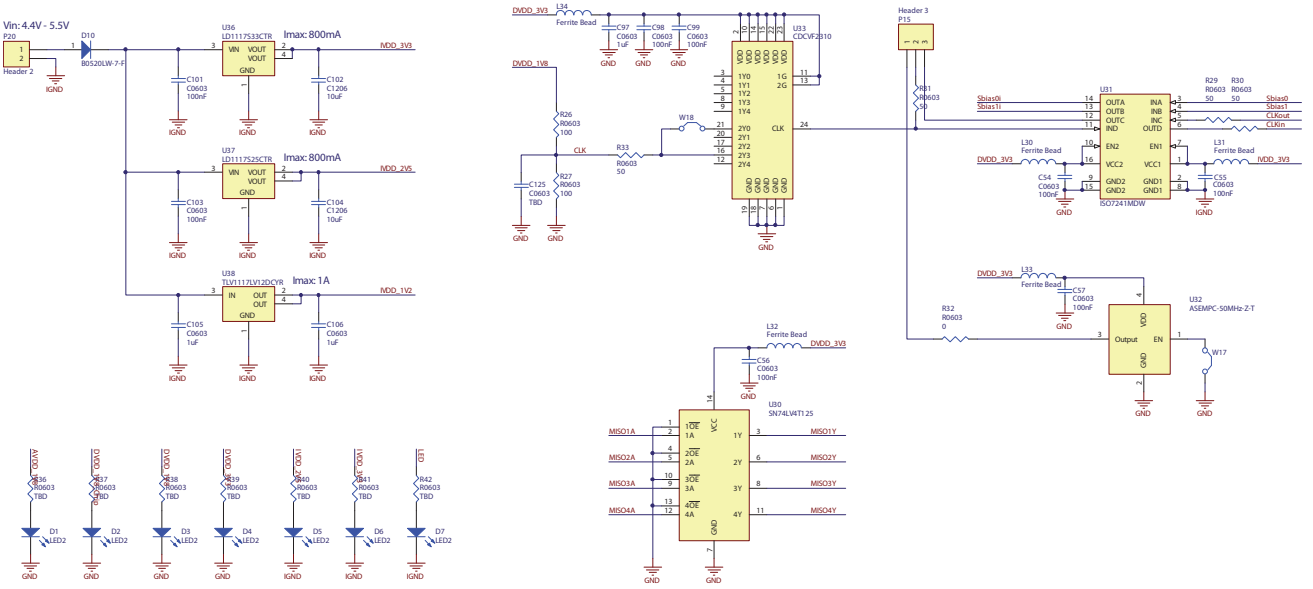

Figure D.3: FPGA supplies, power supply LEDs and chip clock circuit.

