University of Twente

Faculty of Electrical Engineering,
Mathematics & Computer Science

Integrated State-of-Charge Circuit for Rechargeable Batteries

Jasper Velner
MSc. Thesis
August 2008

Supervisors:
dr. ir. R.A.R. van der Zee
prof. dr. ir. B. Nauta
dr. ir. M.W. Kruiskamp
ing. R. Beumer

Report number: 067.3272
Chair of Integrated Circuit Design
Faculty of Electrical Engineering,
Mathematics & Computer Science
University of Twente
P. O. Box 217
7500 AE Enschede
The Netherlands

Acknowledgements

I would like to thank all the people who have supported me in performing this research in any way. In particular, my thanks go to my direct supervisor Ronan van der Zee and my supervisors at SiTel Semiconductor, Wim Kruiskamp and Rene Beumer for their invaluable support and input during the entire cause of this work.
I would also like to thank Bram Nauta and Rein Post for keeping up to date on the state of my research and for several interesting suggestions and remarks from their part.
My final thanks go to all the people at the Hengelo Design Center, which is part of SiTel Semiconductor, for their support. Several of these people have had a particularly profound impact on my work, and I would like to thank them in particular.
The foremost of these people is Lex Lubbers, who originally gave me a document with several project description, including this one. He has also given me invaluable information during the cause of this project, especially in the early design phase.
The second person whom I would like to extend special thanks to is William Aartsen for his insights into charge pump design and device mismatch.
Two other people to whom I owe special gratitude are Chris Smit and Jan Hovius, who have both helped me solve problems I had with Cadence on more than one occasion. In particular I would like to thank Chris for his patience when I was struggling with the AMS simulation (combined analog and Verilog simulation).

Abstract

The aim of this work is to create a low power system that can measure positive and negative voltages varying in magnitude from 100 µV to 100 mV with a relative accuracy of 1% or better. These voltages correspond to a current flow to or from a battery powering a DECT IC. Since these currents are quasi static, no accurate time domain representation is required, as long as the total charge is measured accurately. Since the system is constantly active and battery powered it needs to have a low current consumption. A maximum limit of 100 µA has been specified for this.

The final result is a system that is based on converting the input voltage to a current using a differential pair with a low transconductance and continuously integrating this current. When the output voltage of the integrator reaches a certain threshold, a fixed amount of charge is either added to or removed from the integrator. By counting these charge packages one can measure the total charge that was supplied to or taken from the battery. The system will function correctly if the absolute component values vary, but calibration is required if absolute charge values are to be measured.

A chopping mechanism that can swap the inputs of the system has been added to remove errors caused by offsets and other mismatches in the various building blocks. This mechanism has been functionally tested and the linearity of the system has been verified by manually swapping the inputs and measuring the average output assuming an ideal chopping mechanism. The run time required to test the chopper at its normal operational setting (30 second intervals) is too long for convenient simulations and has been derived mathematically.

With regards to the specifications, the system achieves a 0.24% accuracy without mismatch. With mismatch in the V-I converter added to the simulation the linearity decreases somewhat. The worst accuracy found in this case was 0.7%. Adding mismatch in the charge pump does not cause significant errors. The only issue arises with temperature variations. The measurement accuracy is still much too low when measured over a wide range of temperatures. An error of roughly ± 16 % is achieved here. This is caused by the combination of the charge pump that is used to create the charge packages, in combination with the reference source used to bias said charge pump. Adding a different (NTAT) reference for the charge pump would solve this problem.

The average current consumption is 8.0 µA, which is well within the specified limit.

Table of Contents

1. Introduction

1.1 General

Over the last few decades, portable equipment has become an increasingly important factor in daily life. This equipment is (generally) powered by one of many types of rechargeable batteries. An accurate indication of the available charge on this battery is desirable, both for the convenience of the user and the system that is used to recharge it.
However, creating an accurate indication of the usable charge on a battery has proven to be quite a challenge. In the past, quite a few different approaches have been used. There are systems that use the battery voltage, the current flows from and to the battery, predictive algorithms that use the state of a system to predict the current consumption, and so on. These systems can take into account factors like temperature, (dis)charge current, battery age, and many more. For a more detailed work on batteries and predicting or measuring the available charge on them, see [1].

This work focuses on a system that can measure the current from and to a battery over a large dynamic range and with high precision. The ultimate goal is to incorporate it into integrated circuits for the DECT (wireless phone) market. This system aims to (for the better part) remove the inaccuracy of the measurement from the state of charge indication. Other effects that influence the 'efficiency' of a battery, and thus the accuracy of the indication, are not taken into account for this work.

In chapter two, the system that is currently used is outlined shortly, and a number of different alternate solutions are described. From these alternate solutions, a choice for a system architecture is made. The different building blocks required for this architecture are described, developed and tested in chapter three. Chapter four focuses on chopping of the input terminals, which is used to increase the accuracy and reduce offsets. Finally, in chapter five, the complete system is described and evaluated.

1.2 Specifications

The system must be fully integrated using a 0.18 µm CMOS process. Since it is part of a battery operated piece of equipment, the power consumption must be as low as possible, with a maximum current consumption of 100 µA. For the implementation, a supply voltage of 1.8 V is assumed, and a 1.152 MHz clock is available at all times.

The battery currents are converted into a voltage using an external 0.1 Ω resistor. Depending on the mode of operation of the IC, these currents can vary between 1 mA and 1 A. This results in input voltages between 100 µV and 100 mV, which must be measured with an accuracy of within 1%. Also, because the battery can both be charged and discharged, the system must be able to measure both positive and negative voltages.

2. Comparison of Solutions

2.1 Current Solution

At present time the state of charge (SoC) is measured using a system that integrates the input voltage until a certain threshold is reached. When this occurs, the input terminals are exchanged. Every time this occurs, the SoC counter is updated. An simplified schematic of this system is depicted in figure 2.1.

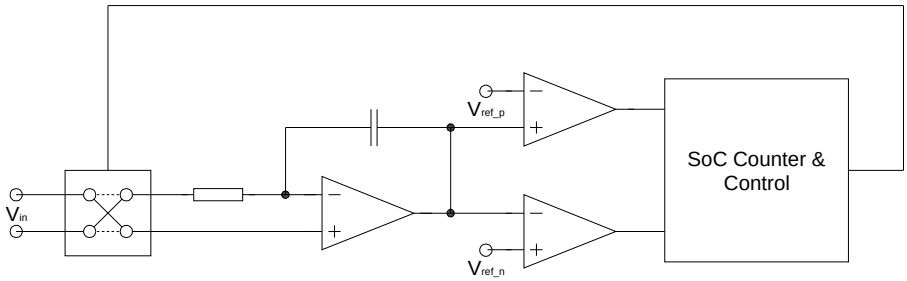

Figure 2.1: Simplified schematic of current system

This system suffers from several drawbacks. Firstly, when the inputs are exchanged, the reference level of the integrator changes. This gives rise to a step in the output voltage of the integrator, for which compensation must be applied.

A second problem comes from the offset voltage of the integrator. Theoretically, this voltage can be equal in magnitude, but opposite in sign, to the input voltage. If this happens, the output voltage of the integrator will not change, and the system will lock up. To prevent this, a timeout is built into the system, which will force the inputs to toggle if the SoC counter is not updated for an extended period of time.

Finally, an issue arises when very small inputs are measured. As the integrator output grows, so will the voltage at its inverting input (which is equal to the output voltage divided by the gain of the amplifier). This will reduce the voltage drop over the resistor, and hence the current that flows into the capacitor.
Since the minimum input voltage is 100 µV, the difference between the inverting and non-inverting inputs must remain smaller than 1 µV at all times. With a supply voltage of 1.8 V and assuming a full scale swing, the gain of the amplifier must be at least 900 mV / 1 µV = 900,000 or 119 dB. This is possible, but not trivial.

Because of issues like the ones mentioned before, the current system achieves an accuracy of about ±3 %, where the goal is ±1 %. Therefore, a different solution is desirable.

2.2 Straightforward Solutions

The most obvious choice for a solution might appear to be a ‘straightforward’ ADC, this is however not the most desirable approach for the following reason. The dynamic range of the system has to be 60 dB, and even at the smallest input the accuracy has to be within 1%, which adds another 40 dB. Also, since both positive and negative voltages have to be measured, another 6 dB is required. This sums up to

106 dB, or an ENOB of 17.3. Though this might be possible, it is needlessly difficult, since the desired accuracy is relative to the measured value, and not to full scale. Therefore, the absolute error at high inputs can be quite large without compromising the specifications. In fact, an error of 1% at maximum input voltage is already ten times larger than the actual minimum input voltage.

One could apply logarithmic compression before measuring, and then logarithmic expansion in the post processing. This operation favors small signals and reduces the required dynamic range. However, to do this a very clean logarithmic device is needed, and the best option for such a device is a bipolar transistor, which is not an easy device to create in a CMOS process. Furthermore, such a log conversion will depend strongly on the absolute parameters of the device in question, which are impossible to control with good accuracy. Therefore signal compression has not been further considered in this work.

Direct conversion, such as in flash or charge redistribution converters suffers from another drawback. For these converters to function properly, a sample and hold circuit is required. These circuits suffer from effects such as charge injection and voltage droop. Apart from this the sample rate is a major issue in a traditional converter. Though the signal is quasi DC, there can be fast transitions (for example when the class D amplifier of the DECT chip is enabled). In order to correctly represent these signals, the sample rate has to be very high to satisfy the Nyquist criterion.

Another possibility is to sample at a random frequency, that is below the Nyquist rate. As long as this frequency is not correlated to any other frequency in the system, the input will be sampled correctly. This is merely possible due to the fact that the only information of interest is in the average DC value of the input. More information about this subject can be found in [2].
The biggest problem with this scheme however is generation of a clock that is not related to any other clock in the system. For starters there are quite a few components that operate at different frequencies and secondly, all these frequencies are derived from a single master clock. Deriving a frequency that is not correlated to any other frequency in the system from this same clock is tricky at best.
If a system can be found that does not actually sample the input these issues can be avoided. By using an integrator as the core of the system, a charge can be build up that is proportional to the time integral of the input voltage and every so often an amount of charge can be removed to prevent the system from clipping.

Another method that can help to achieve the desired accuracy is to create a converter that has its output range segmented into several sub ranges for different input magnitudes. Several converters that employ such a mechanism have been considered and are treated in the next subsection.

2.3 Possible Implementations

2.3.1 Segmented Sigma Delta Converter

This implementation is essentially a multibit sigma delta converter. A conceptual schematic (made in SwitcherCAD III by Linear Technology) is displayed in figure 2.2. Instead of one comparator, as would be the case in a single bit sigma delta, there are three. These comparators control charge pumps that provide a feedback current.

For small signals, the output of the integrator will not change very rapidly, and only the middle charge pump will be activated. If the input becomes larger, the other comparators will trigger as well, doubling the magnitude of feedback current. This increases the speed of the system to compensate for the higher dV/dT at the output of the integrator. Because of this, the time constant of the integrator can be altered to allow a larger swing for small input signals, without risking clipping at large input signals. Such a larger swing reduces the effect of noise, hysteresis, etcetera on decision thresholds in the converter. For this example, the reference voltages are divided equally over the supply voltage.

A few notes on this schematic are in order. For starters there is the switch that is connected in parallel to the integration capacitor. This switch is purely there for simulation purposes, to set the initial charge of the capacitor to zero.

Secondly there is the 1 mΩ resistor in the feedback path. This too was added for simulation purposes only, because the feedback current can be measured at one of its terminals.

Thirdly, this system has not been adapted for input voltages around zero volts yet. Instead, a virtual zero at 900 mV is used. This makes feedback with charge pumps easier, since in order to sink current from the integrator, the supply voltage of the charge pump has to be below the reference level.

Finally there is the time constant of the system. This has been derived using some simple math. The feedback current with only one charge pump enabled is ± 1 µA. This current is in the same order of magnitude as the input current into the system, as will be shown later on. Since the minimum input voltage corresponds to a battery current of 1 mA, and the current into the system should not noticeably affect the result, a factor 1,000 between these two figures has been chose. With a 1 MHz clock and zero input, the swing at the output of the integrator will then be:

$$V_{Int} = \frac{1}{1.5\,pF} \int_{t=0}^{t=1\mu s} 1uA \cdot dt \approx 667mV \qquad \textbf{2.1}$$

Which fits nicely around 900 mV, leaving enough headroom to the supply rails (the swing will be between 0.45 V – 1.35 V). Note that the 1 MHz clock frequency is not quite equal to the 1.152 MHz clock in the system, but it is a bit easier for by-hand calculations.

The input resistor is such that, with maximum input, the current through it is 1.5 times the feedback current from one charge pump. This allows the integrator output to grow beyond the references at large voltages. This way only one charge pump will be active for small inputs, while two charge pumps may be activated for larger signals.

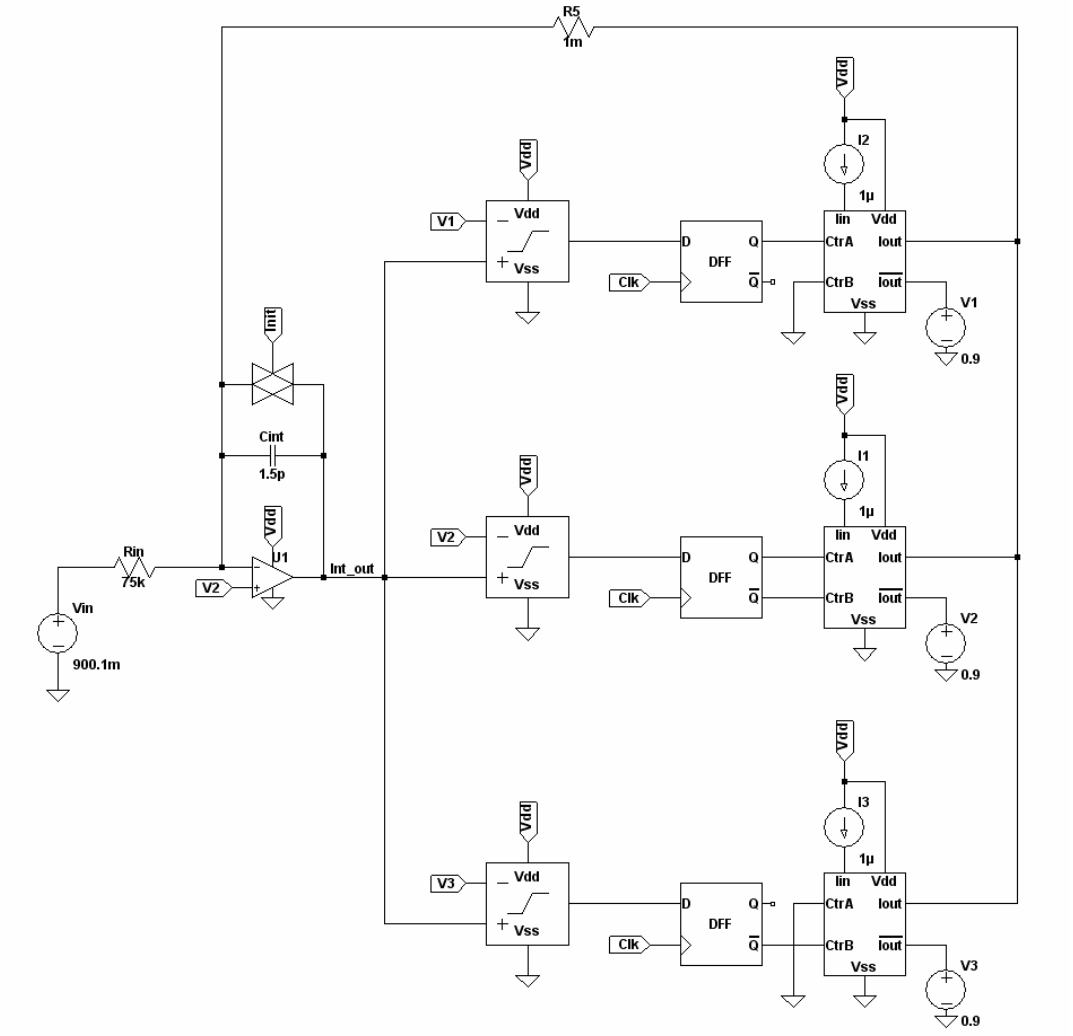

Figure 2.2: Conceptual segmented sigma delta converter

Figures 2.3 and 2.4 show the feedback current and integrator output voltage at respectively 100 µV and 100 mV inputs.

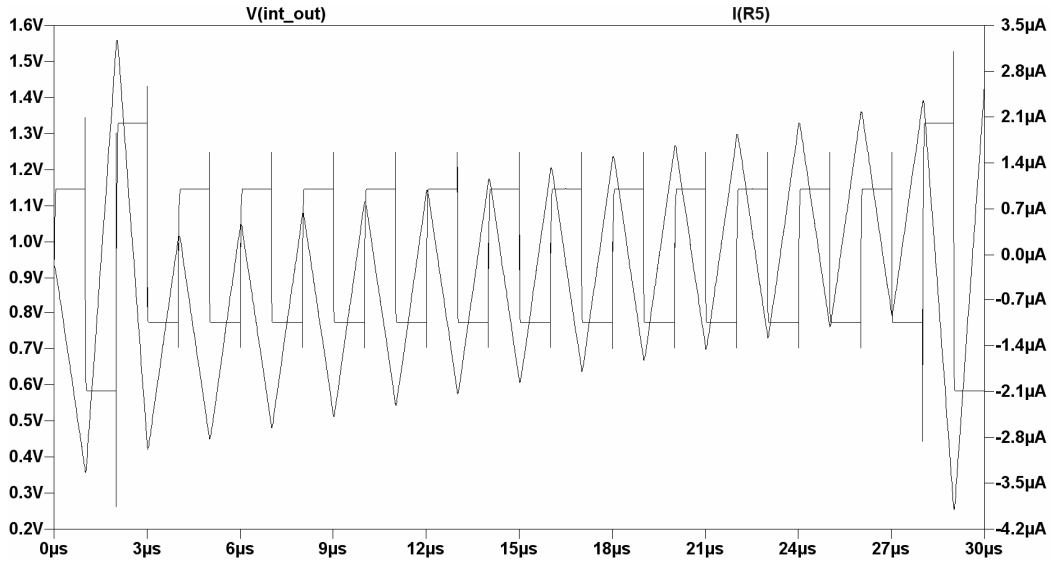

Figure 2.3: V_{in} = 100 µV

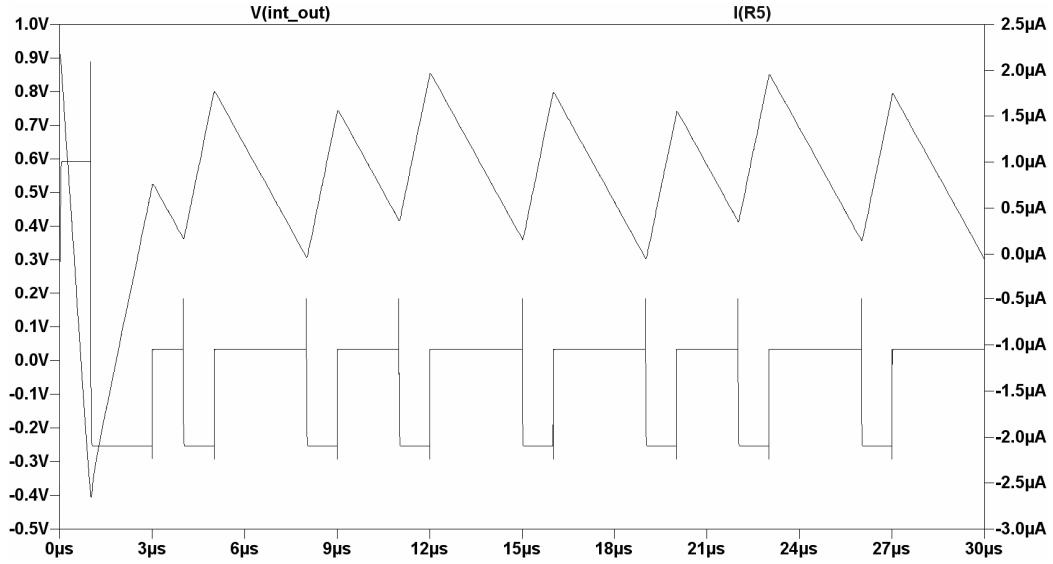

Figure 2.4: V_{in} = 100 mV

As can be seen, the average with a very small input is approximately zero. There is a steady rising trend in the signal, which occasionally triggers the second charge pump, which causes the average to be slightly above zero, as desired. The second plot displays the result at a larger input signal. Here the second charge pump is active for a much larger portion of the time, and the center charge pump never switches from sink to source mode. This means that, conceptually, the system works.

2.3.2 Segmented Dual Slope Converter

A normal dual slope converter is based on the principle of integrating the input for a fixed time, and then discharging the integrator back to zero using a fixed reference. The input can then be determined using the following relationships:

Integration phase:

$$V_{Int} = V_{In} \cdot t_{Int} \cdot \frac{1}{RC} \qquad \textbf{2.2}$$

Discharge phase:

$$V_{Dis} = V_{\mathrm{Re}\,f} \cdot t_{Dis} \cdot \frac{1}{RC} = -V_{Int} \qquad \textbf{2.3}$$

Combined:

$$V_{In} \cdot t_{Int} \cdot \frac{1}{RC} = V_{\mathrm{Re}\,f} \cdot t_{Dis} \cdot \frac{1}{RC} \Rightarrow \frac{V_{In}}{V_{\mathrm{Re}\,f}} = \frac{t_{Dis}}{t_{Int}} \qquad \textbf{2.4}$$

The idea of the segmented dual slope converter is to make the duration of the integration phase variable, or to be more precise an integer multiple of a fixed duration. The discharge phase will then only occur if the integrator output has passed a certain threshold. This way the integration is allowed to continue longer for smaller signals, allowing for better filtering of noise, fewer switching events and lower time and voltage resolution in the decision circuitry. The input/output relationship is now as in formula 2.5.

$$\frac{V_{In}}{V_{\mathrm{Re}\,f}} = \frac{t_{Dis}}{n \cdot t_{Int}} \qquad \textbf{2.5}$$

A conceptual schematic is displayed in figure 2.5.

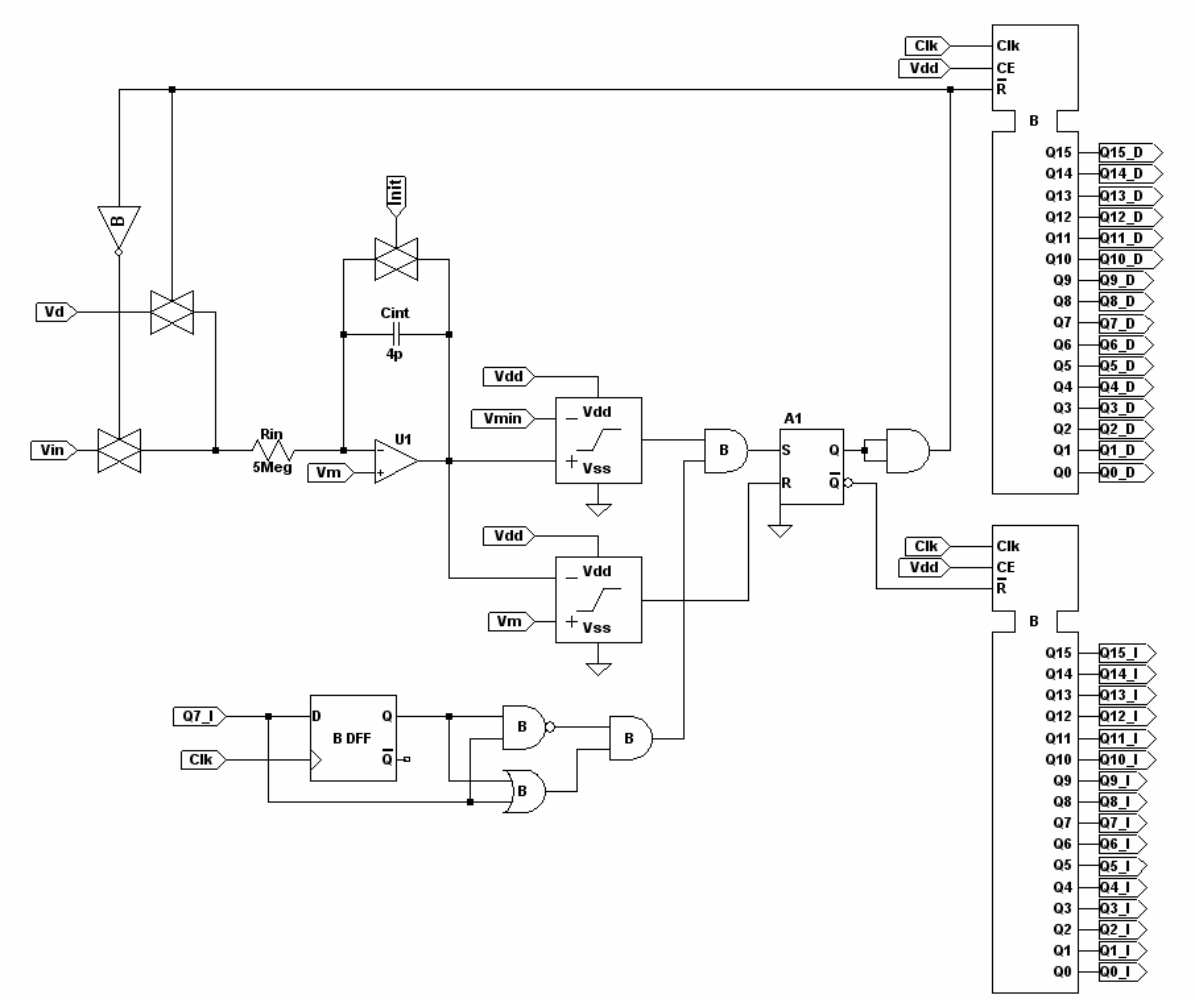

Figure 2.5: Conceptual segmented dual slope converter

This implementation again uses 900 mV as 'ground. Also, it only works for 'negative' voltages. This has been done for simplicity, while still providing enough functionality to prove the concept. Finally, the system works continuously, rather than having a fixed sample rate. This too has been done for simplicity of the test circuit.

The integrator in this example is comparable to the one used in the current system in terms of component values. On the right hand side of the schematic one can see the counters. The bit of logic on the bottom allows the latch that controls the input switches to be set by the comparator every time bit 7 of the integration counter toggles. That way the integration phase can end after an integer multiple of 128 clock cycles. The control latch is reset when the discharge phase ends (detected by the lower comparator).

Runs have been performed with an input of 100 µV and 100 mV again, and the results can be found in figures 2.6 and 2.7. Note that in the system, there are not yet any buffers at the outputs of the counters (to keep the image more clear). For the simulation these have been added, so the value is not reset until the next sample starts.

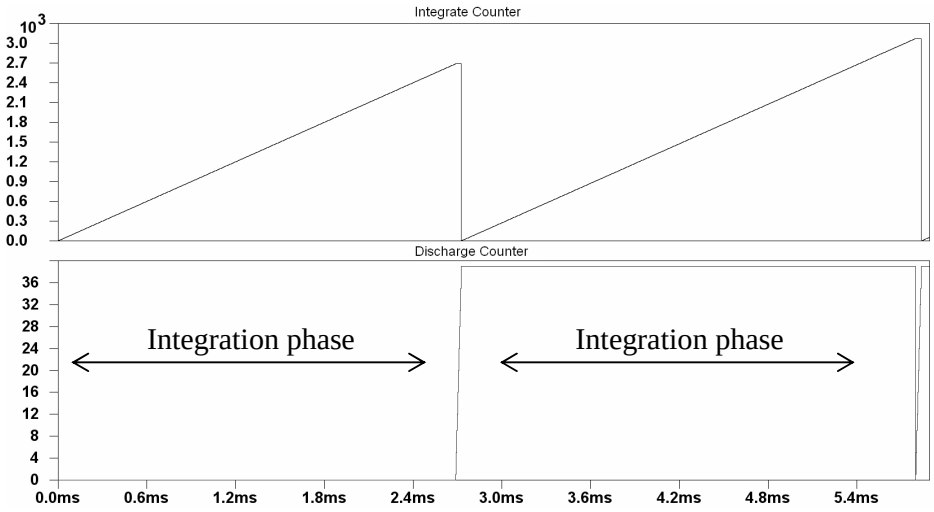

Figure 2.6: V_{in} = 100 µV

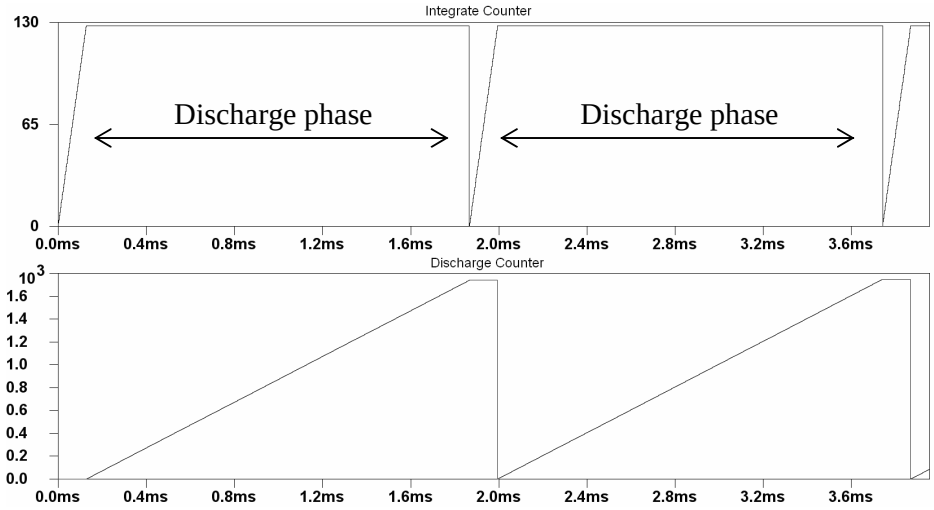

Figure 2.7: V_{in} = 100 mV

As can be seen, in case of the 100 µV input, the integrate counter is active quite long, and the discharge phase only takes a very short period. The integration phase takes 3072 clock cycles, and the discharge phase takes 39 clock cycles, leading to a quotient of approximately $12.69 \cdot 10^{-3}$.
On the other hand, with an input of 100 mV the opposite is the case. The integration step only takes 128 clock pulses (the minimum) and the discharge phase takes a lot longer than before, 1748 steps to be precise, leading to a quotient of about 13.66. The factor between these two quotients is slightly more than expected, but it is close. Therefore, since there have not been any optimizations to this conceptual design, the overall outcome looks promising.
Note that the first cycle from both runs should not be considered, since the initial conditions are not exactly equal to the conditions after the system has settled.

