A Multi-Path Sigma Delta ADC

For use in battery management systems

L.A. Loopik

Technische Universiteit Delft

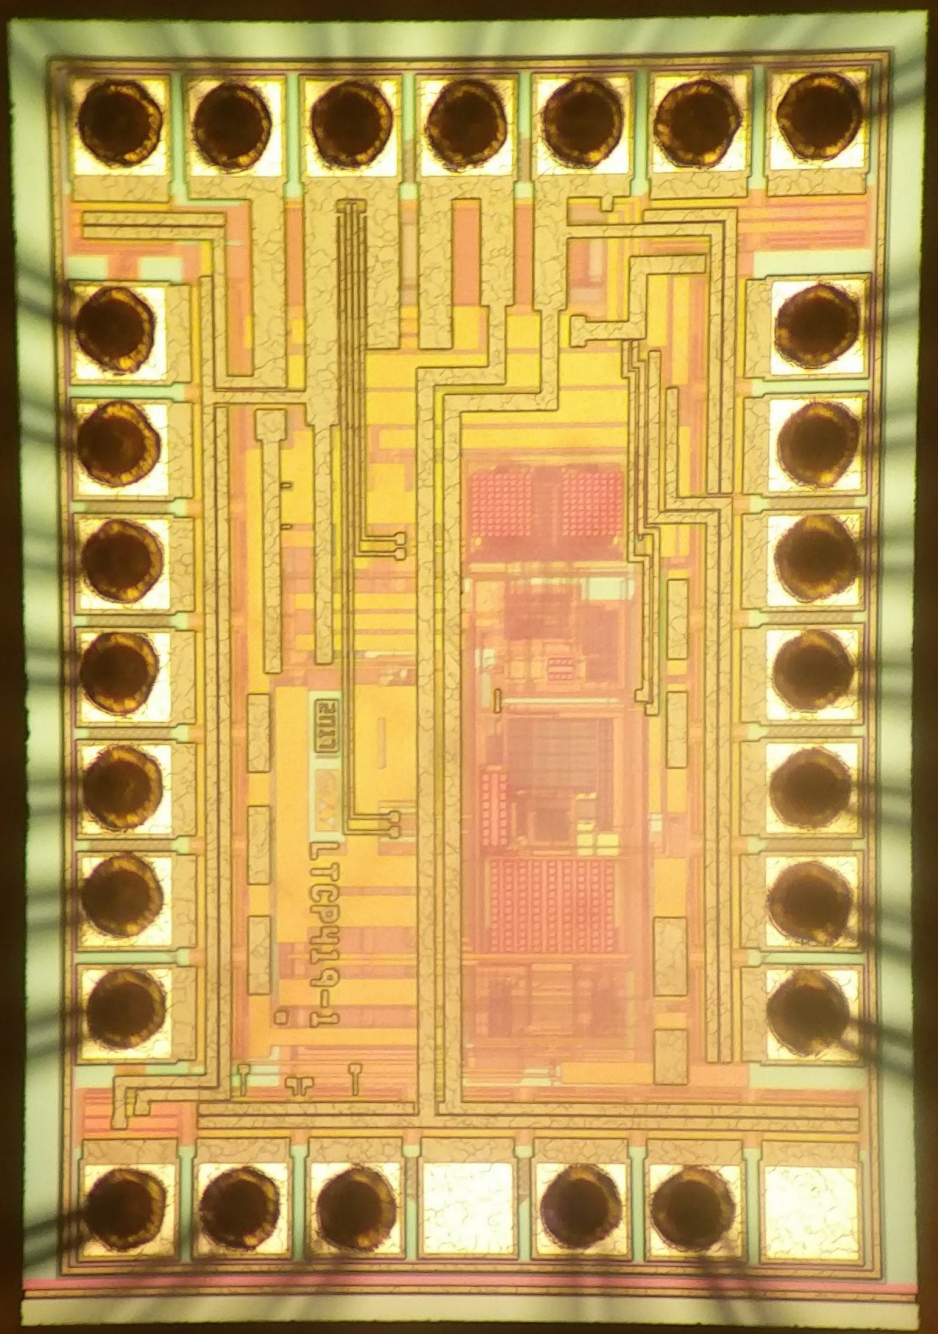

TU Delft Delft University of Technology

Challenge the future

A Multi-Path Sigma Delta ADC

For use in battery management systems

by

L.A. Loopik

in partial fulfillment of the requirements for the degree of

Master of Science
in Electrical Engineering

at the Delft University of Technology,
to be defended publicly on Monday December 18, 2017 at 14:00.

Thesis committee:	Prof. dr. ir. K.A.A. Makinwa,	TU Delft, supervisor
	Dr. F. Sebastiano,	TU Delft
	Dr. C. P. L. van Vroonhoven,	Analog Devices

This thesis is confidential and cannot be made public until December 18, 2020

TU Delft Delft University of Technology

Abstract

This thesis presents the design and testing of a multi-path analog-to-digital converter for current sensing in battery management systems. The specifications of this converter are (1) continuously integrate and convert a signal at 125 Hz with a resolution of 20 bit and (2) perform on demand fast conversions of 20 µs at a resolution of 14 bit. Both conversions should be done with a maximum 1 LSB offset and 0.1 % gain error. The realized system consists of a discrete time 1st order sigma delta converter, followed by a 10 bit charge redistribution extended counting ADC. This design uses the same analog front end for both conversions, thereby saving both power and chip area. To achieve the lowest offset possible, a system level chopping scheme is implemented. A working prototype has been made in TSMC 0.18 BCD technology, using an area of 570 µm × 150 µm with a 5 V supply voltage. The realized chip proves that the concept of combining the analog front end works and achieves a 15 bit linearity. At the 20 bit level, there are still significant issues and despite serious efforts the root cause of those issues has not yet been isolated. Next design and testing steps have however been identified.

Acknowledgements

Let me start by thanking everyone from the TU Delft and Linear Technology/Analog Devices for making it possible for me to reach this point in my life. First of all, many thanks to Prof. Kofi Makinwa for creating the opportunity for me to do my master thesis project with Linear Technology on this interesting and relevant topic. This way I was able to spend a fruitful year in Munich, learn a lot, both professionally and personally.

At Linear Technology, now merged with Analog Devices, my special thanks goes out to Christoph Schwoerer for inviting me and supporting me throughout my project and for making me a job offer! My biggest thanks goes to Caspar van Vroonhoven, my daily supervisor for his continuous support. He was always available to help me on any topic: anywhere from housing to the most detailed technical issues. I greatly enjoyed our technical conversations and thank you for your trust in me.

I also received great help and advice from Gerd Trampitsch on circuit design and from Axel Klein on my test setup and python. I thoroughly enjoyed being a part of the Munich design team. You are all great colleagues to work with! Jiri Dak, thank you for all your amazing layout work, and your help with all our tape-out troubles. Andreas Beermann, thank you for your help with the test PCB and Patrick Wilhelm, for your support with the FPGA's.

Further, I would like to thank my parents for their constant support throughout my study and their never ending trust in me. My fellow students Electrical Engineering and the Electrotechnische Vereeniging for our times together. My housemates at JvB 163 for always being there for me. And last but not least my student association DSB, especialy my friends from Schranz, for all the fun we had.

Let me finish by saying that I am very excited to be returning to Munich for a permanent position as Design Engineer with Analog Devices in January.

L.A. Loopik
Rijswijk, December 2017

Contents

1 Introduction

This thesis describes the design and testing of a Multi Path Analog to Digital Converter (MP-ADC). The aim of this MP-ADC is to explore the possibility of measuring two different quantities (electrical current and electrical charge) over two different time intervals with the same chip. If this succeeds it opens up the possibility to make smaller (die area) Battery Management System (BMS) chips with lower power consumption.

1.1. Battery Management Systems

The world is using more and more batteries every day. Every day, more appliances and machines are battery powered. This change is going hand in hand with improvements in battery technology. And while battery technology is getting better, it is also getting more complicated. In the past, it was quite easy to determine the state of charge of a battery by measuring the cell voltage. But as battery technology improves, side effects such as dropping cell voltage, are getting smaller and smaller. Also, more complicated batteries require more control circuitry to prevent overcharging, undercharging, overheating and to determine the state of charge and state of health of a battery. This control circuitry is called the Battery Management System, or BMS. Because batteries are expensive it is important to maximize their lifetime. Being able to measure the behaviour of the battery and controlling it based on these measurements can increase the lifetime of the battery, and as such, give a significant competitive edge.

1.1.1. System overview

A modern battery system must measure and keep track of many different properties of the battery it is monitoring. In this thesis, we will focus on two main properties used to determine the state of charge and state of health of a battery: cell impedance and delta charge.

The cell impedance is the internal resistance of the battery, this quantity changes as the battery is charged and discharged, as well as over the lifetime of a battery. Because of this it can be used to determine the state of health, as well as give an estimate for the remaining state of charge.

Another method of determining the state of charge, is to use the delta charge of the battery. This quantity gives an good basis for determining the state of charge of the battery, by essentially measuring all charge going into, and going out of the battery. Of course, this measurement is also not perfect, as it does not include, for example, self discharge of the battery.

1.1.2. Impedance measurement

To measure the cell impedance, both the cell voltage and current are measured. The measured cell voltage is now a combination of the internal cell voltage and the voltage drop over the internal resistance of the cell. When this measurement is done with (at least) two distinct current levels, the internal cell voltage can be cancelled out, and the cell impedance can be calculated.

There are multiple techniques available to measure at two distinct current levels. The easiest method is to modulate the load of the battery (connect an extra load resistance) for a short time during the second measurement. Another option is to make use of the natural variations of the current when

the battery is connected to a non-constant load. It is important that these measurements are done with minimal energy use and minimal interference with normal operation. Therefore, these measurements need to be done as quickly and energy efficiently as possible.

Importance of synchronous V and I measurements for impedance measurement

Because in most situations, the load of the battery is not constant. It is very important to make sure that the voltage and current measurements used for the impedance calculation are done at the exact same time.

1.1.3. Coulomb counting

For measuring the delta charge of the battery, a technique called Coulomb Counting is used. The total charge entering and exiting the battery is measured by continuously integrating the battery current. For this integration to be accurate, the current measurement should be done with a very low offset and gain error. Additionally, an accurate time reference is required for accurate integration. For BMS's, this time reference is usually available in the form of a crystal oscillator clock. In order to correctly integrate the signal, it is important that current is continuously measured, without any gaps between the time windows. The speed of these measurements is less important, as they are mostly used to track the status of the battery over a longer time.

1.1.4. Current-sensing techniques

For both the impedance measurement and the coulomb counting, a current measurement needs to be done. There are multiple techniques in use to measure currents. The most used ones are based on either magnetic sensing or shunt-based sensing.

Magnetic

Magnetic current sensing devices use the fact that an electric current produces an magnetic field proportional to the current (Ampère's law) to measure this current[1][1]. Measuring this magnetic field can be done in multiple ways, ranging from using Hall-effect sensors to current transformers. This technique has the advantage that there is no device in series with the current to be measured. This also means that there is a galvanic isolation between the battery and the measurement circuit. The main disadvantage of this technique is that it is hard to do low offset, high precision measurements due to interference with external magnetic fields.

Shunt-based

The other commonly used technique to measure current is to use a shunt resistor in the current path. This can be either an external shunt[2], or a internal shunt[3][2]. The current will now produce a voltage proportional to the current over the shunt (Ohm's law), which can be measured using conventional ADC techniques. In order to have an accurate measurement, the shunt has to be very accurate as well. This means having a well known resistor value, but also a low temperature coefficient (as the shunt heats up during use). The main advantage of using this technique is the high precision measurements it allows. The main disadvantage is the use of an (expensive) shunt resistor that needs to be in the main current path of the battery.

1.2. The multi-path current ADC

As described above, in a BMS needs to do two types of current measurements: a fast, on demand, medium-precision measurement for use in impedance measurements, and a slower, continuous, high precision measurement, for use in coulomb counting. Existing BMS's have two parallel converters that measure the same shunt voltage to fit both those requirements. In order to save chip area (and thus costs), a investigation is done to combine the two types of measurements into one, multi-path, analog to digital converter. This converter should continuously measure and integrate the current through the shunt. In addition to this it should be able to do a quick conversion of the instantaneous current

[1] Open loop: MLX91206 https://www.melexis.com/en/product/MLX91206
Closed loop: DRV411: http://www.ti.com/product/DRV411
[2] LTC2944 http://www.linear.com/product/LTC2944
[3] LTC2947: http://www.linear.com/product/LTC2947

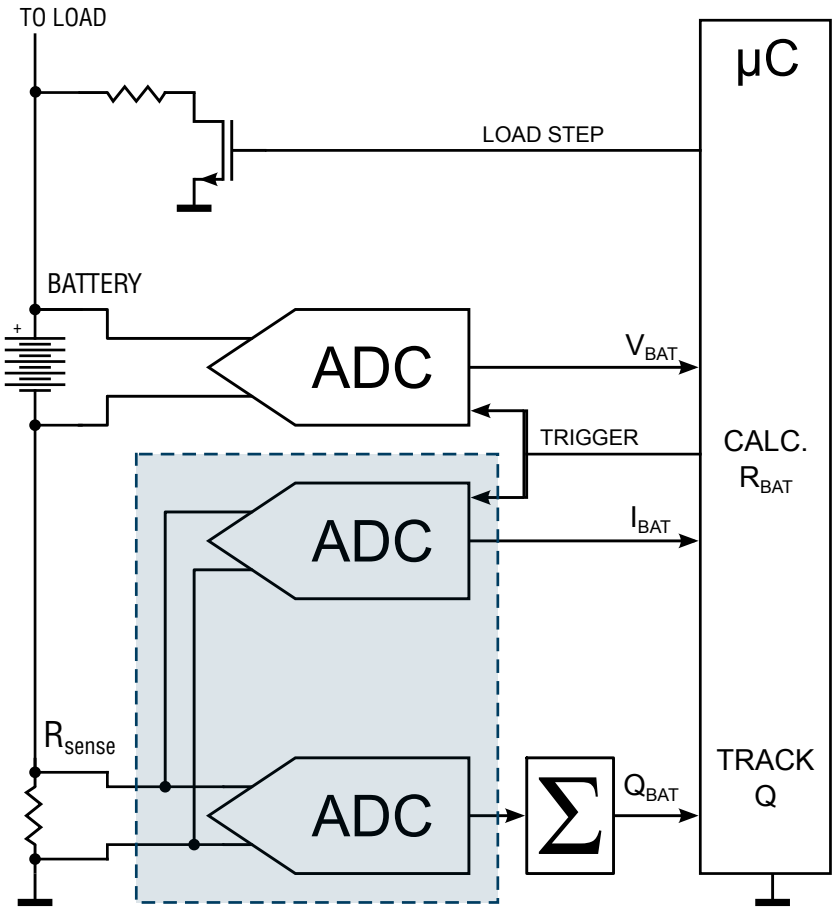

Figure 1.1: A simplified view of the BMS showing the intended combination of converters.

through the shunt, whenever triggered to do so by an external controller. Furthermore, the on demand current measurement should not interfere with the continuous measurements, as to not negatively affect the continuous integration. The targeted combination of converters is shown in figure 1.1.

1.3. Thesis outline

The main research done in this thesis is to design such a converter for use in future BMS's. Chapter 2 will describe the requirements needed for such a converter. Next, chapter 3 looks into different aspects of current-to-digital conversion. With that knowledge, it will be possible to make a system level design, this is discussed in chapter 4. Next the circuit level design is documented in chapter 5. As the final result of this research is produced in silicon, the layout process is described in chapter 6. Chapter 7 describes the testing of the silicon and its results. Finally, in chapter 8 conclusions are drawn and recommendations for further research are made.

2

System requirements

As documented in chapter 1, by combining a continuously integrating converter with a fast, on demand converter, their separate requirements need to be combined. In this chapter the requirements of the two different measurements will be discussed separately.

2.1. Requirements specification

The MP-ADC need to fulfill the requirements in tables 2.1 and 2.2, as given by the sponsor of this project.

In the sections below, some background to the requirements will be discussed

2.2. Coulomb counting

In order to correctly integrate the current through the shunt resistor, and be able to accurately calculate the charge that flowed through the shunt, it is important that the charge measurement is as accurate as possible. There are a few requirements that make a coulomb counting measurement stand out from a usual current measurement. These requirements are: offset and gain stability, continuous measurement and windowing.

2.2.1. Offset and Gain stability

Because the resulting measurements will be integrated indefinitely, it is important to have a very low offset, as an offset in the measurement will result in charge drift in the integrated charge. A gain error will translate directly to a gain error in the calculated charge, so it is also important to keep the gain stable. Linearity of the measurement will also translate to a non-linear error in the charge, but as long as this error is lower then the instability of the gain, this error will not substantially increase the total charge error.

2.2.2. Continuous measurement

To calculate the total amount of charge correctly, even when the current through the shunt is not constant, it is necessary to continuously measure the current trough the shunt, and not to have any "gaps" in the measured data. In other words: as soon as one measurement is complete, the next one has to start, without any delay. This has the effect that the input stage is always busy sampling the input signal, and cannot be time-multiplexed, unless the whole input stage is implemented twice.

2.2.3. Windowing

The final special requirement for coulomb counting is also imposed by the integration. To have a correct integration of the input signal, all input samples need to be weighted equally. This is another way of saying that if you apply a current pulse to the converter, it should always integrate to the same charge (if it is the same pulse), no matter when in the conversion cycle the current pulse is applied to the converter.

Table 2.1: The specifications for the coulomb counter

	Condition	Min	Typ	Max	
Resolution (No missing codes)	2ms Mode	18			Bits
	8ms Mode	20			Bits
Full Scale Input Voltage			±2		V
LSB Voltage of Charge Sense	2ms Mode		15.625		µV
	8ms Mode		3.90625		µV
Noise	2ms Mode		6		µVrms
	8ms Mode		3		µVrms
Gain Error	90% FS			±0.1	%
Offset Voltage	IP=IM=0V			±1	LSB
Integral Non-linearity				±3	LSB
Supply Voltage			5		V
Input Voltage Common Mode Range			2.5		V
Conversion Time	8ms Mode		8		ms
	2ms Mode		2		ms

Table 2.2: The specifications for the on demand current measurement

	Condition	Min	Typ	Max	
Resolution (No missing codes)	125µs Mode	16			Bits
	30µs Mode	14			Bits
Full Scale Input Voltage			±2		V
LSB Voltage of Charge Sense	125µs Mode		62,5		µV
	30µs Mode		250		µV
Noise	125µs Mode		40		µVrms
	30µs Mode		80		µVrms
Gain Error	90% FS			±0.1	%
Offset Voltage	IP=IM=0V			±1	LSB
Integral Nonlinearity				±3	LSB
Conversion Time	125µs Mode		125		µs
	30µs Mode		31		µs
Trigger to measurement delay				1	µs

2.3. On demand current measurement

The requirements for the asynchronous current measurement are different from the requirements for coulomb counting, as it is a single shot measurement. This means it is not constrained by requirements as continuous integration. The main challenge for the asynchronous current measurement is however, to integrate it with the coulomb counter in such a way, that it is almost indistinguishable from two separate ADC's.

2.3.1. Timing constraints

The main purpose of the single shot current measurement is to combine the measurement with a simultaneous cell voltage measurement, to calculate cell impedance. For this calculation to be accurate, the current measurement must be done in the same time frame as the cell voltage measurement, as also explained in chapter 1. Because the voltage measurement is done by a different converter, which is not necessarily synchronized with the multi-path current ADC, the current measurement must be able to start at any time, started by an external trigger signal. This essentially means that the current measurement must be completely independent (in the timing aspect) from the coulomb counter.

2.3.2. Interference with coulomb counting

Apart from being independent of the timing of the coulomb counter, the asynchronous current measurement must also not interfere with the coulomb counter. This is important, as it is important that the charge measured by the coulomb counter should be as accurate as possible, and should not be disturbed by intermediate current measurements.

3

Current-to-Digital converters and their front ends

In this chapter, possible front ends for the ADC's will be discussed. In order to do this, it is first necessary to decide on a converter type. Next, it is then possible to look at different front end designs, and choose the one which best matches the characteristics of this design.

3.1. Converter type

By specification both converters are high precision converters with a low to medium bandwidth and very low offset. To decide what type of converter is best for this type of measurements, a couple of converters are considered.

Flash ADC

The simplest type of ADC is the flash ADC. A flash ADC consists of a way to split the reference into N equal parts, which are then used to compare to the input voltage by using N comparators. Because there is only one thing happening at the moment of conversion (the comparators making their decision), a flash ADC can be incredibly fast. For an n bit converter, however, 2^n comparators are needed, which, at high resolutions, is not a realistic solution. Apart from the sheer number of comparators, the offset and matching of requirements of the comparators become exponentially more challenging with higher resolutions. Because of these properties, a flash ADC is not a realistic ADC type for this system.

Dual slope ADC

Another type of ADC is the dual slope ADC. Here the input signal is first integrated for a fixed amount of time, after which the reference signal with the opposite polarity is integrated until the output of the integrator is back to zero. The time it takes to bring the output of the integrator back to zero is proportional to the input signal. The advantage of this technique is that it integrates the input signal, which suppresses noise. The disadvantage is that during the discharge time, the input signal is not measured. Because the charge conversion requires a continuous input measurement, this is not a viable option for this design.

Successive approximation ADC

The next option are successive approximation ADC's[3], which by trying to match the output of a DAC with the input signal. This is done by trial-and-error for every bit from the MSB to the LSB as follows: First set all bits to 0, then first flip the MSB, if the signal is to high, flip the MSB back and keep it 0, if the signal is still to low, keep the bit 1. When this is continued for every bit, the final DAC value will represent the input signal as close as possible. This means that there is a digital representation of the input signal, and the signal is digitized. The maximum resolution and linearity of a successive approximation ADC is usually limited by matching, and can be increased by extensive trimming and calibration. As extensive trimming is not desired, and getting to 20 bit accuracy with component matching is not possible, the successive approximation ADC is also not a good option for this design.

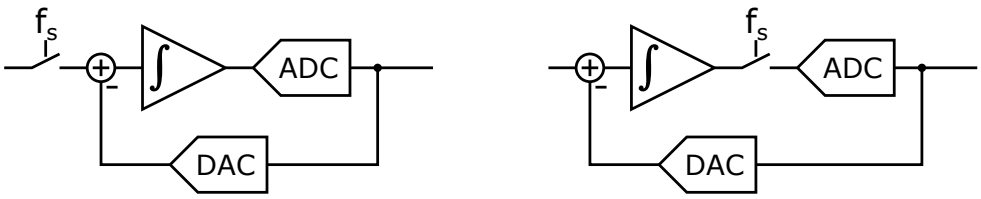

(a) Discrete Time (b) Continuous Time

Figure 3.1: The two main Sigma-Delta sampling architectures

Sigma delta ADC

Finally, there is the option of sigma delta ADC's[4], these work by oversampling and noise-shaping. This means that the ADC samples at a higher frequency than the output frequency. This is done by putting a low resolution ADC and DAC in a feedback loop. While oversampling can be done with any type of ADC, in a sigma delta ADC the feedback loop 'shapes' the quantization noise with a high pass response. Since most of the quantization noise energy is now at higher frequencies, it is possible to remove most of it by passing the signal through a low-pass filter. This way, a low resolution ADC and DAC combination can be used to create a very high resolution ADC without extreme matching requirements. And while the resulting output sample rate is still quite low, the conversions can be combined back to back, without any dead time in between, which is important for the charge conversions. Because of these advantages a sigma delta ADC is thus the ADC type of choice for this project.

3.2. Sampling

A crucial part of any analog to digital converter is sampling. Somewhere in the signal path, the time-continuous analog signal needs to be converted in a time-discrete signal. Especially for sigma-delta converters, there are multiple places in the signal path where this conversion can be done.

3.2.1. CT/DT sigma-delta converters

In a sigma-delta converter there are two conventional places to do sampling: before or after the integrator. This results in two different types of sigma-delta converters: a Continuous Time (CT) sigma-delta converter (see figure 3.1b) and a Discrete Time (DT) sigma-delta converter (see figure 3.1a). The two types vary greatly in implementation and behaviour, but the two most important differences are frequency folding and integrator design.

When using the DT solution, sampling is done at the beginning of the signal path of the converter. This means the input signal (and noise) should already be band-limited to prevent aliasing (and noise folding). This is one of the disadvantages of using a DT system. Another disadvantage is that DT systems will use more power then equivalent CT systems. An advantage of DT systems is that chopping is easier to implement and that a DT integrator is used. DT integrators have as advantage that they can use techniques as auto zeroing (AZ) and finite gain compensation relatively easily. Also the design of a DT DAC is easier then the design of a CT DAC.

A CT solution instead has different advantages and disadvantages. By moving the sampling point behind the integrator, a low pass filtering is applied to the input signal. This means that there is an implicit anti-aliasing filter in a CT system. Another big advantage is that, as mentioned before, CT systems generally use less power for equivalent performance. Disadvantages of a CT system are that chopping becomes more complex, and that a CT (sample-and-hold) DAC needs to be implemented. Also jitter, and other timing inaccuracies become important. And while CT integrators are more power efficient than DT integrators, it is also much harder to get to the same level of precision using techniques like auto zeroing or finite gain compensation.

3.3. Direct input sampling

With a DT system the sampling switch is the input of the system, as long as the input is already band-limited. The easiest way to connect the shunt to the system is thus to add an analog low pass filter to the shunt, and directly sample the voltage. By sampling the shunt directly like this, the input system of the converter is very simple, and chopping can be implemented in a straightforward manner.

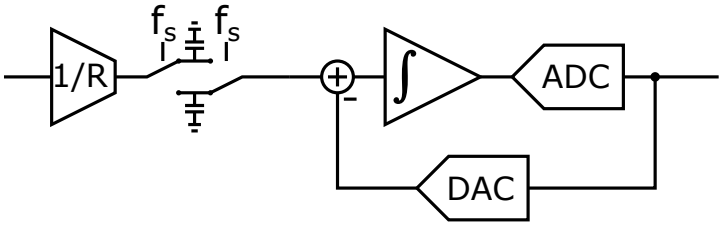

Figure 3.2: The boxcar pre-amplifier scheme

The disadvantages, however, are noise and input impedance. Noise because the shunt voltage will be relatively low (mV levels), and a big sampling capacitor is needed to have low enough kT/C noise. This big sampling capacitor will then lower the input impedance of the converter, unless a pre-charging scheme is used. Depending on the shunt and filter design, a too low input impedance can be a problem.

3.4. Pre-amplifier

The alternative to a direct input sampling system, is to use a pre-amplifier. Using a low noise pre-amplifier, can improve the SNR of the system and thus lower the noise requirements of the rest of the system. Another advantage of using a pre-amplifier is a higher input impedance.

3.4.1. Noise/Bandwidth

When using a DT sigma-delta system, the output of the pre-amplifier will be sampled. Because this causes frequency folding, the bandwidth of the amplifier should be less then the Nyquist frequency of the sigma-delta. But even if this is the case, the output noise of the pre-amplifier will be folded down to frequencies of interest. To prevent this, a low pass filter should be added behind the pre-amplifier.

3.4.2. Boxcar pre-amplifier

Another interesting concept, that can be exploited with the help of the pre-amplifier, is to use a boxcar pre-amplifier. This concept is based on using a linear transconductance amplifier as pre-amplifier[5][6], which then charges a capacitor with it's output current. This capacitor is then used in the switched capacitor DT integrator (see figure 3.2). Because this also produces a low pass response, it makes it possible to combine the inherent anti aliasing filter of the CT system with the precision advantages of using a DT integrator.