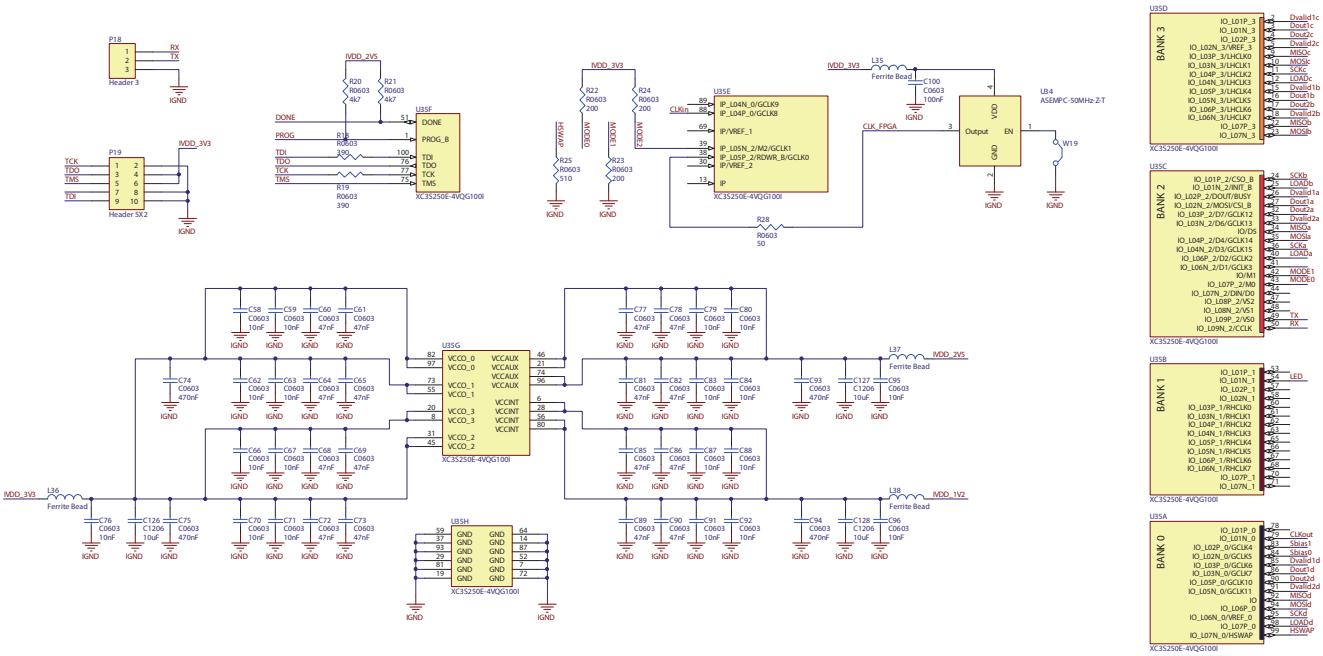

Figure D.4: The FPGA with JTAG and UART header and a clock generator.

Bibliography

[1] K. Makinwa. Smart temperature sensor survey. [Online]. Available: http://ei.ewi.tudelft.nl/docs/TSensor_survey.xls

[2] K. Souri, Y. Chae, K. Makinwa *et al.*, "A cmos temperature sensor with a voltage-calibrated inaccuracy of 0.15 c (3) from 55 c to 125 c," *Solid-State Circuits, IEEE Journal of*, vol. 48, no. 1, pp. 292–301, 2013.

[3] J. Shor and K. Luria, "Miniaturized bjt-based thermal sensor for microprocessors in 32- and 22-nm technologies," *Solid-State Circuits, IEEE Journal of*, vol. 48, no. 11, pp. 2860–2867, Nov 2013.

[4] E. Rotem, J. Hermerding, A. Cohen, and H. Cain, "Temperature measurement in the intel (r) coretm duo processor," *arXiv preprint arXiv:0709.1861*, 2007.

[5] M. Floyd, M. Allen-Ware, A. Buyuktosunoglu, K. Rajamani, B. Brock, C. Lefurgy, A. J. Drake, L. Pesantez, T. Gloekler, J. A. Tierno *et al.*, "Introducing the adaptive energy management features of the power7 chip," *IEEE Micro*, no. 2, pp. 60–75, 2011.

[6] F. Sebastiano, L. J. Breems, K. A. Makinwa, S. Drago, D. M. Leenaerts, and B. Nauta, "A 1.2-v 10-w npn-based temperature sensor in 65-nm cmos with an inaccuracy of 0.2 c (3) from 70 c to 125 c," *Solid-State Circuits, IEEE Journal of*, vol. 45, no. 12, pp. 2591–2601, 2010.

[7] M. Shahmohammadi, K. Souri, K. Makinwa *et al.*, "A resistor-based temperature sensor for mems frequency references," in *ESSCIRC (ESSCIRC), 2013 Proceedings of the*. IEEE, 2013, pp. 225–228.

[8] U. Sonmez, R. Quan, F. Sebastiano, K. Makinwa *et al.*, "A 0.008-mm 2 area-optimized thermal-diffusivity-based temperature sensor in 160-nm cmos for soc thermal monitoring," in *European Solid State Circuits Conference (ESSCIRC), ESSCIRC 2014-40th*. IEEE, 2014, pp. 395–398.

[9] X. Zhang and D. Chen, "An integrated circuit solution of thermal noise thermometer with cascaded preamplifier and 6-bit resolution analog-to-digital converter," in *Circuits and Systems (ISCAS), 2015 IEEE International Symposium on*, May 2015, pp. 2221–2224.

[10] S. Paek, W. Shin, J. Lee, H.-E. Kim, J.-S. Park, and L.-S. Kim, "All-digital hybrid temperature sensor network for dense thermal monitoring," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2013 IEEE International*, Feb 2013, pp. 260–261.

[11] 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," *Solid-State Circuits, IEEE Journal of*, vol. 44, no. 7, pp. 2019–2025, 2009.

[12] P. C. de Jong, G. Meijer, and A. H. Van Roermund, "A 300 c dynamic-feedback instrumentation amplifier," *Solid-State Circuits, IEEE Journal of*, vol. 33, no. 12, pp. 1999–2009, 1998.

[13] L. Vandersypen, R. Hanson, L. van Willems Beveren, J. Elzerman, J. Greidanus, S. De Franceschi, and L. Kouwenhoven, "Quantum computing with electron spins in quantum dots," in *Quantum computing and quantum bits in mesoscopic systems*. Springer, 2004, pp. 201–209.