2.3.3 Charge-Dump Converter

This converter can be seen as a derivative from a sigma delta converter. The idea here is to continuously integrate the input signal, until a certain threshold is reached. When

this happens, a fixed amount of charge is removed from the integrator and the process repeats itself. By counting how many times charge has been removed (or added in case of a negative input), the total amount of charge that has been integrated is known.

The difference between a sigma delta converter and this design is that the 'dump' phase lasts a fixed amount of time. This way, the amount of charge that is removed during said dump phase is always equal and the total charge that has been removed (and thus previously stored) can be measured by simply counting the number of dump phases. A conceptual implementation can be found in figure 2.8.

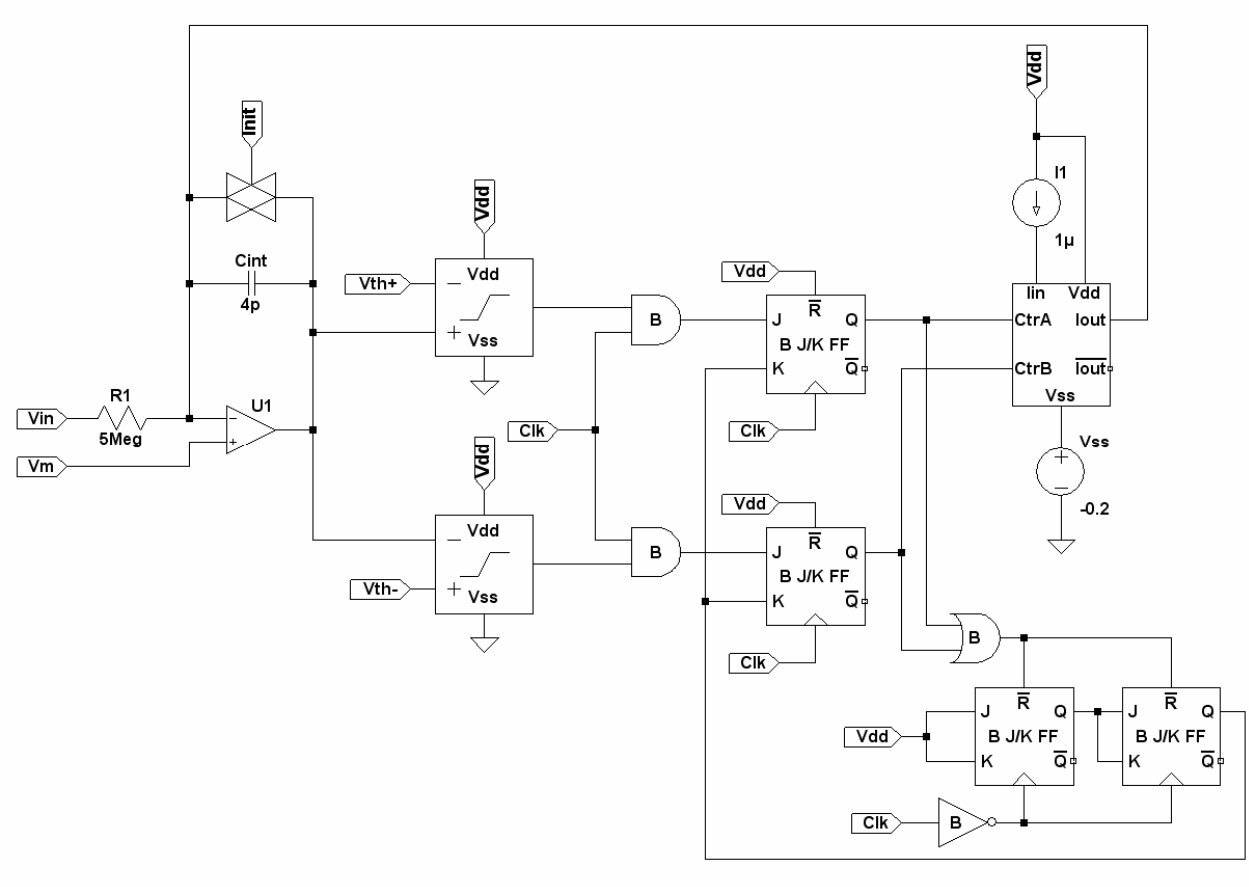

Figure 2.8: Conceptual charge dump converter

This system has only been simulated at inputs of 10 mV and 100 mV, because for smaller inputs the simulation time, and required memory, are excessive when using the SwicherCAD simulator.

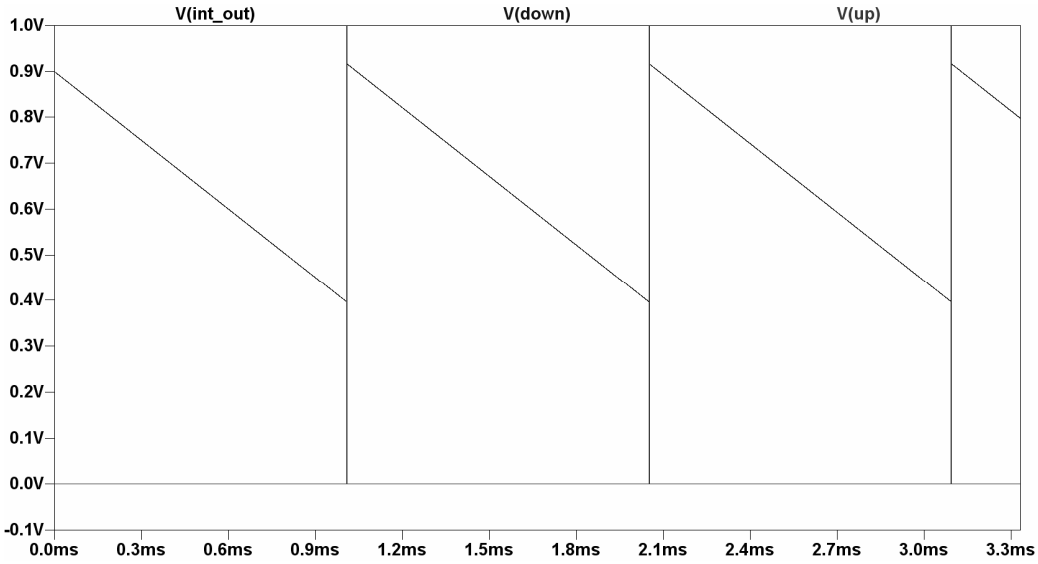

Figure 2.9: V_{in} = 10 mV

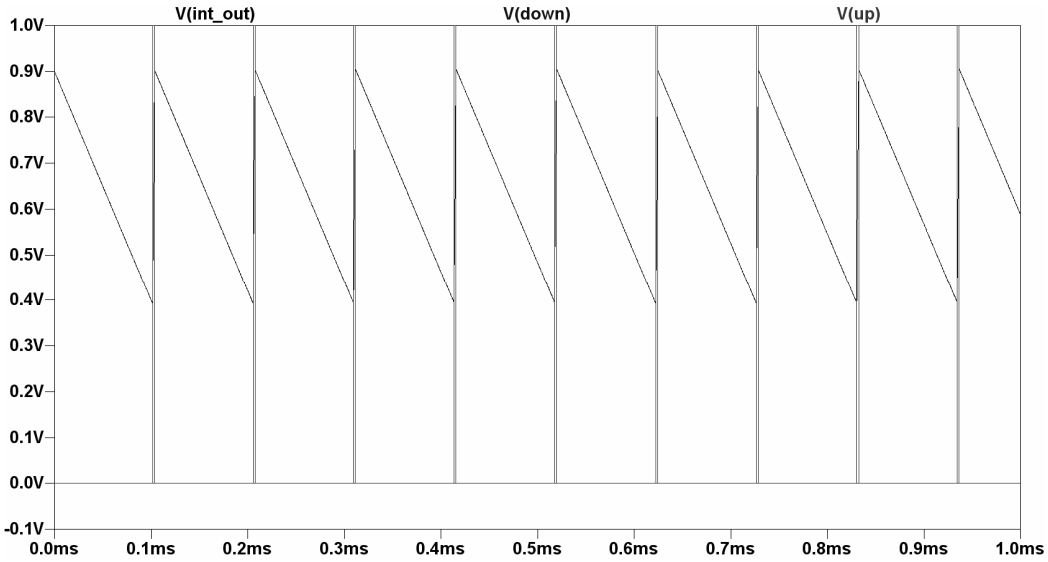

Figure 2.10: V_{in} = 100 mV

As can be seen in these plots, this conceptual system also works as intended. What remains now is to choose to best concept from the three designs that have been treated so far. An extensive comparison can be found in the next sub section.

2.4 Comparison on Non-Idealities

In order to choose the most suitable implementation, the three concepts that have been described so far will be compared on the following aspects:

- Non linearity of components
- Offset (matching)
- Noise
- Absolute tolerances
- Switching speed and other delays
- Capacitor leakage

- Charge injection and clock feed through
- Reference accuracy and stability
- Timing accuracy
- Post processing
- Sampling

2.4.1 Non Linearity of Components

Segmented Sigma Delta:
This is mostly an issue in the input stage. The integration needs to be linear within the specification of 1% for the maximum swing that can occur here. This means that the capacitor and resistor have to be linear, and the OPAMP has to have sufficient gain to suppress its own non-linearity by feedback.

For the feedback path, linearity is not a major issue. There are different possible states (due to the segmented nature of the system) but the accuracy of the feedback will be mostly be due to offsets between the different elements.

Segmented Dual Slope:
The integrator is a key part in this implementation as well, and will thus face the same demands for linearity as the segmented sigma delta.

The switches need not be very linear, because they connect to a very high input impedance, and any non linear resistance added by the switches will be small compared to this impedance.

The rest of this circuit is bi-stable by nature, so non-linearity is not an issue.

Charge Dump:
This system faces the same linearity demands as the segmented dual slope converter.

Conclusion:
The linearity demands for the three systems are similar, so non linearity will not play an important role in system choice.

2.4.2 Offset (Matching)

Segmented Sigma Delta:
The offset at the integration stage is of concern. Any offset that is present here will be integrated and give rise to an erroneous reading. A possible way to reduce this error is by chopping the inputs, but since the system is continuous, chopping without introducing extra errors is not trivial.

Another source of errors in this implementation is the offset in the comparators. With a one bit converter, an offset in the comparator will introduce a DC offset in the output. Offsets in the other converters will yield different quantization intervals, which translates to differential non-linearity.

Finally, the output DAC (be it a voltage DAC or a current DAC as shown in the conceptual design) will need to have it’s components matched to within 1%. If this is not the case, the feedback current will be non-linear and the output value will be distorted.

Segmented Dual Slope:
As before, any offset in the integrator will give rise to errors. Offset in the comparators is not an issue since this will only determine the upper and lower boundaries between which the integrator output can be. As long as these offsets are constant they will not contribute any error after the circuit has settled.
Since the final system will have a fixed sample rate, some dead time can be introduced. This is convenient for chopping the input, which will reduce the error caused by the input stage offset.

A potential problem that has been mentioned before is that the input offset and input voltage cancel each other. In this case a simple overflow detector can be used to reset the circuit and output a sample value of zero. By chopping after such an event, the double value will be obtained (offset + input) and the average value will still be correct.

Offset in the comparators does not matter, as it will only yield a different 'starting point' for the integrator voltage. It is necessary to minimize the hysteresis in the zero crossing OPAMP to minimize the difference between positive and negative input readings. On the other hand, chopping will also reduce these errors.

The final bottleneck caused by offset is in the matching between discharge sources for positive and negative inputs. Any mismatch here will be directly visible in the output reading, as a difference between input signals of opposite symbol but equal magnitude.

Charge Dump:
As with the segmented sigma delta, the fact that the input is always connected to the integrator will make chopping less easy than in the dual slope case. Otherwise the system faces similar issues to the segmented dual slope converter.

Conclusion:
From an offset point of view, the segmented sigma delta converter has the most drawbacks. Between the segmented dual slope converter and the charge dump converter, the segmented dual slope allows for the easiest chopping of the inputs. This makes the segmented dual slope the best solution when looking at offsets.

2.4.3 Noise

Even though noise is an integrate part of any analog system, it is not a large issue here. The system is used to sample quasi DC voltages over very long periods of time, giving it an extremely low noise bandwidth. Since noise has an average of zero by definition, it is safe to assume that it will be filtered out.

2.4.4 Absolute Tolerances

Note that each system has an integrator core. These integrators are made up of components that suffer from both mismatch and process spread. The same goes for the reference sources and feedback paths. Therefore, the absolute output of each of the converters mentioned here is meaningless without calibrating it to a known reference first.

Segmented Sigma Delta:
The only absolute tolerances that are of prime importance in this system are the size of the input resistor and capacitor. These must be such that clipping cannot occur even in the worst case scenario. This means that their values must be chosen such that even in the process corner where both resistor and capacitor values are much smaller than their designed values, the maximum swing out the integrator output leaves enough margin to the supplies.

Segmented Dual Slope:
The same demands as mentioned for the segmented sigma delta converter hold true here.

Charge Dump:
Clipping here will not easily be caused by the dimensions of the input resistor and integration capacitor, but rather by the capacitor and the feedback sources. This is because the time constant from the input to the integrator output is fairly large, and the feedback path can be activated on each clock flank. When activated however, the feedback path is enabled for a fixed period, during which clipping can occur. The feedback can be done using charge pumps, as shown in the concept, or by a voltage source and resistor combination. Since voltage sources and resistors require both these components to be within tolerance, charge pumps are the better option here.

Conclusion:
The systems face similar demands on absolute component tolerances. The charge dump system has a slight edge here, since a reference will be easier to control in absolute magnitude than a resistor value.

2.4.5 Switching Speed and Other Delays

Segmented Sigma Delta:
Comparator delay will not be a major issue most of the time. This is because the comparators are connected to flipflops, which means that errors can only occur if the comparator switches precisely on a clock flank.

The finite switching speed of the feedback DAC can give rise to problems though. Basically what will happen is that a new output code will not just depend on the input, but also on the last output code. If, for example, one of the outputs is high for two clock periods, then low for two clock periods, there will be one rising and one falling flank. If on the other hand the output toggles on every clock flank for four consecutive periods, there will be two rising and two falling edges, yielding a different average value. This is a problem because the output value is used in the feedback loop. This phenomenon is known as inter-symbol interference (or ISI) and is shown graphically in figure 2.11.

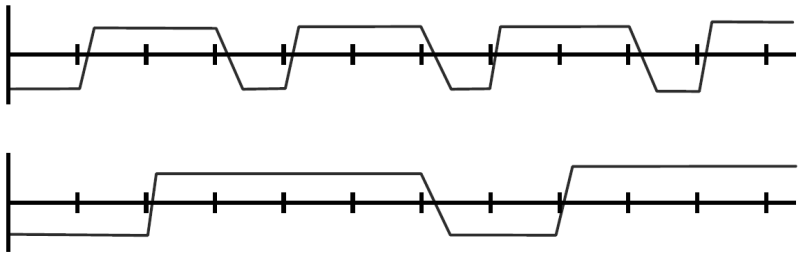

Figure 2.11: ISI displayed graphically

There are several ways to combat ISI, like using a return-to-zero architecture or using PWM rather than PDM. Return-to-zero forces the signal to toggle twice on each sample, which will increase the sensitivity to jitter and drift. PWM 'groups' high and low pulses to reduce the number of flanks. Both solutions do however add to the complexity of the system.

Segmented Dual Slope:

The first potential problem with limited switching (and propagation) speed that springs to mind with this particular architecture is the delay that is introduced by the input switches. However, the total error caused by these can be zero as long as all delays are equally long. This can be easily understood, since both the input and discharge voltage will be connected to the integrator with a certain delay, but will also be disconnected with the same delay. Figure 2.12 shows a graphical representation of this for two input voltages.

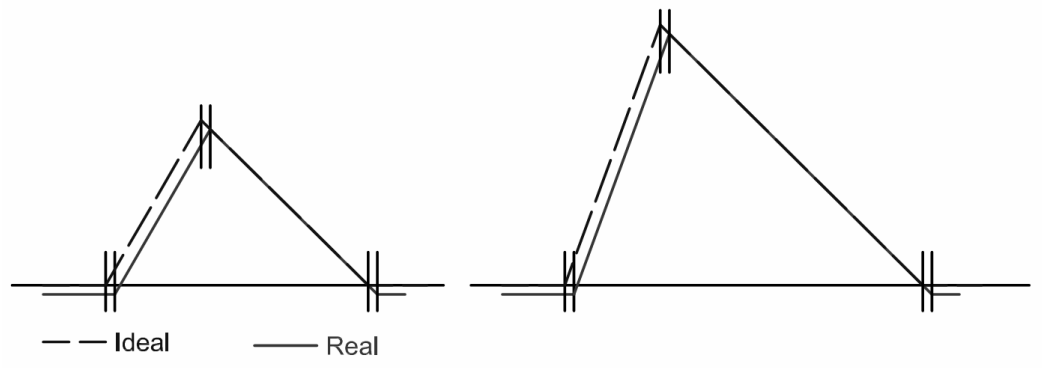

Figure 2.12: Errors due to limited switching speed

Errors may occur if switches have a very non-linear turn on and turn off characteristic, or if both input and discharge voltage are connected simultaneously. To prevent the latter situation, a short dead zone can be implemented.

Charge Dump:

Since the comparators in this system are clocked (in order to synchronize them with the charge pump timer), the integration can last up to one clock cycle too many. However, since the integrator output will continue to grow during this clock cycle, the output after the charge dump will be higher as well. This means the next cycle will take less time, negating the effect. Therefore, this effect will be averaged out over time and no net error will be introduced.

Conclusion:
The segmented sigma delta converter will suffer from ISI, which can be combated but will require attention. The other two systems will not be severely limited by delays.

2.4.6 Capacitor Leakage

Due to the similar input stages, each systems has similar demands on leakage. Essentially, the leakage current of the integrator must be no more than 1% of the minimum input current. How large this number is will depend on the final dimensioning of the system.
Note that, with the current dimensioning, the sigma delta converter has an edge here since it has a lower time constant (to allow the relatively fast switching in said converter). This low time constant is caused by a relatively low input resistance (hence large input current) and small integrator capacitance. Since the input current is higher than with the other converters, the effect of a certain, constant leakage current will be smaller than for the other converters.
On the other hand, capacitor leakage is expected to be very low so no real issues are expected in this area.

2.4.7 Charge Injection and Clock Feed Through

Segmented Sigma Delta:
This effect can occur in the feedback DAC (whether it is voltage or current output). This charge will be added directly into the signal path. Added to that, with this topology, the output will switch a lot when the input is small. This means that more charge injection will occur for small input values, which is of course very undesirable.

Clock feed through will give rise to a similar effect on the system, which holds true for all three systems discussed here.

Note that another option would be to use a switched capacitor integrator. With such an input stage, greater care must be taken to provide proper cancellation of the injected charge, since the switching speed in such an integrator will be higher than the speed at which the input switches are operated, hence more charge will be injected.

Segmented Dual Slope:
With proper compensation, the charge injected by the input switches can be small. If a discrete integrator is used, each switch operates only twice during each sample period, and since a minimum amount of charge is build up (i.e. the integration continues until a certain threshold is exceeded), this effect will in all probability be negligible.

Charge Dump:
This circuit only suffers from charge injection when charge is added or subtracted by the charge pump and the input is chopped. This happens relatively infrequent, hence no problems are expected here.

Conclusion:
The segmented dual slope converter and the charge dump system will suffer the least from charge injection and clock feed through. The segmented sigma delta has an added disadvantage, namely that for smaller inputs, the absolute error due to charge injection will be the largest, due to more frequent switching.

2.4.8 Reference Accuracy and Stability

Segmented Sigma Delta:

As said before, any offset in the comparators will cause offsets and/or non-linearity in the output value. Therefore both the absolute accuracy and the stability of the reference levels is paramount in this circuit.

Segmented Dual Slope:

The accuracy of the reference voltages in this setup is not extremely paramount. This is because absolute value of the charge flow is not important. If the reference voltage of either comparator is slightly off this does not matter either. For the 'zero crossing' comparator, the only thing that will differ is the voltage level at the integrator output after each conversion. The discharge phase will last an integer number of fixed time slots, and variations in threshold can make the difference between one interval more or less, but the final result of the division will not change by this.

The absolute value of the discharge sources is also not important. If their value differs from the designed value, the absolute output value will differ from the predicted value, but the factor between them will be constant and can thus be removed by calibration.

More important than absolute accuracies is the stability of the aforementioned sources. The ratio between input voltage and discharge voltage of a dual slope converter is equal to the ratio between the sample and discharge times. If the discharge voltage varies in time, the output reading will vary linearly with said voltage. Stability demands on this voltage are therefore strictly correlated to the accuracy of the converter.

Charge Dump:

The relative accuracy of the feedback sources is important (as mentioned under matching). The absolute value does not matter, much the same as with the segmented dual slope converter. The reference levels of the comparators will not influence the result, as long as the charge package that is added to or removed from the integrator remains constant, the outcome will be valid. This means that the stability of the references is important for this system.

Conclusion:

The segmented sigma delta is the only system that is sensitive to absolute reference accuracies. The other two systems only rely on reference stability. Therefore the segmented sigma delta is the least attractive system from this point of view.

2.4.9 Timing Accuracy

Segmented Sigma Delta:

The absolute accuracy of the clock is not very important for this topology. However, to the high switching frequency, clock jitter will have an influence on the outcome on a short term basis, where drift will influence the outcome on a longer term basis. Since the signal of interest is quasi static however, both effects will be filtered out over enough time (especially jitter, since it resembles HF noise).

Segmented Dual Slope:
This is paramount for a dual slope converter. Again the absolute accuracy is not of prime importance, since each output will depend on the ratio between two times. Clock drift during conversions is therefore liable to falsify the reading. However, since each conversion takes a relatively large number of clock cycles, any clock drift (which can be thought of as noise) will be low pass filtered, reducing the overall effect. This also means that clock jitter will not be a large issue (increasing the number of clock cycles between events while jitter remains unaffected will decrease the error introduced by said jitter). Added to that, the same filtering over a large number of samples will occur as with the sigma delta converter.

Charge Dump:
The accuracy of the clock will determine how long the charge pump is active. Therefore, drift and jitter will influence the outcome, but the absolute clock frequency will not. Once again filtering over time will also reduce the error caused by both effects.

Conclusion:
Neither system is very vulnerable to timing accuracies.

2.4.10 Post Processing

Segmented Sigma Delta:
For the sigma delta, either a (long) decimation filter, or a pulse counter of some sorts is needed. Note that though a pulse counter may seem simple, it has inherent problems in differentiating between a positive and a negative input (changing the sign of all pulses in a sigma delta changes the sign of the output, but not the number of pulses).

Segmented Dual Slope:
For the dual slope converter, there are two timer values that make up the final conversion value. The result will be the value of the discharge timer (A) divided by the value of the timer that runs during the sample phase (B). Though digital division is not straightforward, it need not be extremely fast, since the sample rate can be relatively low (using uncorrelated undersampling). This means the division can be carried out by successive approximation as follows: $B > A - n \cdot B$
In order to increase the accuracy a final check can be performed to see if the remainder is smaller or larger than B/2 to use proper rounding instead of flooring.

Charge Dump:
The charge dump circuit has by far the simplest post processing. It merely has to increase or decrease a counter whenever charge is added/removed from the integrator by the charge pump.

Conclusion:
From a post processing point of view, the charge dump converter offers the easiest solution.

2.4.11 Sampling

Segmented Sigma Delta:
This converter has its input constantly connected to the signal. This means that no information is lost at any point, and sampling is not an issue. Note the difference

between this system and an audio converter, where the data has to be sampled fast enough for an accurate representation of a certain frequency band.

Segmented Dual Slope:
This system requires a clock to determine the start of a new conversion, and the inputs will be disconnected during a part of the conversion. During this period (and in the dead time after the conversion has finished) the input signal is not sampled, and data is lost. This leads to the same considerations that were mentioned in section 2.2 for a classical converter.

Charge Dump:
Same as with the sigma delta converter, the signal is always connected to the charge dump converter, meaning there is no sampling clock.

Conclusion:
The segmented dual slope converter will be difficult to implement due to the need for an uncorrelated sample clock. The other two systems do not suffer from this drawback.

2.5 Conclusion

A tabular overview, marking which system is strong at which point can be found below in table 2.1. The X's mark the converter that is strongest at that particular aspect. System 1 is the segmented sigma delta converter, system 2 the segmented dual slope and system 3 the charge dump converter.

System: / **Aspect:**	**1**	**2**	**3**
Non linearity of components	X	X	X
Offset (matching)		X	
Noise	N.A.	N.A.	N.A.
Absolute tolerances			X
Switching speed and other delays		X	X
Capacitor leakage	X		
Charge injection and clock feed through		X	X
Reference accuracy and stability		X	X
Timing accuracy	X	X	X
Post processing			X
Sampling	X		X

Table 2.1: Overview

From the previous section and table 2.1, the conclusion can be made that the charge dump converter is the most attractive solution to the problem at hand. Therefore this system will be implemented and tested.

3. Building Blocks

This chapter focuses on the development and testing (by means of simulation in Cadence 5) of the different building blocks that make up the final system. Four main building blocks can be discerned in the conceptual schematic of figure 2.8, namely:

- An input stage comprised of an integrator
- A set of comparators
- Some control logic
- A charge pump

As mentioned before, there will be a chopping mechanism as well. However, this mechanism will be treated separately in chapter 4.

3.1 Input Stage and Reference

One of the issues with the system that is currently in use is the step at the output that arises when the inputs are chopped. Another issue is the high gain that is required for the OPAMP that is used in the integrator. Finally, in order to have input voltages around 0 V, the reference level has to be 0 V as well. This would thus require the charge pump to have a negative supply voltage below GND.

In order to avoid all these problems, a V-I converter will be used instead of a resistor. Such a converter has a high output impedance, which means that the reference voltage of the integrator OPAMP can be arbitrarily chosen (within certain limits), which solves the problem of the negative supply for the charge pump.
Another advantage of a current mode output is that the OPAMP does not need as much gain. If the OPAMP has a gain of over 100 (over its entire output swing), the voltage change at the input will be within 1% of the voltage change at the output. This is easily achieved, and strictly not even necessary. As long as the charge packages that are added or subtracted in the dump phase are fixed the system will function properly. An error caused by low gain will only influence the moment at which the charge pump is activated (much as an offset in either comparator will).
Finally, the problem with the voltage step that occurs when the inputs are exchanged is also solved by using a V-I converter. Swapping the inputs with such a system will only change the direction of the output current, the voltage level is fixed by the reference level of the OPAMP.

Figure 3.1 gives a high level overview of the complete input stage as it will be developed in the following subsections. The PTAT source is used to provide an amount of temperature compensation, as will be shown later on.

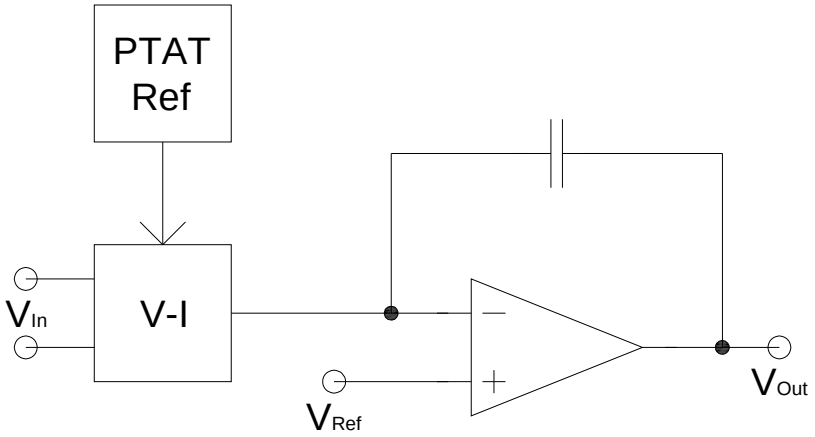

Figure 3.1: Top level of input stage

3.1.1 Theoretical Development of Input Stage

The input stage needs to meet the following demands:

- Linear within 1% over an input range of -100 mV … 100 mV
- Transimpedance in the order of magnitude of 5 MΩ (the same as used at present)
- Low offset
- Stable over a large temperature range (-40 °C … 90 °C was used for testing)
- Functional over all process corners
- High output impedance

The first demand automatically leads to a P-MOST input pair for the V-I converter. Since no negative supply voltages are available, an N-MOST input stage would not turn on.
The second demand leads to large transistors, from which the idea to use a differential pair arises. The linearity of such a pair increases if the length of the devices is increased (as will be shown later on). With long devices, the effect of channel length modulation will also be limited, automatically leading to a high output impedance (and satisfying the last demand).

Since the input voltage is lower than zero volts, a current mirror load, which requires a drain-source voltage higher than V_{TH} for at least one transistor, will cause the input transistors to enter the triode region. This means that a folded cascode is needed. A current mirror can then be used to create a single ended output. This leads to the conceptual schematic in figure 3.2.