3.5. Front-end choice

After having discussed the different possibilities for a front end for the ADC, a front end can be picked. The first important choice is whether or not to use a continuous time or a discrete time integrator. Because a high precision system is designed, and techniques like auto-zeroing are significantly easier to use with discrete time integrators, the discrete time system will be used. Especially because a continuous time system is more error prone, and the limited design time was important factor. The choice for a discrete time system has the disadvantage that it will consume more power then a continuous time system, and it will need an anti-aliasing filter.

In order to keep the sampling capacitor size reasonable and the input impedance high, it was decided to make use of a pre-amplifier. This pre-amplifier will amplify the input voltage $20\times$ from $\pm 100\,\text{mV}$ to $\pm 2\,\text{V}$. To be able to fit within the available time, it is decided to keep the pre-amplifier out of the scope of this thesis. The decision is made to assume a band-limited, $2\,\text{V}$ peak to peak input voltage, and design the rest of the ADC based on this assumption. Further research can be done later to design a fitting (boxcar) pre-amplifier.

4

System level design

In this chapter the design of the charge and the current ADC's are discussed. Both converters are high precision converters with a low to medium bandwidth. As explained in the previous chapter, in such applications sigma delta converters are a good choice. In this chapter the design of the delta sigma converter and the rest of the system will be done on block level.

4.1. Sigma delta analog to digital conversion

Sigma delta ADC's are best known for their ability to achieve high resolution conversions of signals with a low to medium bandwidth.[4] Sigma delta ADC's are composed of two parts, the sigma delta modulator, and the decimation filter. The sigma delta modulator uses oversampling of a low resolution ADC/DAC combination in a feedback configuration with a loop filter to generate a modulated signal. In this signal, the quantization noise of the ADC/DAC is noise-shaped to high frequencies, away from the frequencies of interest. Subsequently, the high frequencies can be filtered out by the decimation filter, leaving a high resolution reproduction of the input signal. This decimation filter will consist of a digital low-pass filter, usually of the SINC^n family. The decimation filter will also down-sample the sample rate (f_s) to the conversion rate (f_c). A block diagram of a sigma delta ADC is shown in figure 4.1.

The ratio between the sampling rate of the sigma delta modulator and the conversion rate is called the oversampling ratio (OSR). Together with the order of the modulator the OSR determines how strongly the quantization noise is suppressed in the overall conversion. The order of the modulator is determined by the number of integrators in the loop-filter. Higher order modulators have higher noise-shaping capabilities, so the quantization noise can be filtered out more effectively.

In order to effectively filter out this more strongly shaped quantization noise, also a higher order decimation filter is needed. In general, the order of the decimation filter should match the order of the modulator. In this case the slope of shaped quantization noise will match the slope of the filter, resulting in a flat quantization noise over frequency.

The disadvantage of using a higher order sigma delta modulators is that these modulators are more sensitive to instability. To compensate for this effect, the ratio between the maximum input voltage

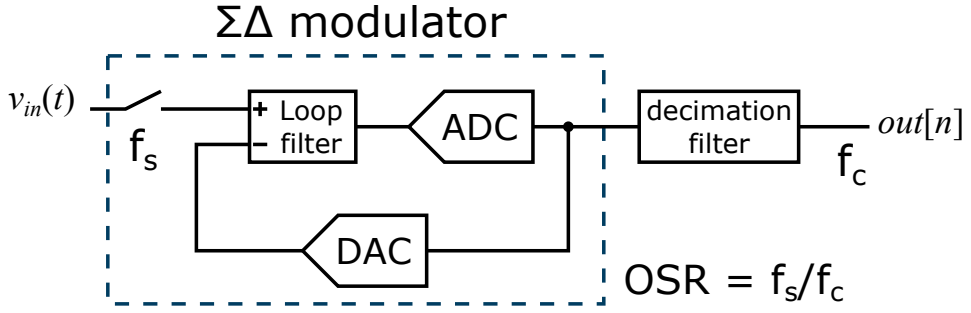

Figure 4.1: The generic block diagram of a discrete time sigma delta converter, with the two different sampling frequencies.

and the reference voltage is reduced.

Sigma delta modulation is based on negative feed-back, which means that as long as the forward gain of the loop is high and the loop is stable, the transfer function of the loop is determined by the feedback path. In the case of the sigma delta modulator, the high forward gain is assured by using integrators in the loop filter.

Because the transfer function of the system is determined by the feedback path, which exists of the internal DAC of the sigma delta modulator, this DAC determines the linearity of the modulator. The resolution of this DAC will need to match the resolution of the internal ADC that produces the quantization noise. While the quantization noise can be dealt with, the linearity of the internal DAC will directly determine the linearity of the system.

While multi-bit ADC's combinations produce less quantization noise, and thus require a lower OSR, realizing very linear multi-bit DAC's is not easy. While techniques exist for creating very linear multi-bit DAC's, most of them rely on complicated extra hardware or complicated switching schemes, such as dynamic element matching (DEM). Because of the added complexity of such techniques, a simpler design has been selected.

This is the use of a single bit ADC followed by a single bit DAC. This has the advantage that a single bit DAC is inherently linear due to having only two possible outputs[1]. Because of this, the overall converter will also be linear, as long as the forward gain of the loop is high enough.

The major disadvantage of using a single bit ADC, however, is that it has very high quantization noise. Which results in higher requirements for the OSR or modulator order in order to push the quantization noise back down to to a low enough level.

Because there are solutions to resolve this disadvantage (see section 4.3), and the specifications ask for a 20 bit linear ADC, a single bit ADC/DAC combination is the best option for this design.

4.1.1. Combining coulomb counting and current measurement

Because the final resolution of the sigma delta ADC is determined by the OSR, and the OSR is only determined at the decimation filter, it is possible to have a system which combines the two sigma delta ADC's into one, with a single sigma delta modulator, and two (digital) decimation filters. The first one needs to have a high OSR to be used for the high resolution, slow charge measurements. The second one needs to have a lower OSR and can be used for the faster, lower resolution current measurements.

Because the split between two converters happens so late in the signal chain and the analog part is the same for both converters, it is possible to first design the converter with the stricter requirements and later fit in the requirements of the second converter. As discussed before, the most important converter is the one used for the charge measurement, because it has the most strict offset and timing specifications. Consequently, in the following sections the main focus will lie on the design of the charge converter.

4.2. Sigma delta design for low offset

Because the result of the charge conversion will be integrated indefinitely, it is very important for the conversion to have extremely low offset, as any offset will result in a increasing error in the integral.

The easiest way to achieve very a low offset circuit is to make use of a switched capacitor circuit in combination with offset cancellation techniques like correlated double sampling (CDS) and block level chopping. These techniques alone are not enough to meet the extreme requirements targeted in this design (section 2.2.1). This is because of residual charge injection. By using system level chopping, however, a sufficiently low offset can be achieved.[7, section III.F]

System level chopping is the technique of chopping the complete circuit from input to output. This is done in a sigma delta converter by doing two complete conversions with opposite input polarity, while also reversing the input of the decimation filter. This is done by including choppers, one directly at the input of the system, and also around all state-holding circuit elements[8][9], like shown in figure 4.2.

4.2.1. State-holding elements

State-holding circuit elements are elements in the circuit that hold a state. In case of the sigma delta ADC, these elements are the capacitors that hold the analog values of the integrators and the digital

[1] The line between two point's is always perfectly linear

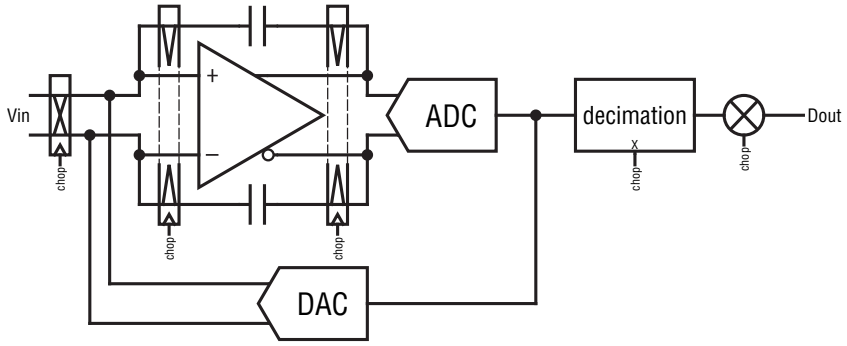

Figure 4.2: An example of system level chopping where all state holding elements are chopped.

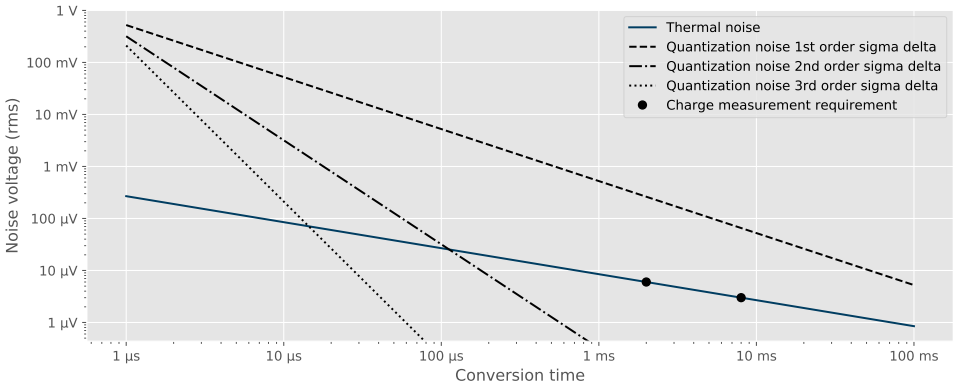

Figure 4.3: The quantization and thermal noise of a n^{th} order sigma delta ADC with an assumed f_s of 4 MHz.

state of the decimation filter. It is not possible to flip the integrator as a whole circuit element, because by doing so the offset of the integrator itself is not canceled. Because of this, only switches around the state holding elements themselves should be added.

Adding these extra switches in the signal path has the unavoidable side-effect that the two conversions, that should be identical (except for their sign), will have slightly different signal paths. The effect of this is that adding more choppers will increase the residual offset. This effect, combined with the fact that flipping around all state holding capacitors will lead to very complex hardware, drives the search for a different solution to implement system level chopping.

4.2.2. Incremental conversions

This different solution has been found by resetting all state-holding elements in the converter between input windows, and completely decouple the different input windows from each other. This mode of operation, also known as incremental conversion[10][11][12], ensures a one-to-one input-output relationship between the input windows and the output samples. This means that every input sample only ever effects one output sample. Because the input windows are now completely independent, converting two input windows while flipping the input chopper in between, will result in two output samples which can be subtracted to produce a single output value with a offset only limited by the resolution of the converter[7, section III.F]. When using incremental conversions, however, equally weighting samples gets more challenging.

4.2.3. Correct integration

In order for the charge measurement to be useful, it should perform correct integration of the input samples (section 2.2.3), meaning all input samples have to be equally weighted. This can be achieved by a very slow anti aliasing filter, by filtering all frequencies higher then half the output sample rate. This filter will, however, interfere with the current measurements, which need to be done at a higher bandwidth.

Another option is to only filter out frequencies higher then half the sampling frequency, and *equally*

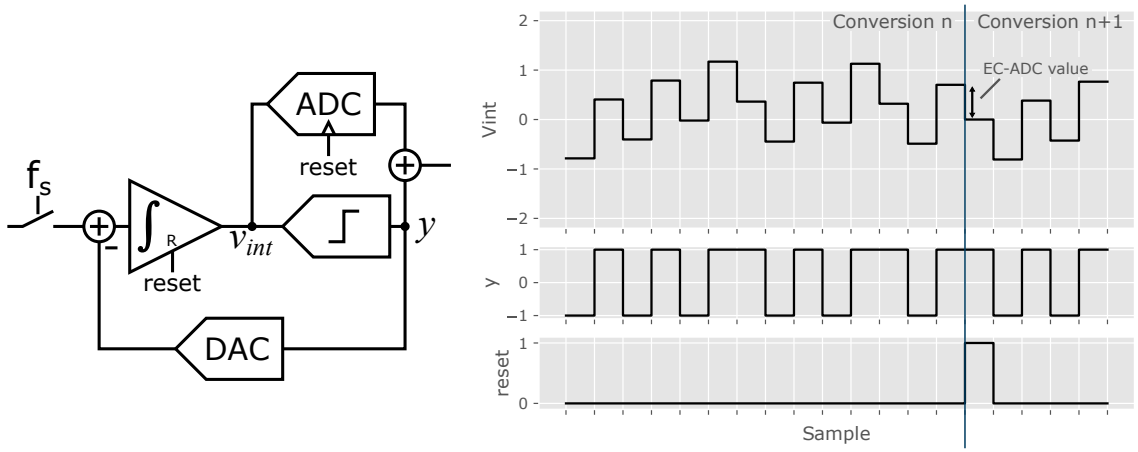

(a) Schematic (b) Wave forms

Figure 4.4: Discrete time 1st order sigma delta modulator with Extended Counting

weighting all samples in constructing the output value. In this way, the integration is partially done by the converter.

By looking at figure 4.3, it can be concluded that a 2nd order sigma delta ADC is necessary, in order reach a thermal noise limited ADC with low enough quantization noise (see table 2.1). Equally weighting samples while doing non-stop incremental conversions however, can only be done by using a first order sigma delta ADC. The reason for this is that higher order sigma delta modulators are only useful in combination with higher order decimation filters. And these higher order decimation filters don't have flat (time domain) responses (like a SINC^1 filter has). This in itself is not a problem, because as long as the higher order SINC filters partially overlap, the overall response will be flat. *When using incremental conversions however, it is not possible to partially overlap the samples just before and just after the reset, because they do not share the same integrator states.*

So, a incremental sigma delta ADC, with a input bandwidth of half the sampling frequency, can only perform correct integration of the input signal, when using a first order design. This means that 2nd order modulator cannot be used in this application, and so, as visible from figure 4.3, another technique is needed to reach the quantization noise level required.

4.3. Extended counting

Because a lower level of quantization noise is required, than a first order sigma delta produces, another technique is needed. Looking at the existing circuit, it can be noted that the quantization noise produced by the system is continuously integrated in the integrator. This means that the value in the integrator at the end of a conversion is exactly the quantization error made by the sigma delta ADC!

Measuring this integrator value, and using it to correct the output of the sigma delta is known as extended counting (EC)[13][14]. The measurement of the final integrator value just before reset is done by using a secondary ADC (see figure 4.4). With this value measured, it is possible to subtract it from the final result. Accordingly, the quantization noise of the converter can be improved by the difference of resolution between the secondary ADC and the quantizer in the sigma delta modulator.

As long as the converter is not yet thermal noise limited, extended counting can thus be used to improve the resolution of the ADC. The effect of adding the extended counting to a first order sigma delta ADC can be seen in figure 4.5.

It should be noted that this technique gives similar improvements to using a multi-bit ADC/DAC combination inside of the sigma-delta modulator, without the challenges of building a linear multi-bit DAC.

In some cases, it is beneficial to combine the circuitry used for the sigma delta modulator for the extended counting ADC[15][16]. This technique cannot be used in this case, however, as the charge measurement must be done continuous. Due to this, the sigma delta circuitry cannot be time multiplexed for secondary use as a extended counting ADC.

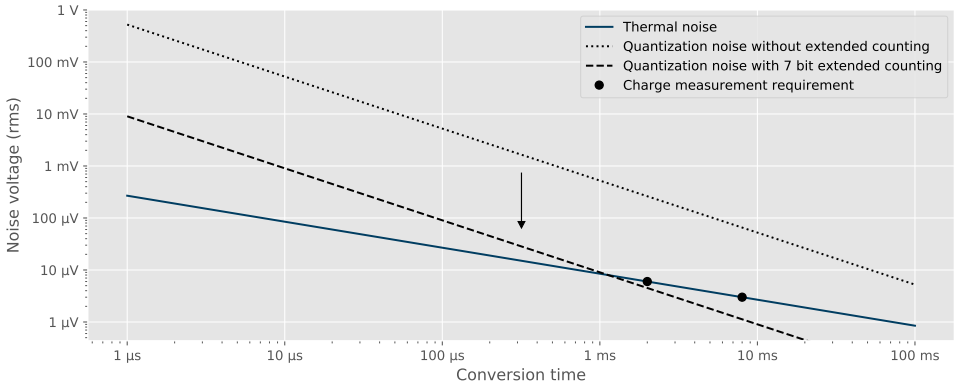

Figure 4.5: The quantization and thermal noise of sigma delta ADC using extended counting.

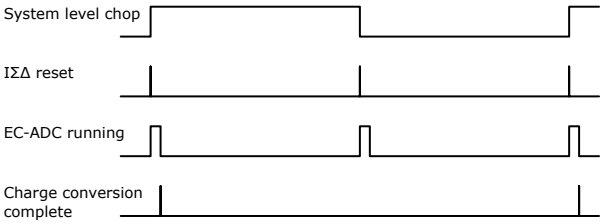

Figure 4.6: The timing diagram of a incremental converter using extended counting.

4.3.1. Timing

With the introduction of a first order sigma delta modulator, system level chopping and extended counting, it becomes possible to define a timing scheme. The converter runs two consecutive conversions with flipped input polarity to implement system level chopping. Between every conversion, the EC-ADC measures the integrator voltage and the integrator is reset. After these two conversions, the results are subtracted, and the charge measurement is complete. A simple timing diagram of these actions can be found in figure 4.6.

4.3.2. Gain matching

When using extended counting, like shown in figure 4.4a, the output of the total converter will be given by equation (4.1). Here is $A_{\text{EC-ADC}}$ the relative gain of the EC-ADC. In order to prevent steps in the transfer function around the hand-off of the sigma delta ADC and the EC-ADC, this gain should be well defined. Because $A_{\text{EC-ADC}}$ might be process dependent (depends on EC-ADC implementation), there can be variations in this factor. because of these variations, a gain trim should be implemented to match the sigma delta ADC and the EC-ADC.

In order to implement a gain trim, it should be taken into account that the resolution of the EC-ADC will also be affected by the gain trim. In order to compensate this, the resolution of the EC-ADC should be around 0.5 bit higher then required by the specifications.

$$\text{ADC}(v_{in}) = \sum_{i=-1}^{\text{OSR}-1} y[i] + A_{\text{EC-ADC}} \cdot \text{EC-ADC}(v_{int}[\text{OSR}]) \tag{4.1}$$

4.3.3. Resolution requirements secondary ADC

In order to find the requirements for the secondary ADC that implements the extended counting, the resolution of the sigma delta ADC needs to be known. This is the case because the EC-ADC essentially provides the extra bits of resolution that the sigma delta ADC 'misses'.

To calculate what the resolution of the EC-ADC should be with a given OSR, a few effects need to be considered:

1. When using a EC-ADC with the same resolution as the ADC/DAC in the sigma delta ADC, no

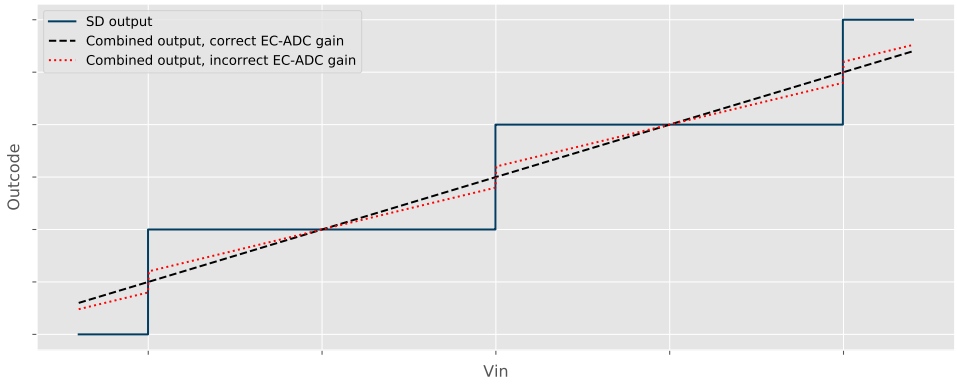

Figure 4.7: The hand-off between the sigma delta LSB and the EC-ADC

improvement will be made, only extra bits help. Because the ADC/DAC combination of the sigma delta modulator is 1b, 1b of the EC-ADC is used to match this.

2. Every charge conversion requires 2 conversions because of system level chopping (see section 4.3.1). Because of this, the EC-ADC needs another 0.5 bit of resolution.

3. Because of the trimming necessity of the gain matching (see section 4.3.2) another 0.5 bits is needed.

In total the resolution of the EC-ADC required by the charge conversion is given by equation (4.2).

$$\text{resolution}_{\text{charge ADC}} = \text{resolution}_{\text{sigma delta}} + \text{resolution}_{\text{EC-ADC}} - 2 \qquad [bit] \quad (4.2)$$

4.4. Adding the current measurement

With the addition of the EC-ADC, a feasible system for the charge measurement is conceived. This can be seen as taking a single point measurement on the combined thermal-noise and quantization noise line from figure 4.5. The next step in the design of the multi-path ADC is to add in the current conversion. This is equivalent to making a second single point measurement on the same line simultaneously. Theoretically this should be possible just by adding a second decimation filter with a different OSR. The difficult question however, is how to do this in a system with extended counting and system level copping.

The first step is to match the extended counting technique to the current measurement, where the OSR is reduced. Additional to the reduced OSR, there is another difference with the charge measurement: in the case of the current measurement, the state of the integrator is unknown at the start and the end of the measurement, because the fast measurement starts in the middle of a charge measurement. To resolve this, the EC-ADC is run twice, once at the start, and once at the end of the current conversion. The transfer of this current conversion started at sample n is now given by equation (4.3).

$$\text{I-ADC}(v_{in}, n) = \sum_{i=n}^{n+\text{I-OSR}} y[i] + A_{\text{EC-ADC}}\big[\text{EC-ADC}(v_{int}[n + \text{I-OSR} + 1]) - \text{EC-ADC}(v_{int}[n])\big] \qquad (4.3)$$

4.4.1. Combined timing

In order to assure that the current measurement does not interfere with the charge measurement, but still can start at any moment in the charge conversion cycle, a more complicated timing scheme is needed.

Starting from figure 4.6, adding the fast current measurement seems quite simple: add a EC-ADC measurement to the begin and end of the current measurement. But because the current conversion is asynchronous with the charge conversion, there are four possible scenarios for the current measurement (also shown graphically in figures 4.8a to 4.8d):

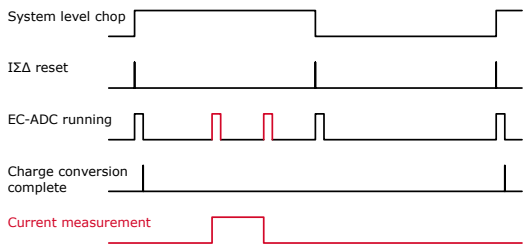

(a) The timing of a current measurement (case 1)

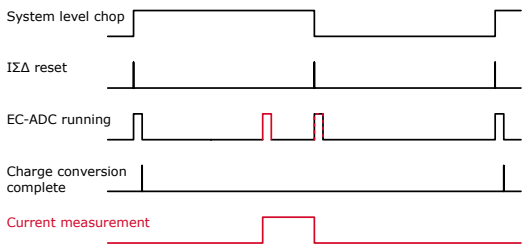

(b) The timing of a current measurement (case 2)

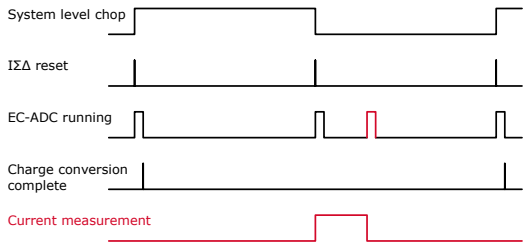

(c) The timing of a current measurement (case 3)

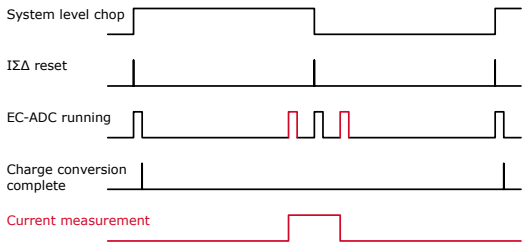

(d) The timing of a current measurement (case 4)

Figure 4.8: Timing diagrams showing the possible different cases for the current measurement

1. The current conversion occurs in the middle of a charge conversion.
2. The end of the current conversion overlaps with the reset of the integrator.
3. The current conversion starts just after the reset of the integrator.
4. The current conversion overlaps with the reset of the integrator.

In the first case the only thing that needs to happen is to add the EC-ADC measurements at the beginning and the end of the current conversion. In the second case, the second EC-ADC measurement overlaps with the measurement that was going to be done for the charge measurement. Because there is nothing different between the two measurements, the result can just be shared. The third case seems the easiest case: the integrator starts from a known (reset) state, so only the EC-ADC measurement at the end of the conversion is needed. While this is correct, any offset in the EC-ADC will not be canceled in this case, to fix this, a shorted input measurement of the EC-ADC is subtracted from the result. The forth case is quite clearly the most complicated case. In order to correctly measure the current when a integrator reset is taking place, the conversion is broken up into two pieces. The first part of the conversion is done as in case 2, and the second part of the conversion is done as case 3. Because of this, there will be 4 EC-ADC [2] conversions done in this case, this means the quantization noise of this case will be $\sqrt{2}\times$ higher than in the other three cases.

4.5. Other considerations

There are a couple of other considerations to look into.

4.5.1. Current conversion quantization noise

After adding the current conversion to the ADC, its resolution will be quantization noise limited. This means that it is relatively power efficient to increase the resolution. In this case, this even needs to be done in order to meet the requirements from table 2.2. To decrease the quantization noise the required amount, the resolution of the EC-ADC is increased. This is shown in figure 4.9.

4.5.2. Digital decimation filter and gain matching

In order to calculate the out-code for a conversion, the gain of the EC-ADC needs to be matched to the gain of the sigma delta ADC to avoid DNL errors (see section 4.3.2). In order to match the two results, there are two possibilities.

[2] the 3 from figure 4.8d plus a shorted input measurement

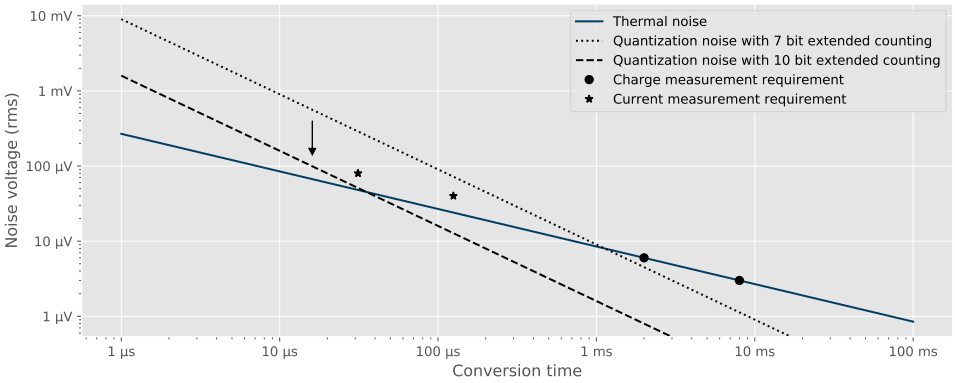

Figure 4.9: The quantization and thermal noise of sigma delta ADC, including the current conversion specifications

1. Adjust the gain of the EC-ADC to the gain of the sigma delta ADC by multiplying the result of the EC-ADC by $A_{\text{EC-ADC}}$. This approach is taken in equations (4.1) and (4.3).