[14] J. Qu, Y. Fu, J. Zhang, H. Rogalla, A. Pollarolo, and S. P. Benz, "Flat frequency response in the electronic measurement of boltzmann's constant," *Instrumentation and Measurement, IEEE Transactions on*, vol. 62, no. 6, pp. 1518–1523, 2013.

[15] T. Blalock and R. Shepard, "Decade of progress in high-temperature johnson noise thermometry," Oak Ridge National Lab., TN (USA); Tennessee Univ., Knoxville (USA). Dept. of Electrical Engineering, Tech. Rep., 1982.

[16] G. Scandurra, C. Ciofi, and A. Gambadoro, "A new approach to johnson noise thermometry based on noise measurements only," *Fluctuation and Noise Letters*, vol. 10, no. 02, pp. 133–145, 2011.

[17] D. Baert and A. Vervaet, "Small bandwidth measurement of the noise voltage of batteries," *Journal of power sources*, vol. 114, no. 2, pp. 357–365, 2003.

[18] C. Borkowski and T. Blalock, "A new method of johnson noise thermometry," *Review of Scientific Instruments*, vol. 45, no. 2, pp. 151–162, 1974.

[19] C. Pickup, "A high-resolution noise thermometer for the temperature range 90–100 k," *Metrologia*, vol. 11, no. 4, p. 151, 1975.

[20] M. Pepper and J. Brown, "Absolute high-temperature johnson noise thermometry," *Journal of Physics E: Scientific Instruments*, vol. 12, no. 1, p. 31, 1979.

[21] H. Klein, G. Klempt, and L. Storm, "Measurement of the thermodynamic temperature of 4he at various vapour pressures by a noise thermometer," *Metrologia*, vol. 15, no. 3, p. 143, 1979.

[22] C. Cannon, "A 2200â° c fuel centerline johnson noise power thermometer," *Nuclear Science, IEEE Transactions on*, vol. 28, no. 1, pp. 763–766, 1981.

[23] M. Imamura and A. Ohte, "A new method of noise thermometry," *Temperature: Its Measurement and Control in Science and Industry*, vol. 5, pp. 139–142, 1982.

[24] L. Crovini, A. Actis, and R. Galleano, "A high temperature noise thermometer for accurate thermodynamic temperature measurements," *Instrumentation and Measurement, IEEE Transactions on*, vol. 42, no. 2, pp. 391–394, 1993.

[25] H. Sakurai, "A cross correlation noise thermometer for thermodynamic temperature measurement at low temperatures," in *SICE 2003 Annual Conference*, vol. 3. IEEE, 2003, pp. 3029–3032.

[26] J. Labenski, W. Tew, S. Benz, S. Nam, and P. Dresselhaus, "A determination of the ratio of the zinc freezing point to the tin freezing point by noise thermometry," *International Journal of Thermophysics*, vol. 29, no. 1, pp. 1–17, 2008.

[27] J. Zhang and S. Xue, "A noise thermometry investigation of the melting point of gallium at the nim," *Metrologia*, vol. 43, no. 3, p. 273, 2006.

[28] L. Spietz, K. Lehnert, I. Siddiqi, and R. Schoelkopf, "Primary electronic thermometry using the shot noise of a tunnel junction," *Science*, vol. 300, no. 5627, pp. 1929–1932, 2003.

[29] D. White, R. Galleano, A. Actis, H. Brixy, M. De Groot, J. Dubbeldam, A. Reesink, F. Edler, H. Sakurai, R. Shepard *et al.*, "The status of johnson noise thermometry," *Metrologia*, vol. 33, no. 4, p. 325, 1996.

[30] A. Savateev, "Compensated thermal noise pulse thermometer," *Measurement Techniques*, vol. 5, no. 2, pp. 114–120, 1962.

[31] T. C. Sepke, "Comparator design and analysis for comparator-based switched-capacitor circuits," Ph.D. dissertation, Citeseer, 2006.

[32] L. Callegaro, M. Pisani, M. Ortolano, V. D'Elia, and F. Manta, "Correlation method errors in johnson noise thermometry," in *CPEM 2010*, 2010.

[33] A. Naderi, H. Mojarrad, H. Ghasemzadeh, A. Khoei, and K. Hadidi, "Four-quadrant cmos analog multiplier based on new current squarer circuit with high-speed," in *EUROCON 2009, EUROCON'09. IEEE*. IEEE, 2009, pp. 282–287.

[34] E. Farshidi, "A current-mode true rms-dc converter based on electronically simulated translinear principle," in *Signals, Circuits and Systems, 2008. SCS 2008. 2nd International Conference on*. IEEE, 2008, pp. 1–4.