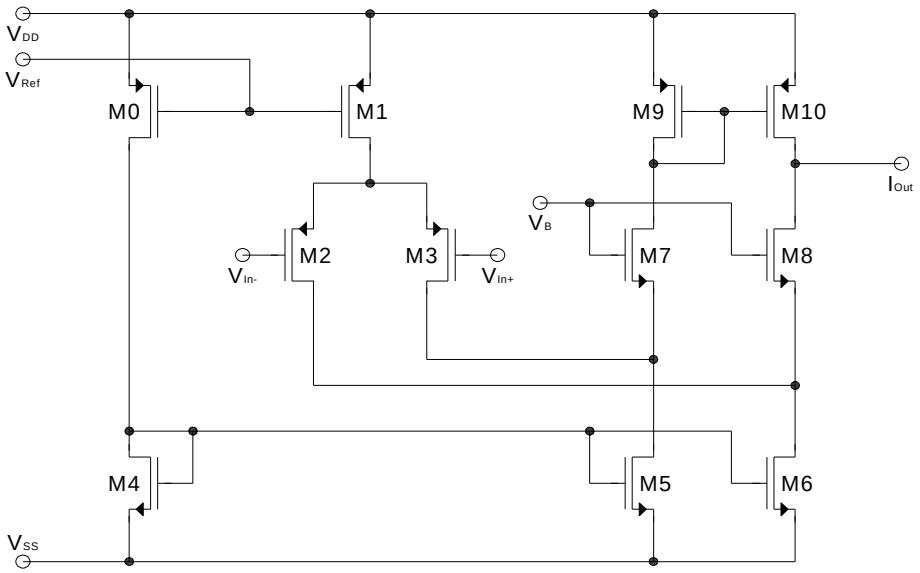

Figure 3.2: Concept for V-I converter

The following relationship holds for a differential pair [3], neglecting secondary effects:

$$\Delta I_D = I_{D,1} - I_{D,2}$$
$$\Delta V_i = V_{i,1} - V_{i,2} \qquad \textbf{3.1}$$
$$\Delta I_D = \frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L} \cdot \Delta V_i \cdot \sqrt{\frac{4 \cdot I_b}{\mu \cdot C_{OX} \cdot \frac{W}{L}} - \Delta V_i^{\,2}}$$

Here $\frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L} \cdot \Delta V$ is linear, so any non-linearity comes from the square root term. In this square root term, there is a constant part and a part that is dependant on the differential input voltage. Therefore, the non-linearity of the transfer function can be viewed as the difference between the square root term at zero and at maximum differential input voltage. If the maximum error is defined as a fraction ε, it can be found using:

$$\frac{\sqrt{\frac{4 \cdot I_b}{\mu \cdot C_{OX} \cdot \frac{W}{L}}}}{\sqrt{\frac{4 \cdot I_b}{\mu \cdot C_{OX} \cdot \frac{W}{L}} - \Delta V_i^{\,2}}} \le 1 + \varepsilon \qquad \textbf{3.2}$$

Using simple mathematics, the following can be derived:

$$\frac{W}{L} \le \frac{4 \cdot I_b}{\mu \cdot C_{OX} \cdot \Delta V_i^{\,2}} \cdot \left(1 - \frac{1}{(1+\varepsilon)^2}\right) \qquad \textbf{3.3}$$

Since the system needs to be highly linear, a good approximation for the differential current can be derived from formula 3.1, by removing the square of the differential input voltage from the square root term, leading to formula 3.4.

$$\Delta I_D = \frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L} \cdot \Delta V_i \cdot \sqrt{\frac{4 \cdot I_b}{\mu \cdot C_{OX} \cdot \frac{W}{L}}} \qquad \textbf{3.4}$$

The transimpedance can then be found from dividing the input voltage by this current, leading to formula 3.5.

$$R_{Eq} = \frac{2 \cdot L}{\mu \cdot C_{OX} \cdot W} \cdot \sqrt{\frac{\mu \cdot C_{OX} \cdot \frac{W}{L}}{4 \cdot I_b}} = \sqrt{\frac{L}{W} \frac{1}{\mu \cdot C_{OX} \cdot I_b}} = \frac{1}{gm_{2,3}} \qquad \textbf{3.5}$$

As can be found in appendix B, the following first order model parameters are used for the MOSFETS:

$\mu_n \cdot C_{OX}$	0.28 mA/V^2
$\mu_p \cdot C_{OX}$	0.048 mA/V^2
$V_{TH,N}$	390 mV
$V_{TH,P}$	-390 mV
L_{Min}	180 nm
L_{Max}	20 µm

Initially, a very low W/L of 500 nm / 100 µm was used for the input devices. To achieve this, five maximum length MOSFETS were placed in series.
In order to achieve a transimpedance of 5 MΩ, the bias current with these values needs to be 167 nA. This results in a linearity of 0.18 % at ±100 mV. To allow for some margin, the error will be considered at ±150 mV as well (and this value will be used as a benchmark). At said value, the linearity error will be 0.4 %, which is well within limits.

The rest of the circuit needs to be dimensioned now. Starting with the reference current sources in the circuit, M0 and M1. As can be seen in appendix B, the effect of the drain-source voltage of the MOSFETS on the drain current, caused by channel length modulation, starts to flatten out at lengths of about 5 µm. In order to achieve a high output resistance for these sources, 4 times this value has been chosen, resulting in maximum length devices.
As will be shown later, the reference generator creates a reference voltage that yields 1 µA when applied to a 1/1 device. Since 167 nA is required, the width of the devices must be 3.33 µm. For simplicity, this has been changed to 4 µm, yielding a slightly higher bias current of 200 nA. This will slightly increase the linearity and slightly decrease the transimpedance (to 4.56 MΩ).

The dimensioning of M4, M5 and M6 is based on matching. Mismatch between M5 and M6 will directly lead to a mismatch in the output current. M4 needs to have the same dimensions as M5 and M6, because these sources both need to conduct the reference current supplied by M1. The length of these transistors, which both

influences K factor mismatch and V_{TH} mismatch, has been chosen at 100 µm, which is the same as that of the input pair. The width has been set to 5 µm.

With these dimensions, and a drain current of 200 nA, the gate-source voltage of M5 and M6 will be:

$$I_D = \frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L} \cdot \left(V_{GS} - V_{TH}\right)^2$$

$$200\ nA = \frac{1}{2} \cdot 0.28\ \frac{mA}{V^2} \cdot \frac{5\ \mu m}{100\ \mu m} \cdot \left(V_{GS} - 0.39\ V\right)^2 \qquad \textbf{3.6}$$

$$V_{GS} \approx 560\ mV$$

This means the drain voltage of M5 and M6 can be no less than 170 mV if these devices are to remain in saturation. Since the drain voltages of M2 and M3 (the input pair) can be one threshold above their gate voltage before entering the triode region, the minimum input voltage is 170 mV – 390 mV = -220 mV which is well beyond the lowest specified input voltage.
In order to keep the input devices in saturation at a -150 mV input voltage, the drain voltage of M5 and M6 can be no more than -150 mV + 390 mV = 240 mV. This means that for all devices to remain in saturation, the drain voltages of M5 and M6 have to be between 170 mV and 240 mV.

The dimensioning of M7 and M8 (which are source followers that maintain a relatively constant voltage at the drains of M5 and M6), and the choice of V_B is based on this demand. With zero input, both M7 and M8 carry a current of 100 nA. Their length has been set to a relatively large value of 10 µm. The reason for this is that shorter devices have been found to conduct 100 nA in moderate inversion, giving an adverse effect on the output resistance and hence the stability of the bias voltage.

With an input swing between -150 mV and +150 mV, the output current lies between -30 nA and +30 nA. Since this current is divided equally between the two branches, and each branch carries 100 nA with zero input, the current through M7 and M8 can vary between 85 nA and 115 nA. The width of these devices has been chosen such that the voltage swing at their sources caused by these current differences is approximately 7 mV. This is about 10% of the range in which all devices remain in saturation. Using formula 3.7, the gate-source voltage as a function of W can be found with a current of 85 nA ($I_{D,Min}$).

$$I_{D,Min} = \frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L} \cdot \left(V_{GS,Min} - V_{TH}\right)^2$$

$$V_{GS,Min} = \sqrt{\frac{2 \cdot I_{D.Min}}{\mu_n \cdot C_{OX}} \cdot \frac{L}{W}} + V_{TH} \qquad \textbf{3.7}$$

$$V_{GS,Min} = \sqrt{\frac{2 \cdot 85\ nA}{0.28\ \frac{mA}{V^2}} \cdot \frac{10\ \mu m}{W}} + 0.39\ V$$

Formula 3.8 relates the minimum and maximum drain-source voltages.

$$I_{D,Max} = \frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L} \cdot \left(V_{GS,Max} - V_{TH}\right)^2$$
$$I_{D,Max} = \frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L} \cdot \left(V_{GS,Min} + V_{d,\max} - V_{TH}\right)^2 \qquad \textbf{3.8}$$

By substituting $V_{GS,Min}$ result into formula 3.8, and knowing that $V_{d,Max}$ is 7 mV, a value for W can be derived as shown in formula 3.9.

$$I_{D,Max} = \frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L} \cdot \left(\sqrt{\frac{2 \cdot I_{D,Min}}{\mu_n \cdot C_{OX}} \cdot \frac{L}{W}} + V_{TH} + V_{d,\max} - V_{TH} \right)^2$$

$$115\ nA = \frac{1}{2} \cdot 0.28\ \frac{mA}{V^2} \cdot \frac{W}{10\ \mu m} \cdot \left(\sqrt{\frac{2 \cdot 85\ nA}{0.28\ \frac{mA}{V^2}} \cdot \frac{10\ \mu m}{W}} + 7\ mV \right)^2 \qquad \textbf{3.9}$$

$$W \approx 3.3\ \mu m$$

A width of 4 µm was used, to be on the safe side. V_B is chosen such that with zero input, the bias voltage is approximately 200 mV, which is roughly in the middle of the 'safe zone'. With a the dimensions found in the previous calculations, an gate voltage of approximately 630 mV is required to achieve this. This voltage is obtained by steering 200 nA through a 250 nm / 10 µm MOSFET diode.

Like M5 and M6, M9 and M10 have been given a length of 100 µm. Their width is smaller than that of M5 and M6, namely 2 µm. The reasoning behind this is that ideally each of these devices will carry 100 nA if the input voltage is zero. With that current, the gate-source voltage will be:

$$I_D = \frac{1}{2} \cdot \mu_p \cdot C_{OX} \cdot \frac{W}{L} \cdot \left(V_{GS} - V_{TH}\right)^2$$

$$100\ nA = \frac{1}{2} \cdot 0.048\ \frac{mA}{V^2} \cdot \frac{2\ \mu m}{100\ \mu m} \cdot \left(V_{GS} - 0.39\ V\right)^2 \qquad \textbf{3.10}$$

$$V_{GS} \approx 850\ mV$$

With a supply voltage of 1.8 V the voltage at the drain of M9 will thus be close to mid scale. This means that if the reference voltage of the integrator OPAMP is set at mid scale, the system will be quite well balanced.

Figure 3.3 shows an implementation of the circuit in Cadence with the device sizes as they have been derived in this section.

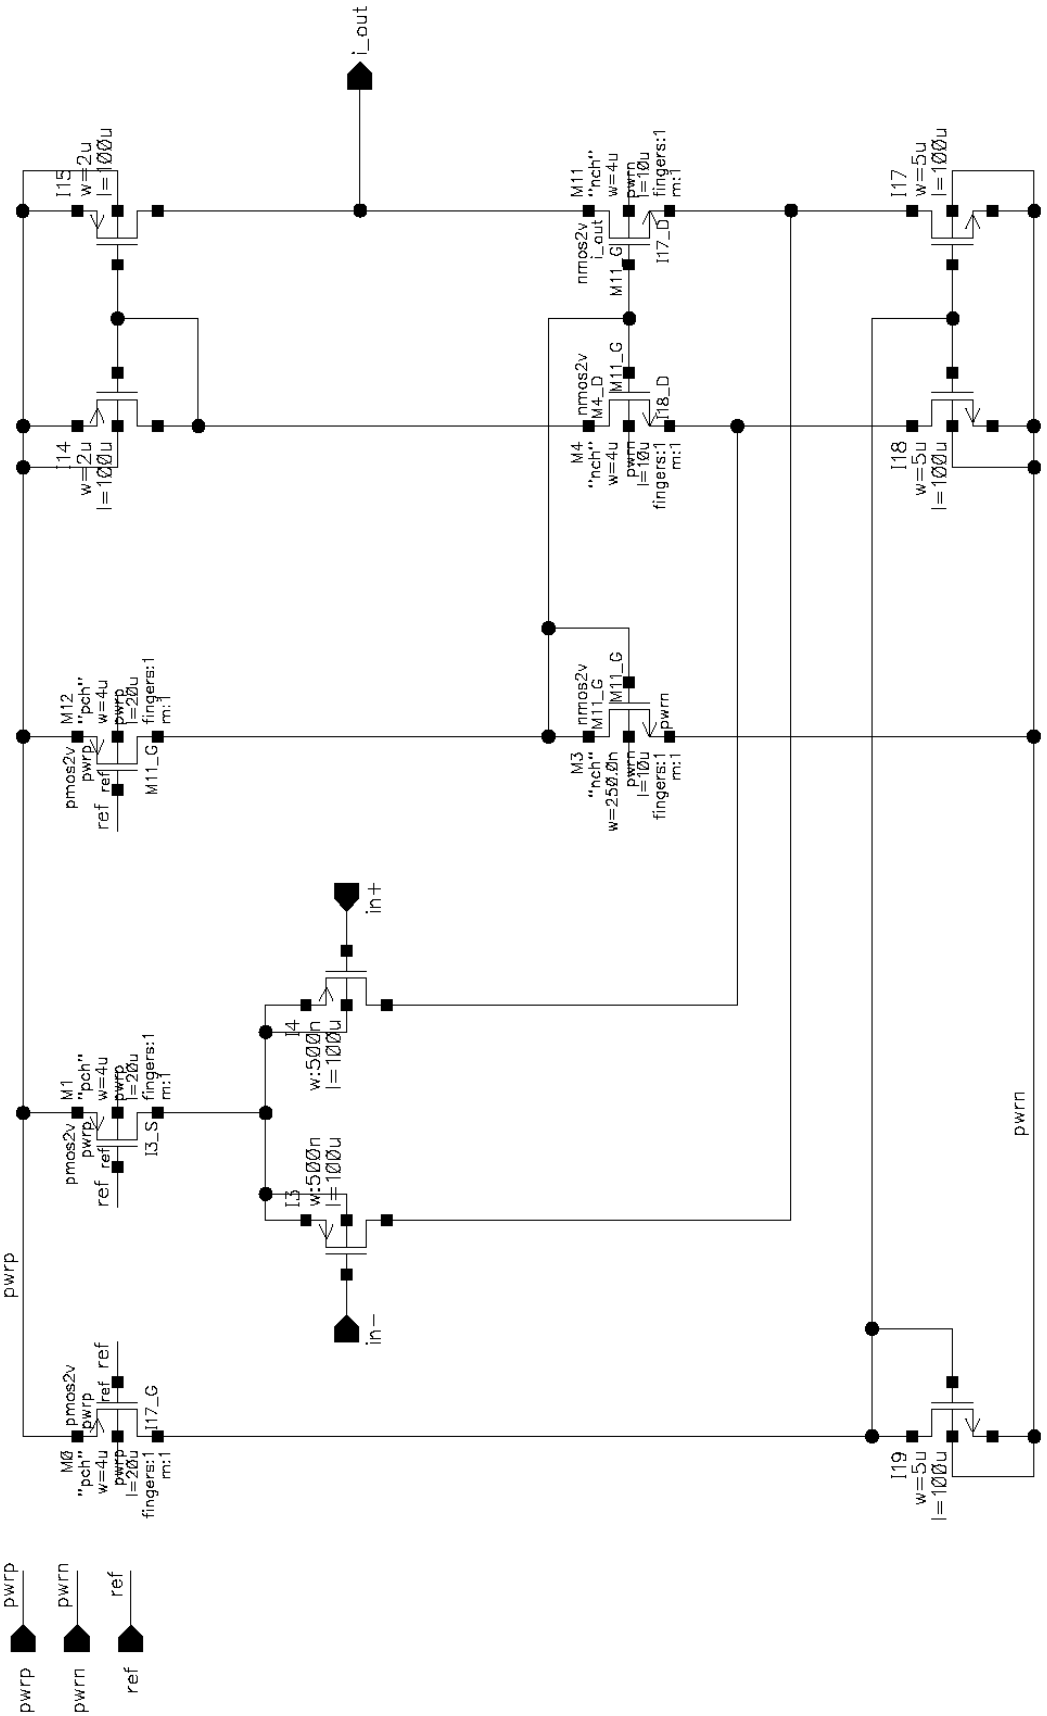

Figure 3.3: Cadence implementation

3.1.3 Current Reference

With an initial system designed, the reference generator has to be developed. Initial tests on an early version of the circuit with a constant current source show a negative dependence between the output current of the V-I converter and the temperature, as depicted in figure 3.4. Though not linear, a fairly good linear approximation can be made. Figure 3.5 shows a possible implementation for a PTAT reference that could negate this dependence. This figure is based on a circuit described in [4].

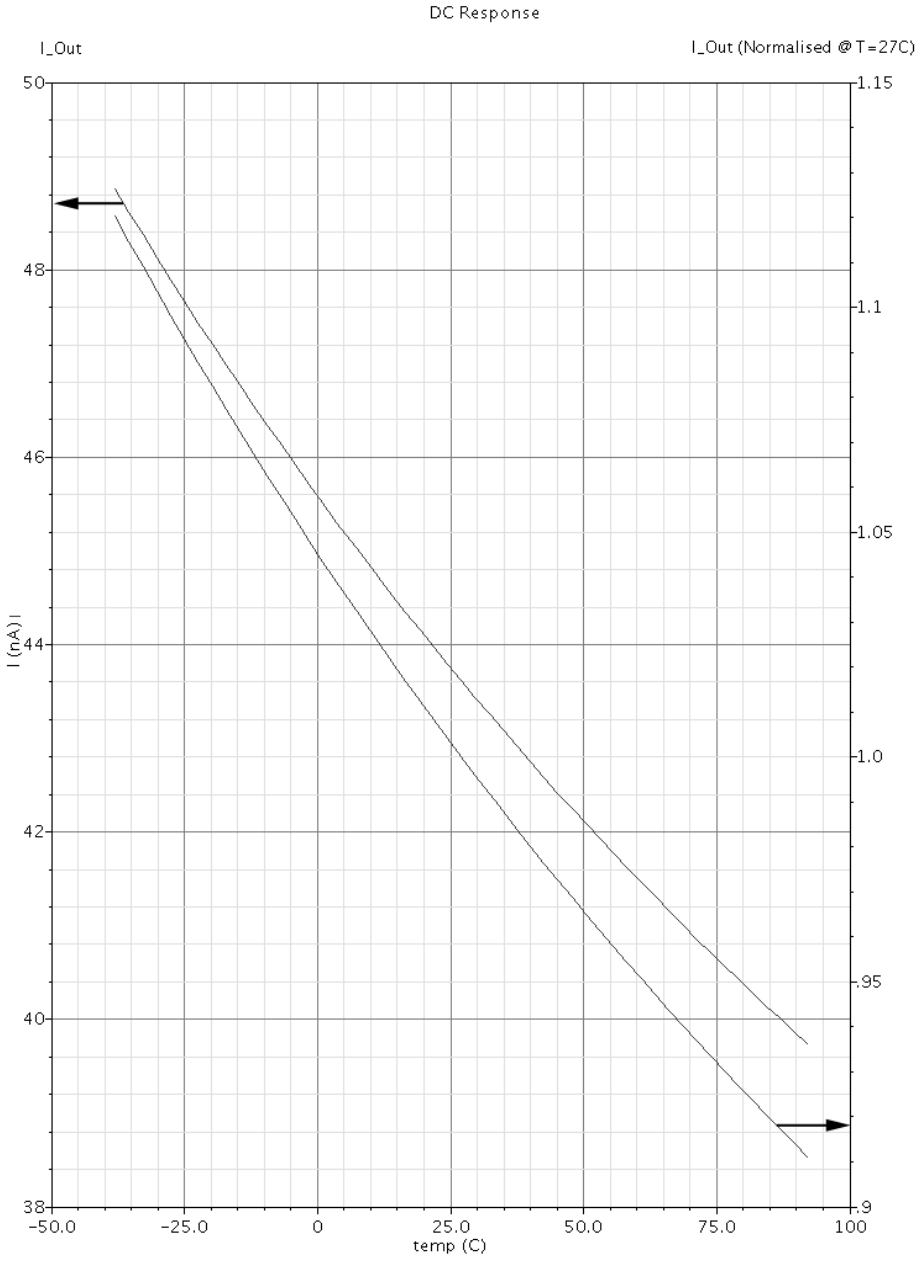

Figure 3.4: Output current versus temperature (V_{In} = 100 mV)

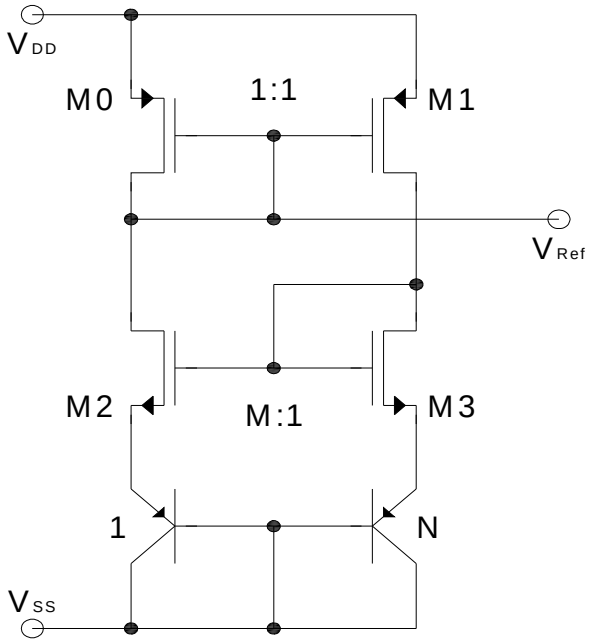

Figure 3.5: Conceptual PTAT source

The reader will probably recognize a supply independent bias generator in circuit 3.5. This circuit has two stable operating points. The first is the trivial solution, where all transistors are off. A startup circuit will be added to prevent this occurrence. The second operating mode has equal, non-zero currents in both branches. An equation that relates the current to the circuit parameters can be derived by equating the currents in both branches. The choice has been made here to choose the value for N and then set M depending on the desired current. Note that the BJTS in this circuit are created using parasitic devices in the CMOS process.

The length of each MOSFET in this reference has been set to 5 µm (as mentioned before, channel length modulation is low at this value, and should be low enough to allow calculations based on the simple quadratic model). N has been set to 5, which is a rather arbitrary value and was selected as a compromise between having a significant difference between $V_{S,M2}$ and $V_{S,M3}$. A reference current of 800 nA has been chosen. This again is rather arbitrary, and was found to be a convenient value. The width of M0 and M1 was set to 4 µm, which means that if the reference is connected to a 1/1 device, a current of approximately 1 µA will flow.

For the B-E current in a BJT the following relationship holds:

$$I = I_S \cdot e^{\frac{V_{BE}}{V_T}} \qquad \textbf{3.11}$$

Where I_S is the saturation current, $2.38 \cdot 10^{-17}$ A for this process and V_T is the thermal voltage, which is 25.9 mV at T = 300 K. For the right branch of the circuit, this current will be N times larger given the same B-E voltage. This can be rewritten to find the B-E voltage, which can then be substituted in the quadratic MOSFET equation as shown in formula 3.12.

$$I_{D3} = \frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L} \cdot \left(V_G - V_T \cdot \ln\left(\frac{I_{D3}}{I_S} \right) + V_T \cdot \ln(N) - V_{TH} \right)^2 \qquad \textbf{3.12}$$

Solving for V_G, which is the gate voltage of both M2 and M3, yields:

$$V_G = -V_T \cdot \ln(N) + V_{TH} + V_T \cdot \ln\left(\frac{I_{D3}}{I_S} \right) + \frac{\sqrt{\frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L} \cdot I_{D3}}}{\frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L}} \qquad \textbf{3.13}$$

Since the current in both branches is equal, $I_{D3} = I_{D2} = I_D$. Formula 3.13 can be substituted in the formula for the current through M2, which then yields equation 3.14.

$$I_D = \frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{M \cdot W}{L} \cdot \left(\sqrt{\frac{I_D}{\frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L}}} - V_T \cdot \ln(N) \right)^2 \qquad \textbf{3.14}$$

From formula 3.14, formula 3.15 can be derived. This formula relates M to the other component parameters.

$$M = \frac{I_D}{\frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L} \cdot \left(\left(\sqrt{\frac{I_D}{\frac{1}{2} \cdot \mu_n \cdot C_{OX} \cdot \frac{W}{L}}} - V_T \cdot \ln(N) \right)^2 \right)} \qquad \textbf{3.15}$$

With L = 5 µm, N = 5 and I_D = 800 nA a feasible solution is W = 1.3 µm and M = 2.5. Setting M to 2 instead allows for the use of two identical devices in parallel to create M2, which is quite elegant from a matching point of view. The current will deviate from the calculated value, but since only first order calculations have been used to arrive at this value, simulations will have to be performed to verify this first.

Using the following temperature dependencies, an estimate of the current vs. temperature can be made (dependence of threshold voltage upon temperature is not taken into account here, which is a reasonable assumption as long as the overdrive voltages are high):

$$\begin{aligned} \mu_n(T) &= \mu_n(300\ K) \cdot \left(\frac{T}{300\ K} \right)^{-1.5} \\ V_T(T) &= 8.625 \cdot 10^{-5}\ \frac{V}{K} \cdot T \end{aligned} \qquad \textbf{3.16}$$

Substituting these values into formula 3.14 and solving for T yields a linear dependency of the current on the temperature, as shown in figure 3.6, which was created with Maple.

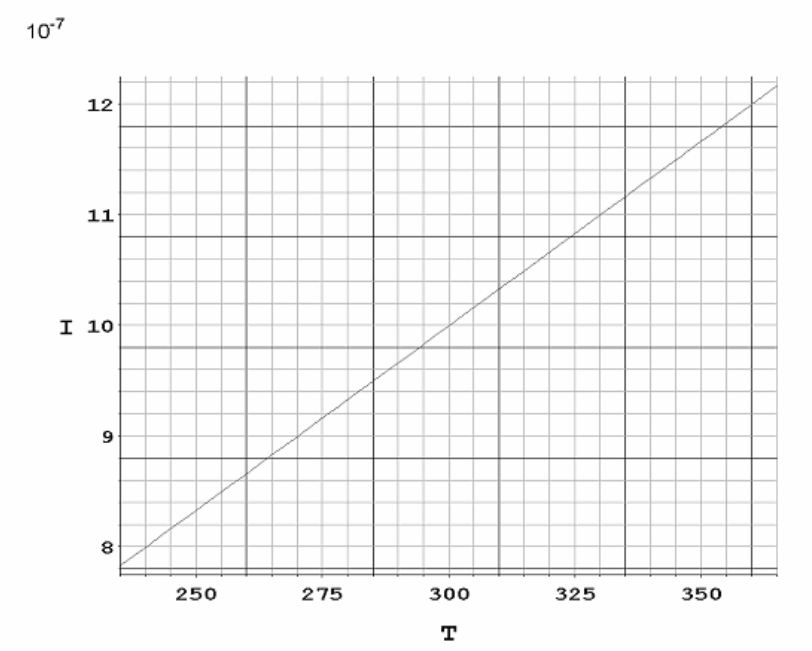

Figure 3.6: Calculated PTAT current vs. temperature

The implementation of the circuit can be found in figure 3.7.

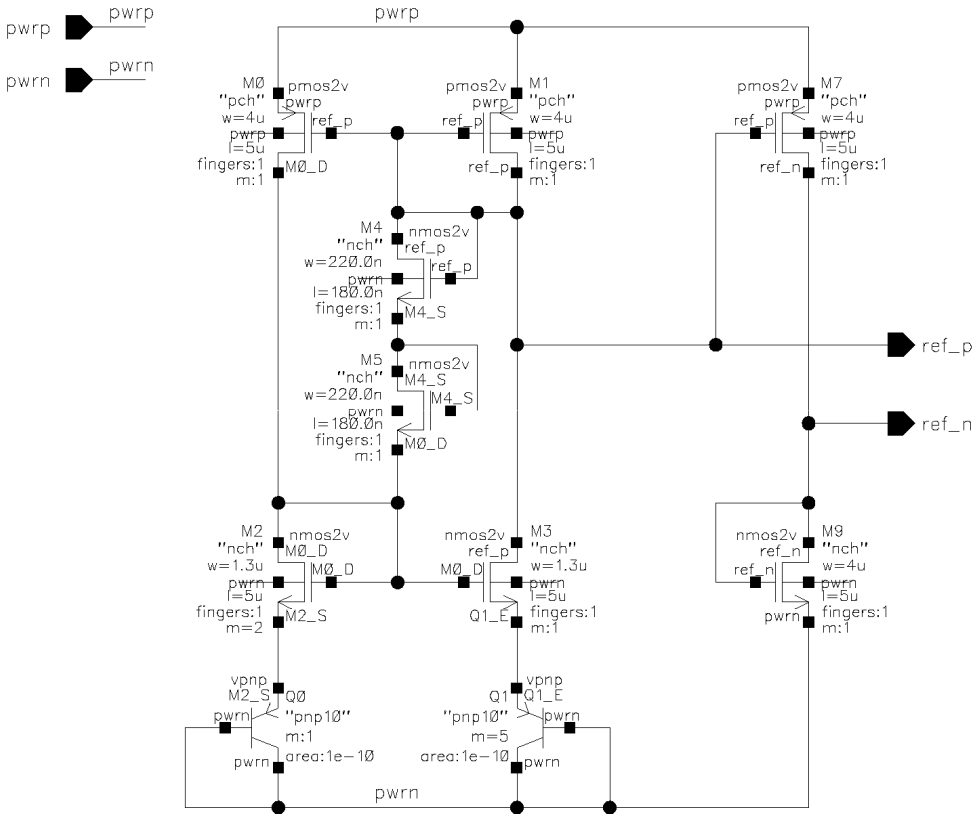

Figure 3.7: Cadence implementation of PTAT source

The two diodes ensure that the circuit will always start up and a generator for an NMOST reference has been added compared to figure 3.5.

A few tests have been performed on this circuit. Figure 3.8 shows the current in both branches as a function of temperature.

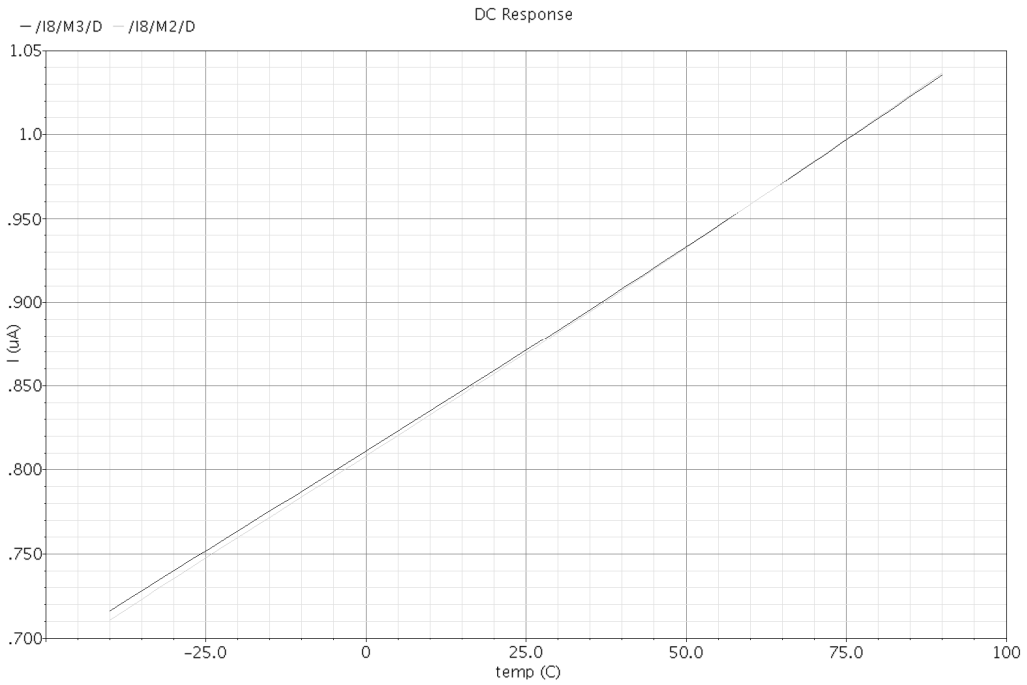

Figure 3.8: Simulated PTAT current vs. temperature

As can be seen, the relationship is close to linear. Also, as expected, the current is slightly higher than the calculated value. Table 3.1 lists both the calculated and simulated derivative of the current to the temperature. This shows that the behavior of the circuit is close to the expected behavior.