2. Adjust the gain of the sigma delta ADC to the gain of the EC-ADC by changing the weights of the bit-stream from 1 to $\frac{1}{A_{\text{EC-ADC}}}$.

Option 1 requires a (expensive) digital multiplier, while option 2 only needs an adder. But, option 2 adjusts the gain of the complete system (which is very stable due to the 1 bit DAC) to the gain of the EC-ADC.

In the intended application, however, the gain of the total system will need to be trimmed anyway, because the value of the sense resistor can vary. This trim will be done by adjusting the reference voltage going to the ADC. This affects both the sigma-delta ADC as the EC-ADC. When this is the case, it is not a problem anymore to adjust the gain of the system to the gain of the EC-ADC. Because of this, option 2 is the better option for this design. Equation (4.4) shows the transfer of the ADC using option 2.

$$\text{ADC}'(v_{in}) = \sum_{i=-1}^{\text{OSR}-1} \left(\frac{1}{A_{\text{EC-ADC}}} \cdot y[i] \right) + \text{EC-ADC}(v_{int}[\text{OSR}]) \tag{4.4}$$

4.5.3. Speed

The speed of the EC-ADC can be split into two independent requirements: the sampling speed and the conversion speed.

The requirement for the sampling speed is set by the speed of the sigma delta modulator. This is because the EC-ADC needs to sample the integrator signal of the sigma delta modulator, at one specific sample. In order for this sampling to be correct, the sampling speed of the EC-ADC should match the sampling speed of the sigma delta loop.

The conversion time of the EC-ADC however, determines the maximum delay time between the external trigger and the start of the current conversion. This is because the EC-ADC can not run two conversions at the same time. This means that the current conversion cannot be started if it would cause two EC-ADC conversions of the charge and current conversion to partially overlap. If they fully overlap, there is no problem, as the conversion result can be shared (case 2 in section 4.4.1). Because all EC-ADC conversions for the charge conversion are done according to a pre-determined schedule, it is possible to schedule all current conversions with a maximum delay of $T_{conv,\text{EC-ADC}}$ without any EC-ADC collisions. As the maximum variation in conversion start is 1 µs (see table 2.2), the maximum conversion time of the EC-ADC is 1 µs.

4.5.4. Integrator leakage

One of the main weaknesses of first order sigma delta ADC's is their sensitivity to integrator leakage, especially when using a single bit DAC. This is caused by the fact that in a first order sigma delta ADC with a single bit DAC, the quantization noise cannot be seen as random, because it is heavily correlated

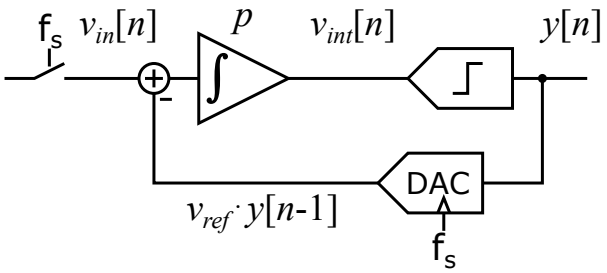

Figure 4.10: The sigma delta modulator as described by equations (4.5) and (4.6)

with the signal. With integrator leakage, this correlation can cause so called limit cycles. Limit cycles are cycles in which a relatively short pattern in the bit-stream is repeated continuously, caused by (near) integer relations between the input voltage and the reference voltage.

In other applications (for example audio), the main problem of these limit cycles is that they cause spurious tones that must be suppressed by the decimation filter. In this application, however, spurious tones are not a problem, because the signal gets integrated, and the integral of the signal stays the same, regardless of the spurious tones.

Limit cycles do have another effect, however: because the same signal continuously repeats, and the quantization noise is correlated with the input signal, the quantization noise in the loop is no longer zero on average. While this quantization error slowly accumulates, it leaks away through the integrator, which causes dead-zones. This effect can been seen from equation (4.5), the main equation of the 1st order sigma delta modulator. In this equation v_{in} is the modulator input voltage, v_{int} is the integrator output voltage, $y[n]$ is the bit stream at sample n and p is the leakage factor of the integrator (see figure 4.10).

$$v_{int}[n] = \sum_{i=0}^{n} \left(v_{in}[i] - v_{ref} \cdot y[i-1] - p \cdot v_{int}[i-1] \right) \tag{4.5}$$

When rewritten as equation (4.7), it is clear how the modulator works as converter. As long as $v_{int}[n]$ is bounded (which it is for a stable converter), and p is very small (near perfect integrator), the sum of the bit stream is a good approximation of the sum of the input voltage. It is also important to note the value's $y[-1]$ and $v_{int}[-1]$ are the reset values of y and v_{int}.

$$v_{ref} \sum_{i=-1}^{n-1} y[i] = \sum_{i=0}^{n} v_{in}[i] - p \sum_{i=-1}^{n-1} v_{int}[i] - v_{int}[n] \tag{4.6}$$

$$\underbrace{v_{ref} \sum_{i=-1}^{\mathrm{OSR}-1} y[i]}_{\text{SD result}} + \underbrace{v_{int}[\mathrm{OSR}]}_{\text{EC result}} = \overline{v_{in}} - p \cdot \overline{v_{int}} \tag{4.7}$$

If the leakage factor of the integrator is not zero, the output of the integrator will be a term in the equation. As long as the modulator is not in a limit cycle, the output of the integrator will be the input signal plus a noise shaped high frequency signal. This means the mean value of the signal is proportional to the input voltage. Because of this, leakage will cause a gain error as long as the modulator does not hit any limit cycles. When the modulator approaches a limit cycle, however, there can be a low frequency component in v_{int}, as shown in figure 4.11. If the leakage of the integrator is bigger than this low frequency component, these components will disappear, and the modulator will 'snap' to simple fractions ($\frac{1}{2}$, $\frac{1}{4}$, etc.) of the full scale, causing the limit cycle. Because the output of the modulator will snap to these fractions as well, it causes non-linearity's known as dead-zones around simple fractions of the full scale.

To prevent this issue, the leakage of the integrator should be small. The maximum leakage is determined by the maximum permitted dead-zone. For the case of a 20b ADC (table 2.1) without a visible dead-zone, the maximum permissible leakage is calculated by equations (4.8) to (4.10). The

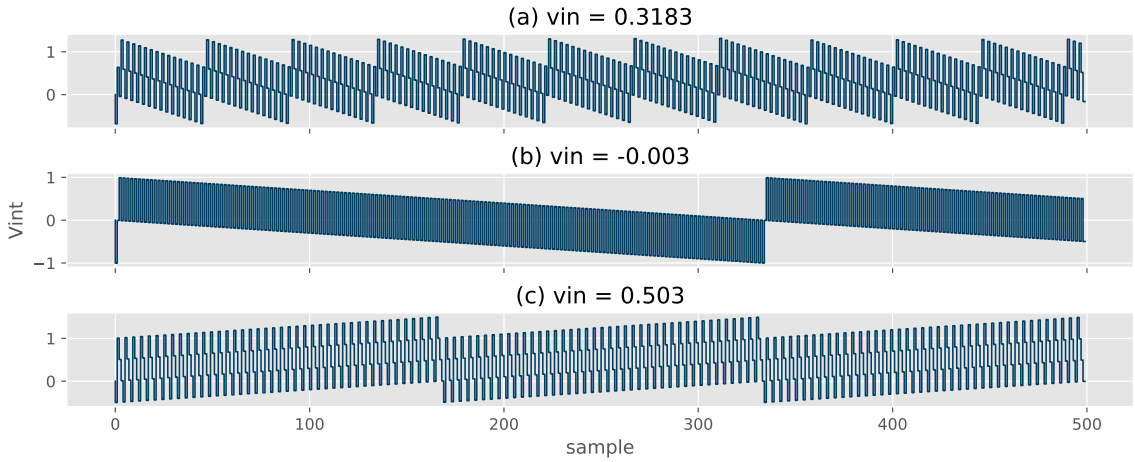

Figure 4.11: The effect of limit cycles on the integrator voltage.

Table 4.1: A table comparing the different possible sampling frequencies and their corresponding OSR, minimum resolution EC-ADC for the Charge ADC and the Current ADC

	Charge ADC		Current ADC	
Fs [*MHz*]	OSR	EC-ADC [*bit*]	OSR	EC-ADC [*bit*]
0.256	2^{11}	11	2^{5}	13.5
0.512	2^{12}	10	2^{6}	12.5
1.02	2^{13}	9	2^{7}	11.5
2.05	2^{14}	8	2^{8}	10.5
4.10	2^{15}	7	2^{9}	9.5
8.19	2^{16}	6	2^{10}	8.5
16.4	2^{17}	5	2^{11}	7.5
32.8	2^{18}	4	2^{12}	6.5
65.5	2^{19}	3	2^{13}	5.5

maximum value of e_{mean}, the difference between $\overline{v_{in}}$ and $\overline{v_{int}}$ is $\frac{1}{2}$ of the input full scale. This is the case at 0 input, when the v_{int} swings between plus and minus v_{ref}[3].

$$e_{leak} = p \cdot \overline{v_{int}} = p \cdot (\overbrace{\overline{v_{in}}}^{\text{gain error}} + \overbrace{e_{mean}}^{\text{non-linear}}) \leq p \cdot (\overline{v_{in}} \pm \tfrac{1}{2} v_{fs}) \tag{4.8}$$

$$e_{leak,non-linear} = p \cdot \tfrac{1}{2} v_{fs} \leq \tfrac{1}{2}\text{LSB} = \tfrac{1}{2} \cdot 2^{-20} \cdot v_{fs} \tag{4.9}$$

$$p \leq 2^{-20} \tag{4.10}$$

4.6. Oversampling rate

With the architecture of the converter determined, the oversampling rate (OSR) can be determined. The OSR is the ratio between the sampling rate of the sigma delta modulator, and the output sample rate. Because the output sample rate is fixed by the specifications, the OSR is completely determined by the sampling rate of the sigma delta.

A main factor in determining the OSR is the distribution of quantization noise between the sigma delta converter and the EC-ADC. In order to determine the optimal OSR, different sampling rates are compared. With the information from tables 2.1 and 2.2, it is possible to generate table 4.1.

[3] All these calculations are done with an integrator and a DAC gain of 1, this means the input full scale $v_{fs} = v_{ref}$. The conclusion of these equations is, however, independent of these gains.

From table 4.1 it is now possible to select an appropriate sampling frequency. In order to choose between the different versions, a few selection criteria are chosen.

1. Because a very high speed design is not desired, the maximum sampling frequency is chosen to be 10 MHz.
2. In order to not go to extremes for the EC-ADC, the maximum EC-ADC resolution is chosen to be 10 bit, as it is relatively straightforward to achieve this level of linearity by only relying on capacitor matching, and not on trimming, calibration or other error correction schemes.

After applying these selection criteria, two possible sampling frequencies remain: 4.1 MHz and 8.2 MHz. Initially, a sampling rate of 8.2 MHz was chosen, but due to process options, the final sampling rate of the system is 4.1 MHz (for more information see appendix B.2). For the EC-ADC, a resolution of 10 bit is chosen.

5

Circuit level Design

In this chapter, a circuit level design, based on the system level design in chapter 4 and the specifications from chapter 2, will be made. First, the sigma delta modulator will be designed, after which the EC-ADC is designed, and finally some timing generation circuitry is discussed.

5.1. Sigma delta modulator

The block level of the sigma delta modulator as described in chapter 4, will be implemented using a switched capacitor integration scheme. The goal of this section is to design a delta sigma modulator which is as small and power efficient as possible, while still meeting the noise, leakage and offset requirements. The first thing to look at is the switching scheme of the integrator, as this determines many of the other factors of the design.

5.1.1. Integrator switching scheme

One of the most strict requirements that need to be fulfilled is the leakage requirement of the integrator calculated in section 4.5.4, equation (4.10). Consider the case of a typical switched-capacitor integrator, as depicted in figure 5.1a. In this system, the DC gain of the amplifier used needs to be quite high, as shown by equation (5.3), in which G is the gain of the integrator[1], A is the DC gain of the amplifier used to implement the integrator and V_{os} is the offset voltage of the amplifier.

$$V_{int}[n] \approx G\left(-V_{os} + V_{in}[n] + \frac{V_{out}[n-1]}{A}\right) + V_{int}[n-1] \qquad \text{(App. C.2)} \quad (5.1)$$

$$p \approx \frac{G}{A} \qquad (A \gg 1) \quad (5.2)$$

$$A \geq \frac{G}{p} = \frac{G}{2^{-20}} = G_{dB} + 120.4\,dB \qquad (5.3)$$

The first issue with this basic switching scheme is that the offset voltage of the amplifier is directly integrated. And while system level chopping is implemented for the charge measurement, it is not helping to reduce the offset in the current measurement. In order to prevent this, a auto zeroing switching scheme can be used, as shown in figure 5.1b. In this switching scheme, the sampling capacitor samples the offset voltage in both phases, and in this manner, prevent the offset from getting integrated.

Another issue is that with reasonable values for G (−15 - 0dB), a DC gain of at least 105 dB is needed. To get to this level of DC gain, either a two stage amplifier or gain-boosting is required. While this is not a big issue, a solution is available.

To improve the leakage characteristics of the integrator, the scheme shown in figure 5.1c is used. This is a variation on the finite gain compensation scheme in [17, figures 30 and 31]. The main advantage of using this switching scheme is that the leakage of the integrator is now proportional to the squared gain of the op-amp (appendix C.5).

[1] $V_{int}[n] = V_{int}[n-1] + G \cdot V_{in}[n]$

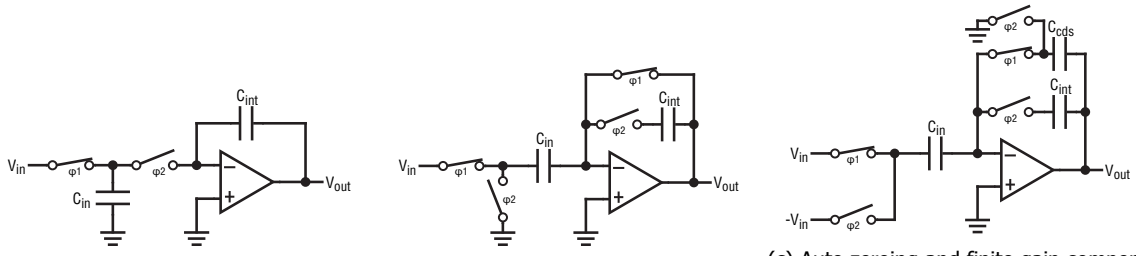

(a) Simple switching scheme (b) Simple auto zeroing switching scheme (c) Auto zeroing and finite gain compensation with double sampling

Figure 5.1: Switched capacitor integrators

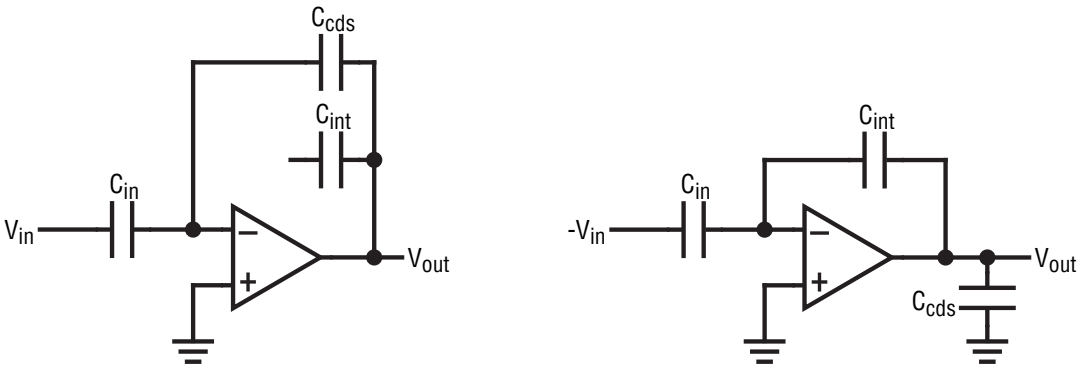

(a) Phase 1 (b) Phase 2

Figure 5.2: The two phases of the finite gain compensating integrator.

Finite gain compensation scheme

The working of the finite gain compensation scheme from figure 5.1c is based on first stepping the output of the integrator in the opposite direction using C_{cds} in ϕ_1, and sampling the step at the virtual ground produced by this output step. After this, in ϕ_2, C_{int} is connected again, and this time, the output steps in the correct direction. Because the step produced at the virtual ground is already added to the signal in ϕ_1, the overall effect is that the error produced by the finite gain of the amplifier is suppressed by another factor A. This can also been seen in mathematical form in equations (5.4) to (5.28).

ϕ_0 is defined as the final state of the previous ϕ_2 phase:

$$V_{vg}[\phi_0] = -\frac{V_{out}[\phi_0]}{A} + V_{os} \tag{5.4}$$

$$V_{C_{in}}[\phi_0] = V_{vg}[\phi_0] + V_{in}[\phi_0] \tag{5.5}$$

$$V_{C_{int}}[\phi_0] = V_{vg}[\phi_0] + V_{in}[\phi_0] \tag{5.6}$$

$$V_{C_{cds}}[\phi_0] = V_{out}[\phi_0] \tag{5.7}$$

Now the equations for the transition to ϕ_1 can be set up (see figure 5.2a):

$$V_{vg}[\phi_1] = -\frac{V_{out}[\phi_1]}{A} + V_{os} \tag{5.8}$$

$$V_{C_{in}}[\phi_1] = V_{vg}[\phi_1] - V_{in}[\phi_1] \tag{5.9}$$

$$V_{C_{int}}[\phi_1] = V_{C_{int}}[\phi_0] \tag{5.10}$$

$$V_{C_{cds}}[\phi_1] = V_{out}[\phi_1] - V_{vg}[\phi_1] \tag{5.11}$$

$$C_{in}(V_{C_{in}}[\phi_1] - V_{C_{in}}[\phi_0]) = C_{cds}(V_{C_{cds}}[\phi_1] - V_{C_{cds}}[\phi_0]) \tag{5.12}$$

Solving for $V_{out}[\phi_1]$ gives the output voltage at the end of ϕ_1:

$$V_{out}[\phi_1] = \frac{A\big(C_{cds}(V_{os} + V_{out}[\phi_0]) - C_{in}(V_{in}[\phi_0] + V_{in}[\phi_1])\big) + C_{in}V_{out}[\phi_0]}{C_{cds}(A+1) + C_{in}} \tag{5.13}$$

$$V_{out}[\phi_1] \approx V_{out}[\phi_0] + V_{os} - \frac{C_{in}}{C_{cds}}(V_{in}[\phi_0] + V_{in}[\phi_1]) + \frac{C_{in}}{C_{cds}A}V_{out}[\phi_0] \qquad (A \gg 1) \tag{5.14}$$

As the output voltage changes due to the sampling of the input voltage on C_{in}, two things happen in equation (5.14): first, the offset voltage is added to the output, such that it will not be sampled when C_{int} is connected; and second, the leakage from ϕ_0 due to the difference in output voltage is added to the output. This means that when C_{int} is connected again, this leakage will be sampled by C_{int}, but because the step in ϕ_1 is of opposite polarity as in ϕ_2, the leakage component will be canceled out.

Similarly ϕ_2 looks as follows (see figure 5.2b):

$$V_{vg}[\phi_2] = -\frac{V_{out}[\phi_2]}{A} + V_{os} \tag{5.15}$$

$$V_{C_{in}}[\phi_2] = V_{vg}[\phi_2] + V_{in}[\phi_2] \tag{5.16}$$

$$V_{C_{int}}[\phi_2] = V_{out}[\phi_2] - V_{vg}[\phi_2] \tag{5.17}$$

$$V_{C_{cds}}[\phi_2] = V_{out}[\phi_2] \tag{5.18}$$

$$C_{in}(V_{C_{in}}[\phi_2] - V_{C_{in}}[\phi_1]) = C_{int}(V_{C_{int}}[\phi_2] - V_{C_{int}}[\phi_1]) \tag{5.19}$$

Next, the equations are rewritten in terms of the samples $n-1$ and n:

$$V_{C_{int}}[\phi_0] = V_{int}[n-1] \tag{5.20}$$

$$V_{C_{int}}[\phi_2] = V_{int}[n] \tag{5.21}$$

$$V_{in}[\phi_0] = V_{in}[n-1] \tag{5.22}$$

$$V_{in}[\phi_1] = V_{in}[\phi_2] = V_{in}[n] \tag{5.23}$$

Now the integrated voltage can be calculated (see also appendix C.5):

$$V_{int}[n] = G'\left(V_{in}[n] + \frac{1}{A}\left(\frac{V_{os}}{2} + V_{in}[n] - \frac{C_{in}V_{in}[n-1]}{2C_{cds}} + \frac{C_{in}V_{in}[n]}{2C_{cds}}\right) - \frac{V_{out}[n-1]}{2A^2}\right) + V_{int}[n-1] \tag{5.24}$$

$$G' \approx 2\frac{C_{in}}{C_{int}} \qquad (A \gg 1) \tag{5.25}$$

$$p \approx \frac{G'}{2A^2} \qquad (A \gg 1) \tag{5.26}$$

$$A \geq \sqrt{\frac{G'}{2p}} = \sqrt{\frac{G'}{2 \cdot 2^{-20}}} = \tfrac{1}{2}G'_{dB} + 57.2\,dB \tag{5.27}$$

$$V_{os} = \frac{V_{os,amp}}{2A} \tag{5.28}$$

From the resulting equation (5.24) is is visible that indeed the leakage ($V_{out}[n-1]$) is suppressed by a factor $\frac{1}{A^2}$. Note that these equations assume A is constant over both phases. The consequences of this assumption will be discussed further in section 7.4.2.

Double sampling

Another technique that is added to this switching scheme is double sampling[18]. This technique works by connecting the input capacitor to $-V_{in}$ instead of to signal ground in ϕ_2. This is easy to accomplish in a fully differential implementation. The effect of this double sampling is that the signal charge going into the integrator is doubled for the same amount of noise. This technique thus improves the signal to noise ratio for a given C_{in} (or gives a smaller C_{in} for a given SNR).

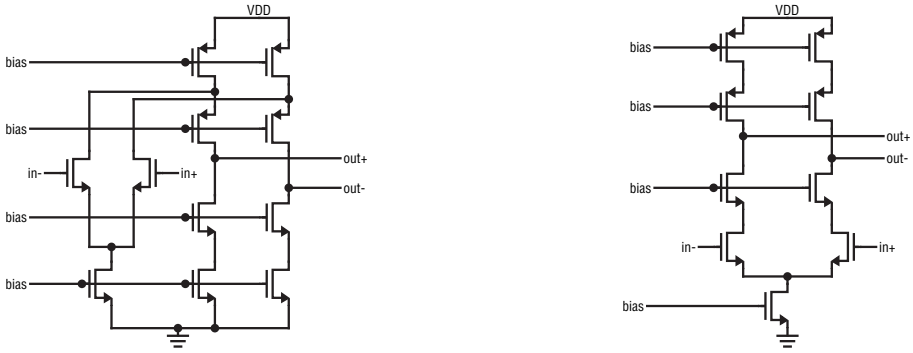

(a) Folded Cascode

(b) Telescopic

Figure 5.3: Two common OTA architectures

5.1.2. Thermal noise

With the OSR and the switching scheme determined, it is possible to calculate the capacitor sizes needed to achieve low enough thermal noise levels. A main factor in determining the capacitor size requirements is the OSR. This is important because bigger OSR's reduce the thermal noise by integrating. Taking this into account, the thermal noise produced by the switching scheme shown from section 5.1.1, can be estimated.

For a switched capacitor circuit with a single stage amplifier, the total input referred noise can be estimated using a factor of $\frac{kT}{C}$[19].

To determine this factor, figure 5.1c will be analyzed. Every time the input capacitor is sampled, the DAC feedback capacitor is also sampled. Additionally, because this is a fully differential circuit, there are two capacitors that will both be sampled. In total this means there are 4 capacitors that will be sampled every phase.

Because there are also switches and an OTA in the loop, the noise will also be higher. For a first order estimate, a factor of 4 will also be taken for the OTA and switch noise. The complete switched capacitor noise for this circuit can thus be estimated by a factor of 16 for preliminary calculations, which leads to equation (5.29). With this information, it is possible to calculate the needed capacitor sizes if $\frac{3}{4}$ of the total noise budget is used for the input thermal noise. Not the full budget is available, as there will also still be some quantization noise in the final result.

$$v_{n,rms,in}^2 = \frac{16 \cdot kT}{C_{in} \cdot \mathrm{OSR}} \tag{5.29}$$

$$C_{in} = \frac{16 \cdot kT}{v_{n,rms,in}^2 \cdot \mathrm{OSR}} \tag{5.30}$$

$$C_{in} = \frac{16 \cdot 1.38 \cdot 10^{-23} \cdot 300}{(3.0 \cdot 10^{-6})^2 \cdot \frac{3}{4} \cdot 2^{16}} = 149.8 \quad [fF] \tag{5.31}$$

5.1.3. OTA design

Because of the finite gain compensating switching scheme, a one stage OTA without gain-boosters is sufficient to reach the DC gain required by the integrator. The OTA design is based on output swing, power consumption, speed and noise.

Architecture

In order to choose the best OTA architecture for the integrator, the advantages and disadvantages of some common architectures should be considered.

Because a single stage OTA is preferred due to its simplicity and noise behaviour, the two architectures that are considered are a folded cascode OTA, and a telescopic OTA (see figure 5.3).

The main advantage of the folded cascode architecture is it's greater output swing, while the telescopic architecture is more power efficient.

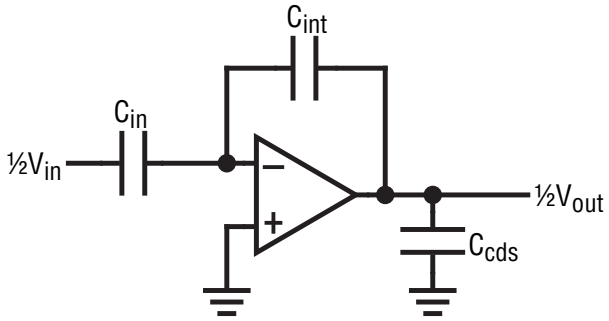

Figure 5.4: A single-ended view of the OTA in phase 2

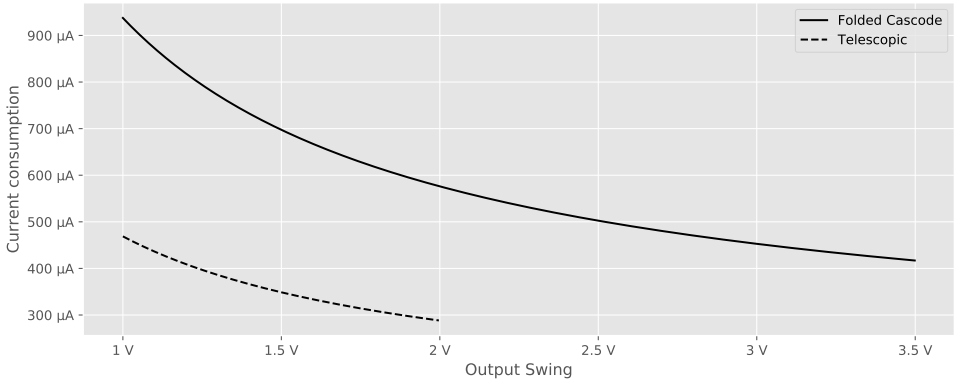

Figure 5.5: The required supply current needed to settle in time with different output swings for a telescopic and a folded cascode amplifier.