[35] A. Gerosa, M. Soldan, A. Bevilacqua, and A. Neviani, "A 0.18-μm cmos squarer circuit for a non-coherent uwb receiver," in *Circuits and Systems, 2007. ISCAS 2007. IEEE International Symposium on*. IEEE, 2007, pp. 421–424.

[36] S. C. Li and K.-L. Lin, "A±1.5 v cmos four-quadrant analogue multiplier using 3 ghz analogue squaring circuits," in *Circuits and Systems, 1998. ISCAS'98. Proceedings of the 1998 IEEE International Symposium on*, vol. 2. IEEE, 1998, pp. 347–350.

[37] S. J. Gift, "A high-performance full-wave rectifier circuit," *International Journal of Electronics*, vol. 87, no. 8, pp. 925–930, 2000.

[38] A. Monpapassorn, K. Dejhan, and F. Cheevasuvit, "A full-wave rectifier using a current conveyor and current mirrors," *International Journal of Electronics*, vol. 88, no. 7, pp. 751–758, 2001.

[39] M. Kumngern, "Cmos current-mode precision full-wave rectifier with improved bandwidth," in *Digital Information and Communication Technology and it's Applications (DICTAP), 2012 Second International Conference on*. IEEE, 2012, pp. 283–286.

[40] A. Virattiya, B. Knobnob, and M. Kumngern, "Cmos precision full-wave and half-wave rectifier," in *Computer Science and Automation Engineering (CSAE), 2011 IEEE International Conference on*, vol. 4. IEEE, 2011, pp. 556–559.

[41] A. Brodskii and A. Savateev, "A new method of absolute temperature measurements," *Measurement Techniques*, vol. 3, no. 5, pp. 397–402, 1960.

[42] C. M. Ralph, "Thermal noise investigation," May 4 1965, uS Patent 3,181,365.

[43] A. Ziel, "Noise," 1954.

[44] P. Kittel, "Comment on the equivalent noise bandwidth approximation," *Review of Scientific Instruments*, vol. 48, no. 9, pp. 1214–1215, 1977.

[45] S. O. Rice, "Mathematical analysis of random noise," *Bell System Technical Journal*, vol. 23, no. 3, pp. 282–332, 1944.

[46] M. J. Pelgrom, *Analog-to-digital Conversion*. Springer, 2010.

[47] B. Razavi, B. Wooley *et al.*, "Design techniques for high-speed, high-resolution comparators," *Solid-State Circuits, IEEE Journal of*, vol. 27, no. 12, pp. 1916–1926, 1992.

[48] G. Xu and J. Yuan, "Performance analysis of general charge sampling," *Circuits and Systems II: Express Briefs, IEEE Transactions on*, vol. 52, no. 2, pp. 107–111, 2005.

[49] T. Sugiyama and K. Yamaguchi, "Pulsewidth modulation dc potentiometer," *Instrumentation and Measurement, IEEE Transactions on*, vol. 19, no. 4, pp. 286–290, 1970.

[50] J. Robert, P. Deval, and G. Wegmann, "Very accurate current divider," *Electronics Letters*, vol. 25, no. 14, pp. 912–913, 1989.

[51] K. B. Klaassen, "Digitally controlled absolute voltage division," *Instrumentation and Measurement, IEEE Transactions on*, vol. 24, no. 2, pp. 106–112, 1975.

[52] P. C. de Jong and G. Meijer, "Absolute voltage amplification using dynamic feedback control," *Instrumentation and Measurement, IEEE Transactions on*, vol. 46, no. 4, pp. 758–763, 1997.

[53] K. S. Lee and Y. M. Lee, "Switched-capacitor cyclic dac with mismatch charge compensation," *Electronics Letters*, vol. 46, no. 13, pp. 902–903, 2010.

[54] P. Chen and T.-C. Liu, "Switching schemes for reducing capacitor mismatch sensitivity of quasi-passive cyclic dac," *Circuits and Systems II: Express Briefs, IEEE Transactions on*, vol. 56, no. 1, pp. 26–30, 2009.

[55] R. Kapusta, H. Zhu, and C. Lyden, "Sampling circuits that break the kt/c thermal noise limit," *Solid-State Circuits, IEEE Journal of*, vol. 49, no. 8, pp. 1694–1701, 2014.

[56] J. Kim, B. S. Leibowitz, J. Ren, and C. J. Madden, "Simulation and analysis of random decision errors in clocked comparators," *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 56, no. 8, pp. 1844–1857, 2009.

[57] A. Graupner, "A methodology for the offset-simulation of comparators," *The Designers Guide Community, www designers-guide org*, 2006.