	I (nA) @ T = -40 °C	**I (nA) @ T = 90 °C**	**dI/dT (nA/K)**
Calculated	635	980	2.65
Simulated	715	1,035	2.46

Table 3.1: Calculated and simulated dI/dT

Figure 3.9 shows a number of transient runs at different temperatures. These runs have been performed with initial conditions such that both current mirrors are inactive. The start up time at lower temperatures is significantly longer than at high temperatures. The system does however start up within an acceptable timeframe for each temperature (less than 1.5 ms as can be seen). With this reference source in place, extensive simulations on the V-I converter can be performed.

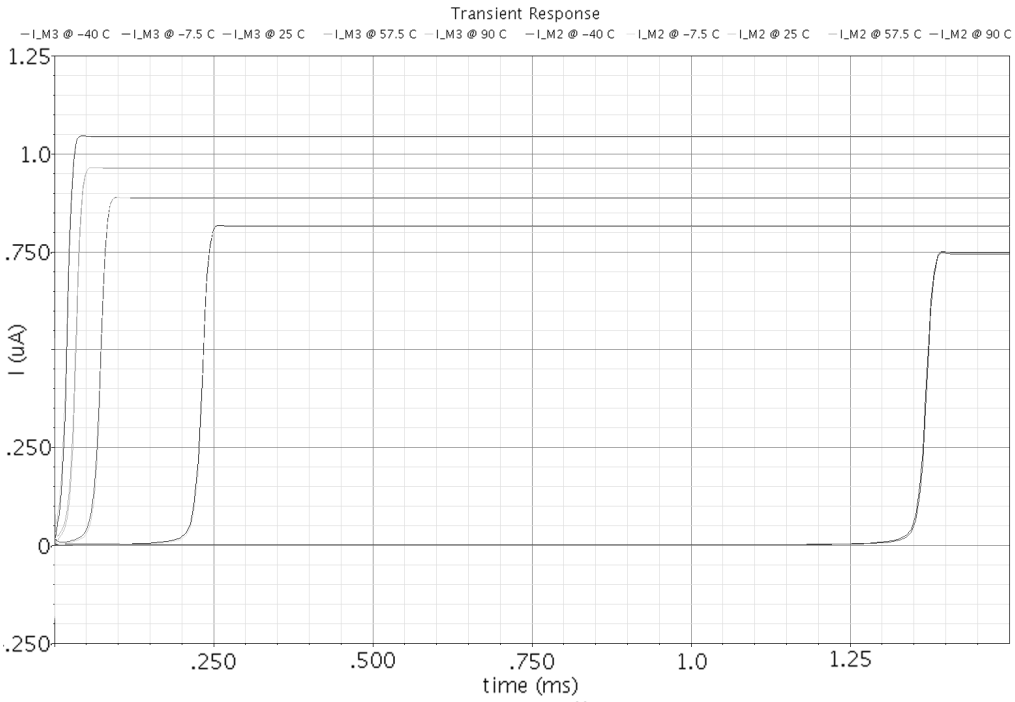

Figure 3.9: Transients at different temperatures

3.1.4 V-I Converter Tests

The first test that was performed on the V-I converter is a DC sweep to test the linearity. The output was connected to an ideal OPAMP transimpedance amplifier, which mimics the integrator that will later be connected to the output. No integrator was used yet, because such a circuit has a very high DC input impedance, which can cause convergence problems. The testbench is shown in figure 3.10. Figure 3.11 shows the output current and figurer 3.12 shows the absolute and relative errors. These error plots are made by comparing the output current to the ideal output current. Said ideal output current is found taking the derivative of the output current to the input voltage at mid-scale and multiplying it with the input voltage. Any offset in the output current is removed first.

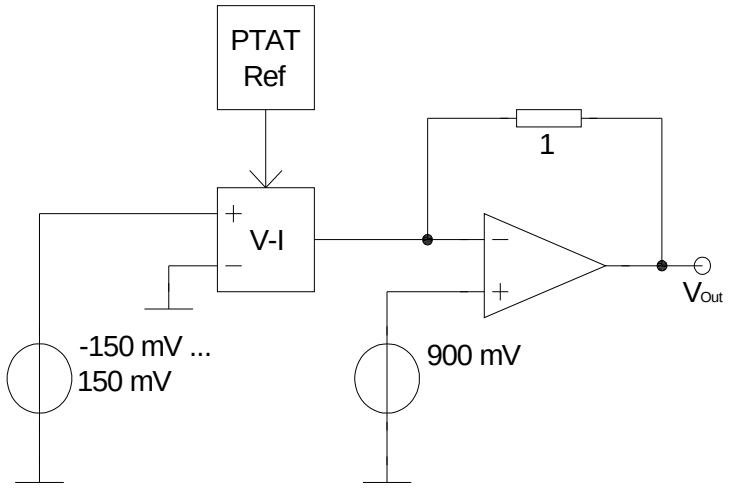

Figure 3.10: V-I converter testbench

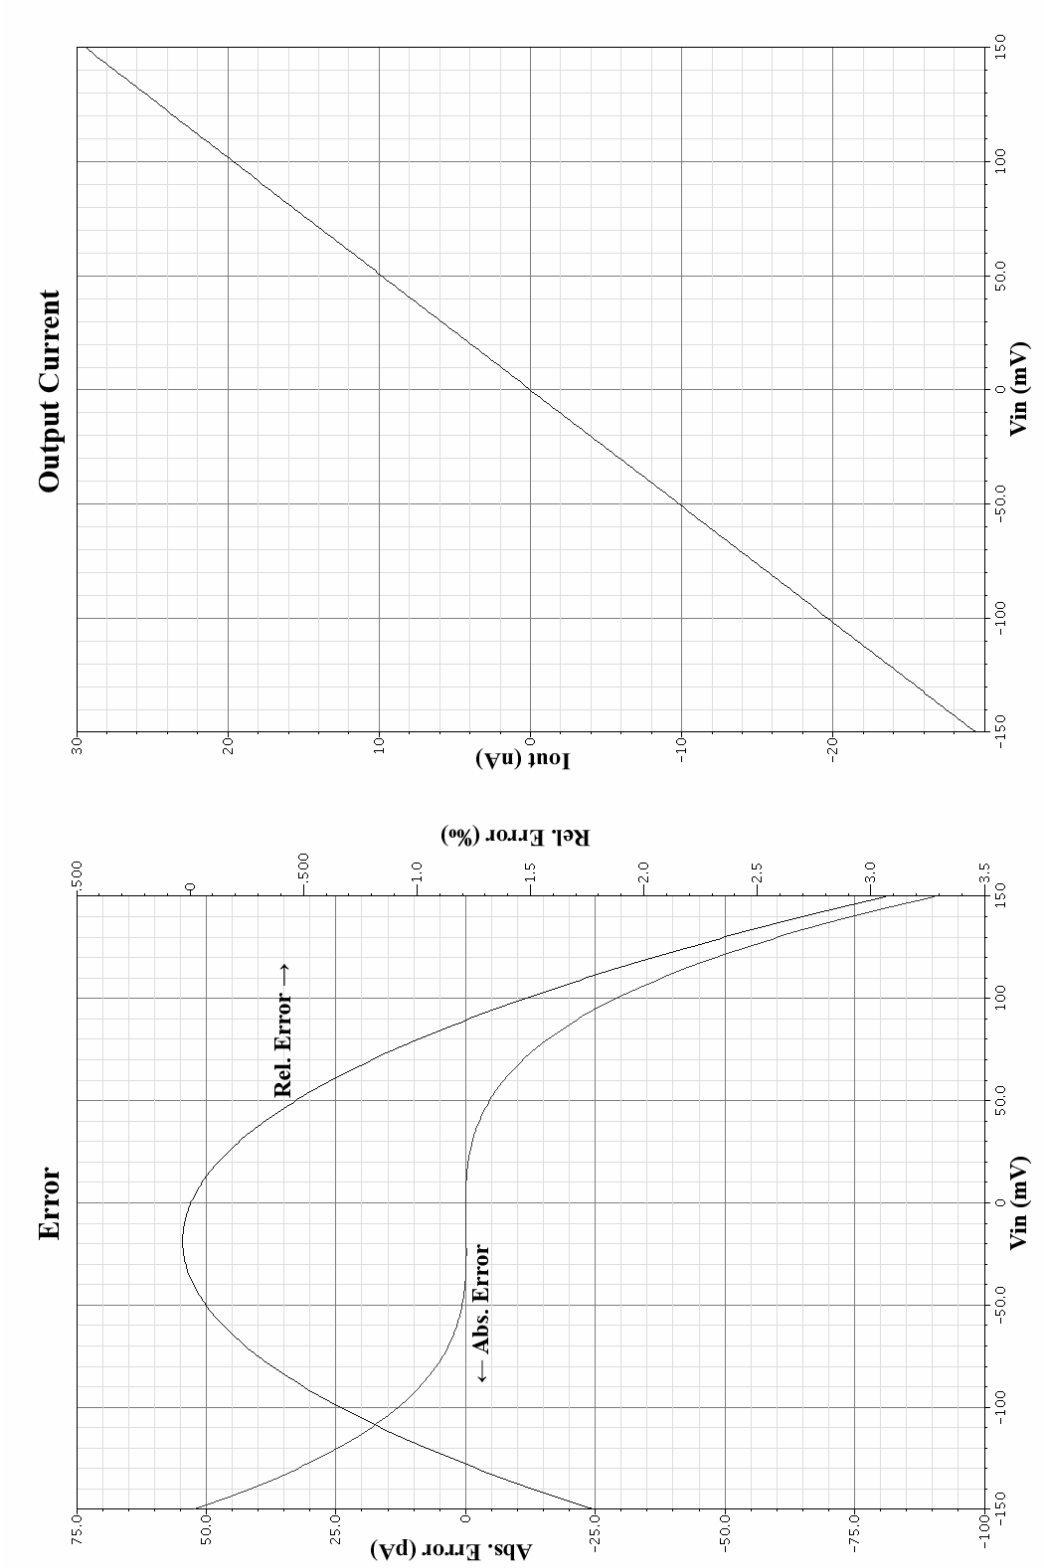

Figure 3.11: Output current and linearity of V-I converter

As can be seen, the linearity is within limits, but is not exactly symmetrical round 0 V. It turns out that the bias current source (M1 in figure 3.3) is pushed out of saturation

for high input voltages. Process corner simulations with otherwise the same settings show this effect even better, as can be seen in figure 3.12.

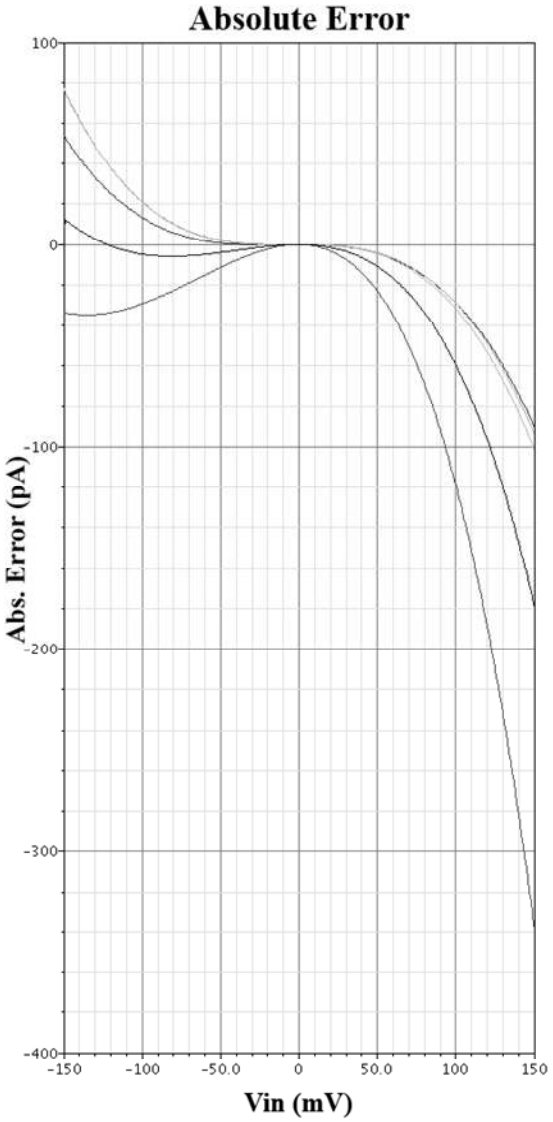

Figure 3.12: Linearity of V-I converter over corners

Figure 3.14 shows that the maximum relative error reaches more then 300 pA under certain conditions, which is over 1%. More specifically, this happens at low supply voltages. This makes sense, as less headroom is available. Also, all symmetry in the error plot is lost.
This problem has been solved by doubling the width of the input transistors. From formula 3.4 one can see that this will increase the differential current by a factor $\sqrt{2}$, but $\mu \cdot C_{OX} \cdot W/L$ of the transistors is increased by a factor 2, meaning less overdrive voltage is needed. This leads to the schematic shown in figure 3.13.

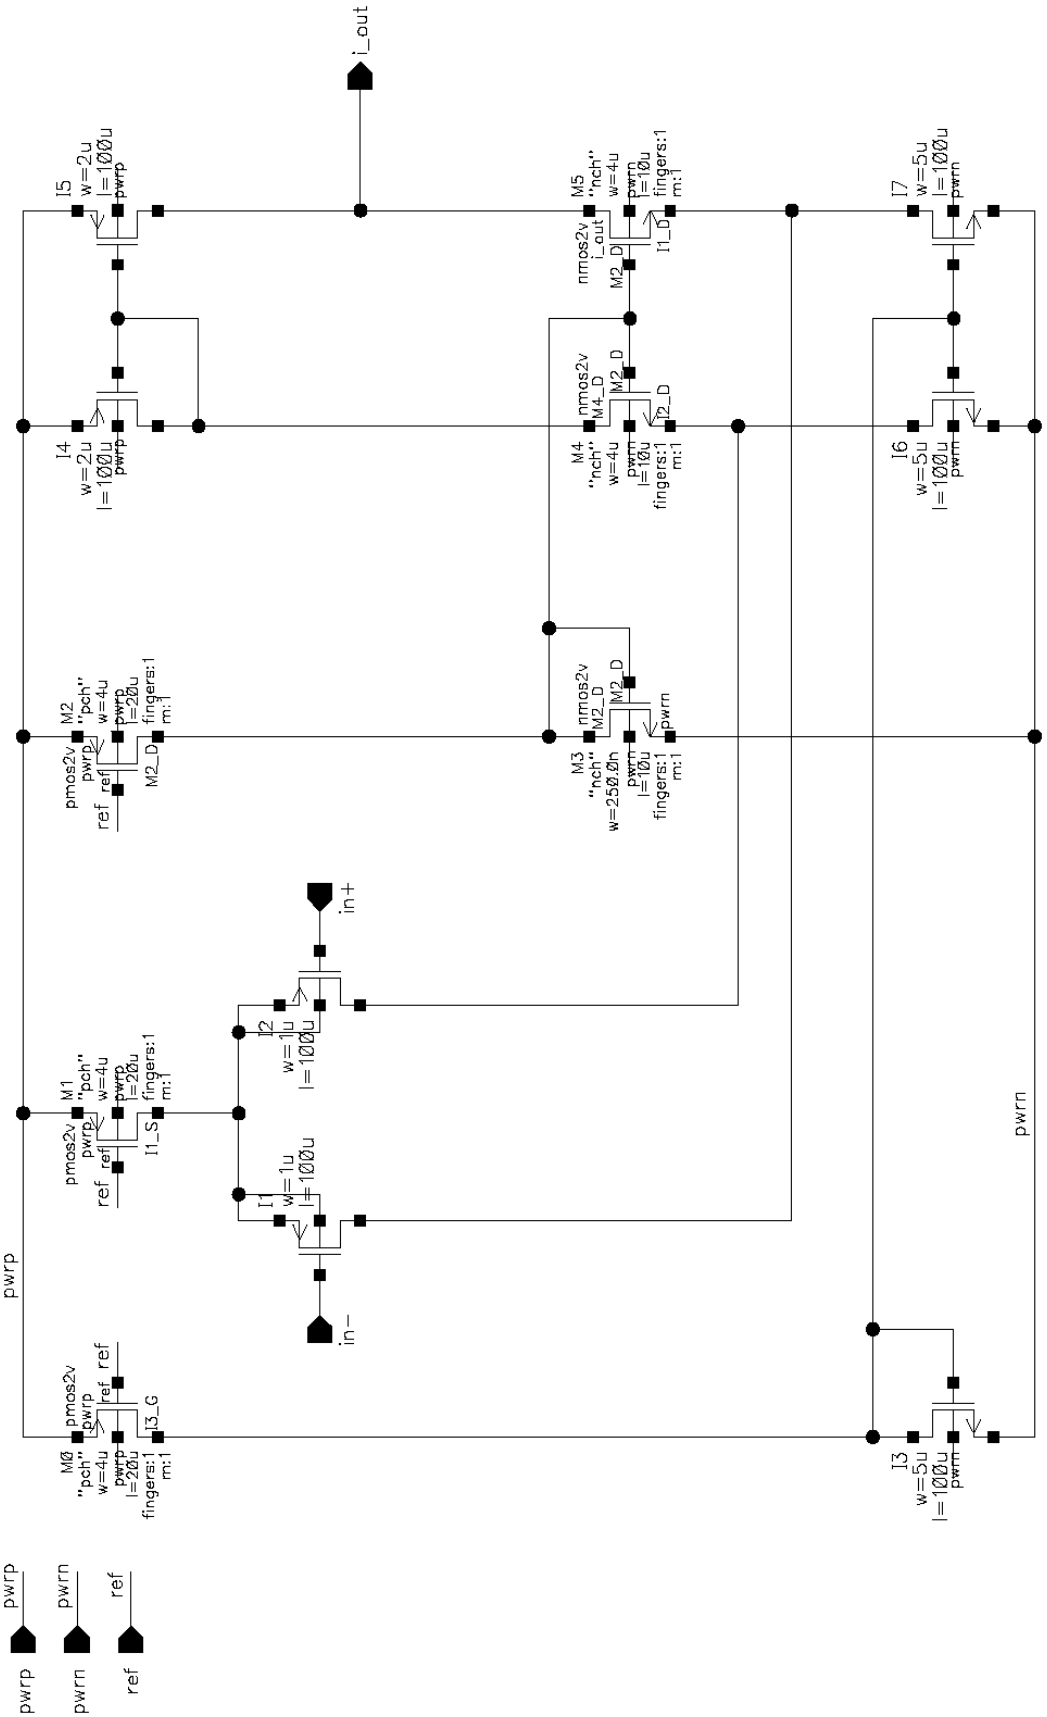

Figure 3.13: New V-I converter

As said before, the transimpedance of the circuit will decrease, by a factor √2. The original circuit had an output current range of approximately -29 nA … 29 nA, which is close to the expected value (-33 nA … 33 nA at 4.56 MΩ). The new circuit is thus expected to have an output current range of -41 nA … 41 nA. Also, the linearity will change somewhat. At an input voltage of 150 mV, a linearity of 0.68 % is expected (formula 3.2).

The same tests have been repeated with this schematic. The linearity test can be found in figure 3.14. Figure 3.15 shows the corner analysis. The output current range has indeed increased to -41 nA … 41 nA. This means that the transimpedance has dropped to 3.66 MΩ. This will slightly increase the speed of the system, but should not cause any issues. The linearity is even better than predicted, about 0.48 % at maximum input magnitude. Both the typical and the corner runs show a nice symmetric behavior now and the linearity error remains well within limits for each process corner.

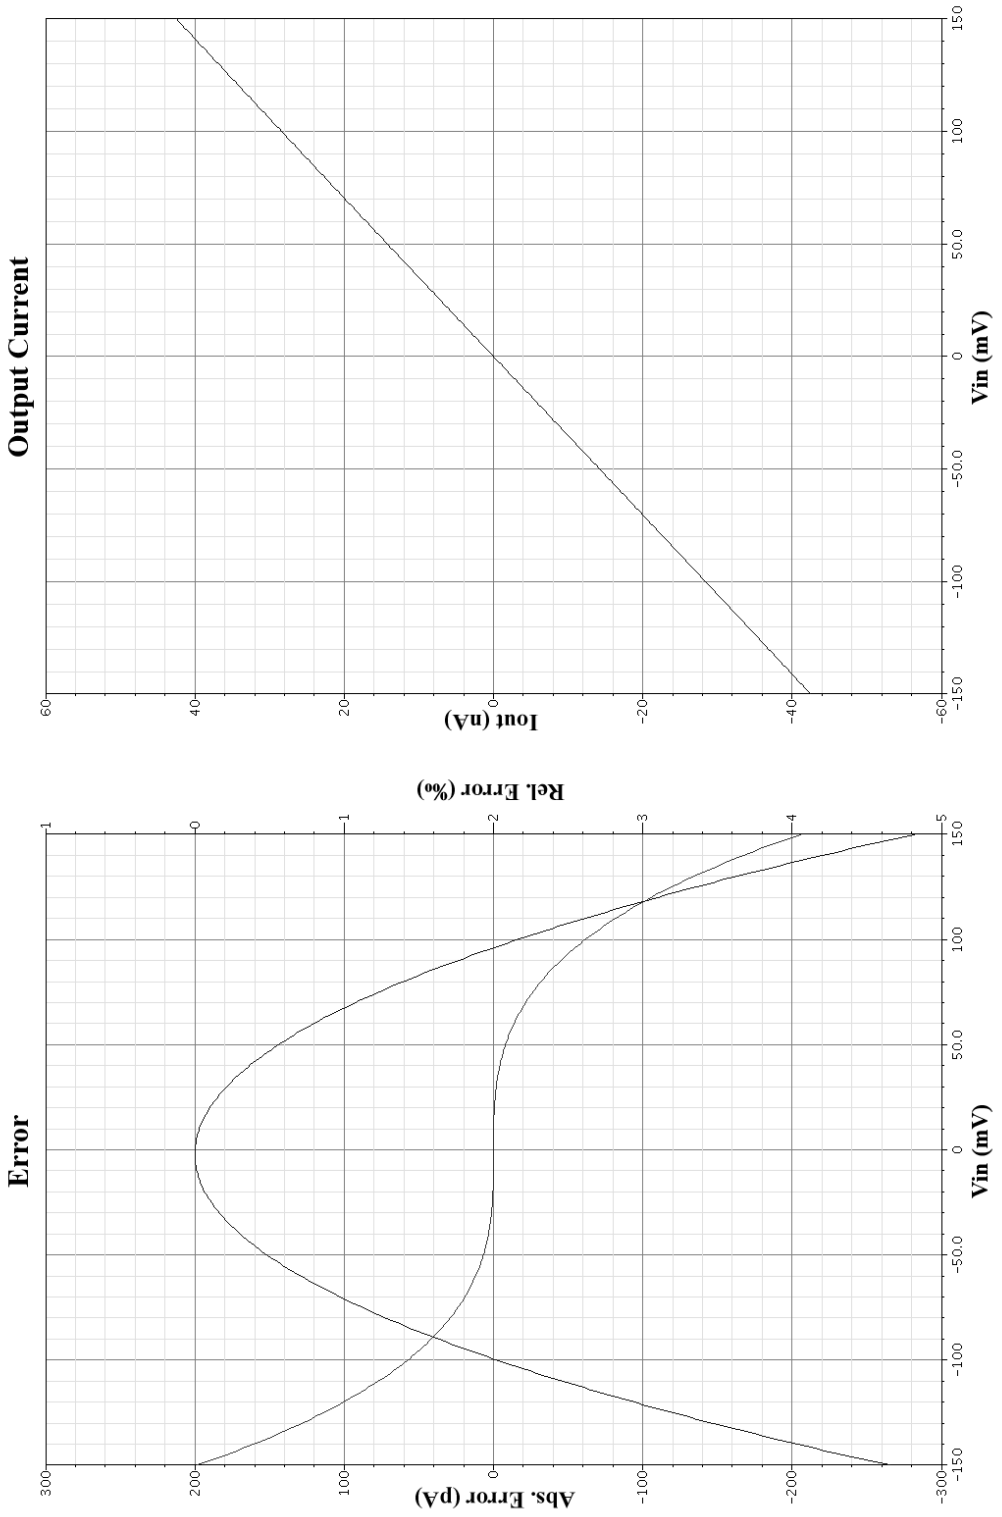

Figure 3.14: New V-I converter linearity test

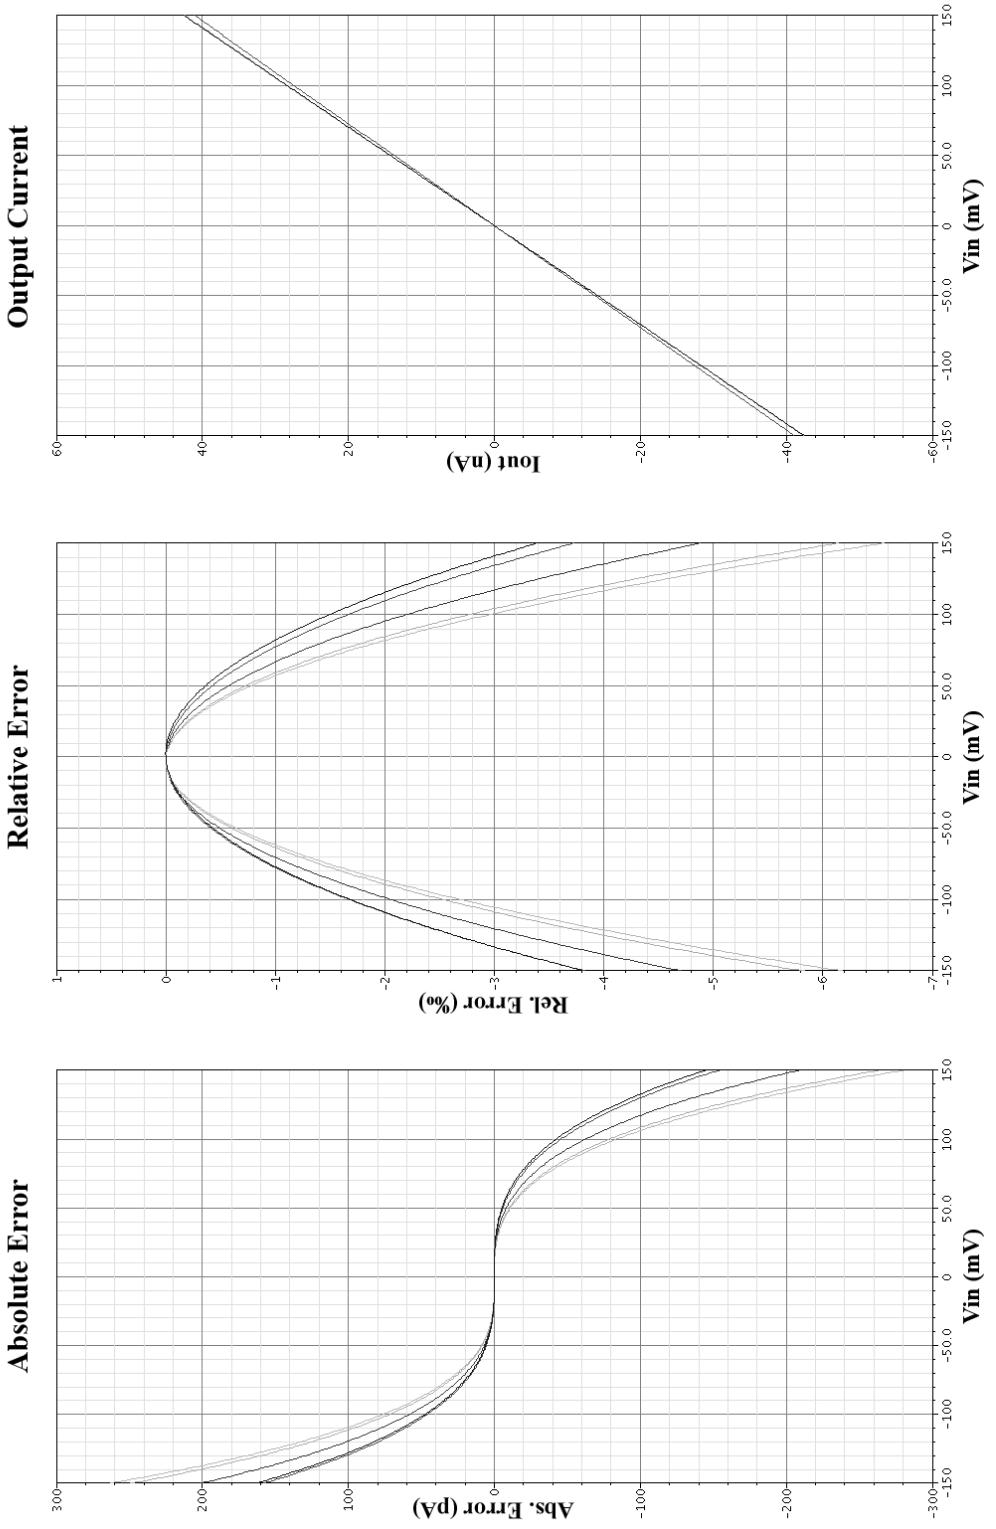

Figure 3.15: Linearity of new V-I converter over corners

The next important test is related to mismatch. The same linearity tests have been repeated in a Monte Carlo analysis, over 100 runs. Mismatch models were used in

both the converter and the reference generator for these simulations. As can be seen in figure 3.16, the transimpedance varies quite a bit over this range. This is mainly caused by mismatch in the reference generator and the bias transistors. The offset has been removed from this plot to better display the difference in transimpedance. A spread of about 10 % can be seen, which is better than what can be expected from a poly resistor.