Because the size of the input capacitor is fixed by noise, the size of the integration capacitor is decided by the output swing of the OTA. Additionally, when using a bigger integration capacitor, a greater transconductance is necessary to meet the speed requirements.

For the OTA to work correctly in the discrete time system, it should settle completely within each phase. In the used switching scheme, the most demanding phase is phase 2, as in phase 2 C_{cds} acts as a load capacitor. To calculate the more efficient architecture, the performance of the architectures is compared during this phase of the switching scheme. For simplicity, the single-ended version of the switching scheme in phase 2 is shown in figure 5.4.

First, the capacitor sizes have to be determined. While the absolute sizes of the capacitors is determined by noise, the relative values determine the behaviour of the system. The gain of the system is determined by the ratio of the input capacitor and the integration capacitor. The gain of the integrator determines the output swing required by the amplifier. To minimize capacitor size, the gain will be chosen as such to fully use the output swing of the amplifier.

While size of the feedback capacitor C_{cds} is not important for the transfer function of the integrator, it is still determines the output swing of the amplifier in ϕ_1. C_{int} is already chosen to fully use the output swing, so $C_{cds} \geq C_{int}$. As C_{cds} forms the load capacitance in ϕ_2, it is desirable to choose as small as possible. The optimal value for the feedback capacitor is thus to have the same value as the integration capacitor.

Because the load capacitance (and thus the required transconductance of the amplifier) is dependent on the output swing, it is not trivial to determine if the telescopic or folded cascode architecture is more efficient. Thus, in order to determine the most efficient amplifier, a simple numerical calculation is done. From this calculation[2], visualized in figure 5.5, it is visible that, with similar headroom, the telescopic architecture is the most efficient architecture. For this reason, a telescopic architecture is chosen for the OTA.

The complete schematic for the finished OTA can be found in appendix C.6.

[2]In this calculation certain values are assumed for the input capacitor size and the settling time, but these assumptions do not affect the ratio between the two curves.

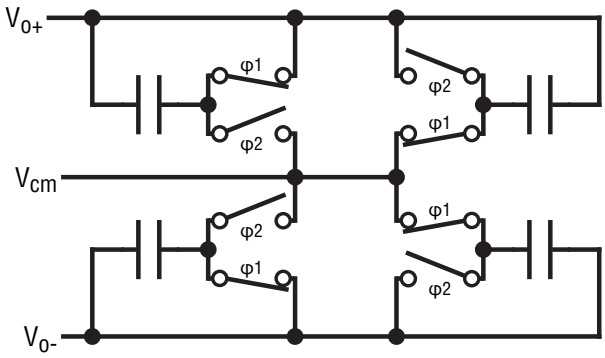

Figure 5.6: The switched capacitor voltage divider, used for the common mode feedback of the OTA.

Isolated input devices

Due to the many switches connected to the virtual ground of the OTA (see figure 5.10b), there will be a lot of common mode charge injection to this node. For this reason, it is very important to keep the common mode rejection ratio (CMRR) as high as possible. In order to increase the CMRR, the V_{sb} of the input transistors needs to be kept constant. In this design this is done by having the input transistors in a separated P-Well, connected to the sources of the input transistors. This way, the whole input stage will float up and down with common mode variations.

Common mode feedback

The common mode feedback of the amplifier is realized by using a switched cap voltage divider as shown in figure 5.6. This is done to keep the output impedance of the OTA as high as possible, while still have accurate common mode measurements with high output swing. The common mode circuitry inside the OTA consists of a simple one stage amplifier that controls the tail current source of the OTA.

Corner Simulations

To assure that the amplifier works as expected in all conditions, corner simulations of the amplifier is done. In these simulations, of which the results are plotted in figures 5.7 and 5.8, the worst case results are:

$$A_{dc} \geq 74.63\,dB$$
$$\text{Phase margin} \geq 69.26°$$
$$\text{Gain margin} \geq 22.38\,dB$$
$$\text{Settling in } 55\,\text{ns} \geq 16.7\,bit$$

These results show that the settling behaviour is not as good as it should be for the correct working of the amplifier (20 bit) in the worst case steps. Because of design time limitation, however, these results where excepted as they do not fundamentally limit the design: by running the design at lower speeds, this limitation can be avoided.

5.1.4. Constant gm current biasing

Because the transconductance of transistors is very temperature dependant, it is important to have a bias circuit that compensates for this change in transconductance. In this circuit, the ideal situation would be to keep the speed of the OTA constant over temperature. To achieve this, the total settling time should remain constant. The total settling time consists of a slewing part and a exponential settling part. The slewing part is proportional to the bias current of the input pair, while the exponential settling part is proportional to the transconductance of the input pair. As the transconductance of the input pair reduces with temperature, this will need to be compensated by a higher bias current. An ideal biasing current for this design would consist of a slightly positive temperature coefficient biasing current.

In this design however, only part of this current was implemented, for simplicity reasons. This biasing circuit is designed to keep the transconductance of the input pair constant over temperature. This is implemented by the circuit in figure 5.9 by fixing the transconductance of a transistor to a

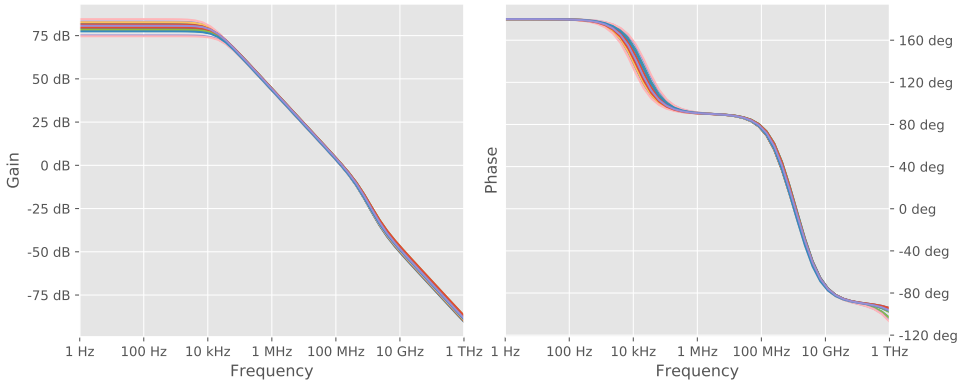

Figure 5.7: The gain and phase response of the OTA over PVT corners

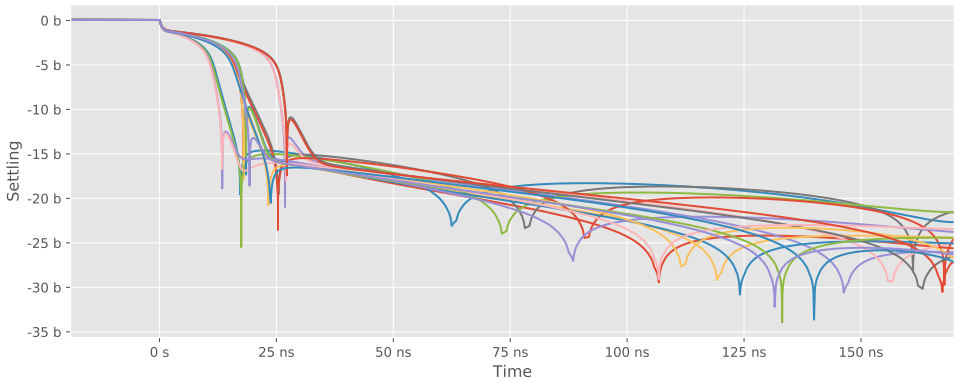

Figure 5.8: The logarithmic worst case step response settling behaviour of the amplifier over PVT corners

resistor value. This is done by having the same current flow though two transistors with different $\frac{W}{L}$ values[20], while the bigger transistor is degenerated by a resistor.

In order to keep the drain voltages of the two transistors approximately equal, to reduce the effects of channel length modulation, a third transistor is added to the loop, so all drain voltages will be about 1 V_{gs}. In order to keep this loop stable, C_{stab} is added.

Because of the relatively high supply voltage, it is possible to make use of Sooch cascode biasing[21] for all current mirrors in the design.

5.1.5. Fast chopping

In order to reach extremely low offset for the charge conversion, system level chopping is used. System level chopping, however, does not effect the current conversion. Because of this, the offset of the sigma delta modulator already needs to be very low without system level chopping.

This is especially important because the residual offset is chopped by the system level chop, making it essentially into a random signal for the current conversion[3].

To combat this 'noise', multiple steps are taken to lower the residual offset before the system level chopping. The first of these steps is to use a switching scheme with auto zeroing. In this way, the offset coming directly from the amplifier is suppressed below the half LSB level.

The switches of this system, however might have mismatch between them, which will cause an offset due to the difference in charge injection. Another cause for differential charge injection is the switches at virtual ground, that have a different source-gate voltage due to the offset of the amplifier. Both of these effects are reduced by introducing nested choppers[22][23] around all the switches and the input pair (the main offset introducing component of the amplifier), as shown in figure 5.10b.

This chopping will remove the offset caused by the mismatched switches, the chopping switches themselves, however, can also be mismatched. The total effect of this is that the residual offset is

[3]The signal is random if the timing of the current conversion is asynchronous and not correlated to the timing of the charge conversion

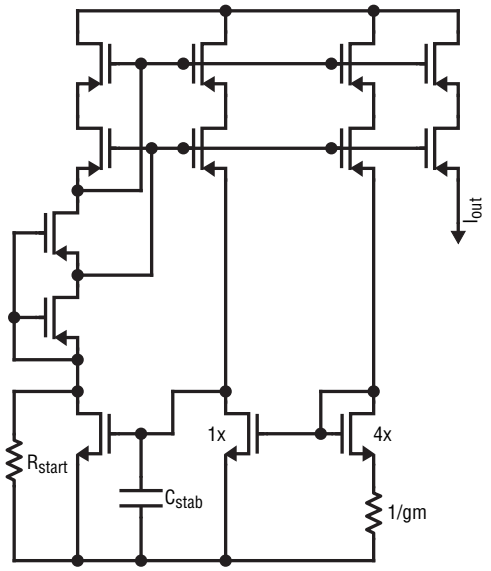

Figure 5.9: The constant transconductance biasing circuit.

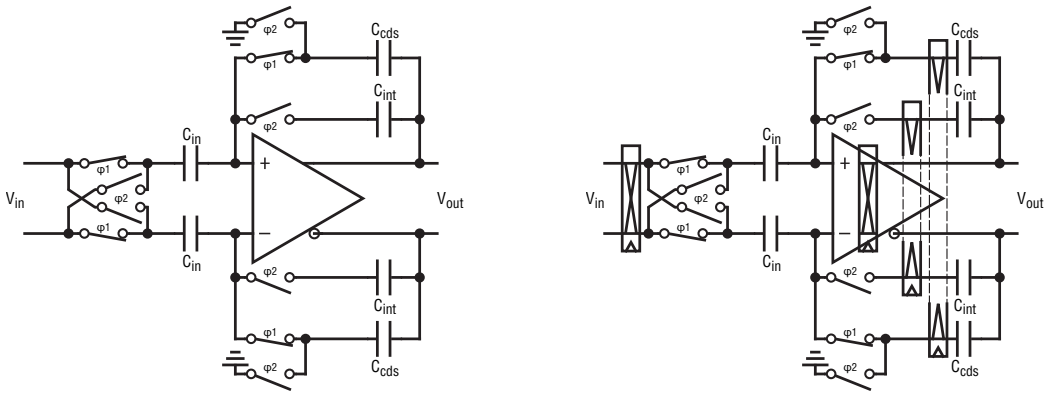

(a) Without fast chopping

(b) With fast copping

Figure 5.10: Fully differential switching scheme

reduced by the ratio between the sampling and the chopping frequency $\frac{f_s}{f_{chop}}$.

As this is the case, it is desirable to keep the chopping speed as low as possible. The lowest possible chopping period which still cancels out the offset for the current conversion, is the conversion time of one current conversion (31 µs).

5.1.6. Comparator design

The 1 bit ADC of the sigma delta loop is implemented as a comparator. Because the errors the comparator makes are noise shaped, the precision of the comparator is not very important. Because the comparator needs to make a decision just before ϕ_2 ends, it is very important for the comparator to have very little kick-back. This is important because if any kick-back of the comparator is not yet settled by the time ϕ_2 ends, it will (partially) end up in the integration capacitor, and thereby adding to the integration error.

The design of the low-kick back comparator is done as a open loop 3 stage amplifier. The first of these stages is a clamped diode amplifier with a gain of about $4.5\times$ in order to reduce most of the kick-back. After this, the second stage increases the gain of the amplifier, and the output turns the output signal into a rail-to-rail signal. A schematic of the comparator is visible in figure 5.11. To reduce the meta-stability of the comparator, a Schmitt trigger is added to the output.

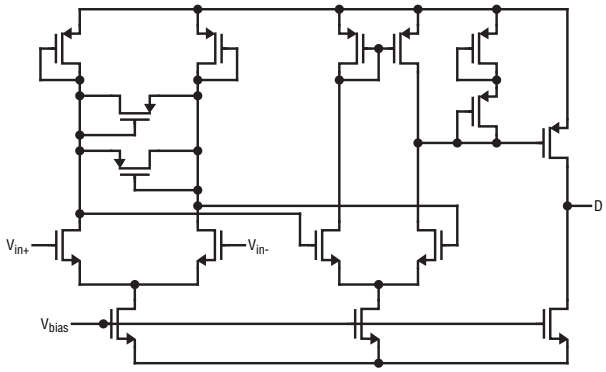

Figure 5.11: The comparator used in the sigma delta modulator

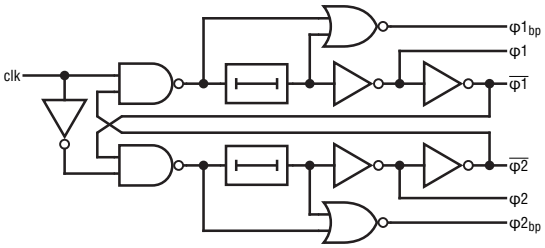

Figure 5.12: A simplified schematic of the clock generator used.

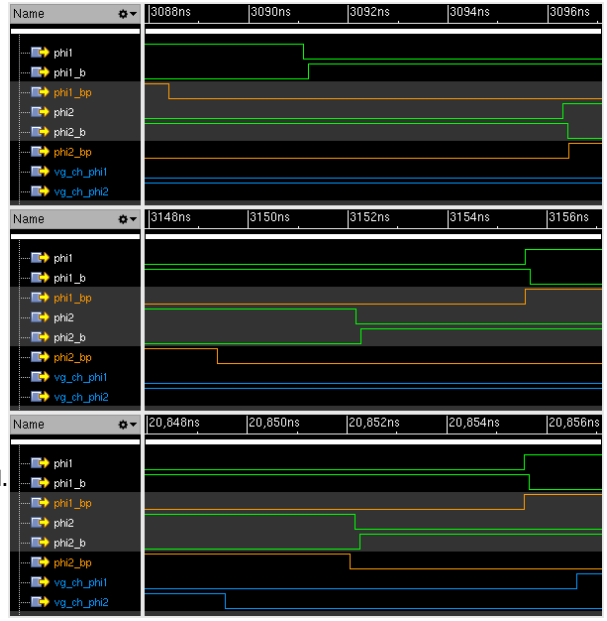

Figure 5.13: A few of the different clock signals during $\phi_1 \rightarrow \phi_2$, $\phi_2 \rightarrow \phi_1$, and $\phi_2 \rightarrow \phi_1$ while chopping

5.1.7. Timing generation

In order for the switched capacitor circuit to work correctly without losing charge because of overlapping switch times, a non-overlapping clock is used. Using a non-overlapping clock assures that all switches open first, before others are closed.

Further, to limit signal dependent charge injection to a minimum, bottom plate sampling is used. With bottom plate sampling, the current into all capacitors is interrupted by first opening a switch with a known (signal independent) voltage. In the case of this design, this is implemented by opening all switches connected to virtual ground first, before doing the rest of the switching.

When the fast choppers are switching (see figure 5.10b), it is important that the chopper switches don't inject signal dependant charge into the integration capacitors. In order to prevent this, the chopper switches open first in the case of a fast chop edge, and the virtual ground switches are delayed.

To generate these clock phases, a non overlapping clock generator is used, similar to the one shown in figure 5.12. The actual circuit is more complex, because it also has to deal with chopping and the DAC signal.

To final clock phases as generated by the clock generator are simulated and can been seen in figure 5.13. PVT corners are also simulated to assure a robust design.

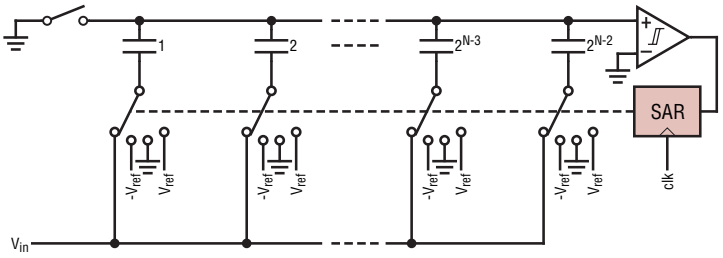

Figure 5.14: A single ended N-bit version of the SAR ADC

5.2. Extended Counting ADC

With the design for the sigma delta modulator complete, the extended counting ADC can be designed. As discussed in chapter 4, the requirements for the extended counting ADC are:

$$\begin{aligned} \text{Resolution} &= 10\,bit \\ T_{sample} &= 1 \text{ sigma delta phase} \\ T_{conv} &\leq 1\,\mu s \end{aligned}$$

5.2.1. Architecture

The goal of the extended counting ADC is to do a relatively fast single shot conversion. The conversion time itself can be longer then the sampling time, and the linearity should be around 10 bit without any trim.

The SAR (successive approximation register) architecture seems to be a very good fit for these requirements, because:[24]

- The Sample time is decoupled from the conversion time.
- The linearity is determined by capacitor matching (can be 10 bit without trim).
- Conversion speed can be high without the need of a fast OTA.

The SAR architecture consists of a capacitor array on which the input voltage is sampled. After this sampling, the input voltage is approximated by successive approximation.

In the type of SAR that is designed (single ended version shown in figure 5.14), the input voltage is first connected to the bottom plate of the capacitors while the top plate is shorted. After this, the charge is transferred to the top plate of the capacitors by switching the bottom plate switches to ground. From here on, every clock cycle, one capacitor will be switched to either $+V_{ref}$ or $-V_{ref}$ depending on the output of the comparator, starting with the biggest capacitor going left.

After $N-1$ cycles, all the bottom plates of the capacitors are connected to either $+V_{ref}$ or $-V_{ref}$, and the top plates at some error voltage close to 0. The past $N-1$ comparator decisions are now the MSB's of the result, and the final comparator decision is the LSB of the result.

5.2.2. SAR Buffer

Because the complete capacitor array can be quite large, it will be quite difficult to settle this signal within one sigma delta clock phase. Directly connecting the SAR array to the sigma delta integrator is also not desirable, as it will slow down and interfere with the integration significantly. Because the SAR is run with a very low duty cycle, it is not worth it to increase the speed of the integrator. A better solution is to have a buffer amplifier in between the integrator and the SAR capacitor array. This buffer amplifier can be turned on only when the SAR is about to run, and thus only consume power when needed.

A other advantage of decoupling the sigma delta integrator and the SAR with a buffer amplifier is that the SAR can have a different input full scale from the integrator output, by having a gain in the buffer amplifier. In this design, a gain of 2 is chosen, to bring the full scale input voltage close to the reference voltage.

The SAR buffer is implemented as a folded cascode switched capacitor amplifier. The capacitors used are small (66 fF), in order to disturb the integrator of the sigma delta as little as possible. Other

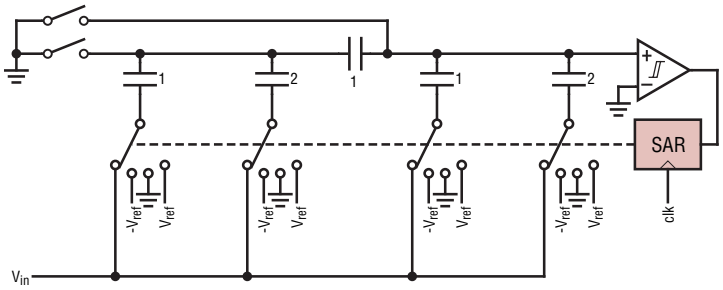

Figure 5.15: A single ended 4-bit SAR ADC with bridge capacitor

features of the buffer amplifier are that it has a power-down mode and it is able to output a wide range of common mode voltages (see section 5.2.4 for why this is necessary).

5.2.3. Bridge capacitor

In a SAR ADC, the minimum capacitor size is usually given by matching instead of by noise. Because, in a N bit SAR, there are $2^{N-1}-1$ unit capacitors (2^N-1 when implemented deferentially), this can become quite a large array.

In order to reduce the size of the array significantly, a bridge capacitor is added (see figure 5.15). This capacitor splits the array into two parts, the LSB part, and the MSB part. The MSB part is the part closest to the comparator. Because the bridge capacitor attenuates the charge transfer from the LSB part to the comparator by connecting it in series, all capacitors on the MSB side can be smaller.

In total, the addition of a bridge capacitor reduces the fully differential capacitor array for a 10 bit SAR ADC from 1022 unit capacitors to 94 unit capacitors.

5.2.4. Further optimization's

In the SAR ADC design proposed, every capacitor needs four switches. Two to plus and minus V_{ref}, one to V_{in} and one to common mode. In order to reduce the complexity of the circuit, some changes are made to these switches.

First, because there is no reliable source of $-V_{ref}$ available, the reference voltages will be changed to $[\text{gnd};\ \frac{1}{2}V_{ref};\ V_{ref}]$. The generation of $\frac{1}{2}V_{ref}$ is less of a problem, because it is a common mode voltage on both sides of the differential circuit, and thus will not be present in the transfer function.

Furthermore, this $\frac{1}{2}V_{ref}$ input is combined with the V_{in} input. This $\frac{1}{2}V_{ref}$ voltage can then be generated by the buffer amplifier, by shorting it's outputs after the top plate switches of the array open. Now, using the common mode feedback circuit from the buffer amplifier, the $\frac{1}{2}V_{ref}$ can be generated.

The steps a conversion now are as follows:

1. The top plate switches of the capacitor array close and the buffer amplifier powers up.
2. The buffer amplifier samples the integrator voltage from the sigma delta.
3. The buffer amplifier amplifies the signal and loads it into the capacitor array.
4. The top plate switches of the capacitor array open.
5. The output of the buffer amplifier is shorted, and the output common mode is changed to $\frac{1}{2}V_{ref}$.
6. In 2^{N-1} steps, the integrator voltage is digitized.
7. Everything resets, and the buffer amplifiers powers down.

5.2.5. Comparator

For the design of the comparator of the SAR ADC, speed, noise, offset and kickback are taken into account. Because the SAR ADC has to run at the highest speed possible in order to minimize the Δt_0 of the current ADC, the speed of the comparator needs to be as fast as possible.

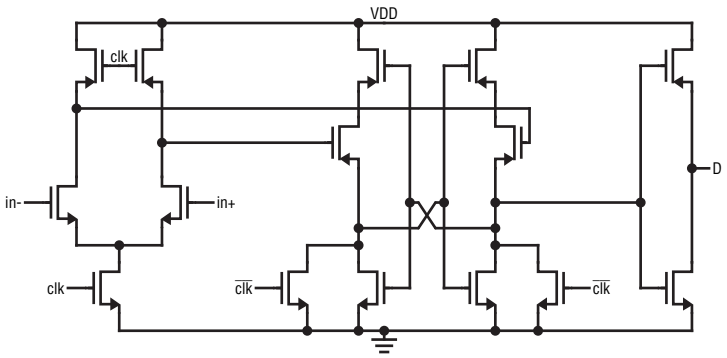

Figure 5.16: The clocked comparator used in the SAR ADC

The noise of the comparator should be lower then the LSB level, to make the design quantization noise limited. This is not the most power efficient way of designing the SAR ADC, but because the duty cycle of the SAR ADC is so low, this is not a problem.

Because the offset of the SAR ADC is canceled by always using the difference of two measurements in the final outcode generation, and the offset of the SAR ADC is directly the offset of the comparator, the offset of the comparator is not of great importance.

Finally, because the charge on the top plates of the capacitor array is always conserved, the kick-back of the comparator is not a problem, as long as no charge leaves the top plate.

In order to satisfy these requirements, a clocked comparator design is chosen (see figure 5.16). By using a clocked design, a high current can be consumed by the input pair while deciding, resulting in low noise, while still not using any current when there is no decision to be made.

5.2.6. SAR and SAR buffer timing

The timing and control of the SAR ADC is kept as simple as possible. The buffer amplifier and sampling switches are controlled by a state machine. This state machine is clocked from the outside of the chip to make the timing as flexible as possible. To make sure the state machine does not go out of sync with the outside world, the state machine is reset on power-down of the SAR. A few timing critical things, however are done on-chip.

One of the on chip timing circuits is the circuit that controls the sampling of the sigma delta integrator voltage into the buffer amplifier. This circuit must run completely synchronous with the sigma delta phases, and for this reason, must be re-timed on-chip.

Another timing optimization that is done on-chip, is the clocking of the successive approximation register. The timing of the register is done by the comparator, clocking it as soon as it made a decision. If the comparator has not made a decision on the falling edge of the SAR-ADC clock, it is forced into a fixed output. The effect of this, is that following an easy decision (large input), there is ample settling time, for a possible next difficult decision (small input). After a difficult decision, however, there will be less settling time, but after a difficult decision always follows a easy decision, as there is always a step made to the input voltage. When the decision is so slow, that the forced output is needed, the input signal was so small (far under the LSB level) that the output of the comparator is not important for the accuracy of the conversion.

5.3. Complete design

With the extended counting ADC implemented, almost all of the circuit level design is done. The result of this can been seen in figure 5.17.

In order to verify the design, multiple simulations are done. Some top level simulations showing the overall performance at low OSR[4] with ideal extended counting are shown in figure 5.18. From these simulations it is visible that the system does not yet have the performance that was intended, but because of time pressure, it was decided to continue with the design as is. This was also done because this is only a prototype, and a proof of concept for the combined current and charge measurement.

[4]A low OSR was used to keep simulation times within reason.

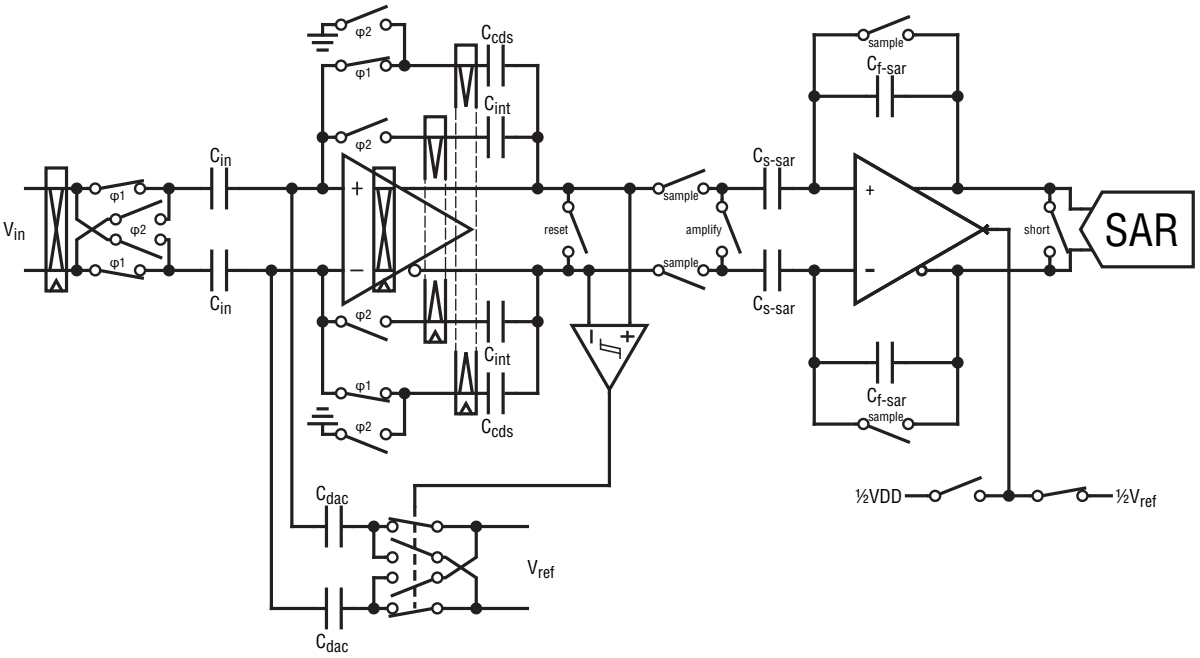

Figure 5.17: A simplified schematic of the complete ADC

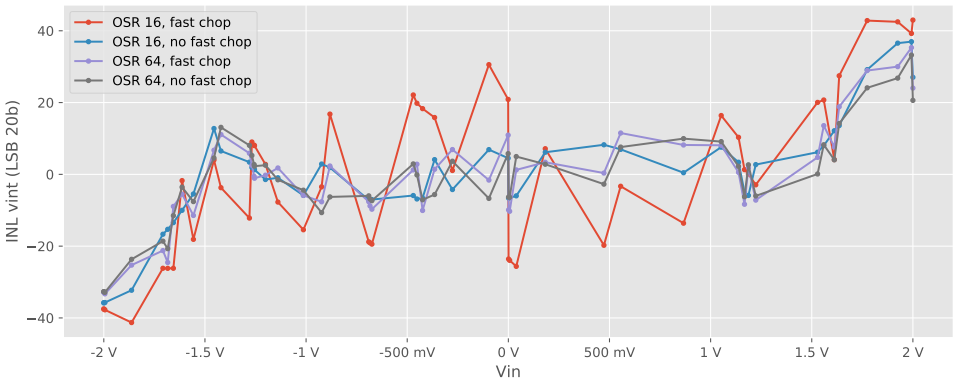

Figure 5.18: The simulated INL of the system at a low OSR with ideal extended counting.

It is also visible that at low OSR's the fast chopping has a big effect on the performance of the system. But this effect seems to reduce with higher OSR's.

5.4. Test modes

In order to verify the workings of (parts) of the design, certain test modes are added to the design. These test modes control specific parts of the chip, in order to help characterize these parts. The test modes added in this design are listed in table 5.1.

5.5. Digital back-end

In order to keep the circuit as flexible as possible, the digital control circuitry needed for the system is not implemented on-chip. Instead, a FPGA is used to control all components. All timing sensitive signals, however, are generated from a master clock on-chip. But this is done in such a way that the clock duty cycle, OSR, chopping frequency etc. are still controllable externally.

Also care is taken to not use resource intensive calculations on the FPGA, in order to make it possible to later implement these digital calculations on-chip in a reasonable chip area.

Table 5.1: The different test modes implemented in the design

Test mode	Effect
I_shutdown	Shut down the constant-gm biasing current.
sar_input	Connect the input of the SAR ADC buffer amplifier directly to the inputs of the chip.
sd_fix_dac	Fix the DAC feedback to a specific value, instead of using the comparator output.
sd_dac_D	The value to which the DAC is fixed, in case of sd_fix_dac.
sd_phi2	Bring out a clock phase from the clock generator, to characterize delays and synchronize. external signals
sd_chop_bp	Turn off the bottom plate sampling for the fast chopping switches (see section 5.1.7).
id	Connect internal resistors to the digital pins, depending on the chip version (see appendix B), to make electrical identification possible.
I_bias (analog)	Output of the bias current, to monitor the biasing circuit.
test_current (analog)	Current flowing into this node will be added to the bias current, to be able to control the bias current.

6 Layout

With the circuit level design complete, it is possible to turn the design into a layout. The chip designed in this thesis will be produced in TSMC 0.18 µm BCD[1] technology. In this technology, high performance analog components like high density MiM-capacitors[2] and thin film resistors are available.

6.1. Sigma delta modulator

In order to mitigate external interference's as much as possible, the layout of the sigma delta converter is done in a completely differential way. As there are a lot of switches around the virtual ground node, there are a lot of connections that need to be made in a fully differential way. This is done by routing all signals through a center channel, with on both sides of this channel all switches, as can been seen from figure 6.1.

In order to shield all wires and signals from each other, every signal runs though a 'tunnel' which consist of a metal sheet above and below the wire, and walls running along each wire. These shields are connected to the analog ground.

The capacitor array of the sigma delta is also designed fully differentially, with the two integration capacitors on either side of the center, and in the center the DAC and input capacitors. This is visible from figure 6.2.

6.1.1. Post-layout simulation

In order to verify if parasitic capacitance's do not cause any unexpected effects, post-layout simulations are run. Similar to the simulations done in section 5.3, these simulations are done at low OSR and with

[1] TSMC 0.18 µm Bipolar-CMOS-DMOS: http://www.tsmc.com/english/dedicatedFoundry/technology/power_ic.htm

[2] Metal-Insulator-Metal Capacitors: horizontally stacked capacitors with a thin layer of high-k dielectric

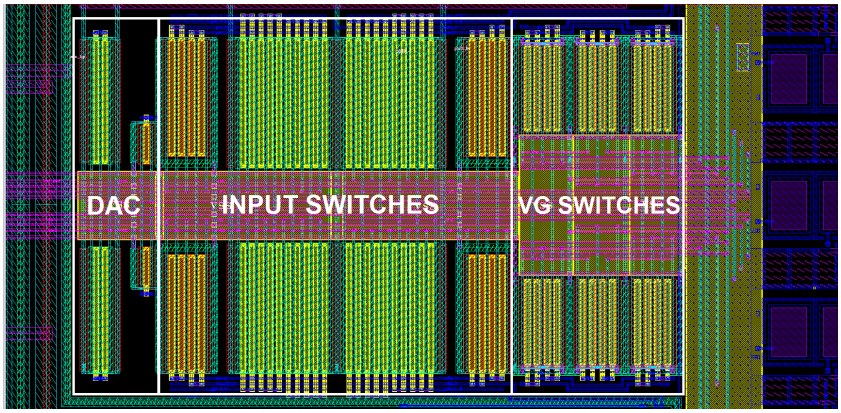

Figure 6.1: The input bus and the switches of the sigma delta ADC

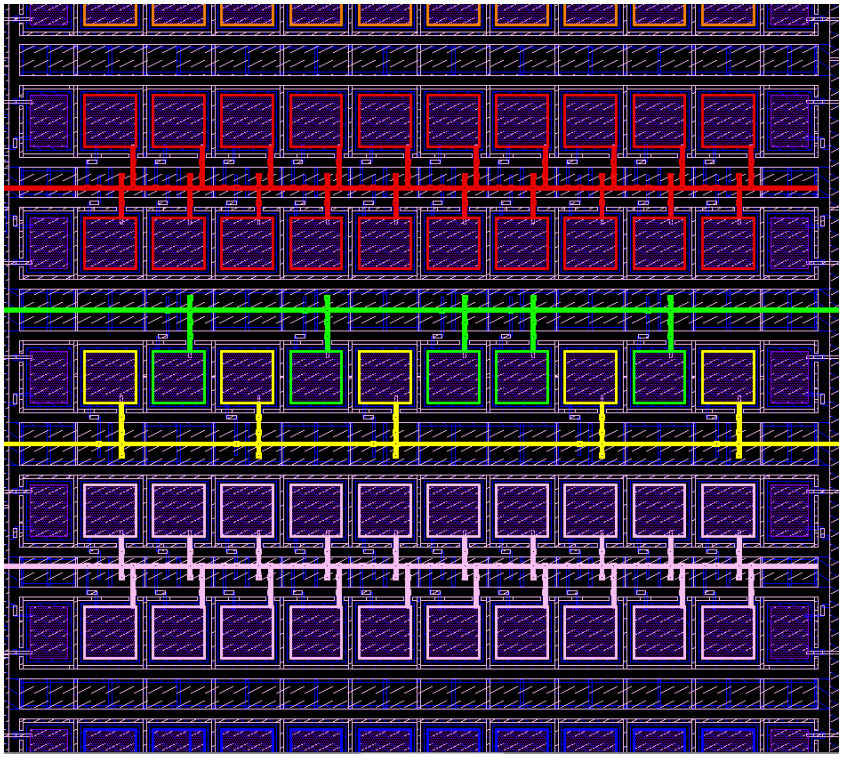

Figure 6.2: Sigma delta capacitor array

ideal extended counting in order to keep reasonable simulation times.

The simulation results are visible in figure 6.3. The results are very similar to figure 5.18, which suggests that the added parasitic capacitance's did not break the design. The results are still not completely within specification, but because of the same reasons mentioned in section 5.3, the prototype chip will be fabricated from this layout.

6.2. SAR

The design and performance of the SAR ADC is very dependent on the layout of the ADC. This is because the linearity of the SAR ADC is determined by the matching of the capacitor array and affected by parasitic capacitance's. Furthermore, to keep the size of the ADC reasonable, the switches and connections should be done in a efficient way.

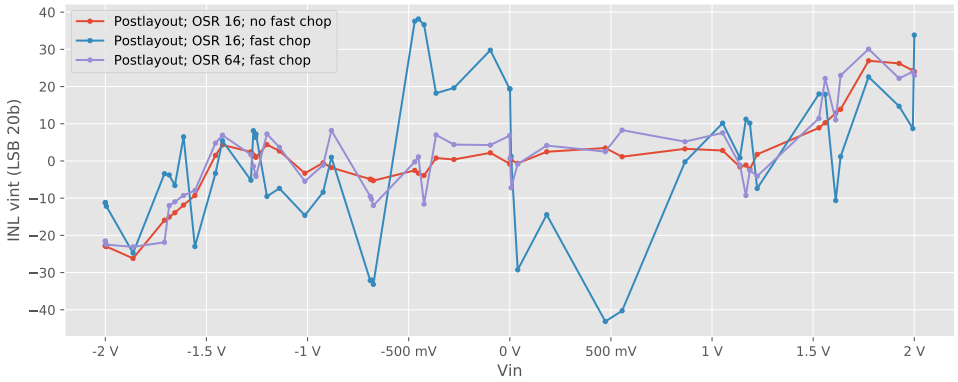

Figure 6.3: The post-layout simulated INL of the sigma delta converter at different OSR's with ideal extended counting.

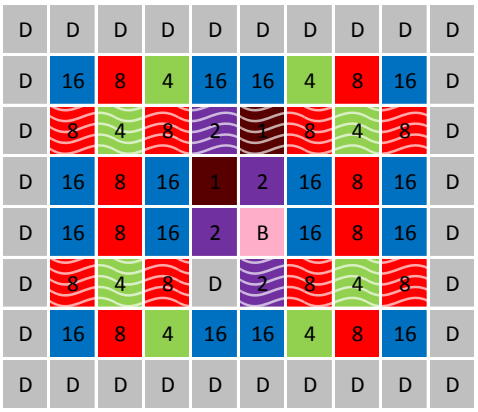

Figure 6.4: The capacitor array layout used for the SAR ADC

6.2.1. SAR Capacitor array

In order to maximize the linearity of the SAR ADC, the matching inside the capacitor array should be maximized.

In the design of the ADC a fully differential design is used, in which there are two capacitor arrays. These two arrays never share any charge, but only redistribute the charge internally. Afterwards, the resulting voltages are compared. Because of this, the capacitor's inside each array need to match, but the arrays themselves don't have to match with each other.

Within each array, the best matching will be achieved by reducing the effects of the most prominent sources of mismatch[25]. In order to cancel gradients in the dielectric, a common centroid layout will be used. On top of this, to reduce the effects of local differences, the capacitors are maximally dispersed. Finally, a row of dummy devices is added around the capacitor array to assure all capacitors have the same environment.

To keep the amount of connections to a minimum, it is also important to minimize the amount of different cells in one row, as the connecting wires will run horizontally, underneath the capacitors.

Unit cells with switches

To keep the size of the SAR ADC as small as possible, the switches connected to the bottom plate of each capacitor are positioned directly underneath the capacitors. Accordingly, every unit is a combination of the switches in the active and lower layers, and the capacitor in the upper layers. The middle layers of the metal stack are occupied by the busses for the control connections, and the different potentials (V_{ref}, gnd, V_{in}).

In order to keep the connection busses small enough to fit underneath a unit capacitor with shielding, the amount of connections to each cell needs to be reduced. In the naive switching scheme of one C-MOS switch for the V_{in} connection, and a P-MOS and N-MOS switch for the V_{ref} and gnd connections, a total of 4 controlling connections need to go to each cell.

As there is still space in the active layer for a inverter, one connection can be removed, and the C-MOS switch will be able to generate the complementary signal locally.

The signals can be optimized even further when it is taken into account that the C-MOS switch needed for the sampling of the input voltage, is no longer needed when the buffer amplifier switches V_{in} to $\frac{1}{2}V_{ref}$. At this low voltage, a pure N-MOS switch suffices. When this is combined with an extra P-MOS switch in the V_{ref} path, the total number of unique wires per cell can be reduced to two. This can been seen in figure 6.5b and table 6.1.

Bridge capacitor

As the SAR capacitor array is split into two parts (see section 5.2.3), there is a most significant part and a least significant part. The least significant part is in series with the bridge capacitor, which attenuates the amount of charge that every switching operation adds or subtracts.

Parasitic capacitance's added to the most significant part of the array attenuate the swing of this part, but the comparator only compares the sign of the voltage, so a attenuation does not change the result of the conversion.

(a) Normal version (b) Optimized version

Figure 6.5: Switch implementation for the SAR ADC capacitor array

Table 6.1: The switching of the signals from figure 6.5. The switching from phase 1 to phase 2 is global, while the switching from phase 2 to phase 3 is unique per cell.

Phase	D0	D1	S0	S1	D2	S2	S3
1. Sampling	1	0	0	1	0	0	1
2. Before switching	1	0	0	1	0	1	1
3. After switching	D	D	1	0	D	1	0

Parasitic capacitance's added to the least significant part, however, will change the result of the conversion. This is because they also attenuate the swing of the least significant part, which in turn reduces the charge transfer to the most significant part via the bridge capacitor. In order to counteract this this effect, the bridge capacitor needs to be a little bit bigger then one unit capacitor. In order to completely cancel out the attenuation caused by the added parasitics, the bridge capacitor should grow with the same factor as the sum of all the capacitors on the least significant side. Because there are 15 unit capacitors on the least significant side, the bridge capacitor should be:

$$C_{bridge} = C_{unit} + \frac{1}{15}\sum C_{par} \tag{6.1}$$

By extracting the expected parasitic capacitance's after the layout of the array. A metal-metal capacitor was created in the same metal layers as where most parasitic capacitance's originating with $\frac{1}{15}$th of the capacitance. By using the same metal layers, the matching between the parasitics should be as accurate as possible, even if the extraction was not completely accurate. The result of the added MoM[3] capacitor can been seen in figure 6.6.

[3]Metal oxide Metal

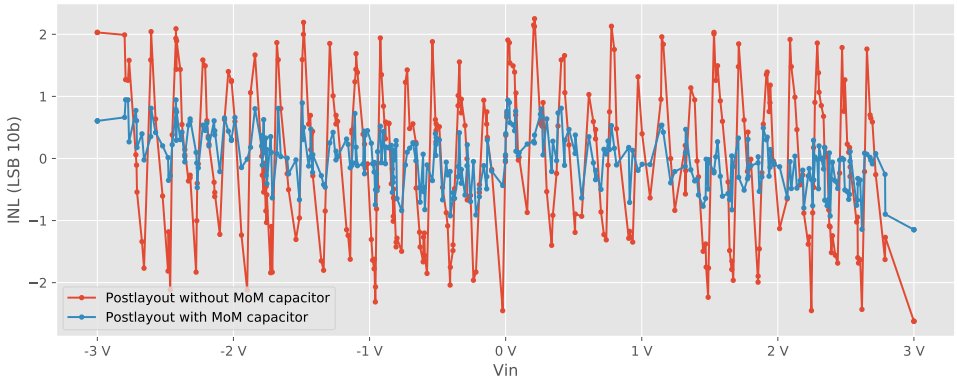

Figure 6.6: The simulated post-layout performance of the SAR ADC, with and without the MoM capacitor added to the bridge capacitor.

7

Measurement Results

In order to fully characterize the produced silicon, a systematic testing approach is taken. First, general properties of the system are tested, such as whether it consumes the correct amount of power, and if it responds as expected to the digital input signals. After that, the parts of the signal chain are tested independently, starting with the SAR, then continuing with the sigma delta converter, and finally the whole system together.

7.1. Measurement setup

In order to be able to correctly test the designed system, it is necessary to design a test environment for the system. For the analog part, a precise input voltage is needed, with as little noise as possible. For the digital part, all digital calculations and control for the ADC is needed, that was excluded from the silicon for flexibility reasons. This digital part also needed designing, implementing and testing of its own.

7.1.1. Analog test system

The analog part of the test system consists of two parts, the dynamic input generation system, and the static reference, biasing and test mode system.

The dynamic input generation system consists of a shunt resistor to convert a input current into a input voltage, a single-ended to differential converter to produce a good quality differential signal and a low pass anti-aliasing filter to prevent frequency folding.

The static part of the analog system consists of voltage references, voltage dividers and buffers. This way all analog input voltages for the chip can be produced with a low output impedance. For the test modes in which the bias current can be adjusted, a v-to-i and i-to-v converter is implemented to make this easy.

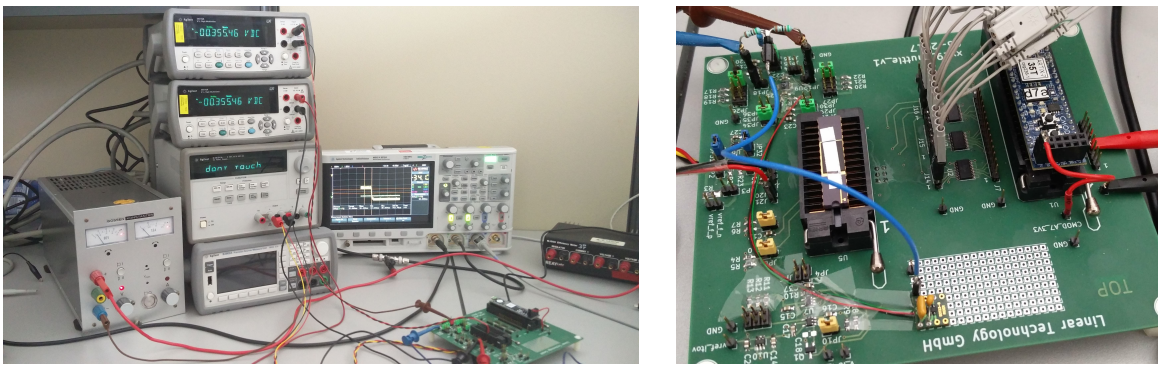

(a) The instruments

(b) The test board

Figure 7.1: The measurement setup used.

7.1.2. Digital test system

Because the ADC is designed to be very flexible, the digital control of the circuitry is completely done by the test system. This means that a big part of the design of the test system is actually still part of the ADC. This causes the digital (test) system to become quite complicated and in need of its own test.

The implementation of the digital system is done on a Cmod-A7 FPGA module[1], powered by a Xilinx Artix-7 FPGA (XC7A35T-1CPG236C). This module can be programmed over USB and also communicates with the test program on the computer using UART over the USB connection.

Tasks of the digital system

The digital system is quite complicated because it performs a lot of tasks that are not completely independent on each other. These tasks are:

- Generate the clock frequency.
- Generate all the chopping signals.
- Reset the integrator at the end of every conversion.
- React on the current conversion trigger, and schedule a current conversion into the chopping scheme.
- Run the SAR ADC at the correct moments for both the charge and the current ADC.
- Decimate the sigma delta modulator output.
- Combine the sigma delta modulator output with the correct SAR ADC outputs.
- Communicate with the test program running on the computer.

Verifying the digital system

In order to verify the digital system, the analog part of the ADC is emulated on a second FPGA. The reason for this is twofold, first it allows debugging of the digital code, with absolute control over the behaviour of the emulated ADC, and secondly, it allowed the digital system to be designed and tested while the ADC silicon was still in production.

The ADC is emulated by implementing a first order approximation of the analog ADC. This is done in fixed point math, with 32 bit precision (4 binary digits + 28 binary decimals). Important parameters as the leakage of the integrator, the offset of the integrator and the input voltage can be adjusted on the fly, software controlled via USB. The rest of the emulated ADC is controlled by the inputs it receives from the digital test system. The digital system, connected to the emulated ADC can be seen in figure 7.2.

7.1.3. Mixed signal design

Because the test board caries both digital signals with a high swing, and high precision analog signals, it is important the board layout is done correctly to avoid interference. On-chip, tunnels for the analog signals could be created due to the many metal layers, but outside of the chip, this is not practically possible.

In order to still maximize the separation between the analog and digital signals, the board is split into two parts. The analog part can only carry analog signals, and the digital part only digital signals. The supply plane is separated on the border of these two parts. The ground plane is not separated, but in some places still has a slit between the two parts, to make sure also low frequency return signals will flow though the correct part of the board. In this way, the digital and analog signals never meet.[26] The final mixed signal text board can be seen in figure 7.1b. The layout used to produce the PCB can be seen in figure 7.3, the purple lines are the cuts in the supply plane, and the blue lines are the slits in the ground plane.

[1]Cmod A7: http://store.digilentinc.com/cmod-a7-breadboardable-artix-7-fpga-module

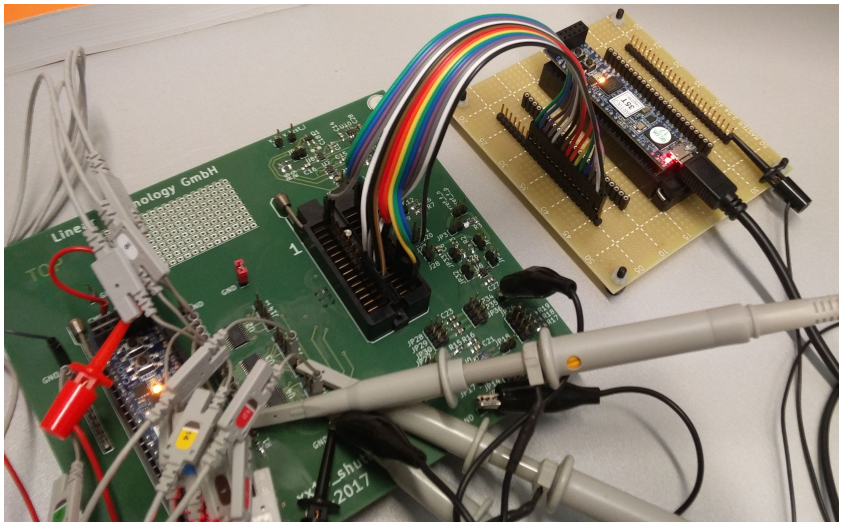

Figure 7.2: The test setup for the digital system with the controller on the left, and the emulator on the right.

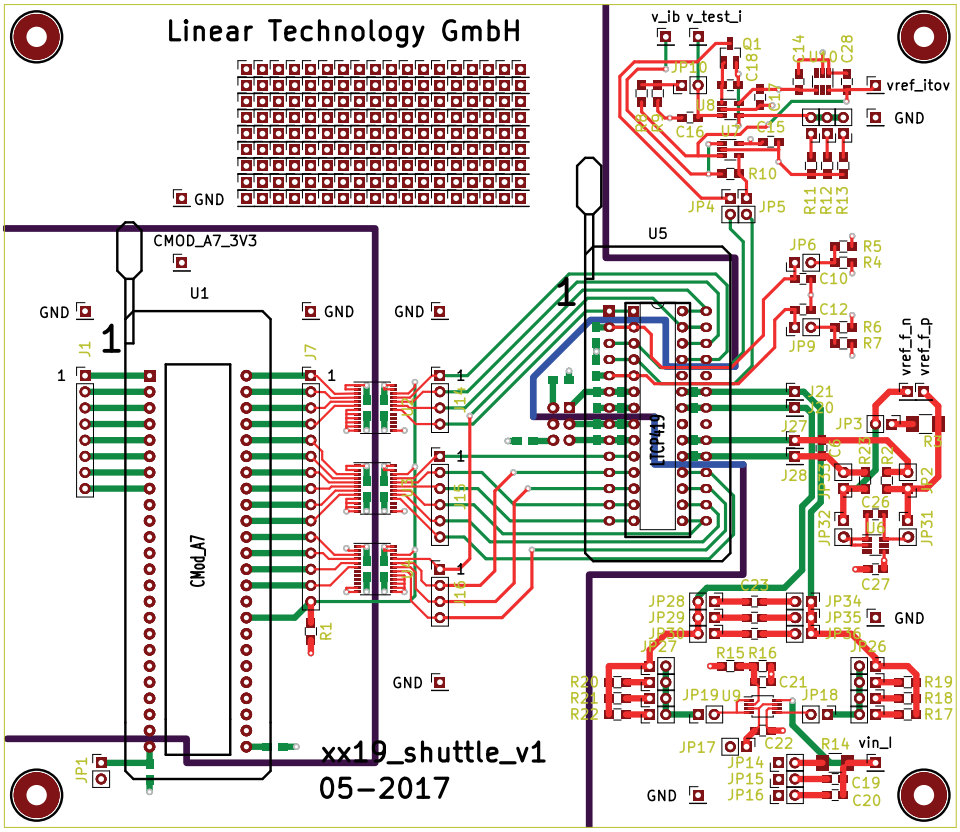

Figure 7.3: The layout of the mixed signal test PCB

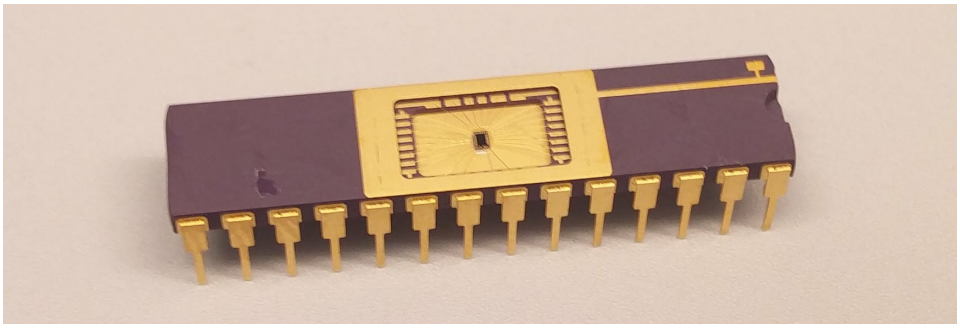

Figure 7.4: The side-braze packaging with the silicon die.