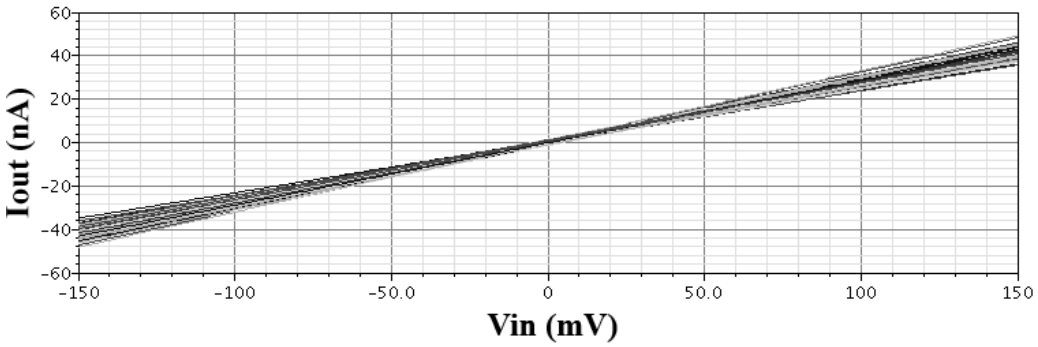

Figure 3.16: Input-output relationship with mismatch (offset removed)

Figure 3.17 shows the maximum relative error under mismatch. The $3 \cdot \sigma$ value lies at 0.73 %. This is well within limits, especially considering that the input range that was used here is 150 % of the specified input range.

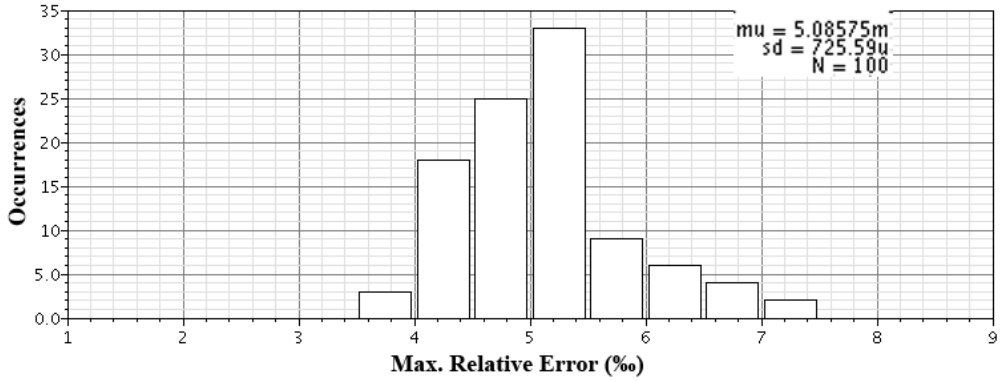

Figure 3.17: Maximum relative error under mismatch

The offset current caused by mismatch is shown in figure 3.18. A spread of 1.73 nA is found within $3 \cdot \sigma$. Since the (ideal) transimpedance is 3.66 MΩ, this corresponds to an input offset voltage of approximately 6 mV.

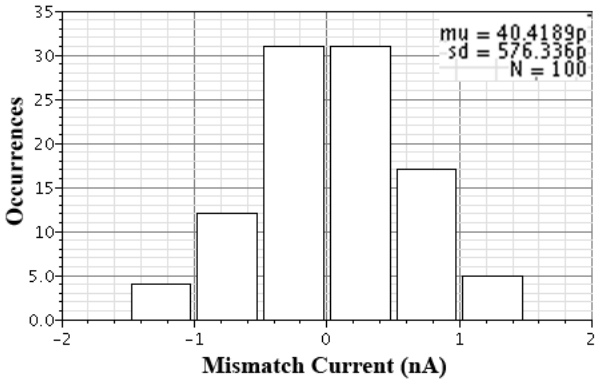

Figure 3.18: Output current offset due to mismatch

The offset will be removed by the chopping mechanism, but the system still has to be linear within 1 % for all input voltages. In order to test this, a DC voltage source of +6 mV has been added to the input. Subsequently, a range of input voltages have been supplied to both the inverting and non inverting inputs and the results have been averaged to mimic the effect of chopping. The results can be found in table 3.2. Note that the results obtained by applying the signal to the inverting input have been negated to compensate for the inverting character of the circuit under this condition.

V_{In}:	$I_{Out,Inverting}$ (nA):	$I_{Out,Non-Inverting}$ (nA):	$I_{Out,Average}$ (nA):
-100 mV	-30.082	-26.686	-28.384
-10 mV	-4.5537	-1.1348	-2.8443
-1 mV	-1.9940	1.4251	-0.28445
-100 µV	-1.7380	1.6811	-28.450 pA
100 µV	-1.6811	1.7380	28.450 pA
1 mV	-1.4251	1.9940	0.28445
10 mV	1.1348	4.5536	2.8442
100 mV	26.683	30.080	28.382

Table 3.2: V-I converter behavior under offset

The largest deviation from the ideal behavior exists at the highest input voltages. A factor of 997.6 is found between the smallest and largest positive inputs, and a factor 997.7 between the smallest and largest negative inputs. This is an error of 0.25 % and thus well within limits. This suggests that with proper chopping, the offset can be removed and the linearity will still be sufficient.

Another important test is the behavior of the system under influence of temperature variations. The offset and maximum relative error have been simulated over a range of temperatures. The results, which can be found in figure 3.19, show that both are well behaved over the entire temperature range. Figure 3.20 shows the behavior of the output current over temperature, at a fixed input voltage of 100 mV. The result is much better than that of the V-I converter without PTAT reference, but there is still an error of -2.4 % at -40 °C and +2.1 % at 90 °C. This is something that can thus still be improved, but the temperature dependence of the charge pump will also play a part in the behavior of the complete system over temperature. Therefore no final conclusion on this issue will be drawn here.

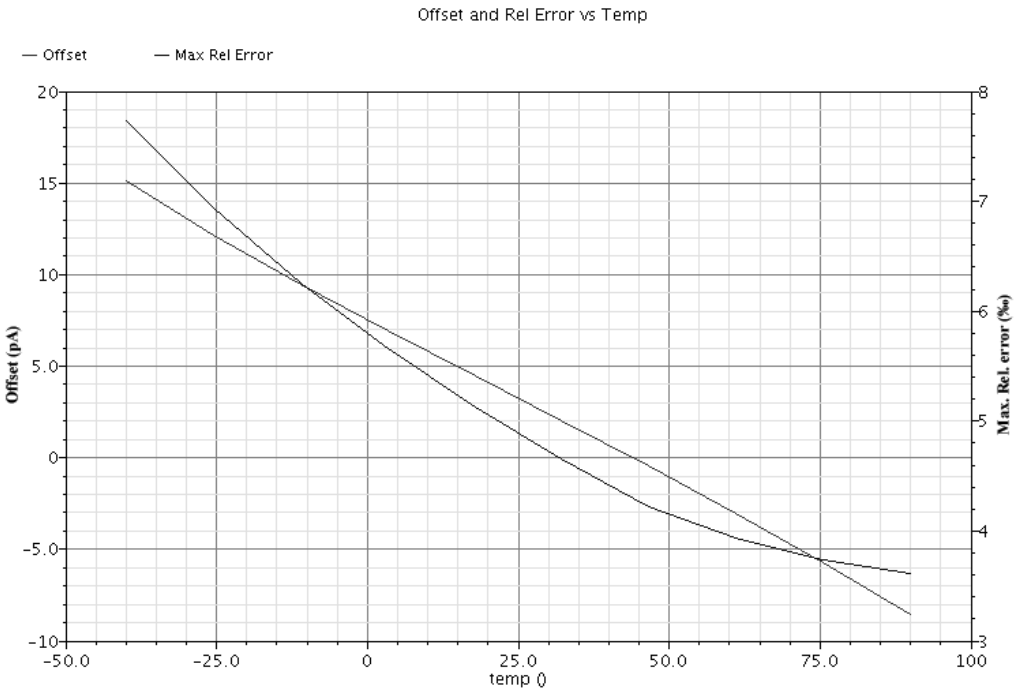

Figure 3.19: Offset and maximum relative error over temperature

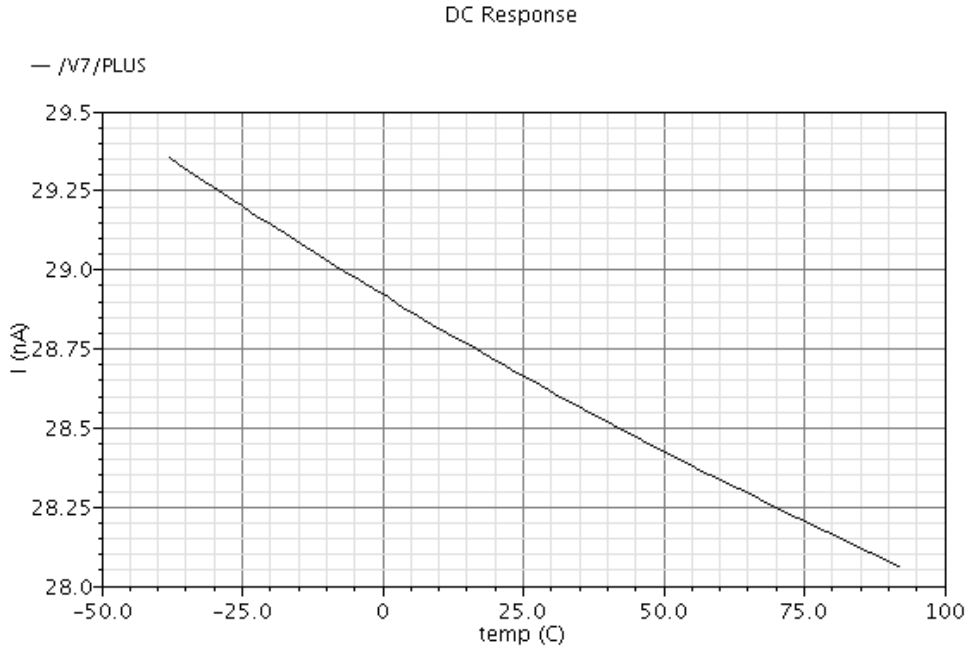

Figure 3.20: Output current at 100 mV input over temperature

The final test verifies the bandwidth of the circuit. This is not extremely important, since only the average of the input is of importance, but it is good practice to compare the bandwidth of the circuit to that of a 'normal' integrator with a resistor as input. Since the transimpedance of the circuit is 3.66 MΩ and the integration capacitor will be 4 pF (based on the current design), the bandwidth of an equivalent R-C integrator would be:

$$BW_{-3\,dB} = \frac{1}{2 \cdot \pi \cdot R \cdot C} \approx 10.9\ kHz \qquad \textbf{3.17}$$

An AC sweep of the circuit has been performed with an AC input of 100 mV. The result can be found in figure 3.21. The bandwidth of the converter is 42.5 kHz, which is well above the bandwidth of the complete integrator, hence the circuit will not have a large influence on the AC performance of the system.

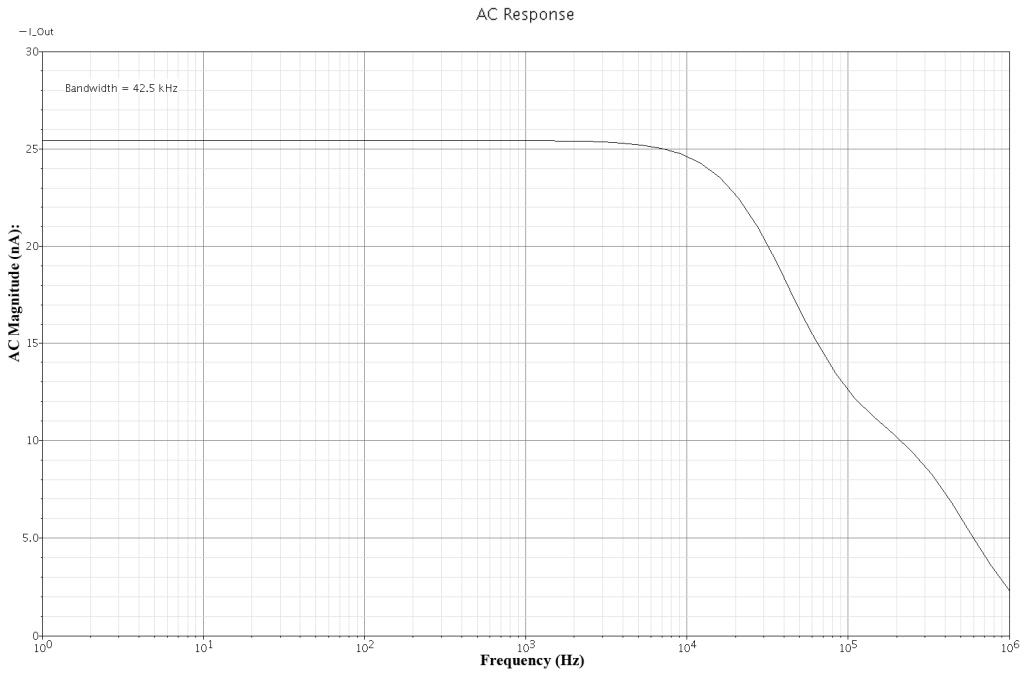

Figure 3.21: AC sweep

3.1.5 Integrator OPAMP

The last component that needs to be designed for the input stage is the OPAMP. As said before, the demands on this block are quite relaxed due to the current mode output of the V-I converter. In order to have predictable behavior for this block between different samples, a gain of at least 1,000 is desired over the entire output range. With that much gain, the feedback loop will always keep the swing at the input node small compared to the output node and the system can be thought of as an ideal integrator. A relatively high phase margin is also desirable, to limit overshoots when the charge pump is (de)activated. Finally, the output stage has to be able to sink and source approximately 40 nA to prevent slew rate limitations.

Figure 3.22 shows the schematic for the OPAMP. As can be seen, it is a simple two stage device. Moderate length devices are used in the second stage and in the input pair and current mirror in the first stage. This yields some degree of matching and a fair amount of gain, without requiring excessive die area. The differential pair is biased at 500 nA and the output pair is biased at 1 µA. Since this is significantly higher than the maximum current from the V-I converter, slew rate will not cause any issues here. Note the two anti-parallel Miller capacitors. These devices are so called

CAPA's which have a highly non-linear capacitance, which depends on their voltage level (see appendix B). By placing two of them anti-parallel, this non linear relation is linearized to some degree, which stabilizes the gain margin over the entire output voltage range.

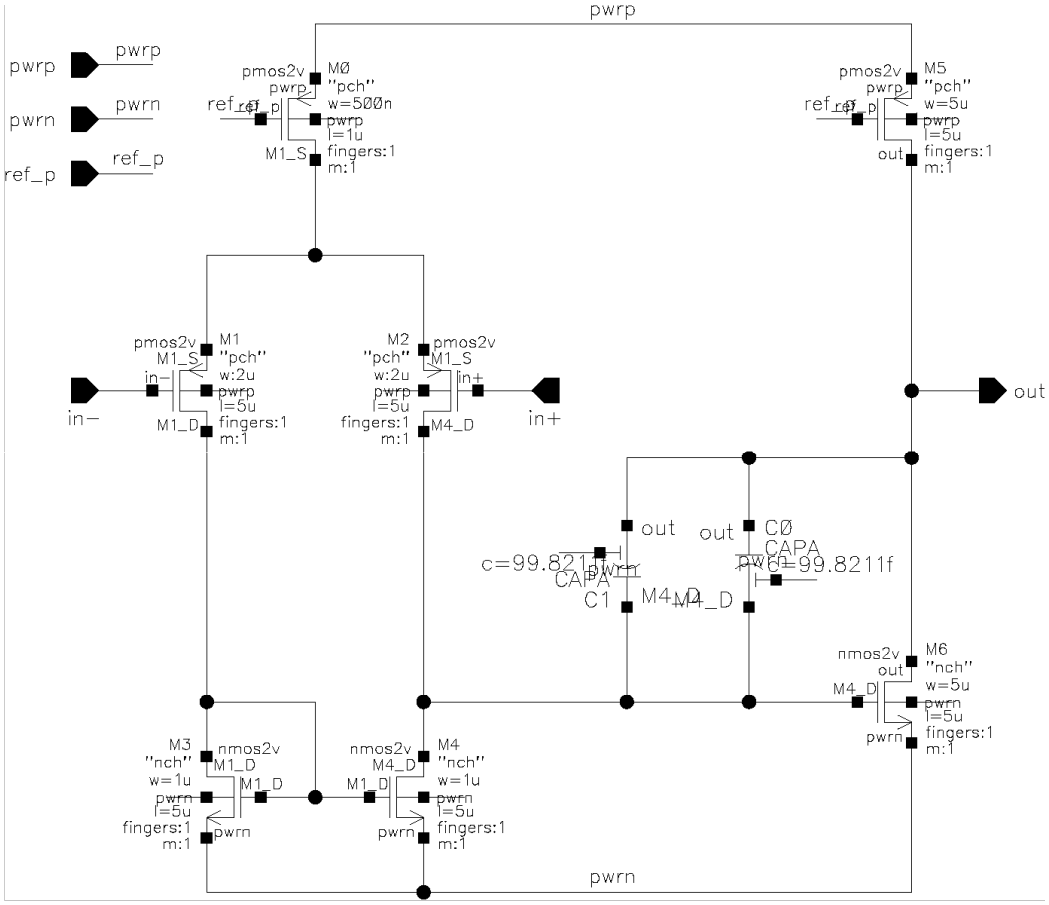

Figure 3.22: OPAMP

In order to test this OPAMP, the testbench in figure 3.23 was used.

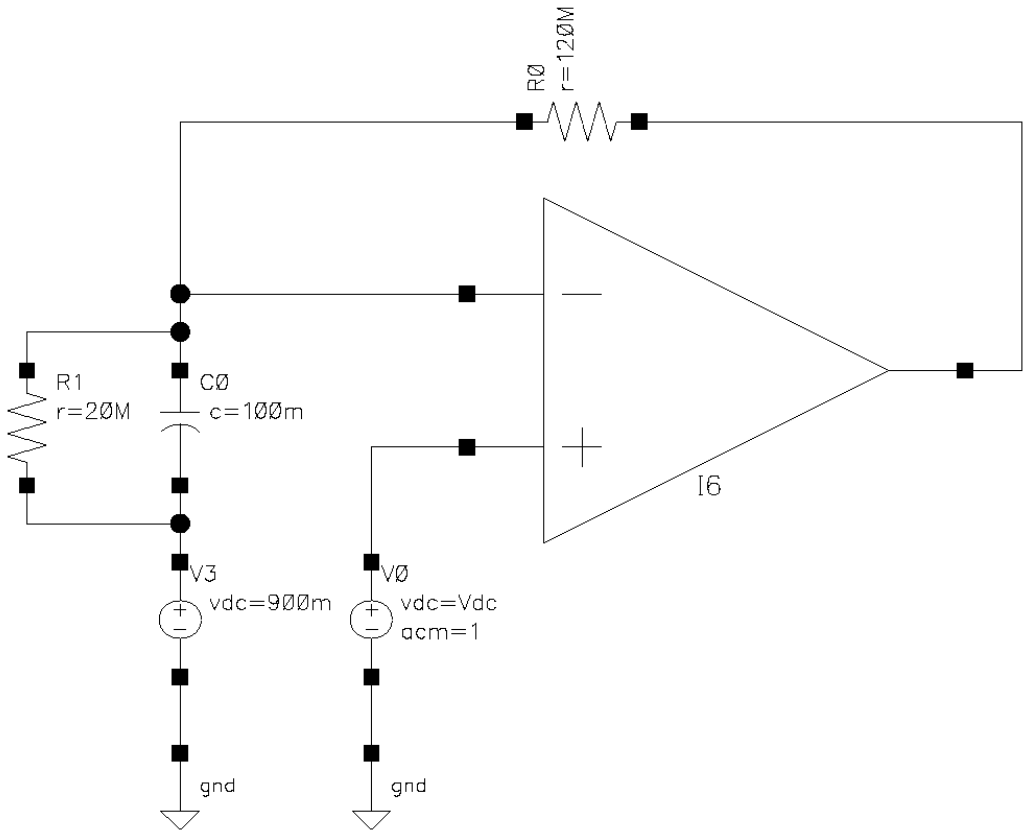

Figure 3.23: OPAMP testbench

The capacitor in this schematic has an infinite impedance at DC, yielding a DC gain of 7. This way the DC voltage at the output can be set between 200 mV and 1.6 V, by setting Vdc between 800 mV and 1.0 V. For high frequencies, the capacitor represents a short circuit, allowing the open loop gain to be simulated. A run at different output voltages has been performed using otherwise ideal conditions. The result can be found in figure 3.24

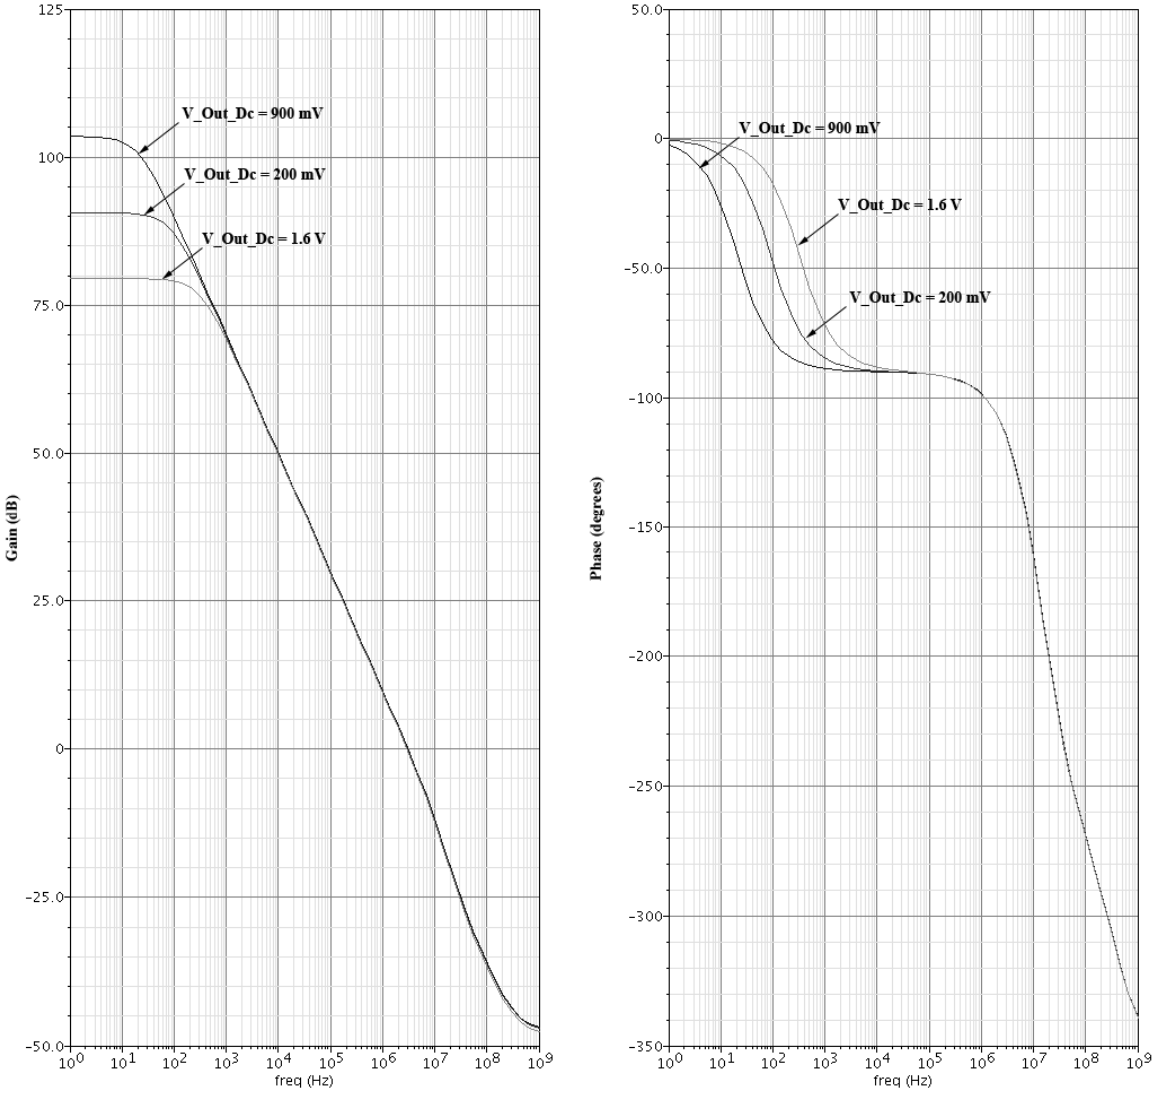

Figure 3.24: Gain and phase versus output DC level

As can be seen, the gain drops as the output voltage approaches the supply rails, but remains well above 60 dB for all levels. The phase margin is roughly 65°, regardless of output voltage. Figures 3.25 and 3.26 show that the gain and phase margin behave well over process corners and under mismatch conditions as well. Note that for these simulations, the DC voltage of the output was set to 900 mV.

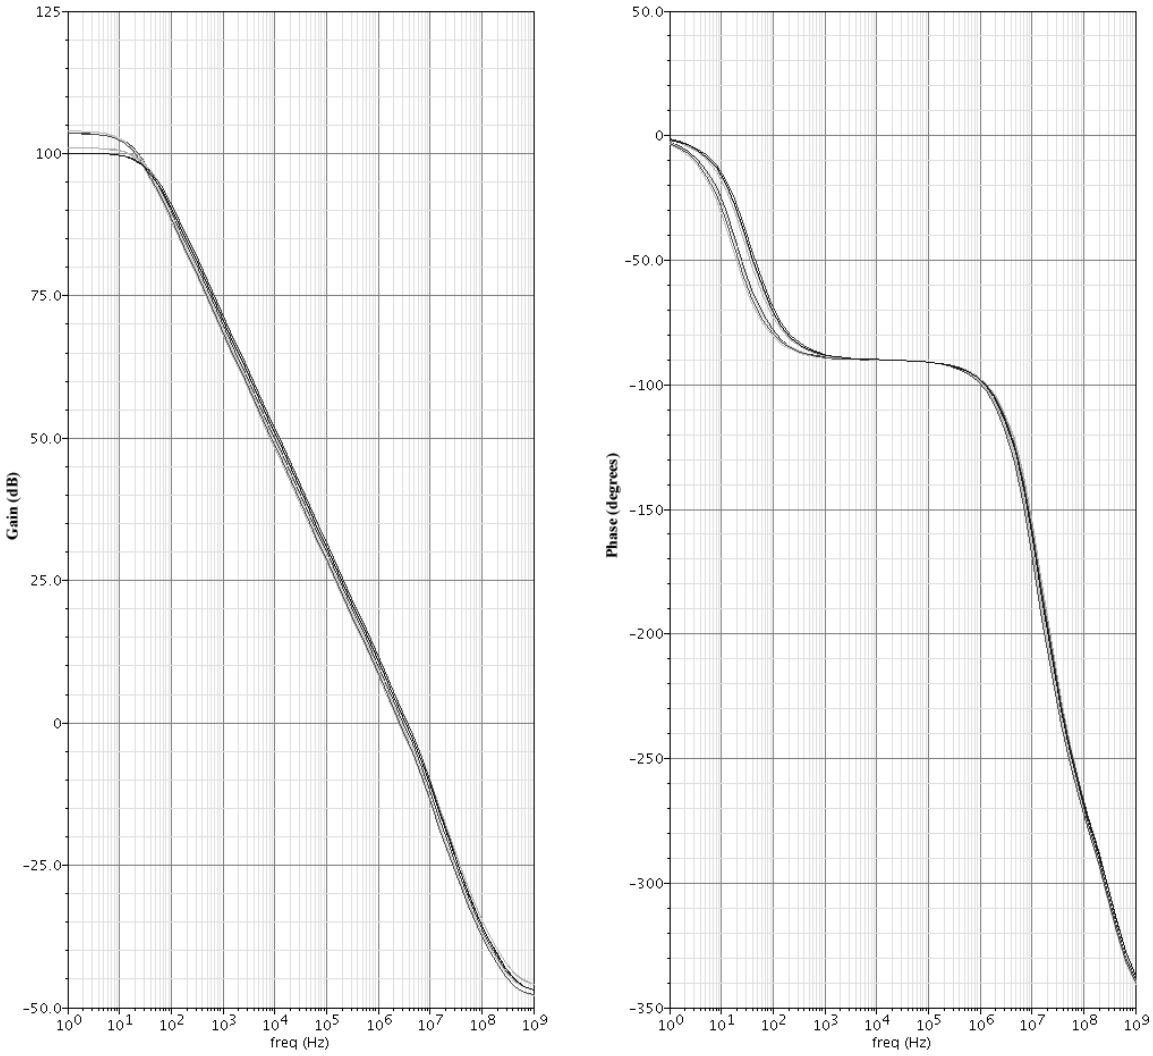

Figure 3.25: Gain and phase over corners

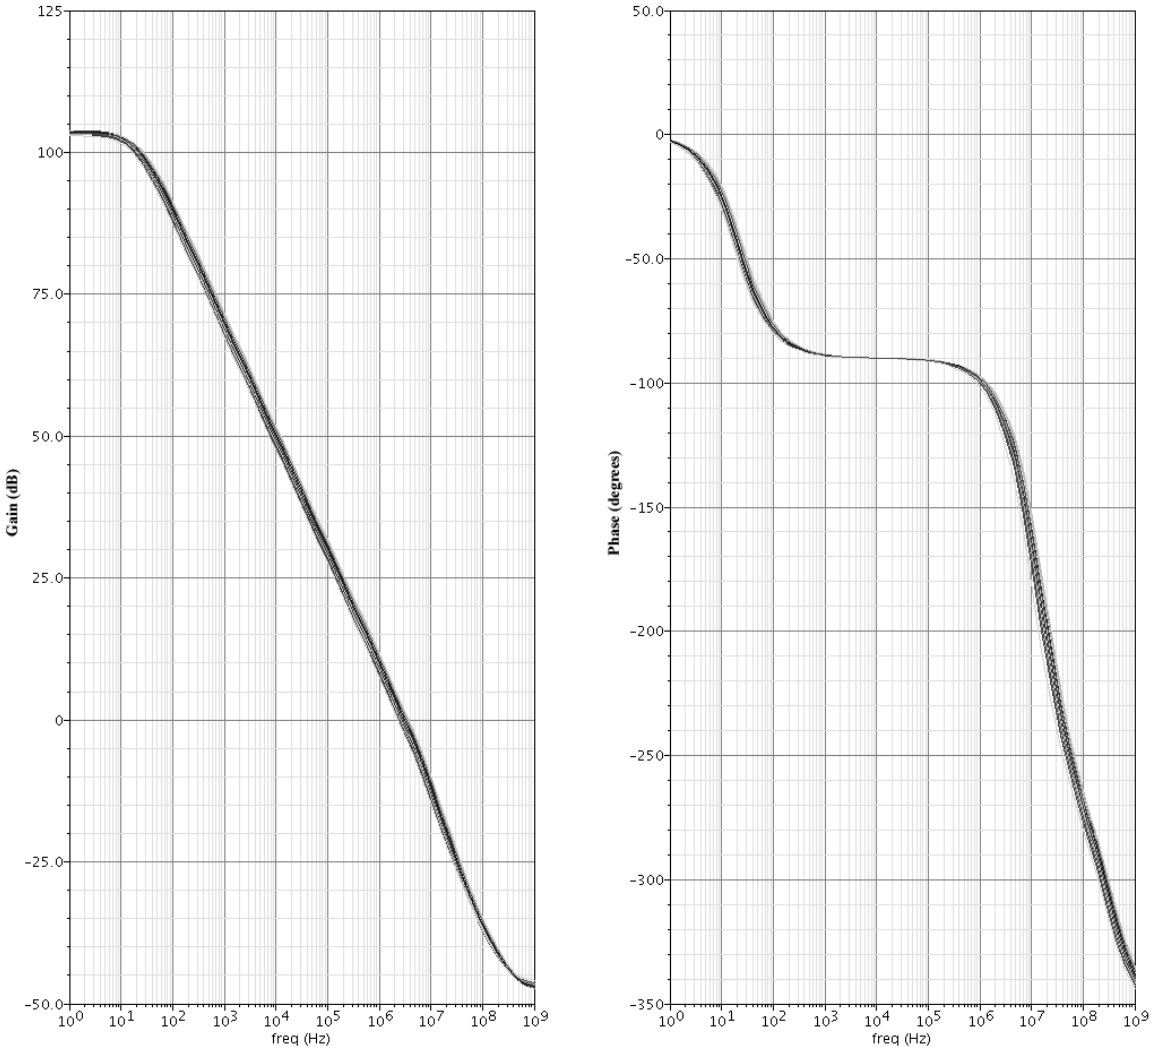

Figure 3.26: Gain and phase with mismatch

A last test op the OPAMP has been preformed to find the offset. As mentioned before this is not of prime importance, but it is good to have an idea of what to expect in practical circuits. The result can be found in figure 3.27. This plot was obtained by performing a DC analysis using unity gain feedback. The offset was found to be approximately 11 mV in $3 \cdot \sigma$. This should not cause any issues.

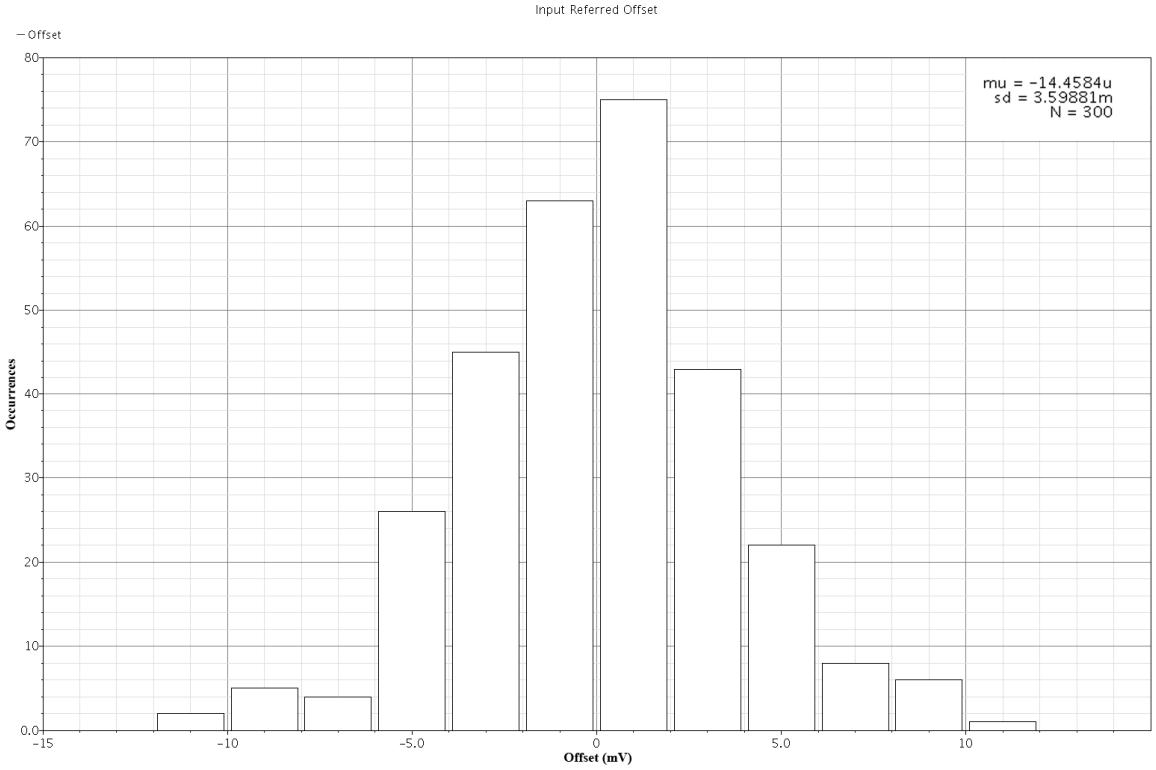

Figure 3.27: Input referred offset

3.2 Comparators

The comparators in the system have fairly mild specifications. They need not be very fast and their input offset does not impact system performance either. One important issue is power consumption, since the system should use as little power as possible. Another important aspect is that the charge pump needs to be activated in sync with the clock (to ensure that there is no difference in dump phase durations). This, combined with the low power demand makes latched comparators a nice option. Figure 3.28 displays a schematic of such a comparator with an NMOS input for high voltages.

In this design, bias transistor M11 is activated when the clock goes high. At the same time M4 and M8 are deactivated, yielding a high impedance at the drains of M6 and M7. These two devices make up a latch, that yields infinite gain for the input voltage difference. The right half of the system is another latch. This latch is forced to a low level at both outputs when the clock is low. When the clock goes high, the outputs are released from GND and either M12 or M16 is activated (depending on the input voltage). This forces one output high, and due to M19 and M20, the other output low. When this occurs, the output of the NOR gate goes low, and the bias current is turned off. M13, M14, M15, M17 and the inverters ensure that the output state of the latch remains fixed when the bias is disabled. When the clock goes low the circuit is reset. This way, the output is in sync with the clock, and the bias current is only consumed during the comparison phase. During this comparison phase, a bias current of 1 µA is consumed (which is an arbitrary value).

The PMOS equivalent of the circuit, which is used for low input voltages, can be found in figure 3.29. It is basically the inversed version of 3.28.

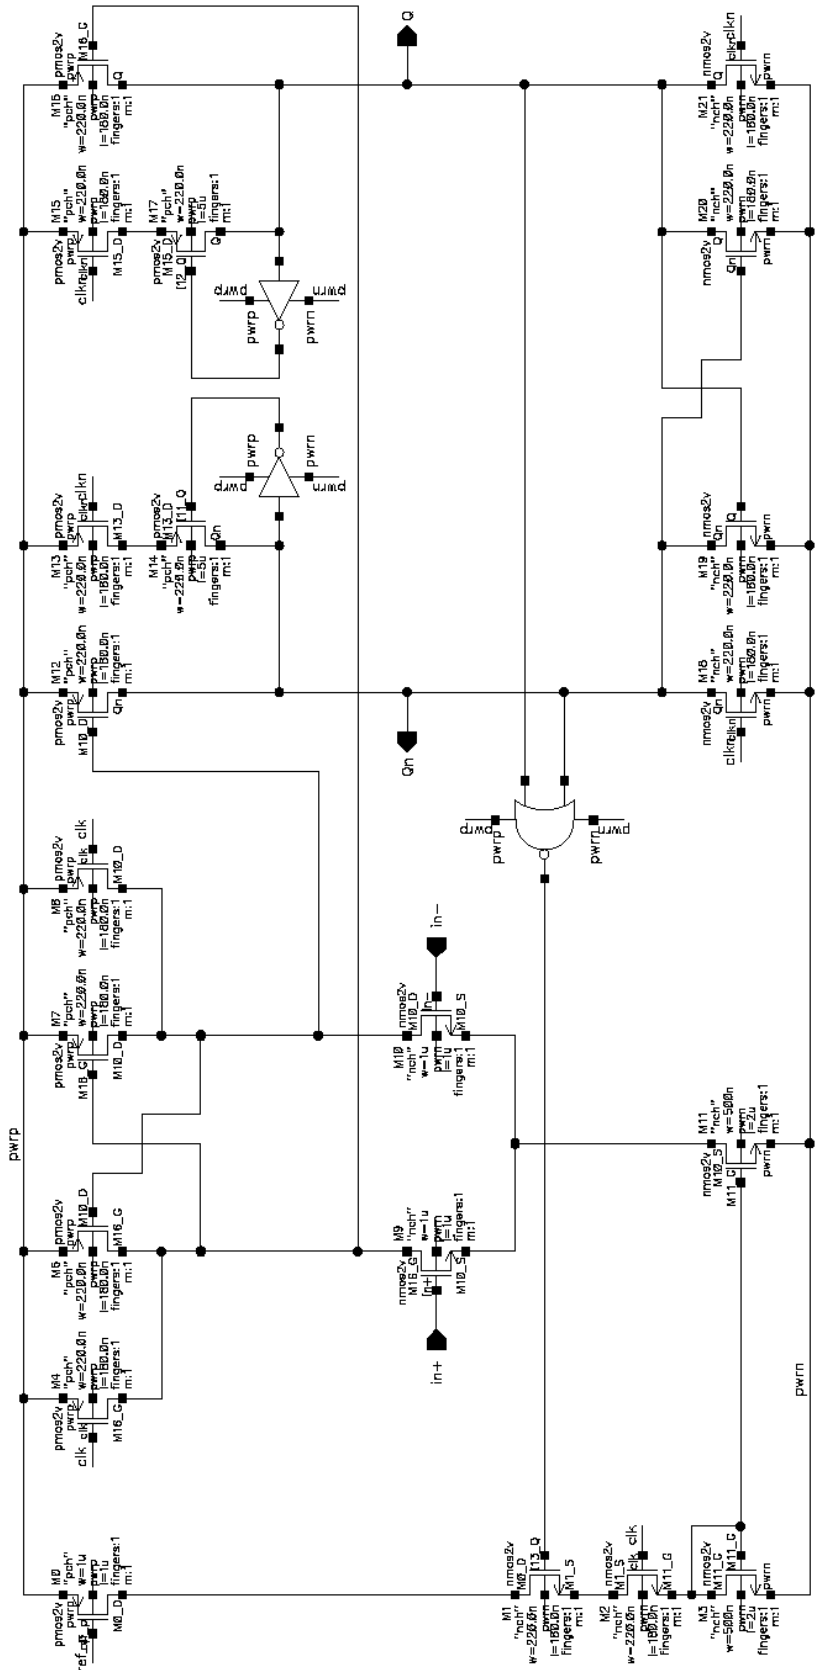

Figure 3.28: N-Type clocked comparator

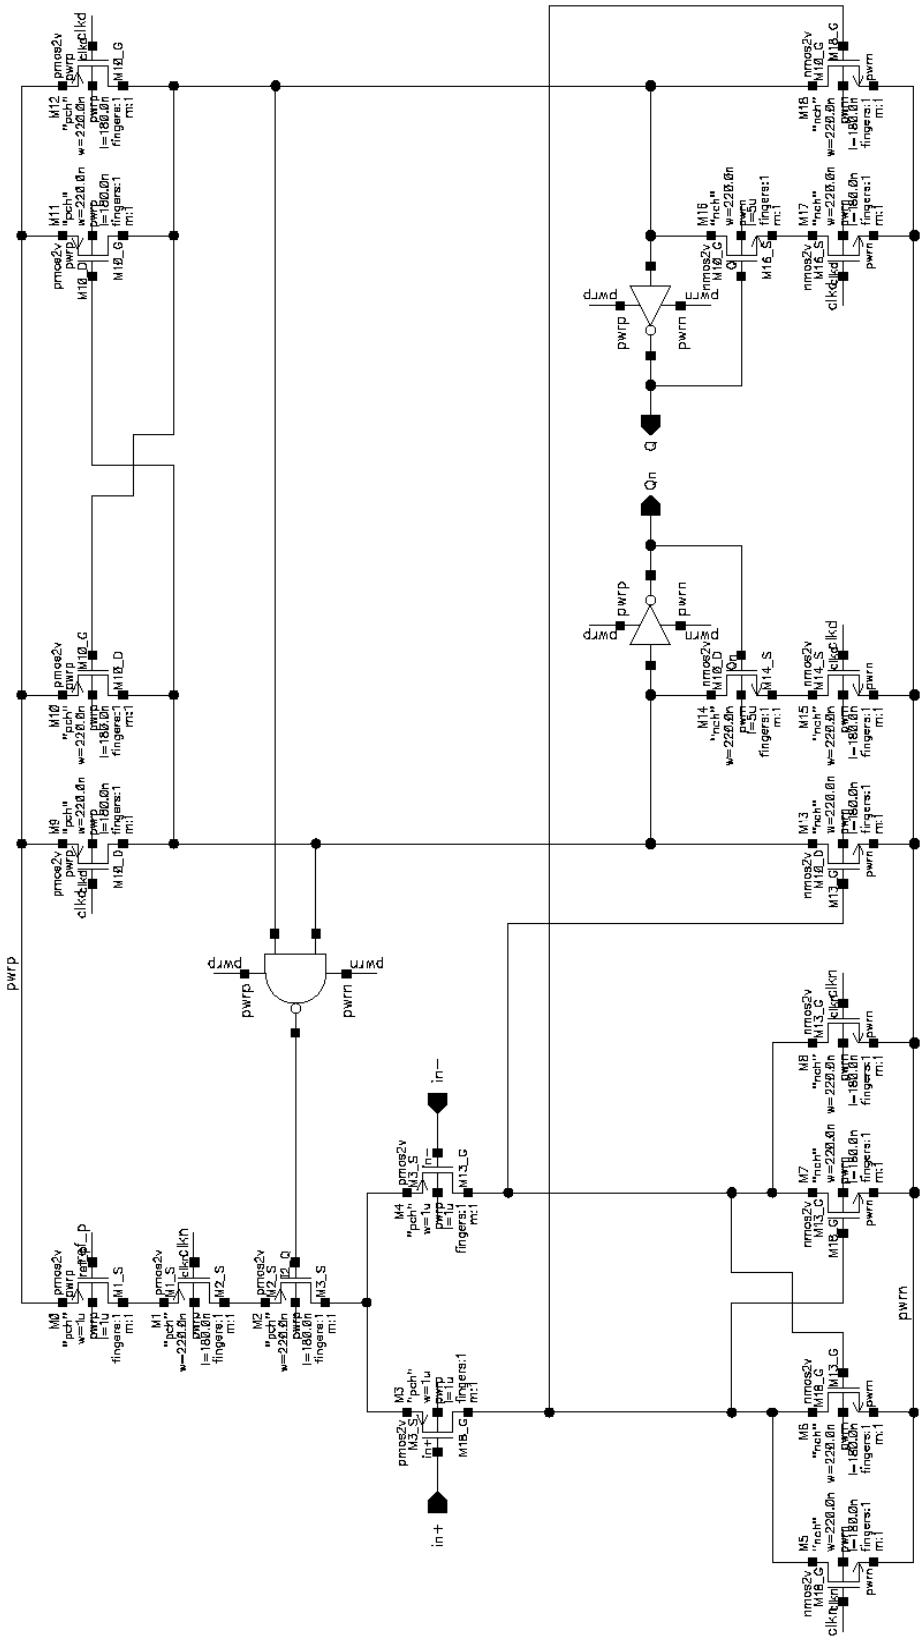

Figure 3.29: P-Type clocked comparator

Note that most of the devices in figures 3.28 and 3.29 are minimum sized, meaning neither circuit will use much die area. Transient test results from both circuits can be found in figures 3.30 and 3.31. In the top plots of these figures the input signals of the comparators can be seen. As can be seen, a sine wave was applied to the non inverting inputs, and a constant voltage to the inverting input. The lower part of both plots shows the output signals, Q and Qn. As expected, the Q output is high when the inverting input is lower than the non inverting input ands vice versa. Note that either output is only high when the clock is high as well.
Figures 3.32 and 3.33 depict the propagation delay of both comparators. As can be seen this is less than 10 ns for both types, when measured at an input voltage difference of 3 mV. Finally, the current consumption has been measured at an average of less than 100 nA.