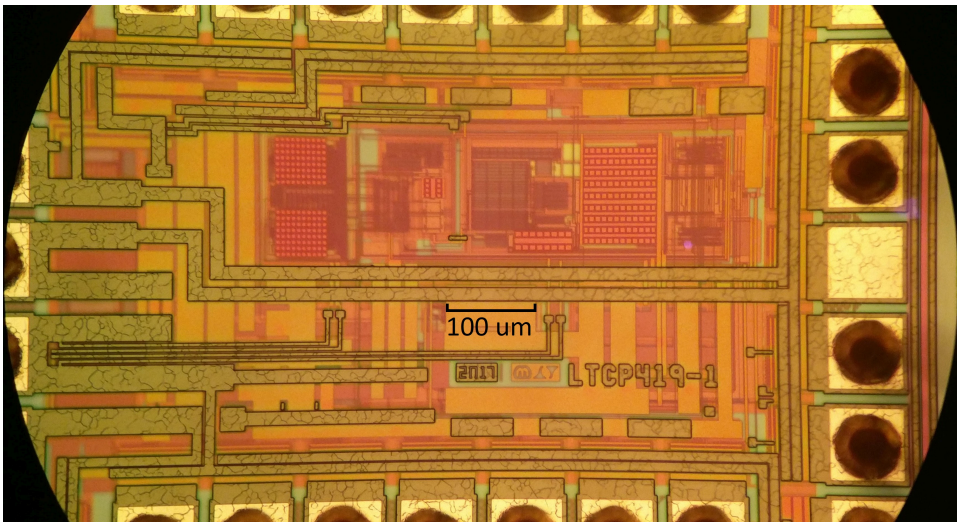

Figure 7.5: A die photograph of the produced silicon.

7.1.4. Measurement instruments

Finally, in order to generate the supply voltage, input signal and measure back the results, a combination of lab equipment was used, visible in figure 7.1a.

For the measurement of the actual input voltage and reference voltage of the ADC, two Agilent 34410A[2] $6\frac{1}{2}$ digit multimeters where used.

To generate the input signal, a Agilent B2902A[3] precision source/measure unit was used. The output current of this current is converted into a voltage through a shunt resistor.

Finally, for the supply voltage, a Agilent E3634A[4] power supply is used. When there was a need for more then one power supply, a Gossen Konstanter power supply was used.

All instruments, except for the Gossen Konstanter, where controlled via software.

7.2. General functioning

With the test system done, and the silicon back from TSMC, the chip can be tested. Unfortunately the chip arrived with a month delay because of a production process error, which limited the available test time. 40 chips in 4 different versions (see appendix B) where available for testing.

A photo of the produced silicon is shown in figures 7.4 and 7.5.

The first thing to test is if the system respond the same to the digital inputs as the emulated ADC. This was the case, and also all test modes where working. So the basic functioning of the chip was as designed!

The next thing to check is the power consumption of the device. The results of this can be seen

[2]Agilent 34410A: https://www.keysight.com/en/pd-692834-pn-34410A
[3]Agilent B2902A: https://www.keysight.com/en/pd-1983585-pn-B2902A
[4]Agilent E3634A: https://www.keysight.com/en/pd-836423-pn-E3634A

Table 7.1: The simulated and measured power consumption of the chip

	Measured [μA]	Simulated [μA]
Analog	327	385
Digital	488	82.2

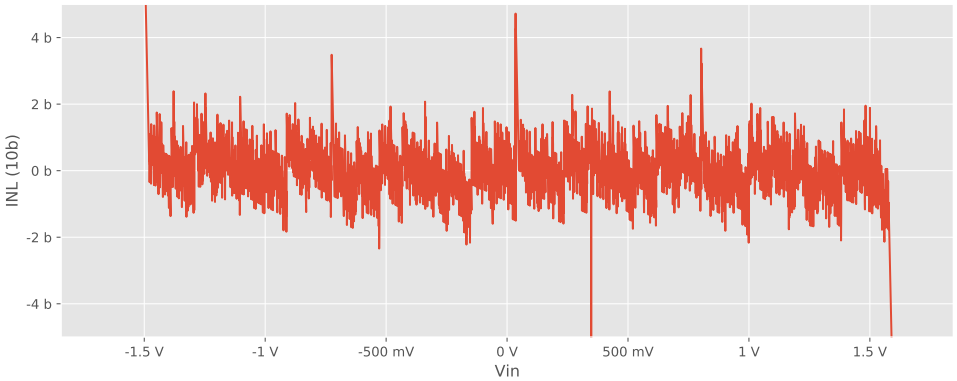

Figure 7.6: The measured INL of the SAR ADC including the SAR buffer amplifier.

in table 7.1. While the analog power consumption is very similar to what was predicted by simulation, the digital power consumption was quite a bit higher. This extra power consumption can, however, be explained by taking into account the digital load capacitance of the outputs of the chip.

As the digital outputs are driving level shifters with 10 pF input capacitance with a swing of 5 V, this consumes quite a big amount of power. In this setup there are three digital output signals: one phase shifted copy of the clock (4 MHz), the bitstream, and the inverse of the bitstream (≈2 MHz). The power consumed by these signals is calculated in equation (7.1).

$$I = VCf = 5\,V \cdot 10\,pF \cdot (4\,MHz + 2 \cdot 2\,MHz) = 400 \qquad [\mu A] \quad (7.1)$$

7.3. SAR and SAR buffer

After verifying that the silicon is on high level working, and reacting as expected to input signals, next the analog performance can be tested. First, the SAR ADC and the buffer that drives it are tested. In order to do this, the chip is put into the test-mode that connects the analog input of the chip directly to the input of the buffer.

Gain

The gain error of the SAR and buffer combination is 1.7 %. This gain error most probably originates from the capacitor matching inside the buffer. As there are no dummies used in the layout of these capacitors. This gain error is not a problem, as the gain of the SAR will be trimmed to the gain of the sigma delta converter.

INL

The measured Integral Non Linearity (INL) of the SAR ADC including buffer is shown in figure 7.6. From this graph it possible to conclude that the SAR ADC, including the buffer amplifier, works as intended. The maximum INL is about 2 LSB. Around 0 V and ±750 mV there are some DNL spikes, caused by the maximum amount of capacitors switching between two outcodes. In this graph there also is a spike in the results around 350 mV, this is probably a measurement error, as it did not show up in further measurements.

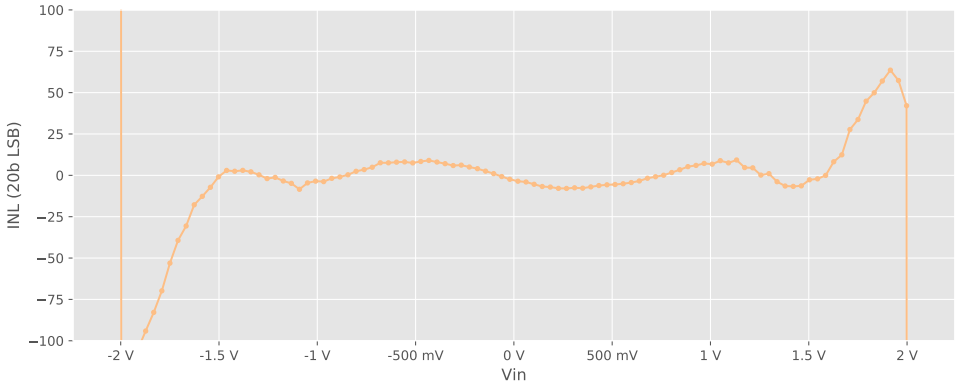

Figure 7.7: The measured INL of the sigma delta ADC without the SAR.

Offset

The offset of the measured SAR ADC was 46.3 mV (15 LSB). This is within the anticipated range, as no efforts where made to reduce the amount of offset in the comparator (as this offset will be canceled out in the overall conversion). In the limited time frame of the measurements, no accurate offset measurements over multiple units where performed for the SAR ADC.

7.4. Sigma delta converter

As the SAR works as expected, the next step is to characterize the sigma delta converter. This is done by running the sigma delta converter with a very high OSR. Because of this, the quantization noise of the sigma delta converter is pushed down to below the thermal noise. Now, the general performance characteristics, like linearity and noise, can be evaluated, without requiring the interaction between the sigma delta converter and the SAR converter yet. All measurements shown here are measured at $\frac{1}{4}$ clock speed, in order to not measure the slow settling effects from figure 5.8.

7.4.1. Linearity

When evaluating the performance of the sigma delta converter, it is directly visible that the sigma delta converter does not perform as well as expected. First the INL of the sigma delta is a lot bigger then simulated during the design. The measured INL can been seen in figure 7.7, and is about ±10 LSB at an input voltage between ±1.5 V, and a lot worse outside of this voltage range. Although in simulations there was also visible degradation at the edges of the full scale, it was a smaller effect, of around ±25 LSB.

The cause of this non-linearity is probably due to the distribution of the charge that is injected by the bottom-plate sampling switches. This distribution is signal dependent due to the fact that the input switches have different impedance's at different input voltage levels. In simulation this is verifiable by replacing the bottom-plate sampling switches by ideal switches. After this is done, the resulting linearity curve is around ±5 LSB, which matches the expected non-linearity from the input capacitors. A way to improve the performance would be to use bootstrapped input switches[27][28] to keep the impedance constant.

7.4.2. Dead-zone

The most serious issue discovered in testing can been seen when zooming into the transfer curve around 0 V, see figure 7.8. There is a quite significant dead-zone. This dead-zone, where a change in input voltage does not cause a change in output code, is caused by integrator leakage, as explained in section 4.5.4. The measured dead-zone is around 100 µV wide, or 25 LSB. While this leakage was not expected, because of the use the finite gain compensating technique from section 5.1.1, it can be verified by simulation. The simulation results, however, are very sensitive to small circuit changes, making finding the root cause difficult.

A lot of time has been spent to find the root cause of this issue, many simulations have been run in order to compare the integrator leakage of different parts of the circuit. The results of the most interesting of these simulations are listed in table 7.2. In order to reach the specifications, a leakage

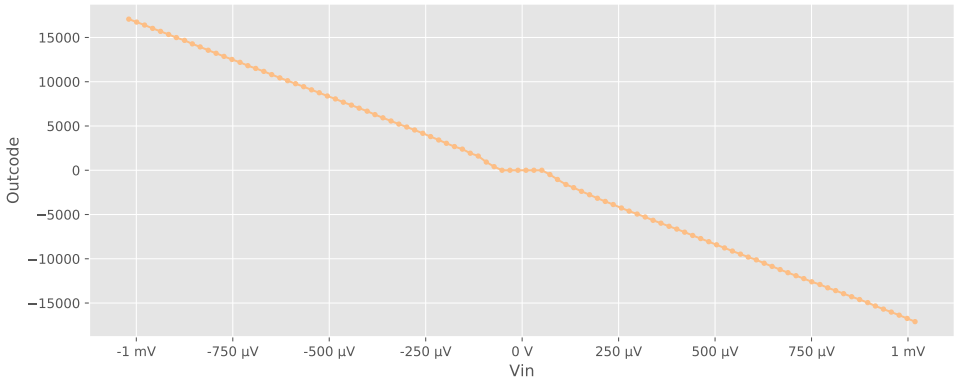

Figure 7.8: A zoomed in portion of the transfer curve of the sigma delta ADC.

Table 7.2: The amount of leaked charge per integration step for different circuit changes, all simulations are done at $\frac{1}{4}$ clock speed

Simulation	Leaked charge [μ%]
Measured result	5560
Postlayout	606
Schematic	3610
Postlayout with ideal OTA	46.7
Postlayout with ideal OTA with input capacitance	73.3
Postlayout with ideal current follower	247
Postlayout without parasitics	3920
Postlayout with higher accuracy settings	606
Schematic with 4.50 V vdd	1560
Schematic with 5.50 V vdd	5590
Schematic with 3.80 pF load capacitors	6.67
Schematic with all sampling switches ideal	100
Schematic with buffered comparator	3330
Schematic with buffered bias	3240
Schematic with gain boosters	3640
Schematic with $\frac{1}{40}$ clock speed	3600
Schematic with correct bottom plate sampling	7410

of less then 250 µ% is necessary.

Simulation results

A couple of simulations from table 7.2 require some extra explanation. First of all, the *'Postlayout with ideal current follower'* fixes the output of the OTA at a fixed voltage, but still carries over the current to the output. This is the same as having cascode transistors with infinite gain in the output stage. In order to keep the gain of the amplifier the same, a load resistor is added with a size of the output impedance of the amplifier. This is also done with the *'Postlayout with ideal OTA'* simulation, to give it approximately the same frequency and DC behaviour as the normal OTA.

In the simulation *'Postlayout with higher accuracy settings'*, all accuracy parameters of the simulator have been increased by $10\times$ in order to verify that this does not not change the result of the simulations.

The simulation *'Schematic with 3.8 pF load capacitors'* adds two quite big load capacitors to the outputs of the integrator. To compensate for the slower settling of this circuit, the clock speed also has been brought down.

In the *'Schematic with buffered comparator'* simulation, the output of the integrator is buffered before it is connected to the comparator, to eliminate any kick-back effects. And in the 'Schematic with buffered bias' simulation, all bias voltages are buffered by a ideal voltage follower, to rule out any really long time constants or weird feedback coming from the bias circuit.

Finally, while doing the simulations it was discovered that one of the switches that did sample a signal was not connected to the bottom plate sampling phase of the clock. This is the switch that connects C_{cds} to common mode in $\phi 2$. Luckily, any signal dependant charge injection that may result from this mistake are suppressed by the gain of the OTA before they reach C_{int}. The last simulation in table 7.2 is a simulation where this mistake is corrected. Unfortunately, this only makes the leakage worse.

Possible explanations

With all the different simulations done it has not yet been possible to arrive at a clear view of what the problem might be. The first investigated potential issue is a settling problem. Increasing the settling time by $40\times$, however, does not change the result of the simulation. This suggests that the leakage is not caused by a settling issue.

As the post-layout simulation gives significantly better results then the schematic simulation, there still seems to be some timing component involved. In order to verify if this is due to some mismatch between the schematic and the layout, all parasitic capacitance's where removed from the net-list of the post-layout simulation. This resulted in a similar leakage as the schematic, which is as expected. So there is no mismatch between the schematic and the layout.

To further look into the why the post-layout simulation performed better then the schematic, a load capacitor was added to the output of the OTA, to make it slower in the schematic. Also the clock speed was reduced to not involve incomplete settling. This became the best simulated result (*'Schematic with 3.8 pF load Capacitors'*). While this seems to suggest that there is something going wrong in the non-overlap/switching timing, these timings are working correctly as shown by the simulation results (all clock phases that should not overlap are non-overlapping).

While performing all calculations from section 5.1.1 in order to determine the DC gain required for the OTA, only 1st order effects where included. When it is taken into account that the DC gain of the amplifier is dependent on the output voltage, a different result applies (for calculation see appendix C.6). The leakage parameter p of the integrator now depends heavily on x, which is the ratio between the OTA gain in $\phi 2$ and the OTA gain in $\phi 1$. If the change in OTA gain is only 1 dB, we still see that the leakage is only improved by a factor 10, or 20 dB, instead of the earlier calculated $A\times$! So this effect should have the following contribution to the dead-zone.

$$p \approx \frac{G}{2A}\left(\frac{1}{x} - 1\right) \tag{7.2}$$

$$p \approx \frac{1}{10}\frac{G}{2A} \qquad \text{x} = -1\,\text{dB} \tag{7.3}$$

$$p \approx \frac{0.015}{A} \qquad \text{G} = 0.3 \tag{7.4}$$

$$p \approx -36.5 - A = -36.5 - 80 = -116.5 \qquad [dB] \tag{7.5}$$

As the necessary leakage factor for a 20 bit ADC is −120.4 dB, there would still be a expected dead-zone of 1.6 bit or 392 µ% of leaked charge. And while this is quite a very nice result, this is with a gain difference of 1 dB. While staying in the middle of the range of the amplifier, the actual gain difference can be much smaller then this, but when reaching the edges of the amplifiers range, the actual gain difference can also become a lot bigger. This effect will thus mostly affect large input signals.

To look for other effects, the DC gain is analyzed. To do this an ideal current follower is added to the output of the amplifier, and the R_{out} is varied. While the ideal current follower does improve the performance (which can be explained by the effect above, as now the gain is stable), increasing the DC gain does not seem to help. Just improving the DC gain by adding gain boosters also does not seem to change the result.

Finally, looking at the combined results of table 7.2, it is clear that touching anything that has to do with charge injection results in changes in the leakage. It would be very interesting to see if changing the supply voltage also has such a big effect on the silicon as it has in the simulations, but unfortunately there was no time left to do this measurements in time for this thesis. Another interesting measurement that can shine more light on this issue is comparing the version with smaller capacitors to this version (see appendix B). While there is already a lot of data gathered about the leakage, further research is necessary to find the root cause.

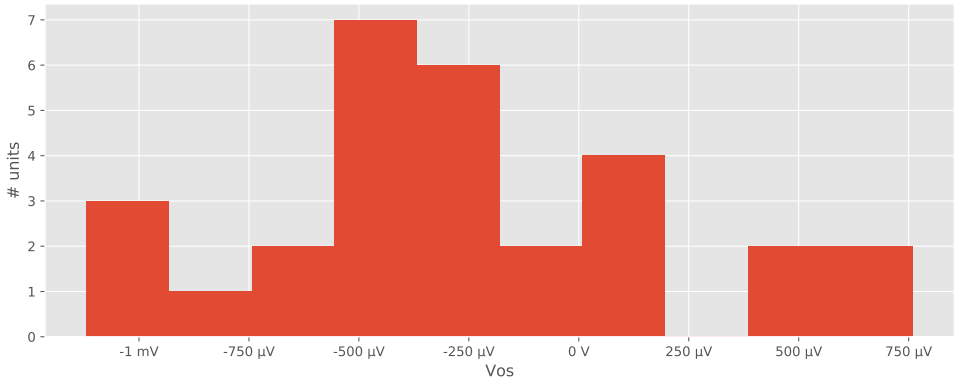

Figure 7.9: The offset of the sigma delta ADC without any chopping.

7.4.3. Offset

Producing a low offset was a key requirement for the design of this ADC, it is interesting to look at the offset resulting from this design. The raw offset of the sigma delta converter over 30 chips is shown in figure 7.9. A small reference voltage (15 mV) was used in order to be able to measure these offsets while not being limited by the integrator leakage discussed above.

Unfortunately, when using the fast-chopping, there is some extra leakage introduced by switching a lot of parasitic capacitance's. During the design this was thought to be accepted, because all transferred charge would end up in C_{cds} and this charge would be suppressed by the gain of the OTA before ending up in C_{int}. But this seems not completely correct. The cause of this is most probably related to the cause of the overall leakage seen in the sigma delta ADC. Because the normal leakage seemed to be a more severe problem, research was prioritized for the normal leakage. For now that means there is no additional data on this effect, so it will need further research in order to resolve.

7.5. SAR to SD matching

The next thing to verify is the matching of the sigma delta converter to the SAR ADC. It is important that the transfer of information works, and that the gain difference between the two converters is stable, so it can be trimmed.

$$\frac{1}{A_{\text{EC-ADC}}} = \frac{V_{fs} G_{sd}}{V_{sar-fs}} \cdot 2^{10} = \frac{2\,V \cdot 0.3}{3\,V} \cdot 1024 = 204.8 \tag{7.6}$$

The theoretical matching ratio is given by equation (7.6). In the post-layout simulations, this is changed slightly due to parasitics to 203. This matches nice with the measured result, where the optimal matching ratio is also 203.

With this matching ratio, the INL of the combined ADC can be measured. In figure 7.10 the INL of the current ADC in 14 bit mode is shown. From the INL it is visible that the current ADC works as intended, even if it is measured in the middle of a charge conversion.

7.6. Timing

In order to verify the working of the combined timing of the charge and current measurements, the trigger point of the current ADC is swept in relation to the charge measurement.

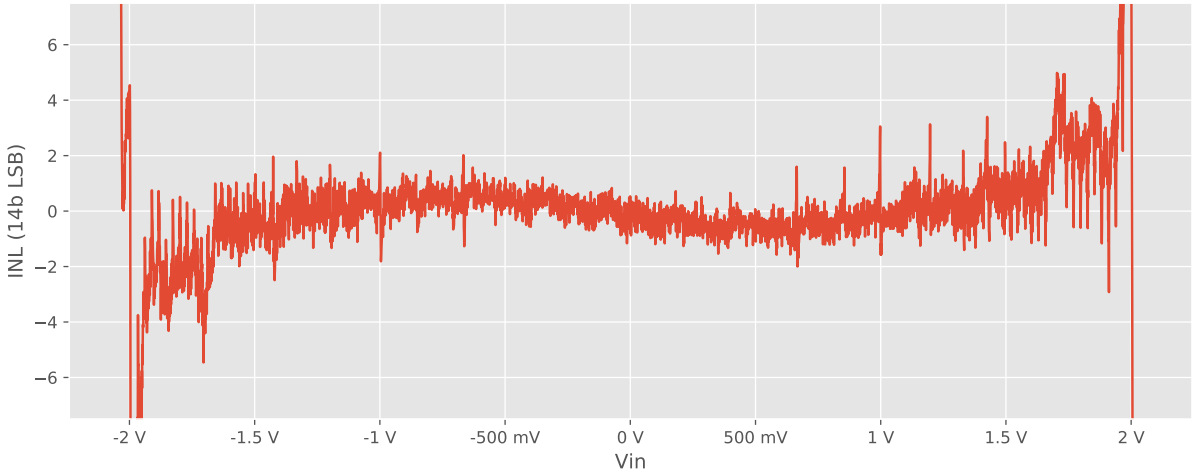

Figure 7.10: The INL of the complete current conversion, measured during a charge conversion.

8

Conclusions and Recommendations

This thesis has looked into the design and testing of a multi-path ADC. At the end of the project the following conclusions and recommendations can be made.

8.1. Conclusions

- It is possible to combine a slow, continuous measurement with a fast, on demand measurement in a single ADC into a more economical design by saving chip area. This is shown by the chip designed, built and tested within this project.
- The 10 bit SAR ADC and its accompanying buffer amplifier work completely as intended. While shifting the common mode in the middle of the conversion was first seen as a difficult and risky task, the buffer amplifier managed to perform excellent at this task.
- In the time frame of one year, designing, taping-out and testing a 20 bit multi-path ADC is quite a challenging task, while a working system has been created, not all specification where met. Big steps are made in the direction of a final solution and a lot has been learned about the work still to be done.
- The current ADC works as intended in 14 bit mode, combining the result of the sigma delta ADC and the SAR ADC with the same gain matching parameter as expected from simulation. It is expected to also work correctly in 16 bit mode, although with some small dead-zones. This is not certain, however, because this has not been conclusively measured.
- The 20 bit charge ADC suffers a lot from integrator leakage, but apart from that, seems to work. There is still room for improvement on the linearity of the charge ADC, but this can most probably be improved by using bootstrapped input switches.
- At the 20 bit level, there are still significant issues and despite serious efforts the root cause of those issues has not yet been isolated, next testing steps have been identified.
- Merging a fast on demand measurement into an existing slow continuous measurement does not cost a lot of extra hardware or precision. Having a very low offset specification on the fast measurement however, introduces some challenges. While a solution for these challenges was attempted, in the end, it resulted in a design with a complicated switching scheme with interesting effects that would benefit from more research.
- From the results from section 7.5 it is possible to see that the two measurements can run in parallel. While the measured results directly show that the slow measurement does not impair the fast measurement, it is not fully possible to conclude for sure that the fast measurement does not influence the slow one. This is because, while no influence could be detected, the slow measurement still suffers from bigger non-idealities that might overshadow possible interference.

- The finite gain compensation switching scheme from section 5.1.1 does not actually square the effective gain for leakage purposes. Instead it produces an improvement related to the gain stability over output voltage of the OTA.
- Even though the main specifications from chapter 2 are not completely met, the design has a lot of potential. It combines a number of circuits and switching schemes, and there is a lot more to learn about the interaction between these circuit elements and there shortcomings.

8.2. Future research

- There are a lot of things that can still be learned from the designed circuit. First of all, it is very interesting to learn what causes the integrator leakage from section 7.4.2. While from the measurements and simulations done, there is not yet a clear root cause. The following measurement steps will give more insight in the problem.
  - Comparing the leakage at different power supply levels.
  - Comparing the leakage between the main version and the version with the smaller capacitor sizes (see appendix B).
  - Comparing the leakage at different bias current levels.
- While the problems that the fast chopping generates seem to be related to the integrator leakage, without any more data on the problem, it is impossible to verify that. Researching this would also make for a interesting continuation of this project.
- Simple single measurements have been done on the system level chopping, unfortunately there was no time to create a full data-set that compares the system offset before and after system level chopping. As offset is a very important requirement for this system, this would be good data to have.
- In general, more detailed measurements on the linearity, noise, gain and offset for the sigma delta converter, the SAR ADC and the two combined would give a lot of insight the system. Useful measurements would for example be:
  - A time versus resolution measurement to compare with figure 4.9 and verify the noise calculations.
  - A offset and mismatch analysis of the SAR ADC over different parts.
  - The gain stability of the sigma delta converter over different parts.
  - The gain stability of the SAR ADC over different parts.
- While in simulations a high temperature range was simulated, only room temperature measurements have been taken, a nice future step would be to characterize the system over temperature.

8.3. Known improvements

- While bottom plate sampling is used for the virtual ground switches, the linearity of the system will probably benefit from using bootstrapped input switches, to make the switch impedance signal independent.
- Even though the designed OTA does work, in light of the target specifications, its performance is marginal, and its settling behaviour can definitely be improved. Because the finite gain compensation does not work as well as expected, some more DC gain is also desirable. Gain boosters on the output cascodes could be a nice way to achieve this.
- While the gm of the OTA will be constant over temperature due to the constant-gm biasing, the settling time will not be constant because there is a significant slewing time. Adding a constant with temperature current to the bias current could thus improve temperature behaviour.
- The digital design for this project was done with possible implementation in mind, but also with a lot of flexibility. Refining the digital design and removing some of the flexibility would be necessary in order to also put the digital part of the ADC into silicon.

- The design of the ADC pre-amp, which was out of scope is also interesting follow up research. In chapter 3 already several options for a pre-amp have been analyzed. The boxcar type pre-amp/filter combination seems an interesting option to further explore.

Bibliography

[1] J. Jiang and K. Makinwa, *A hybrid multipath cmos magnetic sensor with 210µtrms resolution and 3mhz bandwidth for contactless current sensing,* in *2016 IEEE International Solid-State Circuits Conference (ISSCC)* (2016) pp. 204–205.

[2] S. H. Shalmany, D. Draxelmayr, and K. A. A. Makinwa, *A ±36-a integrated current-sensing system with a 0.3% gain error and a 400- μA offset from -55°c to +85°c,* IEEE Journal of Solid-State Circuits **52**, 1034 (2017).

[3] J. L. McCreary and P. R. Gray, *All-mos charge redistribution analog-to-digital conversion techniques. i,* IEEE Journal of Solid-State Circuits **10**, 371 (1975).

[4] M. A. Pertijs and J. H. Huijsing, *Sigma-delta analog-to-digital conversion,* in *PRECISION TEMPERATURE SENSORS IN CMOS TECHNOLOGY* (Springer Netherlands, Dordrecht, 2006) pp. 107–157.

[5] L. Acosta, M. Jiménez, R. G. Carvajal, A. J. Lopez-Martin, and J. Ramírez-Angulo, *Highly linear tunable cmos $Gm-C$ low-pass filter,* IEEE Transactions on Circuits and Systems I: Regular Papers **56**, 2145 (2009).