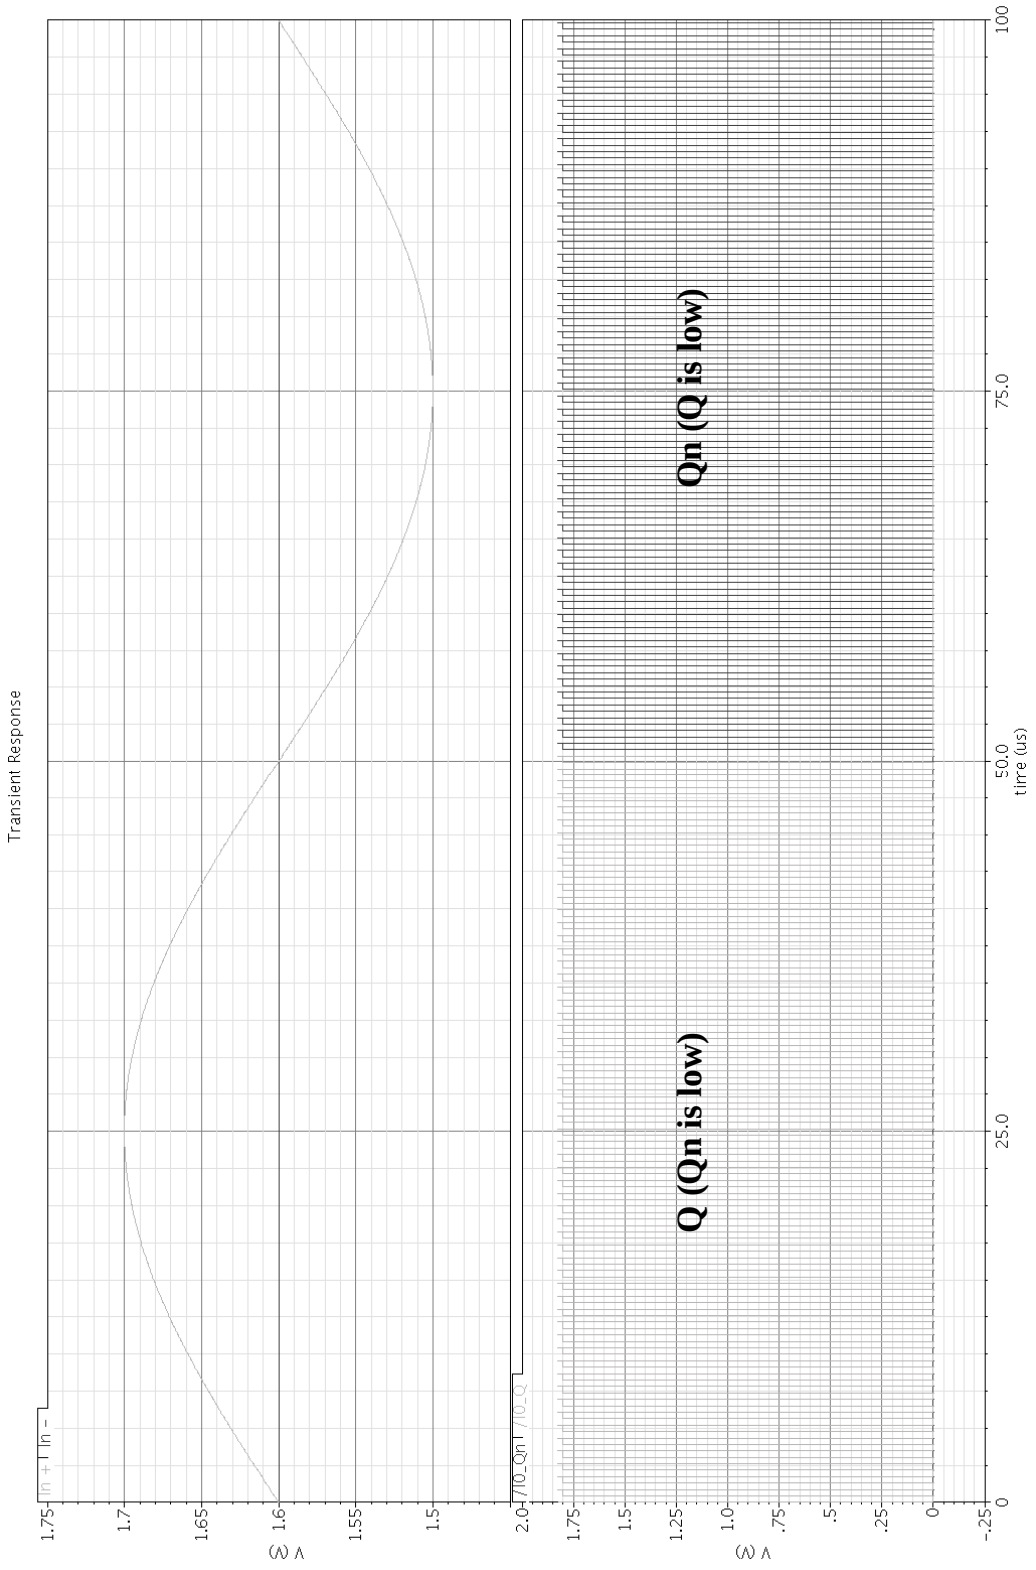

Figure 3.30: N-Type comparator transient

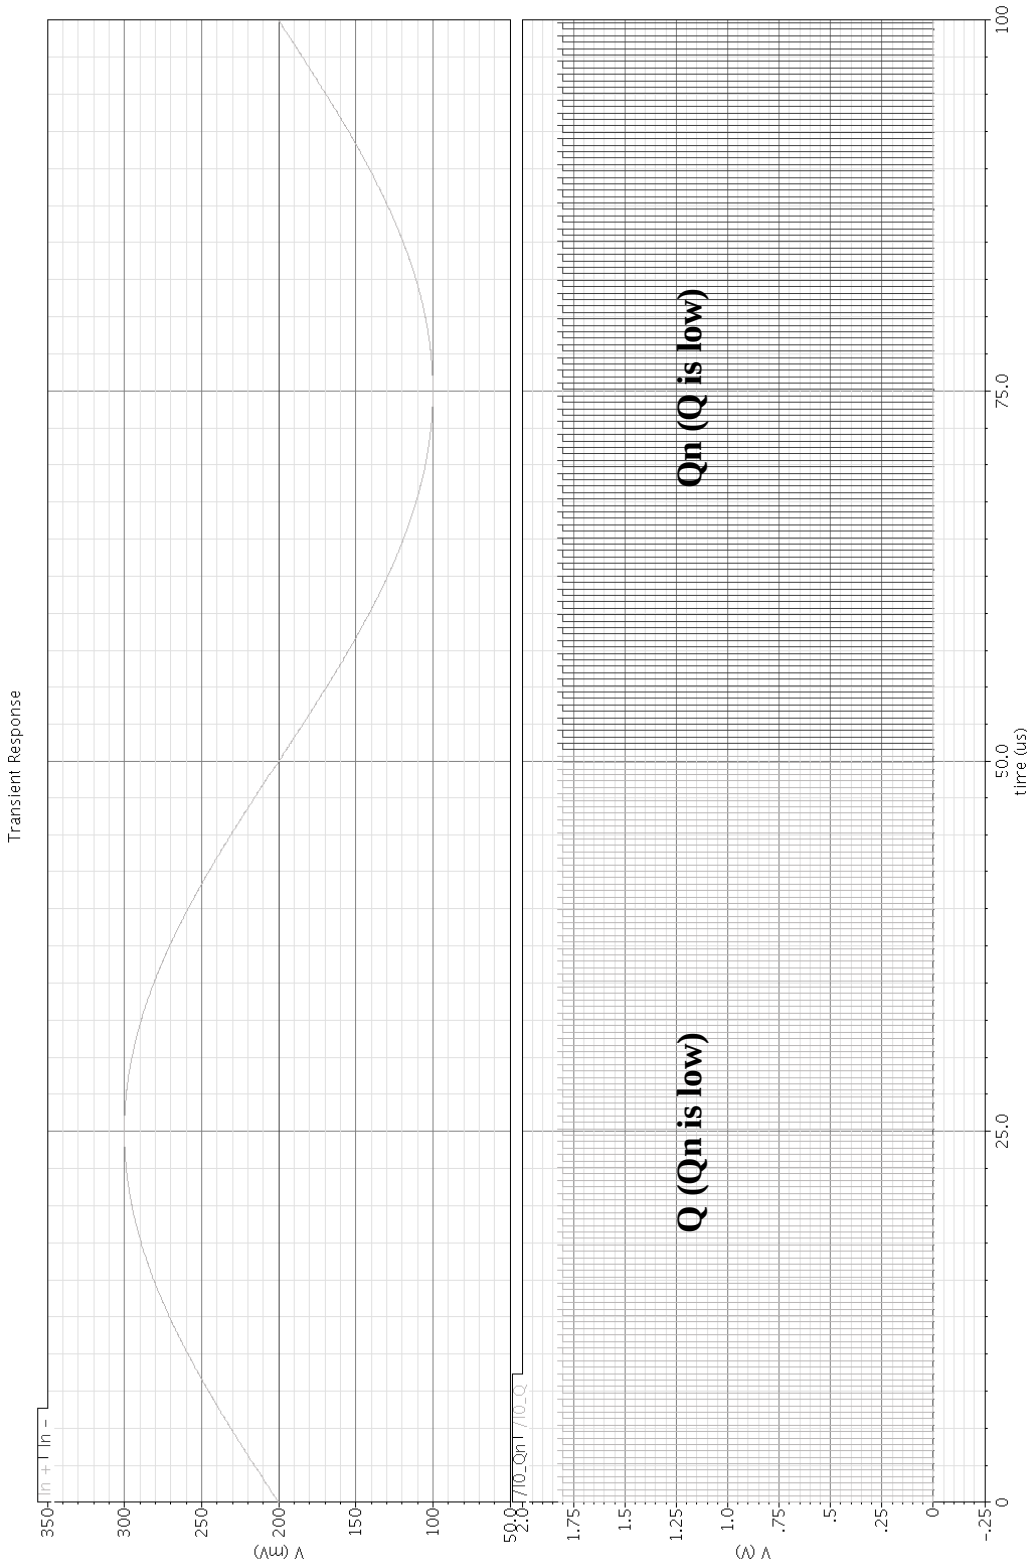

Figure 3.31: P-Type comparator transient

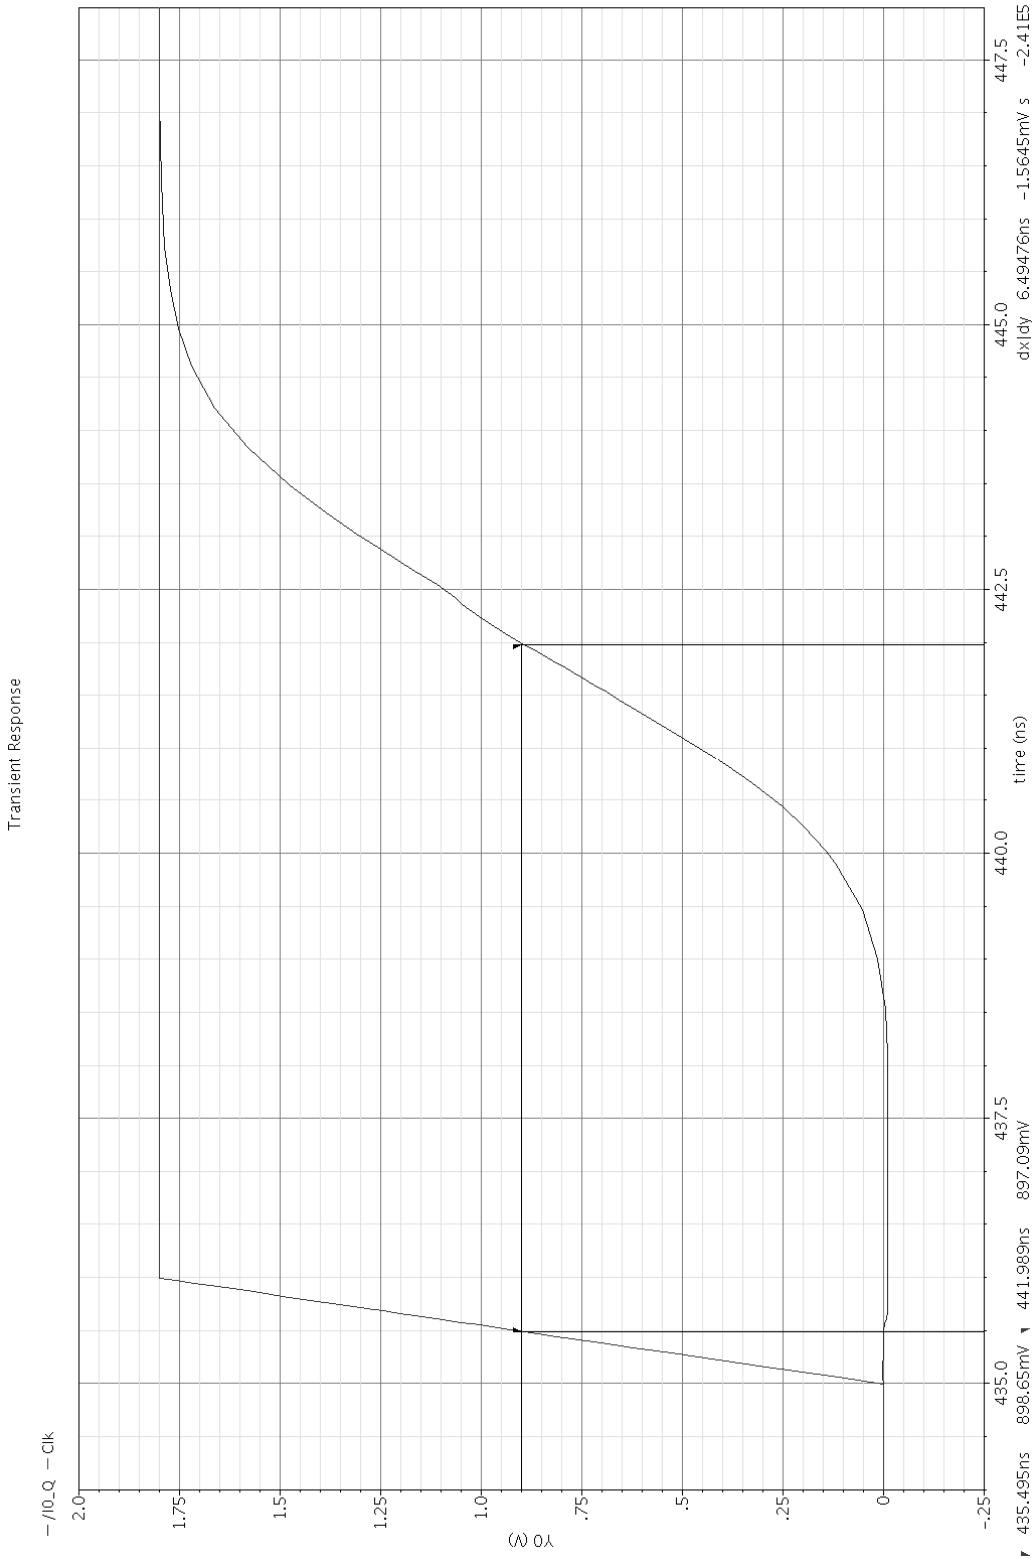

Figure 3.32: N-Type comparator propagation delay

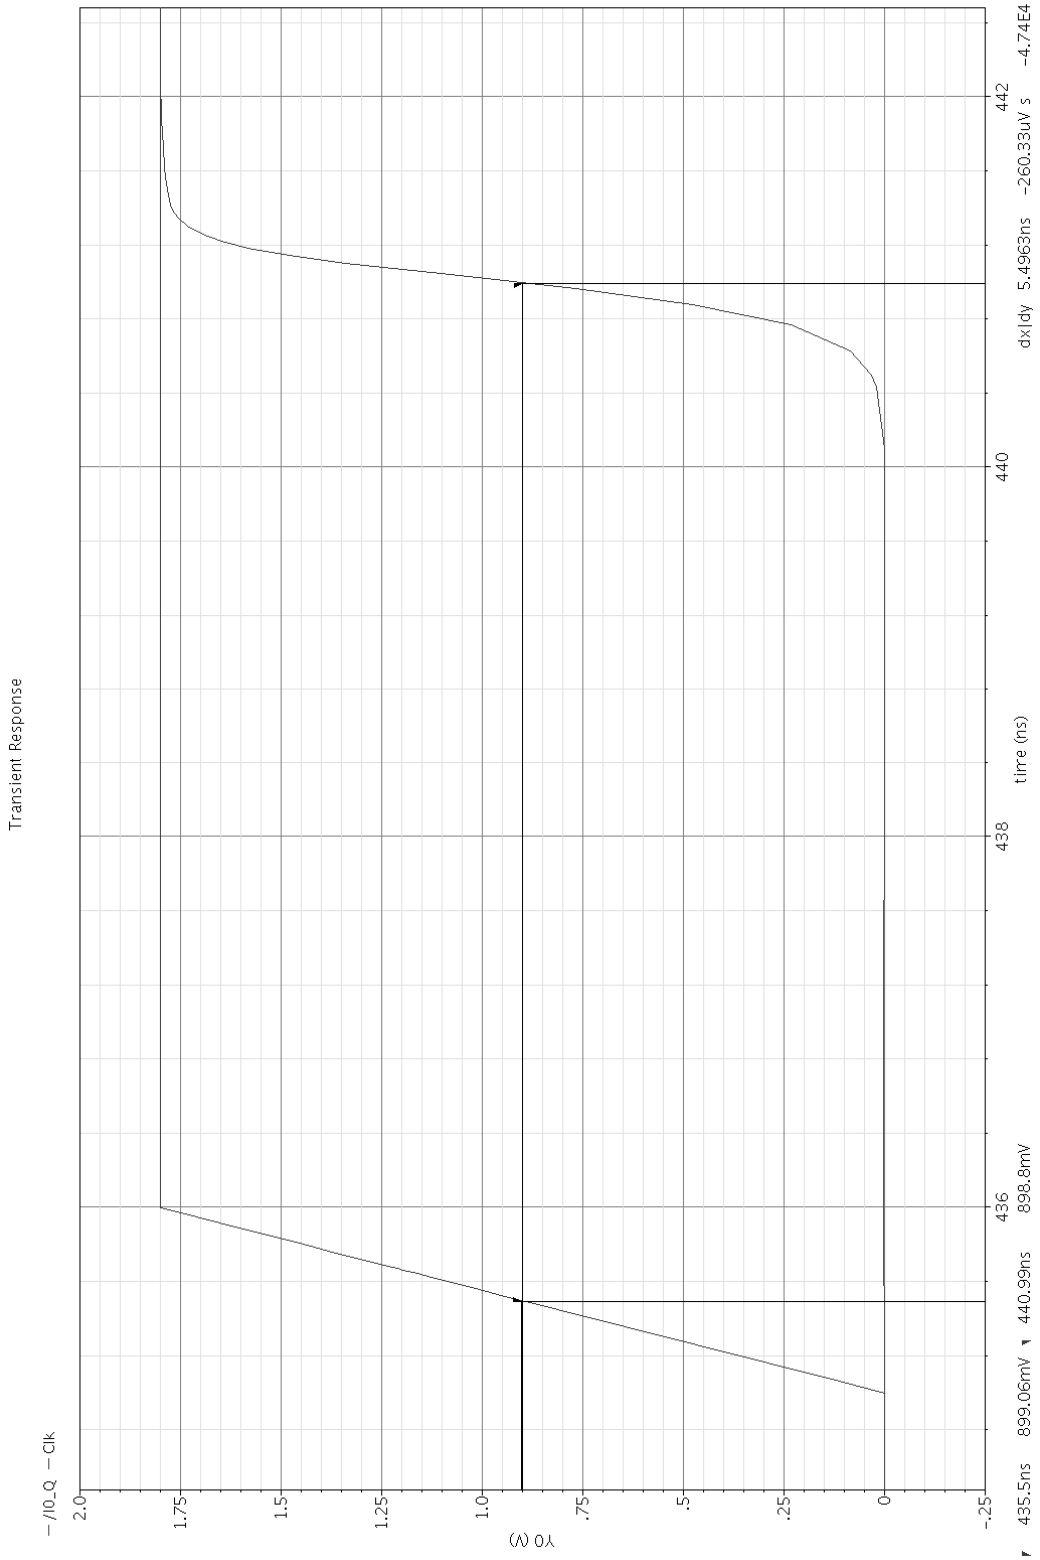

Figure 3.33: P-Type comparator propagation delay

Though these comparators appear fully functional, there are a few drawbacks to them. It turns out that some current is coupled through the gates of the input devices upon

clock flanks. Simulations show that this yields a ripple on the reference voltage when (relatively high impedant) divider strings are used to generate said references.
Also, due to the switching nature of the circuits, the reference current experiences some ripple, which will influence the V-I converter. To combat that effect, another reference source is required. This new reference source and any buffers that may be incorporated to stabilize the references will use a certain amount of power from the supply.
This means that even with the low power use of the comparators themselves, the total current consumption is still significant. All in all a fair amount of complexity is added to the system when these comparators are used without actually gaining much benefit in the end.
A better solution for this system, due to the mild demands on the comparators, is thus to use simple OPAMP like devices that are always active. Synchronization with the clock can be done by performing an AND operation on the clock and the output of the comparators. Figures 3.34 and 3.35 show the schematics of these new comparators.

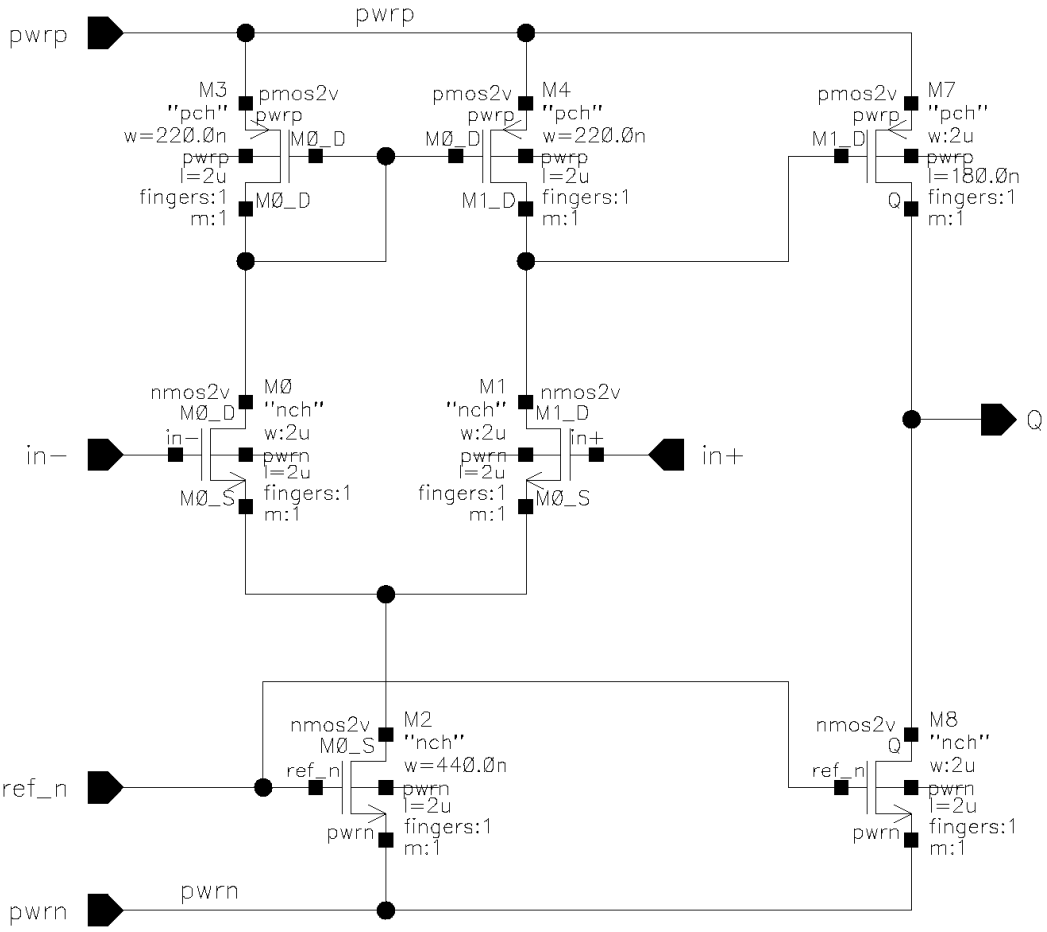

Figure 3.34: N-Type comparator

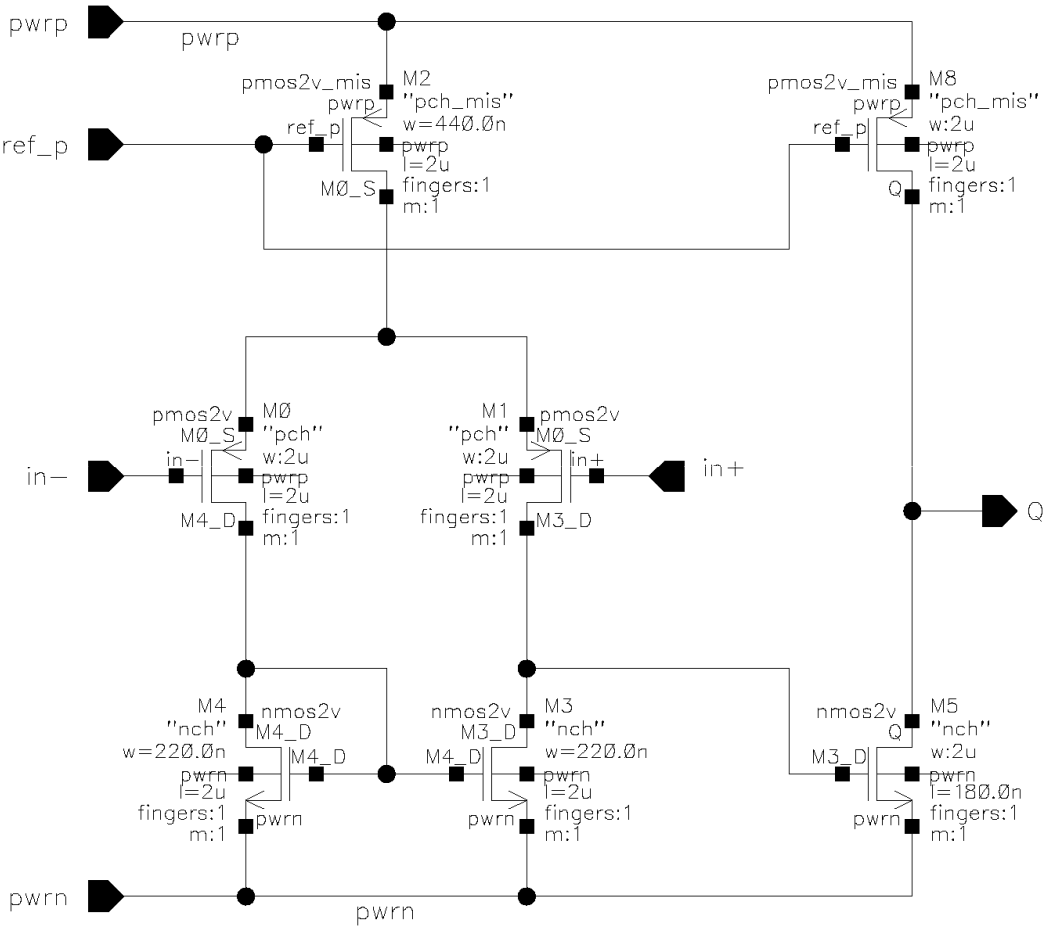

Figure 3.35: P-Type comparator

As can be seen, these systems are simple two stage OPAMPS, using current mirror loads for the input pair. No Miller capacitances are of course required since no feedback will be applied. The bias current of the input stage is approximately 220 nA and moderate size devices are used to have a fair amount of gain and fairly good matching. The output stage is biased at a somewhat higher value of 1 µA to improve the switching speed.

A response time of 149 ns was simulated for the N-type comparator, and 307 ns was found for the P-type. Both response times were measured with an input voltage difference of 3 mV again. Even though this is much slower than the clocked comparators, the system will not be affected. The only possible effect is that the dump phase starts slightly later, if a clock flank occurs during the response time of a comparator. The amount of charge that is removed from or added to the integrator will however not be affected.

Using a DC sweep in a Monte Carlo analysis, the offset of both comparators can be found. The results from those analyses can be found in figures 3.36 and 3.37. The offset of both circuits is roughly 15 mV in $3 \cdot \sigma$. Same as for the propagation delay, this can only affect when a charge dump is initiated, but will not affect the measurement result since the size of the charge package remains unaffected.

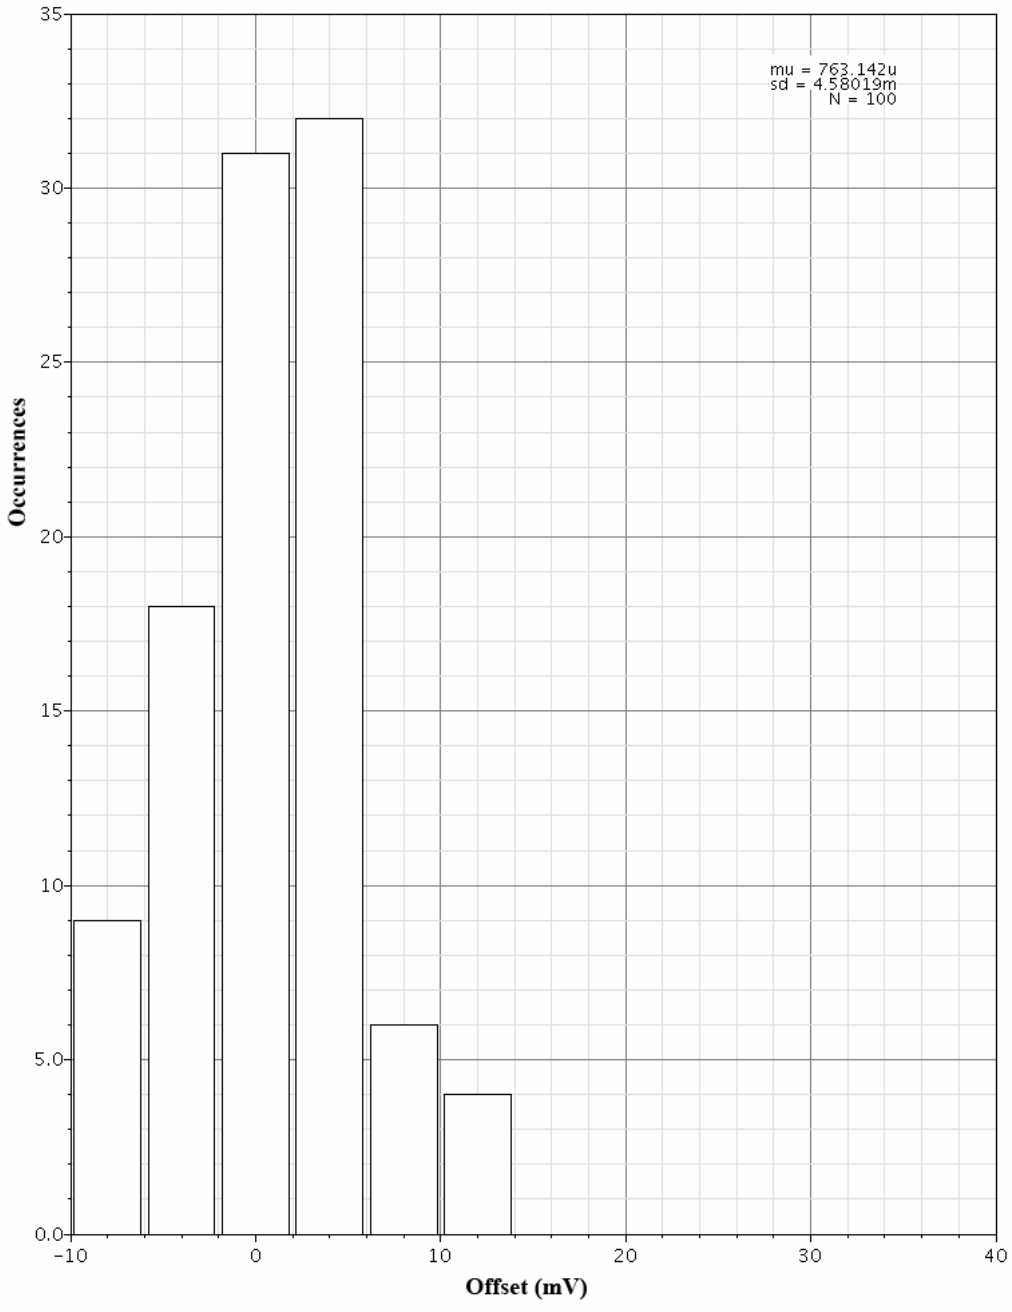

Figure 3.36: N-type comparator offset

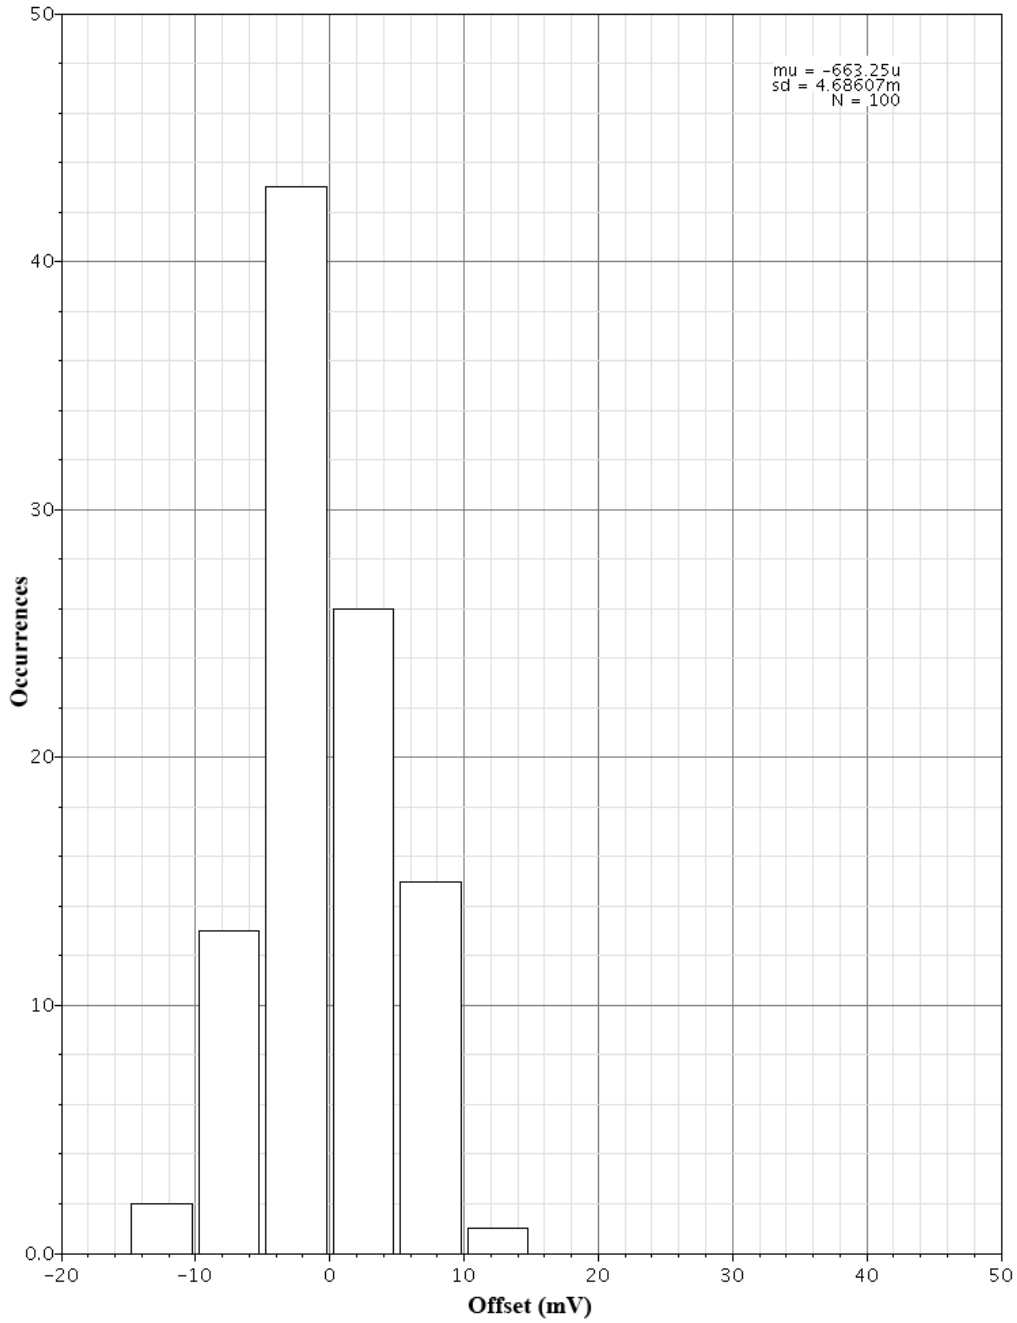

Figure 3.37: P-type comparator offset

3.3 Charge Pump Control Logic

This part of the circuit is responsible for controlling the actual charge packets that are applied to the integrator during the dump phase. The first step in designing this block is to determine the size of those packages, as well as finding a compromise between their magnitude and duration. The following demands are used to make a decision on this front:

- The size of the package has to be such that a significant part of the available supply voltage is utilized. Since this must hold true if the maximum input

voltage is applied to the system, the current must be significantly larger than the 28 nA input current associated with this voltage (figure 3.14).
- Clipping must be prevented under all process corners. In practice this means that the capacitor value can be as small as 85 % of its nominal value, and as large as 115 % (these values are taken from the process corner definitions used by SiTel).
- The duration of the dump phase is preferably as long as possible. This way the total number of switching events is minimized, and there is more averaging of clock jitter. Also, variation propagation delays will play a smaller role in this case, since the total duration of the dump is larger compared to these delays.

Since the current needs to be significantly larger than 28 nA, a value of 100 nA is selected. The factor four between the dump current and maximum input current ensures that a significant part of the supply voltage is utilized when the maximum input is applied, without the risk of clipping when the input current is small.

With the current chosen, the size of the charge package has to be determined. No clipping may occur even with the smallest possible integration capacitor. In this circuit, 'clipping' occurs if the total voltage difference on the capacitor during a dump phase exceeds the range between the comparator thresholds. If this situation should occur, the system will oscillate, since each dump will place the integrator output beyond the threshold of the opposite comparator. This would essentially turn the system into a sigma delta converter with a large hysteresis.
In order to create the threshold values of the comparators, a string of MOSFET diodes may be used (mimicking a resistive divider), or a MOSFET diode biased at a constant current. Either way, the threshold will be at least one V_{Th} removed from either supply rail, which is 390 mV. This makes 400 mV to 1.4 V a convenient voltage range for the integrator output.

In practice, the maximum voltage difference during a dump will occur if the input voltage is minimal. This is because no current from the V-I converter will counter the dump current in that case. Assuming the input voltage is zero (for simplicity) and the capacitor is at 85 % of its nominal value, the size of the charge package can be no more than:

$$\begin{aligned} \Delta Q_{Max} &= C_{Min} \cdot \Delta V_{Max} \\ \Delta Q_{Max} &= 0.85 \cdot 4\ pF \cdot 1V = 3.4\ pC \end{aligned} \qquad \textbf{3.18}$$

Note that, theoretically, the worst case scenario occurs if the input signal changes sign and assumes its maximum amplitude at the exact time the charge dump phase starts. However, if such an event causes the output of the integrator to go beyond the threshold of the opposite comparator, the system will recover by simply executing another charge dump (opposite to the previous one), which will counter the new input.

With a current of 100 nA, or 100 nC/s, the total duration of the dump phase can be no more than 34 μs. A value slightly less than this number is desirable, since offsets in the threshold values and comparator input stages can serve to narrow the acceptable voltage range. Also, mismatches and spread can influence the charge dump current. With this in mind, a duration of about 30 μs appears to be a safe option. Since the

clock operates at 1.152 MHz, this corresponds with 34.56 clock cycles. By using 32 clock cycles instead, a simple 5 bit counter can be used as timer for the dump phase. As can be seen in formula 3.19, this results in a charge package of 2.8 pC, which gives rise to a voltage difference of 820 mV.

$$\Delta Q = I_{Dump} \cdot t_{Dump}$$
$$\Delta Q = 100\ nA \cdot \frac{1}{1.152\ MHz} \cdot 32 \approx 2.8\ pC \qquad \textbf{3.19}$$

The opposite worst case scenario occurs when the input voltage is at its maximum value (countering the charge dump current) and the capacitor is larger than nominal. In this case, the current from the V-I converter can be as large as 30 nA (28 nA from a 100 mV input combined with maximum offset). The voltage difference that is built up during the dump phase will in that case be:

$$\Delta V_{Min} = \frac{\Delta Q_{Min}}{C_{Max}}$$
$$\Delta V_{Min} = \frac{\left(100\ nA - 30\ nA\right) \cdot \frac{1}{1.152\ MHz} \cdot 32}{1.15 \cdot 4\ pF} \approx 423\ mV \qquad \textbf{3.20}$$

This is still a significant swing, meaning the charge dump current and duration derived here are a viable solution.

Note: in the initial design and test phases, the dump phase lasted 38 clock periods. This yields a higher voltage swing between dumps, which was later found to cause problems if the capacitor is at 85% of its nominal value. For the sake of readability these early considerations are not fully treated here, but the initial test results have been obtained with them, hence they are mentioned.

The next task is thus to create a block that can control the charge pump. It is assumed that the charge pump has a control input for both positive and negative currents. If either comparator goes high, the respective control input has to be kept high for 32 clock cycles. A simple way to do this is by connecting SR latches to the control inputs and using the outputs of the comparators to set these latches. The control logic simply has to start counting when either latch is set, and reset both latches after 32 clock cycles. Note that the set input of the latches needs to be synchronized with the clock as well, which can be done by means of a simple AND operation between the clock and the comparator outputs.

A gate level implementation of a circuit that performs this operation can be found in figure 3.38. As can be seen, it is a simple 6 bit counter that is kept in a reset condition unless either input is high. When an input goes high, the counter starts and produces a high level at Q5 after 32 periods. The output can reset the latches, which in turn resets the counter itself.