[6] C. D. Ezekwe and B. E. Boser, *A mode-matching $\Sigma\Delta$ closed-loop vibratory gyroscope readout interface with a 0.004°/s/$\sqrt{}$ hz noise floor over a 50 hz band,* IEEE Journal of Solid-State Circuits **43**, 3039 (2008).

[7] Y. Chae, K. Souri, and K. A. A. Makinwa, *A 6.3 μw 20 bit incremental zoom-adc with 6 ppm inl and 1 μv offset,* IEEE Journal of Solid-State Circuits **48**, 3019 (2013).

[8] M. A. P. Pertijs, K. A. A. Makinwa, and J. H. Huijsing, *A cmos smart temperature sensor with a 3σ inaccuracy of ±0.1°c from -55°c to 125°c,* IEEE Journal of Solid-State Circuits **40**, 2805 (2005).

[9] K. Souri and K. A. A. Makinwa, *A 0.12 mm^2 7.4 μ w micropower temperature sensor with an inaccuracy of $\pm 0.2\,^\circ$c (3 σ) from $-30\,^\circ$c to $125\,^\circ$c,* IEEE Journal of Solid-State Circuits **46**, 1693 (2011).

[10] V. Quiquempoix, P. Deval, A. Barreto, G. Bellini, J. Markus, J. Silva, and G. C. Temes, *A low-power 22-bit incremental adc,* IEEE Journal of Solid-State Circuits **41**, 1562 (2006).

[11] A. Agah, K. Vleugels, P. B. Griffin, M. Ronaghi, J. D. Plummer, and B. A. Wooley, *A high-resolution low-power incremental $\Sigma\Delta$ adc with extended range for biosensor arrays,* IEEE Journal of Solid-State Circuits **45**, 1099 (2010).

[12] J. Márkus, P. Deval, V. Quiquempoix, J. Silva, and G. C. Temes, *Incremental delta-sigma structures for dc measurement: An overview,* in *Custom Integrated Circuits Conference, 2006. CICC'06. IEEE* (IEEE, 2006) pp. 41–48.

[13] C. Jansson, *A high-resolution, compact, and low-power adc suitable for array implementation in standard cmos,* IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications **42**, 904 (1995).

[14] I.-J. Chao, C.-C. Huang, Y.-C. Wu, B.-D. Liu, C.-Y. Huang, and J.-M. Lin, *A 12-bit 4-khz incremental adc with loading-free extended counting technique,* in *Next-Generation Electronics (ISNE), 2013 IEEE International Symposium on* (IEEE, 2013) pp. 29–32.

[15] C.-H. Chen, Y. Zhang, T. He, P. Y. Chiang, and G. C. Temes, *A micro-power two-step incremental analog-to-digital converter,* IEEE Journal of Solid-State Circuits **50**, 1796 (2015).

[16] Y. Zhang, C. H. Chen, T. He, K. Sobue, K. Hamashita, and G. C. Temes, *A two-capacitor sar-assisted multi-step incremental adc with a single amplifier achieving 96.6 db sndr over 1.2 khz bw,* in *2017 IEEE Custom Integrated Circuits Conference (CICC)* (2017) pp. 1–4.

[17] C. C. Enz and G. C. Temes, *Circuit techniques for reducing the effects of op-amp imperfections: autozeroing, correlated double sampling, and chopper stabilization,* Proceedings of the IEEE **84**, 1584 (1996).

[18] J. De Maeyer, P. Rombouts, and L. Weyten, *A double-sampling extended-counting adc,* IEEE Journal of Solid-State Circuits **39**, 411 (2004).

[19] R. Schreier, J. Silva, J. Steensgaard, and G. C. Temes, *Design-oriented estimation of thermal noise in switched-capacitor circuits,* IEEE Transactions on Circuits and Systems I: Regular Papers **52**, 2358 (2005).

[20] J. Chen and B. Shi, *Novel constant transconductance references and the comparisons with the traditional approach,* in *Southwest Symposium on Mixed-Signal Design, 2003.* (2003) pp. 104–107.

[21] N. S. Sooch, *Mos cascode current mirror,* US Patent 4550284 A (1985).

[22] A. Bakker, K. Thiele, and J. H. Huijsing, *A cmos nested-chopper instrumentation amplifier with 100-nv offset,* IEEE Journal of Solid-State Circuits **35**, 1877 (2000).

[23] J. C. van der Meer, F. R. Riedijk, E. van Kampen, K. A. A. Makinwa, and J. H. Huijsing, *A fully integrated cmos hall sensor with a 3.65μt 3σ offset for compass applications,* in *ISSCC. 2005 IEEE International Digest of Technical Papers. Solid-State Circuits Conference, 2005.* (2005) pp. 246–247 Vol. 1.

[24] S. Liu, Y. Shen, and Z. Zhu, *A 12-bit 10 ms/s sar adc with high linearity and energy-efficient switching,* IEEE Transactions on Circuits and Systems I: Regular Papers **63**, 1616 (2016).

[25] C.-W. Lin, J.-M. Lin, Y.-C. Chiu, C.-P. Huang, and S.-J. Chang, *Common-centroid capacitor placement considering systematic and random mismatches in analog integrated circuits,* in *Design Automation Conference (DAC), 2011 48th ACM/EDAC/IEEE* (IEEE, 2011) pp. 528–533.

[26] H. W. Ott, *Partitioning and layout of a mixed-signal pcb,* http://www.hottconsultants.com/pdf_files/june2001pcd_mixedsignal.pdf (2001).

[27] A. M. Abo and P. R. Gray, *A 1.5-v, 10-bit, 14.3-ms/s cmos pipeline analog-to-digital converter,* IEEE Journal of Solid-State Circuits **34**, 599 (1999).

[28] J. Steensgaard, *Bootstrapped low-voltage analog switches,* in *Circuits and Systems, 1999. ISCAS'99. Proceedings of the 1999 IEEE International Symposium on*, Vol. 2 (IEEE, 1999) pp. 29–32.

Glossary

ADC See: Analog to Digital Converter

Analog to Digital Converter A device that takes in an continuous (analog) signal, and outputs an quantized (digital) version of that signal.

Battery Management System An electrical system, designed to monitor the state of charge and state of health of a rechargeable battery.

Bit-stream The output of sigma delta modulator. This is a stream of digital data that still needs to pass though a decimation filter to be usefull.

BMS See: Battery Management System

CDS See: Correlated Double Sampling

Chopper A circuit element is able to switch the polarity of the signal based on a control signal.

Correlated Double Sampling The technique to sample a signal twice, in which the wanted signal switches sign or is absent in one of the two samples. By subtracting the samples, the constant non-idealities can be removed.

Coulomb Counting Measuring the charge going into and coming out of a battery by integrating the battery current.

Decimation Filter A (usually low pass) filter that filters out the band of interest from a over sampled signal and then down samples the signal.

EC See: Extended Counting

Extended Counting The technique where a first order sigma delta converter is improved by measuring the integrator state on reset.

Field Programmable Gate Array A reconfigurable integrated circuit that can programmed to run digital circuit configurations.

FPGA See: Field Programmable Gate Array

INL See: Integral Non Linearity

Integral Non Linearity The integrated non linearity of a ADC. Equivalent to the difference of the output of the ADC and a straight line.

MP-ADC See: Multi Path Analog to Digital Converter

Multi Path Analog to Digital Converter An ADC that has multiple parallel paths from the input to the output(s) of the converter.

OSR See: Oversampling Ratio

Oversampling Ratio The ratio between the sampling rate and the output rate of an ADC.

Quantization Error The difference between the output and input of a ideal ADC. Caused by the quantized output steps of a ADC.

Quantization Noise A approximation of the quantization error as noise. Because the quantization error is not noise, this can lead to wrong results.

SAR See: Successive Approximation Register

SD See: Sigma Delta

Sigma Delta A modulation technique based on a feedback loop that tries to reduce the error to zero. Can be used to build an ADC when combined with a decimation filter.

State of Charge The state of charge of a battery is the amount of energy that is left in a battery, or how 'full' it is.

State of Health The state of health of a battery is how far the battery is in its life cycle.

Successive Approximation Register A ADC type that uses successive steps to approximate the input signal with a DAC.

System Level Chopping A chopping technique where the whole converter, not only parts of it, are chopped to minimize offset to a absolute minimum.

A

Toplevel simulations LTCP419 shuttle

Before PG, a few final top level simulations where done to confirm the working of the circuit as intended.

A.1. Sigma Delta operation

In the first simulation (figure A.1) the $\Sigma\Delta$ circuit integrates two input signals for 16 cycles. First, it settles from any incorrect DC-operating point calculations, then it integrates a very small signal (1 mV) and then a full scale signal (2 V). It is visible that the $\Sigma\Delta$ circuit works as intended: there is negative feedback, the integrator output is reset wile the reset signal is high, and the fast chopping is operational.

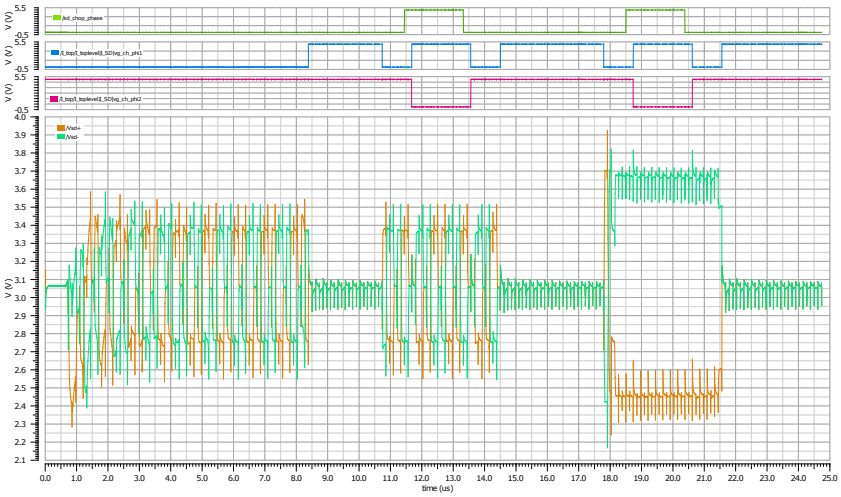

Figure A.1: Two conversions with OSR=16 of the $\Sigma\Delta$ converter with fast chopping

Because in the first test, the system level chopping of the circuit is not tested yet, the functioning of the system level chop signals is quickly tested in a second simulation, seen in figure A.2. As visible, the ch_phi1 signal is the xor of the fast chopping signal, and the system level chop signal.

A.2. SAR operation

The second subsystem of the LTCP419 is the 10b SAR with buffer. Because the SAR generates all its necessary signals itself, based on an externally clocked state machine, it is important to verify it's inner workings.

To be able to control the SAR correctly, it is necessary to have a reset signal to reset the state machine. In case of the LTCP419, this reset signal is the sar_powerdown signal. In figure A.3 the state machine is tested to see if it runs through all states and resets properly.

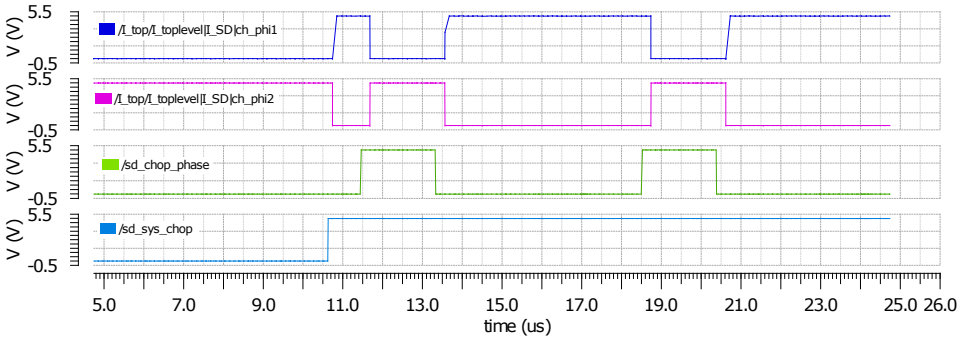

Figure A.2: The important signals for system level chopping

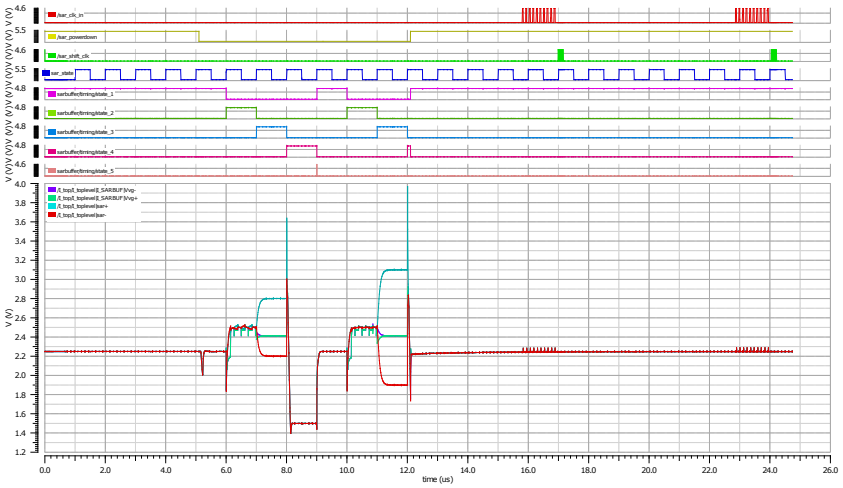

Figure A.3: A testbench to check if the SAR state machine is functioning correctly

Without a functioning SAR, it is not possible to view the state of the integrator in the test-chip. Because this will make the chip very hard to debug in case of errors. It is very important to have a correctly operating SAR. In the simulation shown in figure A.4, the SAR is put into power down mode, woken up, and used for two conversions. As is visible in the zoomed in conversions (figure A.5), the SAR functions correctly. The 5 states of the SAR are:

1. Standby mode
2. Sample the $\Sigma\Delta$ integrator voltage
3. Amplify the signal and load it into the SAR capacitor array
4. Sample the SAR capacitor array and change the common mode voltage of the buffer to $\frac{1}{2}V_{ref}$. During this state, the SAR conversion will be done by clocking the SAR
5. Load the conversion result into the output shift register and reset back to standby mode

After the conversion is done, the conversion result is stored in the output shift register. As visible in figure A.5, the correct data is clocked out of this shift register on request.

A.3. Test mode operation

To be able to gather as much test data as possible from the LTCP419 test chip, a few test modes have been implemented to change the behaviour of the circuit. These test modes are:

th_I_shutdown Shut down the internal constant-gm bias current generator. The bias current can now be completely controlled from the outside.

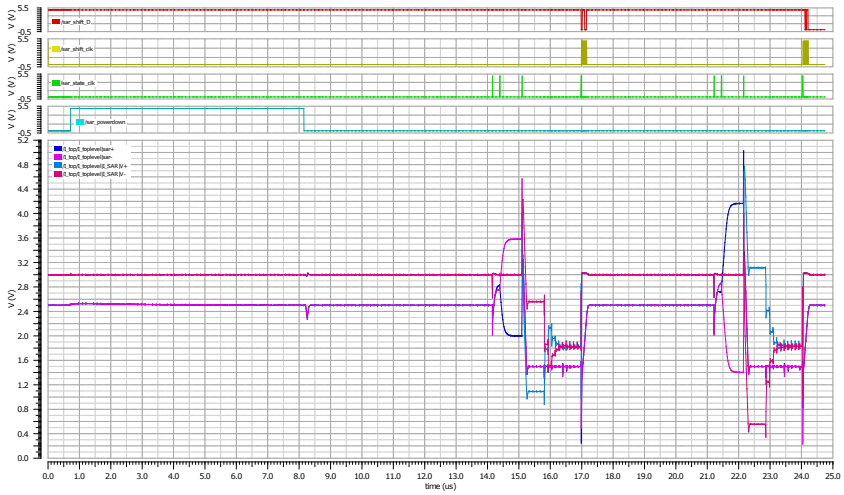

Figure A.4: A testbench showing the intended behaviour of the SAR system

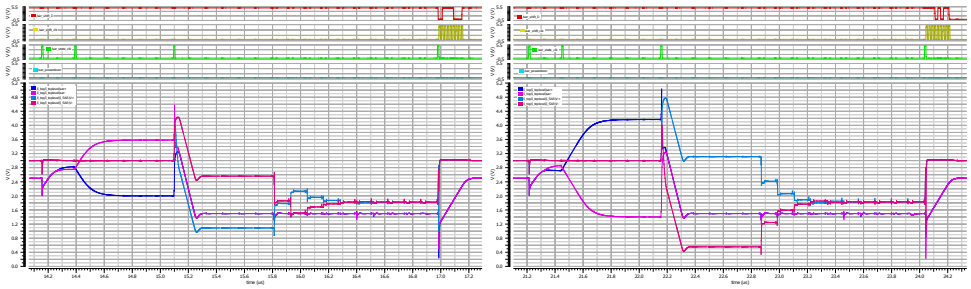

Figure A.5: Zoomed in versions of the conversions in figure A.4

th_sar_input Connect the SAR buffer to the input pins instead of the output of the $\Sigma\Delta$ integrator. This makes sure the SAR can be tested separately from the $\Sigma\Delta$ modulator.

th_sd_fix_dac Do not use the bitstream to control the DAC of the $\Sigma\Delta$, but the **th_sd_dac_D** bit.

th_sd_phi2 Output a copy of the phi2 internal clock signal. This can be used to more precisely synchronise the test system to the internal state.

th_sd_chop_bp Turn off the bottom plate sampling by the virtual ground choppers on chop, used to measure the importance of this measure.

th_id Connect different valued resistors to the testmode input pins. Used to easily identify the different versions of the chip electricaly.

To quickly verify the correct working of the shift register, a data bit is shifted through the different test bits, this can be seen in figure A.6.

The individual workings of the test mode bits have been verified before, and because almost all test bits go directly to digital timing blocks, that have been simulated extensively before, and have experienced no change after that, we wont be discussing them again now.

The th_I_shutdown bit however, is not part of the digital blocks. Figure A.7 shows that this test bit also works. When the th_I_shutdown bit is low, the gate of the first n-mos that generates the bias current, is pulled low. It is also visible, that when the shutdown bit is cleared, the bias current will start again, with a dampened oscillation.

A.4. Power consumption

In order to get an estimate for the power consumption of the chip, a worst case test bench was produced, where the bit stream was toggling constantly, with a load capacitance of 10 pF. The integrated current is shown in figure A.8 and the average current is shown in table A.1

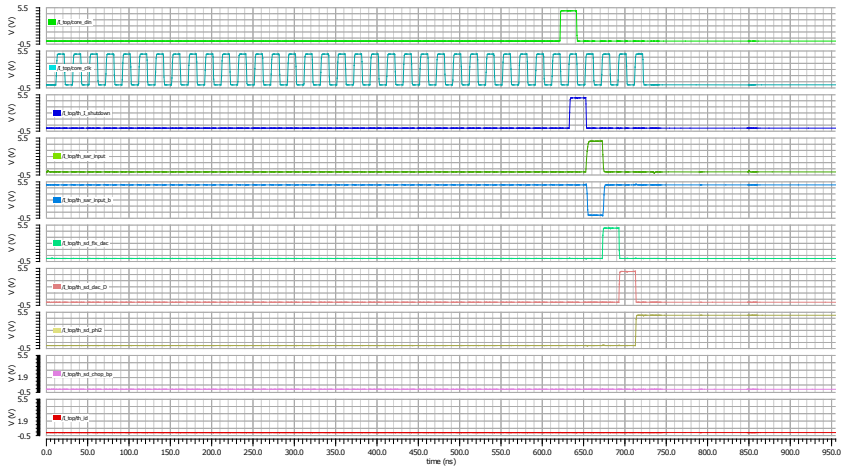

Figure A.6: A quick test of the test mode shift register

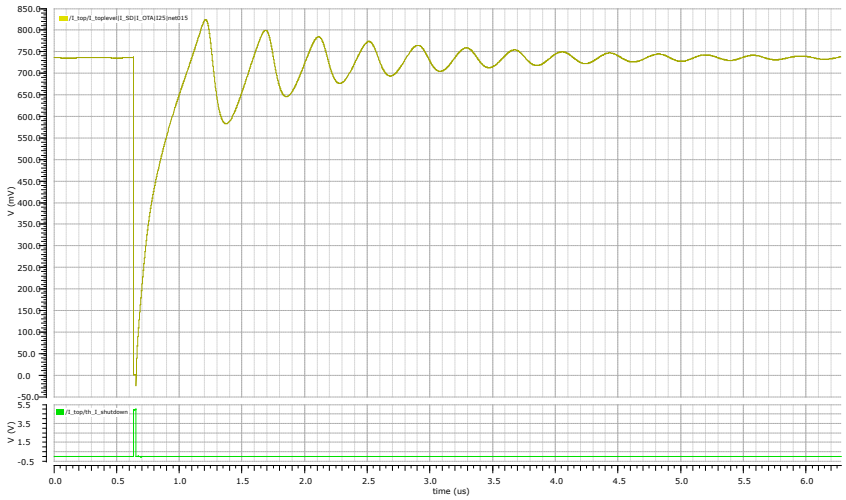

Figure A.7: th_I_shutdown turns off and on the constant-gm bias current generator

A.5. Power on

The final check is based on the power on of the chip. When the supply is turned on, the part should get to a steady state operating voltage, with the correct common mode voltage levels. To check this, a simulation is done where the supply voltage is slowly ramped, after which the reference voltage is ramped, and finally the input voltage. The result of this simulation can been seen in figure A.9 and, zoomed in on the conversion after power on, in figure A.10. The most important thing to notice from this graph, is that all voltages settle to there intended common mode voltage.

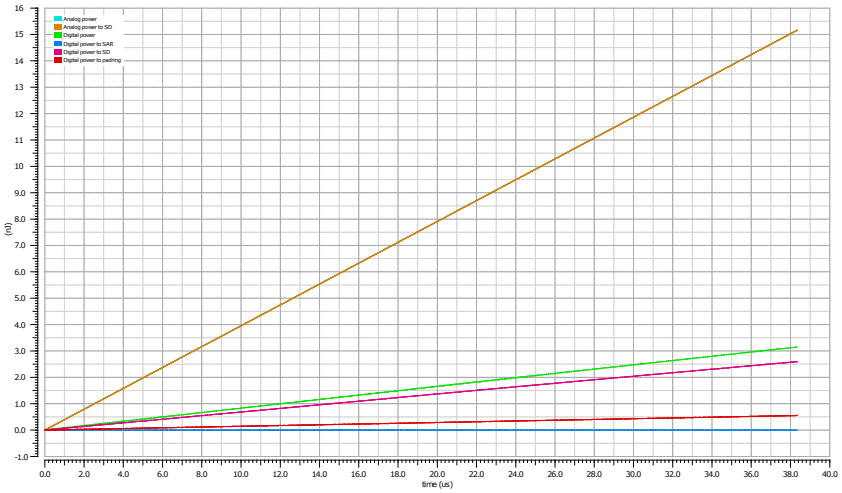

Figure A.8: The integrated current of the test chip with a constantly toggling bit stream

Table A.1: The average current of the chip

Expression	Value (µA)
Analog power	395.4
Analog power to SD	395.4
Digital power	82.18
Digital power to SAR	$4.4 \cdot 10^{-3}$
Digital power to SD	67.66
Digital power to padring	14.51

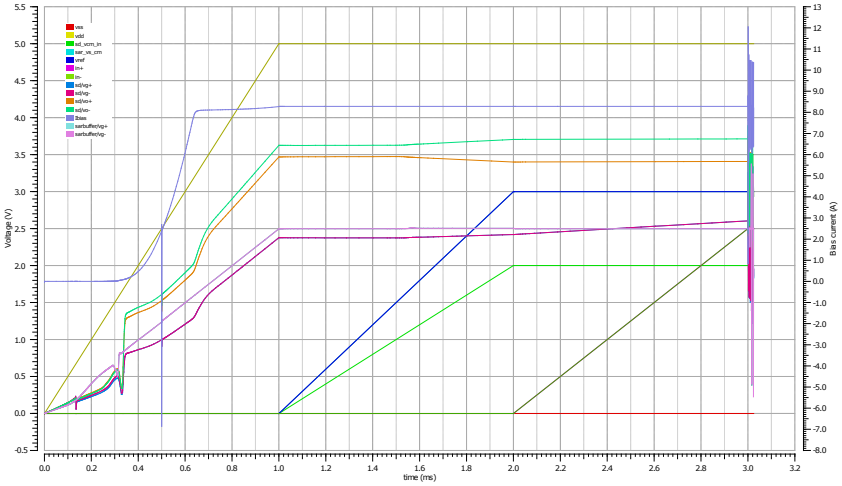

Figure A.9: The critical voltages and currents when ramping up the power

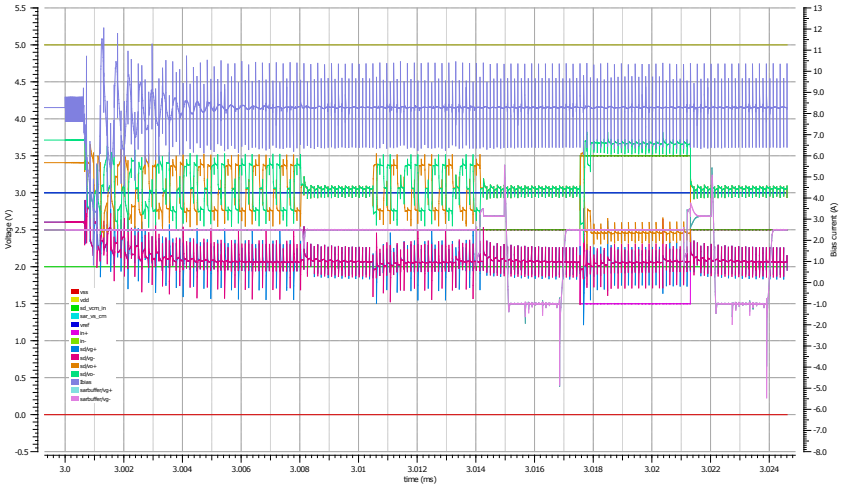

Figure A.10: The critical voltages and currents during a conversion after ramping up the power

B

LTCP419-x version information

On the LTCP419 shuttle, there are four different versions of the LTCP419 circuit. They only differ from each other in a few small details. The four different versions are:

LTCP419-1 The default version.

LTCP419-2 The cascode chopped version.

LTCP419-3 The small capacitor version.

LTCP419-4 The improved input bus version.

The LTCP419-1 is the original version. This is the version on which all verification and extraction simulations are run prior to PG. The first planning was to only produce this version. Later in the design process however, a flaw in the original design was discovered. This was already very late in the layout phase of the design, and it was deemed to risky to change the circuit in a critical way just before PG. Because there was previous experience with putting multiple versions of a chip on a shuttle, it was decided that this was the preferred way to go forward.

The multiple versions of the LTCP419 have different changes to their layout, as specified in table B.1.

Table B.1: A summary of which changes are done to each version

Version	LTCP419-1	LTCP419-2	LTCP419-3	LTCP419-4
Cascode chopping		X		
Small capacitors			X	
Improved input bus		X	X	X

B.1. Cascode chopping

As explained before, a critical change in the circuit was necessary just before PG. This change was the moving of the $\Sigma\Delta$ OTA input chopper from before the input pair to in between the input pair and the cascodes, as can been seen in figure B.1.

The big difference between the two versions is the influence on the offset voltage at the virtual ground node. This offset voltage is primarily caused by the mismatch of the input pair of the OTA. This offset is not chopped around if the chopper is at the virtual ground node, but it is chopped around with the chopper between the input pair and the cascodes. In a first order analysis of the system, this offset voltage does not propagate to the output of the integrator, which is why this architecture was first not deemed necessary. A more careful look however, shows that this offset voltage can reach the output of the integrator trough the voltage dependant charge injection of the sampling switches. Because of this, it is important to move the chopper.

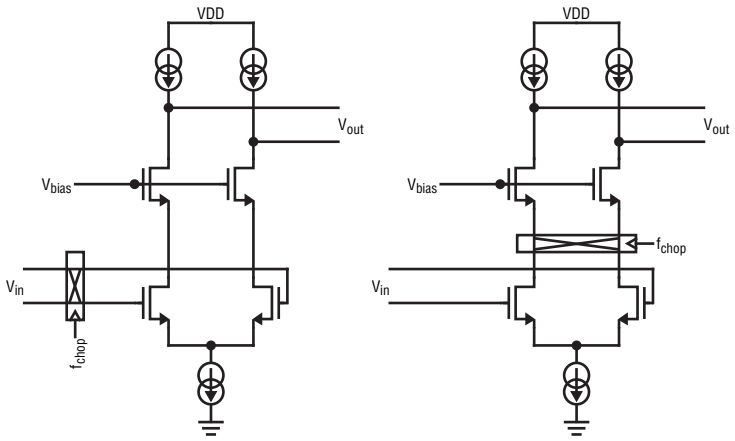

(a) Chopper at Virtual ground (b) Chopper at cascode node

Figure B.1: The OTA with the chopper at different nodes

Expected results

The expected results of moving the chopper from the virtual ground node to the cascode node are reducing the offset in the fast (current) conversions. It will not have a big effect on the charge conversions, because these conversions use system level chopping to eliminate the offset on a higher level.

B.2. Small capacitors

In the TSMC .18 BCD process, there are three different MiM capacitor options: 1 fF/µm^2, 2 fF/µm^2 and 1 fF/µm^2 HL (High Linearity). Of these three, the HL capacitor has the lowest voltage coefficient, followed by the 2 fF/µm^2. The regular 1 fF/µm^2 has the worst performance. The initial design used the HL MiM capacitors, in order to reach the highest linearity possible. Halfway through the layout, however, it was discovered that the HL MiM capacitor was not available between metal 4 and metal 5 (where the capacitors in this design are located). There where now two options to resolve this issue:

1. Only change the type of the capacitor to 2 fF/µm^2, have all capacitors be 2x in capacitance, and run the circuit at $\frac{1}{2}$ speed (possible due to extra bit in SAR)

2. Change the type of the capacitors, but also reduce their area by 50%. Therefore keeping the original capacitance value.

For the main design, option 1 was chosen, because it required the least design effort, and it reduced the importance of parasitic capacitance's, thereby making the design more robust. In the "Small Capacitor" version of the design (LTCP419-3), option 2 has been implemented, in order to test if the design is capable of running on the original speed.

Expected results

The expected results of using smaller capacitance's are that the noise per sample will be $\sqrt{2}$ more, and the settling should be 2x faster. So when using a 2x OSR, the noise per conversion should stay the same. By running the design at a higher speed, while still using a 10b SAR, the quantization noise will be lower, but all parasitic effects will be bigger, so the performance is probably worse then the design with bigger capacitance's.

B.3. Improved input bus

The last change that is done, is a pure layout change. This change is applied to all versions except for the original version. It is only a minor change in how the input and virtual ground switches are connected to the capacitor array, in order to optimise and equalize the parasitic capacitance's between the different wires, and adding a little bit of extra shielding. This was done as an attempt to find a large

offset in the extracted simulations (which, in the end, was caused by a floating shield). The biggest changes can been seen in figure B.2

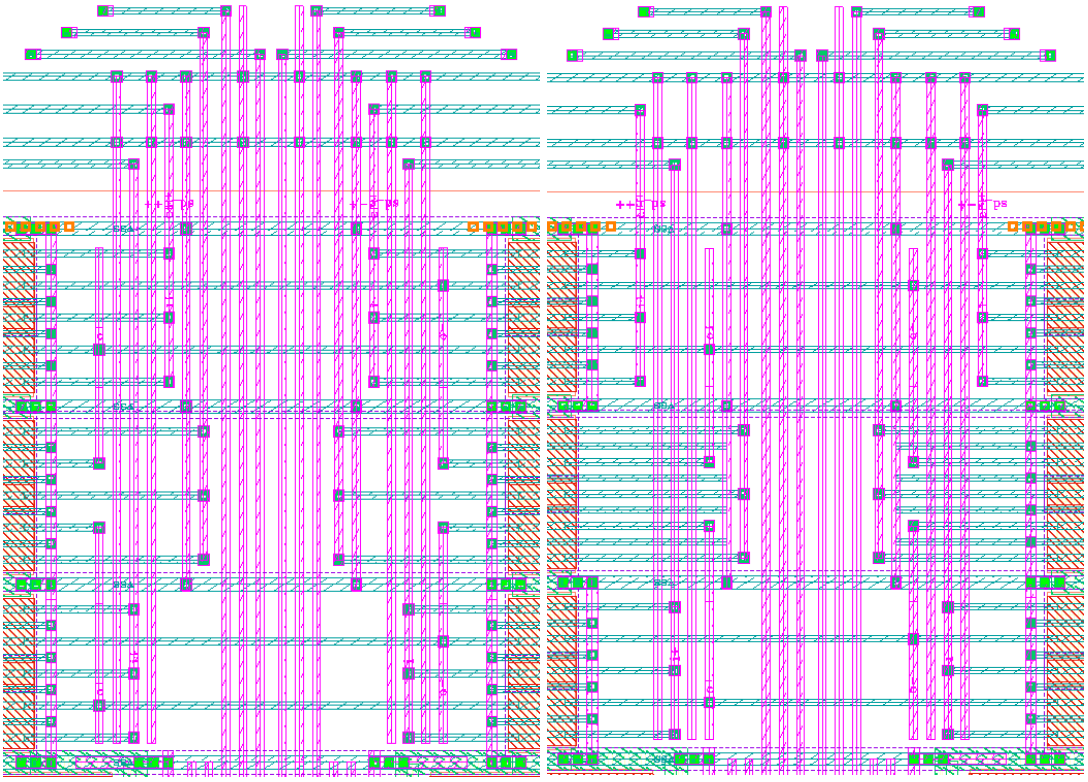

(a) Old input bus (b) New input bus

Figure B.2: The biggest changes of the input bus

Expected results

Because the changes are very minor, the change will probably not have a lot of effect on the performance of the complete circuit. If there is any visible difference with the "default" version, it will probably have a little bit lower offset on the current channel, due to more equal parasitic capacitance's in the choppers.

C

Switched Capacitor Integrator Circuits

In this appendix, the transfer functions of the different switched capacitor integrator circuits will be calculated. This is done using SymPy, a symbolic math library for Python.

From the transfer functions of the integrator, different properties of the amplifier can be derived. The properties that will be needed in this thesis are:

- G: The intended gain of the integrator ($V_{int}[n] = V_{int}[n-1] + G \cdot V_{in}$).
- K: The gain error of the integrator ($V_{int}[n] = V_{int}[n-1] + G \cdot K \cdot V_{in}$).
- $V_{int}[n]$: The transfer function itself.
- p: The leakage factor of the integrator ($V_{int}[n] = (1-p)V_{int}[n-1] + G \cdot K \cdot V_{in}$).

First, in appendix C.1, the code used for the calculations will be explained. After that, in appendices C.2 to C.5, the transfer function of 4 different switched capacitor integrator circuits will be calculated. And finally in appendix C.6, a more complete calculation will be done for the integrator from appendix C.5.

C.1. Used Python code

Below is the python code used for calculating the transfer functions of switched capacitor circuits. It consists of 4 subroutines:

- **solveProblem(equation, lhs)** This function finds the solution to a linear equation, and returns the value of the symbol indicated by lhs.
- **solveFor(var, glbs)** This is a decorator that will solve the set of equations given by a function for the symbol indicated by var. The glbs parameter should contain the *globals()* dictionary.
- **Qredist(C1, V1, C2, V2, phi1, phi2)** This function returns the equation for the charge distribution between two capacitors between two phases.
- **getK(expr, term)** This function is used to find the gain error of the integrator.

All voltages are represented as arrays, the indexes of these arrays give the phase (ϕ_0, ϕ_1, ϕ_2). For more detailed information, please refer to the inline comments of the python script.

```
# -*- coding: utf-8 -*-
"""
Symbolic switched capacitor circuit solver

@author: lloopik
```

```
"""

import sympy

# solve a simple equation for one of its terms
# ex: solveProblem(Eq(y, 3*x), x) = 1/3*y
def solveProblem(equation, lhs):
    solutionset = sympy.solveset(equation, lhs)
    if (len(solutionset) != 1):
        raise Exception('No easy solution')
    solution = solutionset.as_relational(lhs)
    return solution.rhs

# Decorator to make opamp calculations easier,
# add this to a function that returns an equation that must be true
# and it will solve for one other variable.
# Make use of global variables! (Hack hack hack)
# ex:
#  @solveFor('x', globals())
#  def equations():
#      global x, y1, y2
#      y1 = 5*x
#      y2 = 2*x-3
#      return Eq(y1,y2)
#
#   # x will be -1 now, and y1 and y2 wil be -5
def solveFor(var, glbs):
    def runSolver(func):
        # Solve the problem and assign the solved value to var
        exec(var + ' = solveProblem(func(), eval(var))', glbs, locals())
        # Run the function again to get rid of all old dependancies
        check = func()
        # and check if the problem is indeed solved
        assert(check.simplify())
    return runSolver

# Charge redistribution between two capacitors
def Qredist(C1, V1, C2, V2, phi1, phi2):
    return sympy.Eq(C1 * (V1[phi2]-V1[phi1]), C2 * (V2[phi2]-V2[phi1]))

# Declare all variables for all phases
Vout = list(sympy.symbols('V_{out}[n-1], V_{out}[1], V_{out}[2]'))
Vin = list(sympy.symbols('V_{in}[n-1], V_{in}[n]'))
Vin.insert(2, Vin[1])
A, Cint, Cin, Cfb, Vos = sympy.symbols('A, C_int, C_in, C_cds, V_os')
#Vvg = list(map(lambda x: -x/A+Vos, Vout))

Vvg = [None, None, None]
Vcin = [None, None, None]
Vcint = [None, None, None]
Vccds = [None, None, None]


#
# All calculations go here!
# Vint should be the integrated voltage
# G should be the intended gain of the circuit
#

# A helper function to find K
def getK(expr, term=Vin[1]):
```

```
    fact = expr.simplify().expand().cancel().as_ordered_factors()
    assert(len(fact) == 2)
    if(fact[0].is_Pow):
        i = 0
    else:
        i = 1
    assert(fact[i].is_Pow and fact[i].args[1] == -1)
    vin_part = fact[1-i].collect(term).coeff(term,1)
    main_part = vin_part.collect(A).as_ordered_terms()[0]
    return main_part*fact[i]


# Calculate other Coefficients and collect everything
K = getK(Vint/G)
Vint_collected = (Vint/K/G).simplify().expand().collect(A)

# Split coëfficients
A3 = Vint_collected.coeff(A, -3)
A2 = Vint_collected.coeff(A, -2)
A1 = Vint_collected.coeff(A, -1)
A0 = Vint_collected.coeff(A, 0)

# Check if we did not make any errors
# Vint should be equal to our new formula K*G*(A0+A1/A+A2/A^2+A3/A^3)
assert(sympy.Eq(Vint, ((((A3)/A+A2)/A+A1)/A+A0)*G*K).simplify())

# Print everything
Vintn, Vintn1, Gp, Kp = sympy.symbols('V_{int}[n], V_{int}[n-1], G, K')
print('Constant factor: (K)')
display(K.cancel().collect(A))
print('Dependant on:')
display(K.free_symbols)
print('Is approximate:')
display(K.subs({A: 1000, Cfb: 20*64e-15, Cin: 3*64e-15, Cint: 20*64e-15}))


print('Integrator gain (G):')
display(G)
print('Direct:')
display(A0)
print('Suppressed by 1/A:')
display(A1)
print('Suppressed by 1/A**2:')
display(A2)
print('Suppressed by 1/A**3:')
display(A3)

print('Total formula:')
formula = sympy.Eq(Vintn, Gp*Kp*Vint_collected + Vintn1)
display(formula)
```

C.2. Simple SC integrator

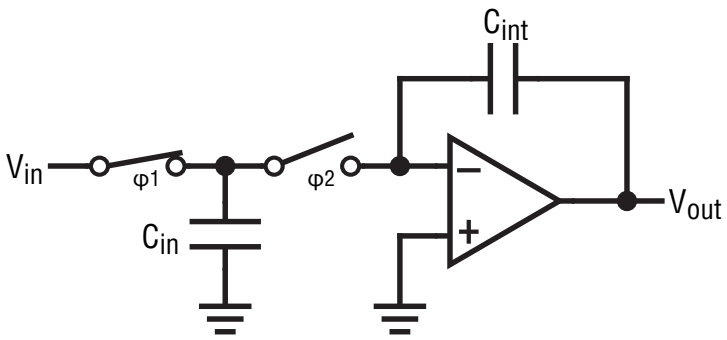

Figure C.1: Simple SC integrator

The first integrator for which to calculate the transfer function is the standard switched capacitor integrator shown in figure C.1. In order to calculate its transfer function, the following equations are derived from the schematic.

```
# Opamp transfer function (in reverse)
def getVvg(Vout):
    return -Vout/A+Vos

G = -Cin/Cint

#phi0
Vvg[0] = getVvg(Vout[0])
Vcin[0] = Vvg[0]
Vcint[0] = Vout[0] - Vvg[0]

#phi1
Vout[1] = Vout[0]
Vvg[1] = getVvg(Vout[1])
Vcin[1] = Vin[1]
Vcint[1] = Vout[1] - Vvg[1]

#phi2
@solveFor('Vout[2]', globals())
def phi2():
    Vvg[2] = getVvg(Vout[2])
    Vcin[2] = Vvg[2]
    Vcint[2] = Vout[2] - Vvg[2]

    return Qredist(Cin, Vcin, Cint, Vcint, 1, 2)

Vint = Vcint[2] - Vcint[0]
```

From this input, the python code calculated the following transfer function. From which the gain error and the leakage parameter are calculated.

$$G = -\frac{C_{in}}{C_{int}} \tag{C.1}$$

$$K = \frac{AC_{int} + C_{int}}{AC_{int} + C_{in} + C_{int}} \tag{C.2}$$

$$K \approx 1 \quad (A \gg 1) \tag{C.3}$$

$$V_{int}[n] = GK\left(-V_{os} + V_{in}[n] + \frac{V_{out}[n-1]}{A}\right) + V_{int}[n-1] \tag{C.4}$$

$$p \approx -\frac{GK}{A} = \frac{|G|K}{A} \tag{C.5}$$

C.3. SC integrator with autozeroing

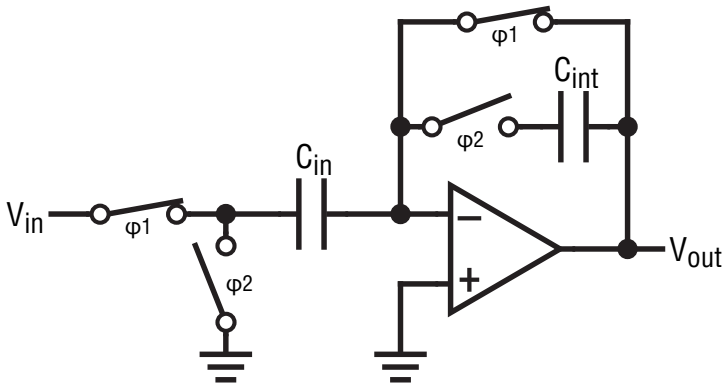

Figure C.2: Auto-zeroing SC integrator

The second switched capacitor integrator uses auto zeroing. Again, the equations are derived from the schematic.

```
# Opamp transfer function (in reverse)
def getVvg(Vout):
    return -Vout/A+Vos

G = Cin/Cint

#phi0
Vvg[0] = getVvg(Vout[0])
Vcin[0] = Vvg[0]
Vcint[0] = Vout[0] - Vvg[0]

#phi1
@solveFor('Vout[1]', globals())
def phi1():
    Vvg[1] = getVvg(Vout[1])
    Vcin[1] = Vvg[1] - Vin[1]
    Vcint[1] = Vcint[0]

    return sympy.Eq(Vout[1], Vvg[1])

#phi2
@solveFor('Vout[2]', globals())
def phi2():
    Vvg[2] = getVvg(Vout[2])
    Vcin[2] = Vvg[2]
    Vcint[2] = Vout[2] - Vvg[2]
```

```
    return Qredist(Cin, Vcin, Cint, Vcint, 1, 2)

Vint = Vcint[2] - Vcint[0]
```

The result of running the code shows the difference between the simple integrator and the integrator with auto-zeroing: in the integrator with auto-zeroing the offset is suppressed by a factor A.

$$G = \frac{C_{in}}{C_{int}} \tag{C.6}$$

$$K = \frac{AC_{int}}{AC_{int} + C_{in} + C_{int}} \tag{C.7}$$

$$K \approx 1 \quad (A \gg 1) \tag{C.8}$$

$$V_{int}[n] = GK\left(V_{in}[n] + \frac{V_{os} + V_{in}[n] - V_{out}[n-1]}{A} - \frac{V_{out}[n-1]}{A^2}\right) + V_{int}[n-1] \tag{C.9}$$

$$p \approx \frac{GK}{A} \tag{C.10}$$

C.4. SC integrator with auto-zeroing and double sampling

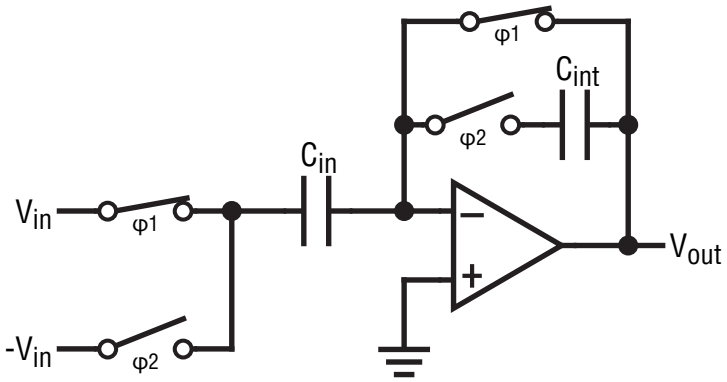

Figure C.3: Auto-zeroing double sampling SC integrator

The only difference between the auto-zeroing integrator and the double sampling auto-zeroing integrator is that instead of shorting the input in ϕ_2, it is connected to $-V_{in}$ (easy to do in a fully differential circuit).

```
# Opamp transfer function (in reverse)
def getVvg(Vout):
    return -Vout/A+Vos

G = 2*Cin/Cint

#phi0
Vvg[0] = getVvg(Vout[0])
Vcin[0] = Vvg[0] + Vin[0]
Vcint[0] = Vout[0] - Vvg[0]

#phi1
@solveFor('Vout[1]', globals())
def phi1():
    Vvg[1] = getVvg(Vout[1])
    Vcin[1] = Vvg[1] - Vin[1]
    Vcint[1] = Vcint[0]
```

```
    return sympy.Eq(Vout[1], Vvg[1])

#phi2
@solveFor('Vout[2]', globals())
def phi2():
    Vvg[2] = getVvg(Vout[2])
    Vcin[2] = Vvg[2] + Vin[2]
    Vcint[2] = Vout[2] - Vvg[2]

    return Qredist(Cin, Vcin, Cint, Vcint, 1, 2)

Vint = Vcint[2] - Vcint[0]
```

From the resulting transfer function, it can be seen that by double sampling, the gain of V_{in} goes up by a factor two, which makes other effects relatively smaller.

$$G = 2\frac{C_{in}}{C_{int}} \tag{C.11}$$

$$K = \frac{AC_{int}}{AC_{int} + C_{in} + C_{int}} \tag{C.12}$$

$$K \approx 1 \quad (A \gg 1) \tag{C.13}$$

$$V_{int}[n] = GK\left(V_{in}[n] + \frac{1}{A}\left(\frac{V_{os}}{2} + V_{in}[n] - \frac{V_{out}[n-1]}{2}\right) - \frac{V_{out}[n-1]}{2A^2}\right) + V_{int}[n-1] \tag{C.14}$$

$$p \approx \frac{GK}{2A} \tag{C.15}$$

C.5. Finite gain compensated integrator

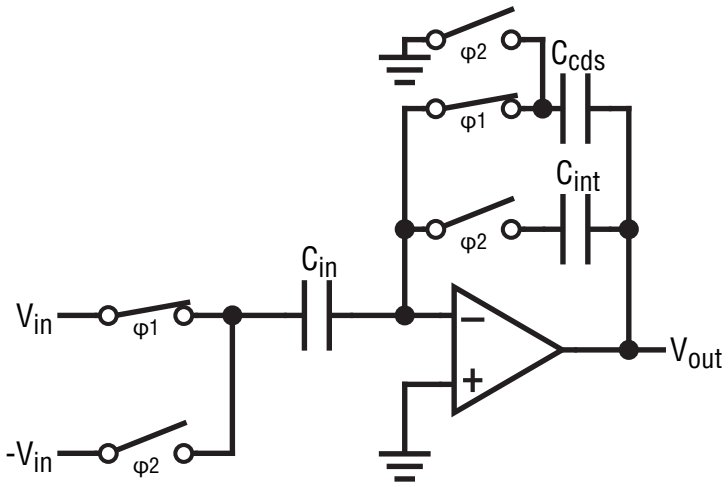

Figure C.4: Double sampling SC integrator with finite gain compensation

In the finite gain compensated integrator, instead of shorting the amplifier in ϕ_1, it is connected via a different capacitor C_{cds}. The input equations then look like this.

```
# Opamp transfer function (in reverse)
def getVvg(Vout):
    return -Vout/A+Vos

G = 2*Cin/Cint

#phi0
```

```
Vvg[0] = getVvg(Vout[0])
Vcin[0] = Vvg[0] + Vin[0]
Vcint[0] = Vout[0] - Vvg[0]
Vccds[0] = Vout[0]

#phi1
@solveFor('Vout[1]', globals())
def phi1():
    Vvg[1] = getVvg(Vout[1])
    Vcin[1] = Vvg[1] - Vin[1]
    Vcint[1] = Vcint[0]
    Vccds[1] = Vout[1] - Vvg[1]

    return Qredist(Cin, Vcin, Ccds, Vccds, 0, 1)

#phi2
@solveFor('Vout[2]', globals())
def phi2():
    Vvg[2] = getVvg(Vout[2])
    Vcin[2] = Vvg[2] + Vin[2]
    Vcint[2] = Vout[2] - Vvg[2]
    Vccds[2] = Vout[2]

    return Qredist(Cin, Vcin, Cint, Vcint, 1, 2)

Vint = Vcint[2] - Vcint[0]
```

The result of this extra capacitor is that the leakage factor of the amplifier goes down. One of the assumptions done in this calculation is that the gain of the amplifier A is constant.

$$G = 2\frac{C_{in}}{C_{int}} \tag{C.16}$$

$$K = \frac{A^2 C_{cds} C_{int} + A C_{cds} C_{int}}{A^2 C_{cds} C_{int} + A\left(C_{cds} C_{in} + 2C_{cds} C_{int} + C_{in} C_{int}\right) + C_{cds} C_{in} + C_{cds} C_{int} + C_{in}^2 + C_{in} C_{int}} \tag{C.17}$$

$$K \approx 1 \qquad (A \gg 1) \tag{C.18}$$

$$V_{int}[n] = GK\left(V_{in}[n] + \frac{1}{A}\left(\frac{V_{os}}{2} + V_{in}[n] - \frac{C_{in} V_{in}[n-1]}{2C_{cds}} + \frac{C_{in} V_{in}[n]}{2C_{cds}}\right) - \frac{V_{out}[n-1]}{2A^2}\right) + V_{int}[n-1] \tag{C.19}$$

$$p \approx \frac{GK}{2A^2} \tag{C.20}$$

C.6. Finite gain compensation with non-constant gain

When the calculation from appendix C.5 is redone, this time with a gain of A in ϕ_1 and a gain of $x \cdot A$ in ϕ_2, the following happens.

```
# Opamp transfer function (in reverse)
def getVvg(Vout):
    return -Vout/A+Vos

G = 2*Cin/Cint

#phi0
Vvg[0] = getVvg(Vout[0])
Vcin[0] = Vvg[0] + Vin[0]
Vcint[0] = Vout[0] - Vvg[0]
Vccds[0] = Vout[0]
```

```
#phi1
@solveFor('Vout[1]', globals())
def phi1():
    Vvg[1] = getVvg(Vout[1])
    Vcin[1] = Vvg[1] - Vin[1]
    Vcint[1] = Vcint[0]
    Vccds[1] = Vout[1] - Vvg[1]

    return Qredist(Cin, Vcin, Ccds, Vccds, 0, 1)

# Change the value of A to A*x
# this means that x is the ratio between the gain
# in phi 1 and phi 2
oldA = A
x = sympy.symbols('x')
A = x*A

#phi2
@solveFor('Vout[2]', globals())
def phi2():
    Vvg[2] = getVvg(Vout[2])
    Vcin[2] = Vvg[2] + Vin[2]
    Vcint[2] = Vout[2] - Vvg[2]
    Vccds[2] = Vout[2]

    return Qredist(Cin, Vcin, Cint, Vcint, 1, 2)

# Change A back for the calculations
A = oldA

Vint = Vcint[2] - Vcint[0]
```

As visible from the transfer function below, the leakage increases significantly if the effect of the non-constant gain of the amplifier is taken into account. The consequences of this effect are discussed further in section 7.4.2.

$$G = 2\frac{C_{in}}{C_{int}} \tag{C.21}$$

$$K = \frac{A^2 C_{fb} C_{int} x}{A^2 C_{fb} C_{int} x + A\left(C_{fb}C_{in} + C_{fb}C_{int}x + C_{fb}C_{int} + C_{in}C_{int}x\right) + C_{fb}C_{in} + C_{fb}C_{int} + C_{in}^2 + C_{in}C_{int}} \tag{C.22}$$

$$K \approx 1 \qquad (A \gg 1) \tag{C.23}$$

$$\begin{aligned} V_{int}[n] = GK\Bigg[&V_{in}[n] + \frac{1}{A}\left(\frac{V_{os}}{2} + V_{in}[n] + \frac{V_{in}[n]}{x} + \frac{V_{out}[n-1]}{2} - \frac{V_{out}[n-1]}{2x} - \frac{C_{in}V_{in}[n-1]}{2C_{fb}} + \frac{C_{in}V_{in}[n]}{2C_{fb}}\right) \\ &+ \frac{1}{A^2}\left(\frac{V_{os}}{2x} + \frac{V_{in}[n]}{x} - \frac{V_{out}[n-1]}{2x} - \frac{C_{in}V_{in}[n-1]}{2C_{fb}x} + \frac{C_{in}V_{in}[n]}{2C_{fb}x} + \frac{C_{in}V_{out}[n-1]}{2C_{fb}} - \frac{C_{in}V_{out}[n-1]}{2C_{fb}x}\right) \\ &- \frac{V_{out}[n-1]}{2A^3x}\Bigg] + V_{int}[n-1] \end{aligned} \tag{C.24}$$

$$p \approx \frac{GK}{2A}\left(\frac{1}{x} - 1\right) \tag{C.25}$$

OTA Schematic