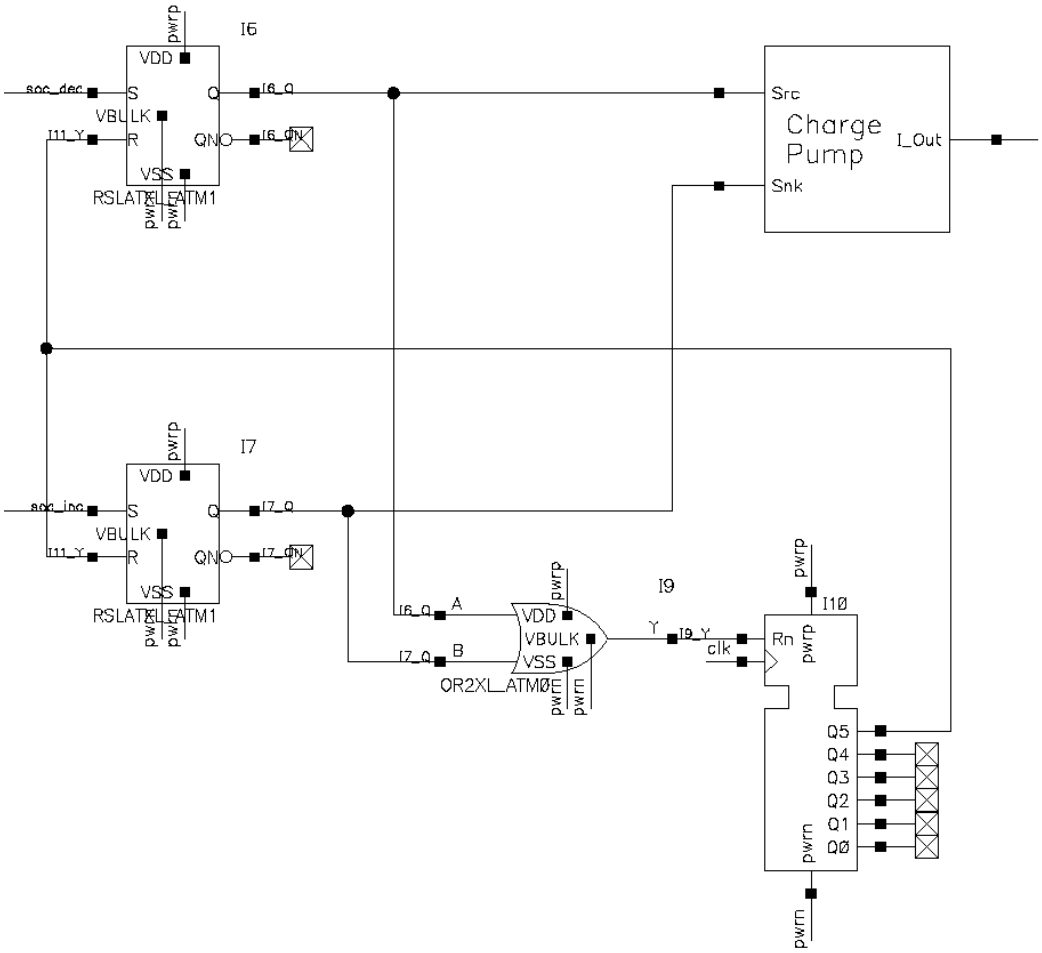

Figure 3.38: Charge pump controller

To speed up the simulation, a behavioral description in Verilog was created as well. This description can be found in code fragment 3.1, and includes the timer and the OR gate that connects to its Rn input. The latches are tested using the Verilog descriptions provided in their model files. Testing of this block will be done as an integral part of testing the whole system.

```
//Verilog HDL for "thesis", "charge_pump_control" "functional"


module charge_pump_control ( Q, A, B, clk, pwrn, pwrp );

  //Inputs & Outputs
  input clk;
  input pwrp;
  input A;
  input B;
  input pwrn;
  output Q;

  //Internal Registers
  reg[5:0] counter;
  reg res;
  reg Q_int;

  //Initial Conditions
  initial
  begin
    Q_int <= 1;
    counter <= 0;
    res <= 1;
  end

  //Count only if reset is inactive
  always @(posedge clk)
  begin
    if (res)
    begin
      counter <= 0;
    end
    else
    begin
      counter <= counter + 1;
    end
  end

  //Perform OR operation and invert reset input
  always @(A or B)
  begin
    res <= !(A || B);
  end

  //Assign output
  always @(counter)
  begin
    Q_int <= counter[5];
  end

  assign Q = Q_int;

endmodule
```

Code Fragment 3.1: Verilog code for charge pump control

3.4 Charge Pump

The final part of the basic system is the actual charge pump. The basic idea behind this block is to create a current dump source and a current sink source from the reference by use of current mirrors and connecting either source to the output by means of switches. The dimensioning of the devices will be based on matching properties and the switches and their drivers are optimized to reduce charge injection and clock feedthrough.

The current sources are created from the reference by means of a number of current mirrors, that will be cascoded for increased precision. This way the effects of channel

length modulation are minimized, creating a good match between the sink and source currents. This is of course only true when mismatch is not considered.

For the analysis of the current mirrors, they are split up in two sections: the current sources and the cascode devices. The dimensioning of the current sources is based on matching. The source that creates the actual reference current needs to have a W/L of 1/10, to conduct 100 nA (the reference voltage yields 1 µA in a 1/1 device). Since this source provides the current for both the dump and sink sources, its matching properties are not very important. To create a fairly high output resistance, the length of this source is 5 µm, yielding a fairly small width of 500 nm.

For the other sources, matching plays an important role. The aim is to create a $3 \cdot \sigma$ mismatch between source and sink currents that is in the order of magnitude of the maximum measurement error. This error will still be reduced by the chopping mechanism, but since the error in the input source can already be very high compared to the input signal, it is desirable to minimize the other error sources in the system.

For finding the mismatch between the source and sink currents, the circuit in figure 3.39 is used. In this circuit, any deviation in the properties of M0 from the designed values influence I_1 and I_2 in the same manner, and M0 is thus ignored in this derivation. In order to find the mismatch between I_1 and I_3, the mismatch between I_1 and I_2 is quadratically added to the mismatch of I_2 and I_3. This assumes all mismatches to be uncorrelated. The result is then divided by two, since a certain total mismatch between the two sources yields half that mismatch as average current.

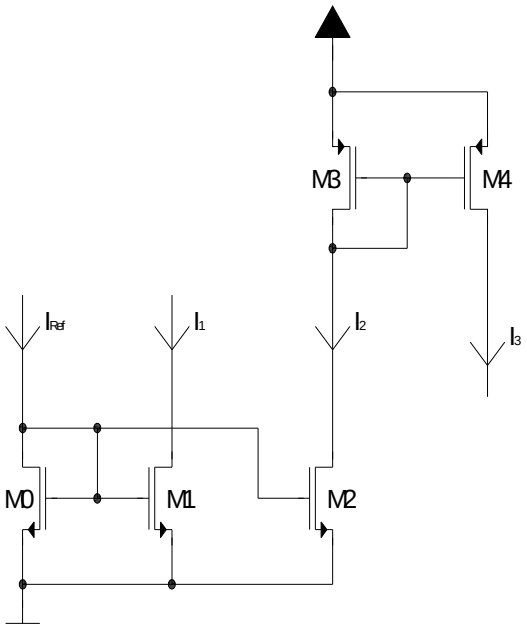

Figure 3.39: Circuit used for mismatch estimation

Assuming the mismatches are small compared to the component values, their influence on the current can be found by taking the derivative of the drain current with respect to the parameter of interest, and multiplying the result by the absolute value of the mismatch. Since mismatch can work in either direction, the absolute value is taken as well. For the threshold voltage mismatch this results in formula 3.12.

$$I_D = \frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L} \cdot (V_{GS} - V_{Th})^2$$
$$\frac{dI_D}{dV_{Th}} = -\mu \cdot C_{OX} \cdot \frac{W}{L} \cdot (V_{GS} - V_{Th}) \qquad \textbf{3.21}$$
$$|\Delta I_D| \approx \mu \cdot C_{OX} \cdot \frac{W}{L} \cdot (V_{GS} - V_{Th}) \cdot \Delta V_{Th}$$

This can be worked out further by substituting the formulas for $V_{GS} - V_{Th}$ and ΔV_{Th} as in formula 3.22.

$$V_{GS} - V_{Th} = \sqrt{\frac{I_D}{\frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L}}}$$
$$|\Delta I_D| \approx \mu \cdot C_{OX} \cdot \frac{W}{L} \cdot \left(\sqrt{\frac{I_D}{\frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L}}} \right) \cdot \Delta V_{Th}$$
$$|\Delta I_D| \approx \sqrt{2 \cdot I_D \cdot \mu \cdot C_{OX} \cdot \frac{W}{L}} \cdot \Delta V_{Th} \qquad \textbf{3.22}$$
$$|\Delta I_D| \approx \sqrt{2 \cdot I_D \cdot \mu \cdot C_{OX} \cdot \frac{W}{L}} \cdot \frac{A_{VTH}}{\sqrt{W \cdot L}}$$
$$|\Delta I_D| \approx \sqrt{2 \cdot I_D \cdot \mu \cdot C_{OX}} \cdot \frac{A_{VTH}}{L}$$

For the K-factor mismatch, a parameter of $\Delta K/K$ was simulated (appendix B), hence the result of the derivation to K has to be multiplied by K again to use this factor. This simply yields the quadratic equation back, as shown in formula 3.23.

$$I_D = \frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L} \cdot (V_{GS} - V_{Th})^2$$
$$\frac{dI_D}{dV_K} = (V_{GS} - V_{Th})^2 \qquad \textbf{3.23}$$
$$|\Delta I_D| \approx \frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L} \cdot (V_{GS} - V_{Th})^2 \cdot \frac{\Delta K}{K}$$

Substitution of the formulas for $\Delta K/K$ and $(V_{GS} - V_{Th})^2$ yields:

$$|\Delta I_D| \approx I_D \cdot \frac{A_K}{\sqrt{W \cdot L}} \qquad \textbf{3.24}$$

With these formulas an estimate for the mismatch currents can be made. In order to find the total mismatch for one device, formulas 3.22 and 3.24 are added

quadratically. The result is then multiplied by √2 to find the mismatch between two devices.

$$\left|\Delta I_D\right| \approx \sqrt{2} \cdot \sqrt{2 \cdot I_D \cdot \mu \cdot C_{OX} \cdot \left(\frac{A_{VTH}}{L}\right)^2 + \frac{\left(I_D \cdot A_K\right)^2}{W \cdot L}} \qquad \textbf{3.25}$$

As can be seen, the effect of increasing the length is more profound than that of increasing the width. Since at this time the equation is based on two unknown variables (L and W), a choice for either value has to be made. Since both influence the result, but L more so than W, L is simply set to 2 · W. Note that this is not necessarily the optimal solution, but it is an educated guess that is used to yield an equation that can be solved. This was done to reduce formula 3.25 to a formula with only one unknown variable. This changes formula 3.25 into the set of formulas found under 3.26.

$$\begin{aligned}
\left|\Delta I_D\right| &\approx \sqrt{2} \cdot \sqrt{2 \cdot I_D \cdot \mu \cdot C_{OX} \cdot \left(\frac{A_{VTH}}{2 \cdot W}\right)^2 + \frac{\left(I_D \cdot A_K\right)^2}{W \cdot 2 \cdot W}} \\
\left|\Delta I_D\right| &\approx \sqrt{2} \cdot \sqrt{\frac{I_D \cdot \mu \cdot C_{OX} \cdot A_{VTH}{}^2}{2 \cdot W^2} + \frac{\left(I_D \cdot A_K\right)^2}{2 \cdot W^2}} \qquad \textbf{3.26} \\
\left|\Delta I_D\right| &\approx \frac{1}{W} \cdot \sqrt{I_D \cdot \mu \cdot C_{OX} \cdot A_{VTH}{}^2 + \left(I_D \cdot A_K\right)^2} \\
W &= \frac{1}{2} \cdot L
\end{aligned}$$

Since the matching parameters of both NMOS and PMOS devices are in the same order of magnitude, the mismatch in the PMOS mirror is expected to be smaller than that of the NMOS mirror (due to the lower μ_p). Therefore, for the choice of W and L, the error is assumed to be mainly in the NMOS part of the circuit. The average current is a measure of the matching between the upper and lower halves of the circuit.

An average current of |2 nA| within 3 · σ seems like a reasonable value. With such a value, the error caused by this block remains significantly lower than that caused by the input stage under all conditions. Since the average current is found by adding the source and sink currents and then dividing by two, the offset can be a factor two higher than the maximum error in the average, or |1.33 nA| in 1 · σ. For this, a width of 14 µm is required, and hence a length of 28 µm. For convenience, a width of 10 µm and a length of 20 µm was chosen for all devices. A slightly higher offset may thus be found, but one has to bear in mind that approximations have been used to estimate the offset in the first place.

The next step is to dimension the cascode devices. These devices serve to minimize the adverse influence of a second order effect, and need not be excessively large. A length of 5 µm has been used, which still gives a fairly high output resistance (see appendix B). With such a length, and a current of 100 nA, the required gate source voltage can be found by means of formula 3.27.

$$V_{GS} = \sqrt{\frac{I_D}{\mu \cdot C_{OX} \cdot \frac{W}{L}}} + V_{Th}$$

$$V_{GS,N} = \sqrt{\frac{100\ nA}{0.28\ \frac{mA}{V^2} \cdot \frac{W}{5\ \mu m}}} + 390\ mV \qquad \textbf{3.27}$$

$$V_{GS,P} = -\sqrt{\frac{100\ nA}{0.048\ \frac{mA}{V^2} \cdot \frac{W}{5\ \mu m}}} - 390\ mV$$

Using a width of 1 µm for both N and P type devices, gate source voltages of no more than 490 mV are required. Since the actual current sources need less overdrive, each cascode will not consume more than half the supply voltage, leaving enough margin to ensure that all devices are in saturation.

The next important part of the charge pump are the switches. These are designed to minimize the amount of current injected during the switching phases. In order to do this, a few demands have been stated:

- The switches are as short as possible (180 nm). Their width is also small, but not too much so to ensure a low on-resistance.
- The current in each branch is never switched off. Instead, the current is steered to either the output of the system or a dummy node which is created by copying the voltage level at the output.
- The switches are not driven into the triode region. When a MOSFET enters triode, its parasitic capacitances change, which makes compensation for clock feedthrough and charge injection by dummy devices difficult.

To meet the first demand, each switch consist of two devices of the same type. Controlling the switch therefore requires complementary signals. The switch drivers have to be designed such that these signals match each other as closely as possible during the transitions. The NMOS switch can be found in figure 3.40 and the PMOS switch can be found in figure 3.41. The complete charge pump can be found in figure 3.44, which shows how these switches are connected.

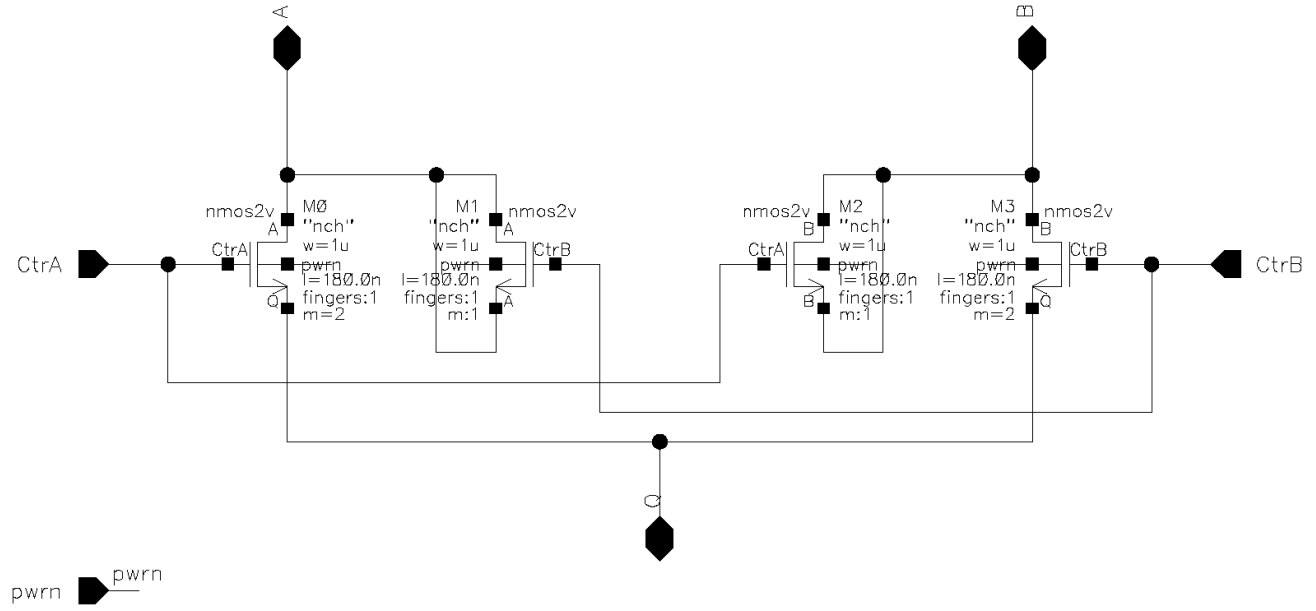

Figure 3.40: NMOS switch

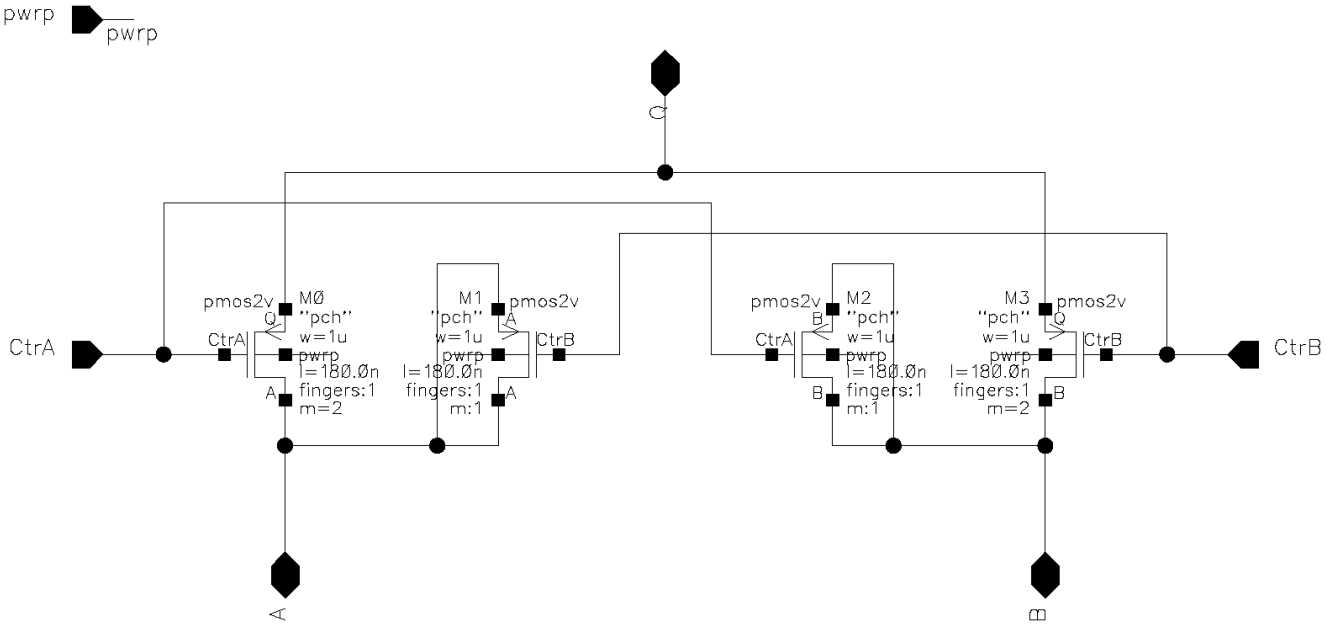

Figure 3.41: PMOS switch

As can be seen, the switch devices are twice as large as the dummy devices. This has been done, because the dummies have their drain and source shorted, which doubles the capacitances. Compensation at the sources of the switching devices is performed by the complementary operation of said devices (i.e. M0 is compensated by M3 and vice versa).

Avoiding deep triode operation can be achieved by clipping the switch control signals before they reach either supply rail. Therefore, the switch driver outputs switch between 400 mV and 1.4 V. The voltage at the output of the charge pump is 900 mV (due to the reference level of the integrator). This means the devices will still be driven into the triode region, but not as deep as would be the case if the entire supply

range was used. Since an active switch acts as a source follower, the level at the source of an active NMOS switch will be close to one threshold below its gate voltage. This same level appears at the source of the inactive switch, resulting in a negative gate source voltage for said device which ensures a high off-resistance.

The schematic of the switch drivers can be found in figure 3.42. The top half, which creates Q, consist of a string of inverters. The first inverter is simply a minimum size block, with a PMOS that is wider than the NMOS to compensate for the difference between μ_n and μ_p. The second and third inverters also have this compensation, but it is achieved by scaling both W and L. This has been done in such a way that the total device area of NMOS and PMOS devices is roughly equal as well. The CAPAS at the output make up a non-linear, but symmetrical capacitive load, that is large compared to the input capacitance of the switches. This way the resemblance between positive and negative transitions is improved, and the amount of clock feedthrough is reduced by decreasing the high frequency components of the switching signal.
The additional CAPAS in the lower half of the circuit, which generates Qn, serve to add some propagation delay, since this stage lacks one inverter compared to the top half. Note that the output stages have different supply voltages than the other inverters.

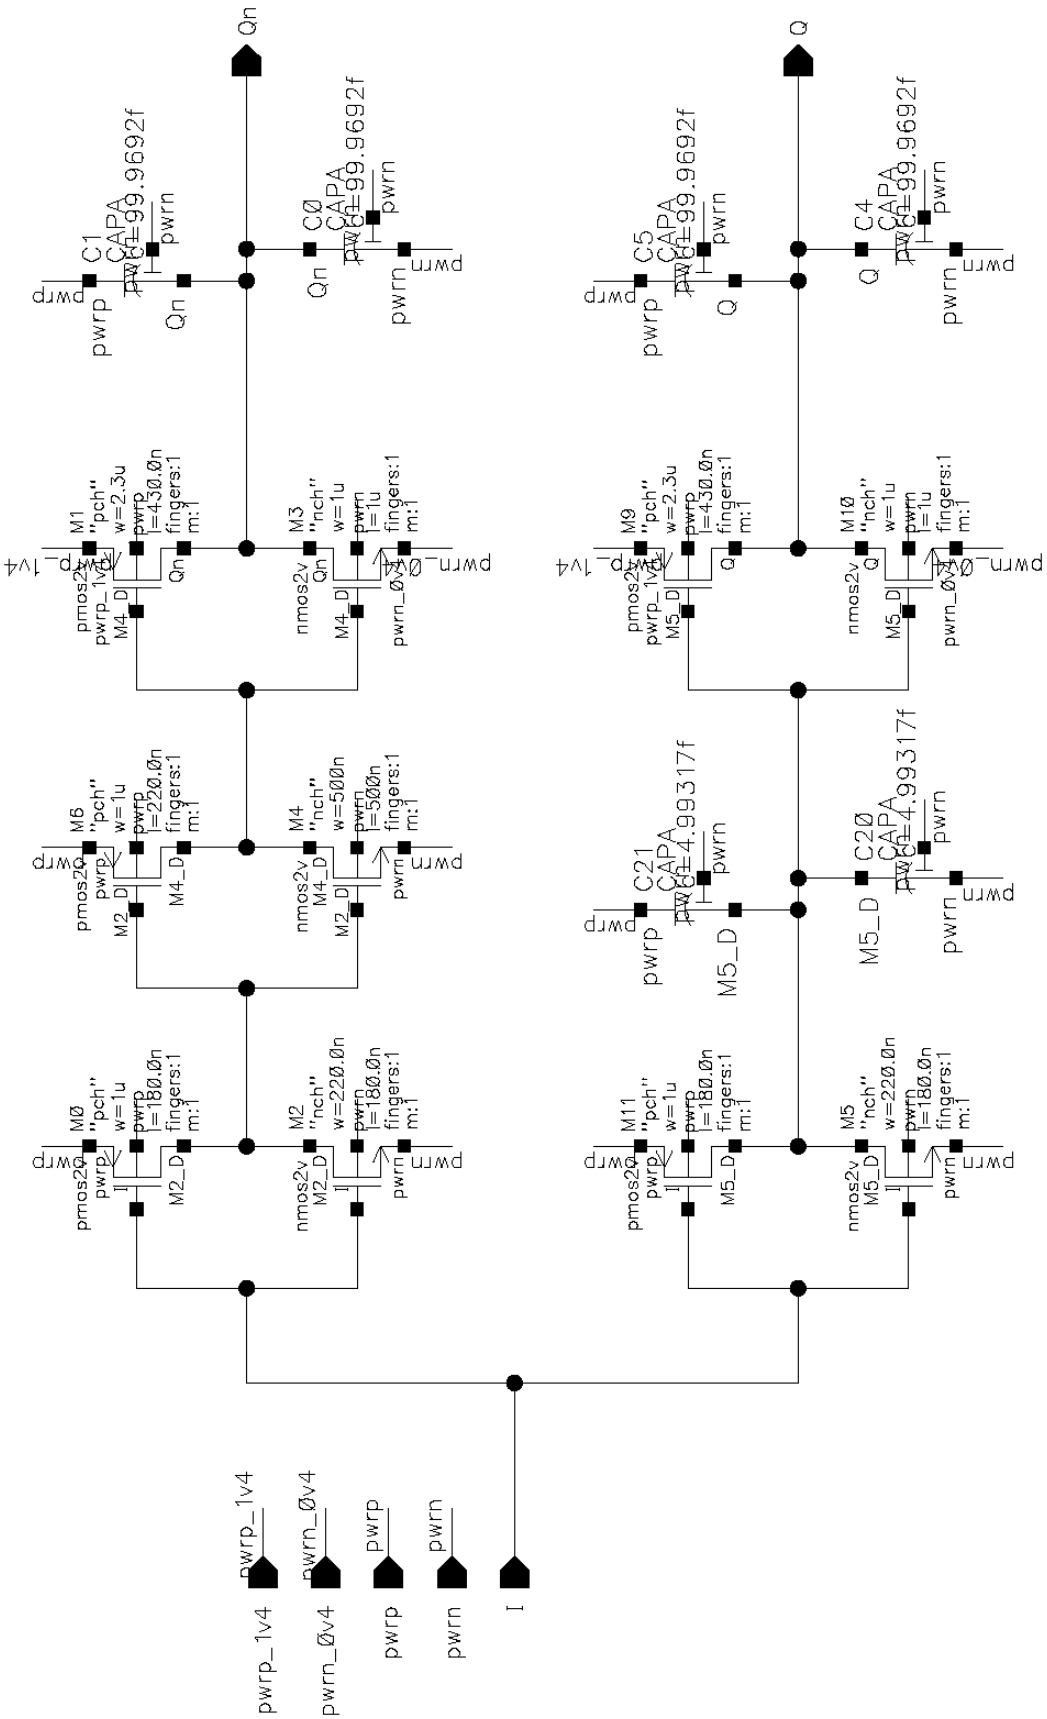

Figure 3.42: Switch driver

Figure 3.43 shows a transition in the outputs of the switch driver. A close match between positive and negative flanks can be seen.

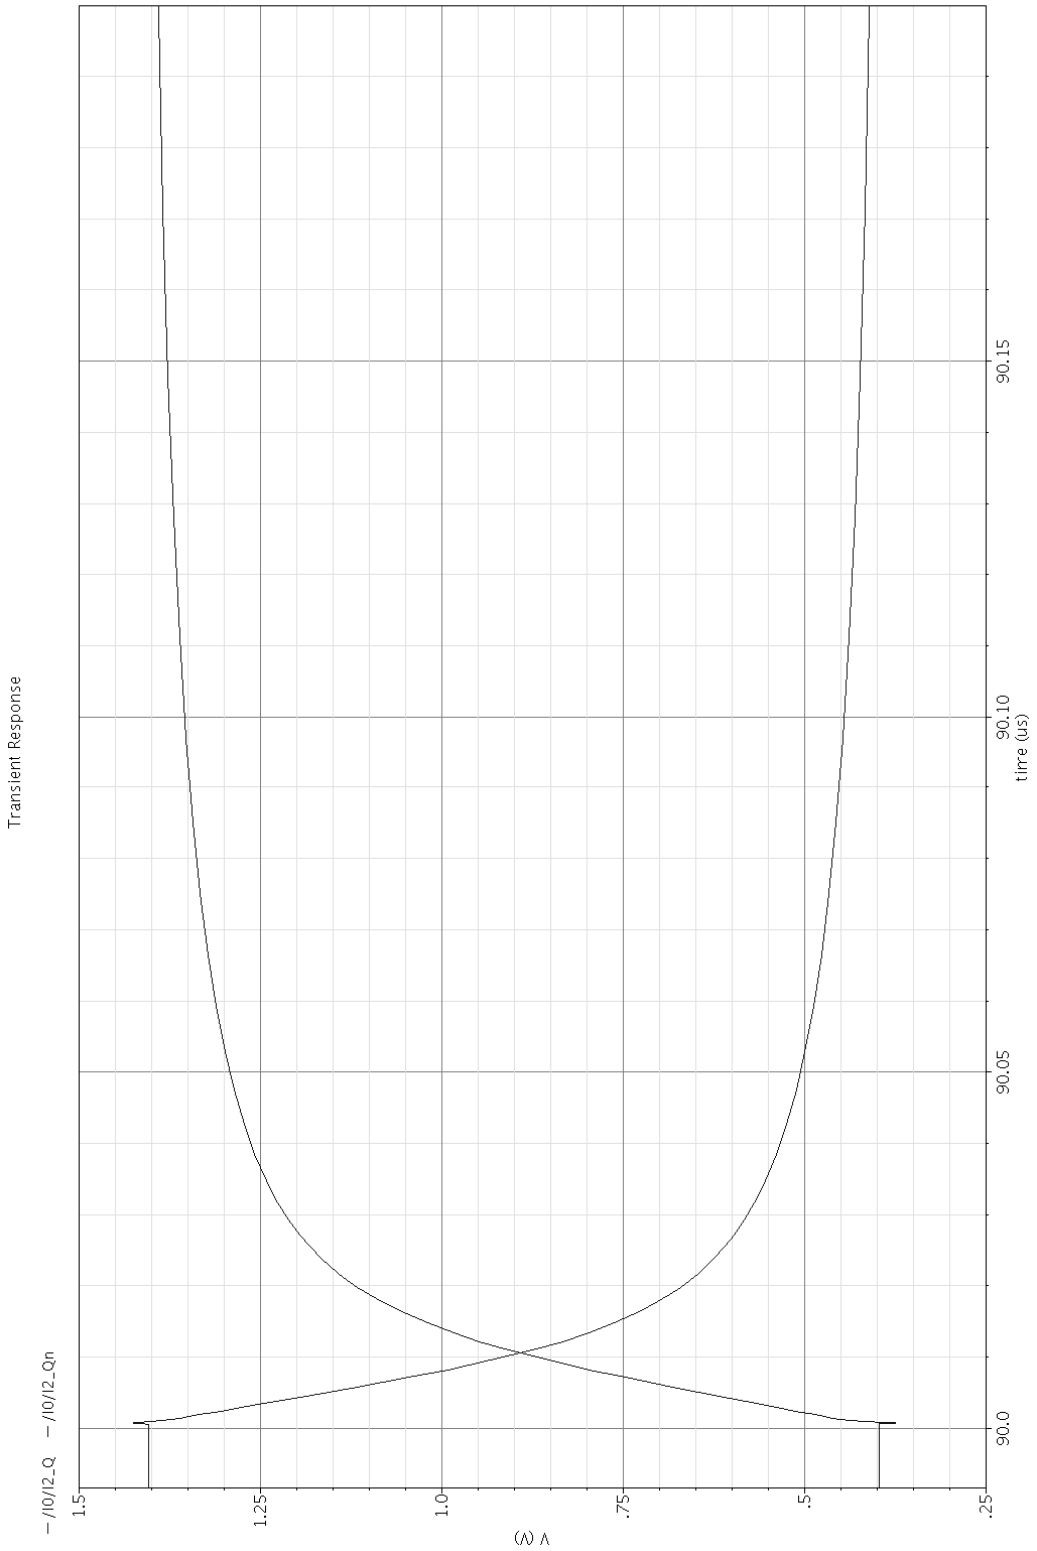

Figure 3.43: Switch driver transition

With each component defined, the total charge pump circuit can be created. The schematic can be found in figure 3.44. Apart from the current mirrors and the switch drivers and switches, a unity gain buffer can be seen that creates the dummy output voltage. Also, a divider string is added to create the 400 mV and 1.4 V references.

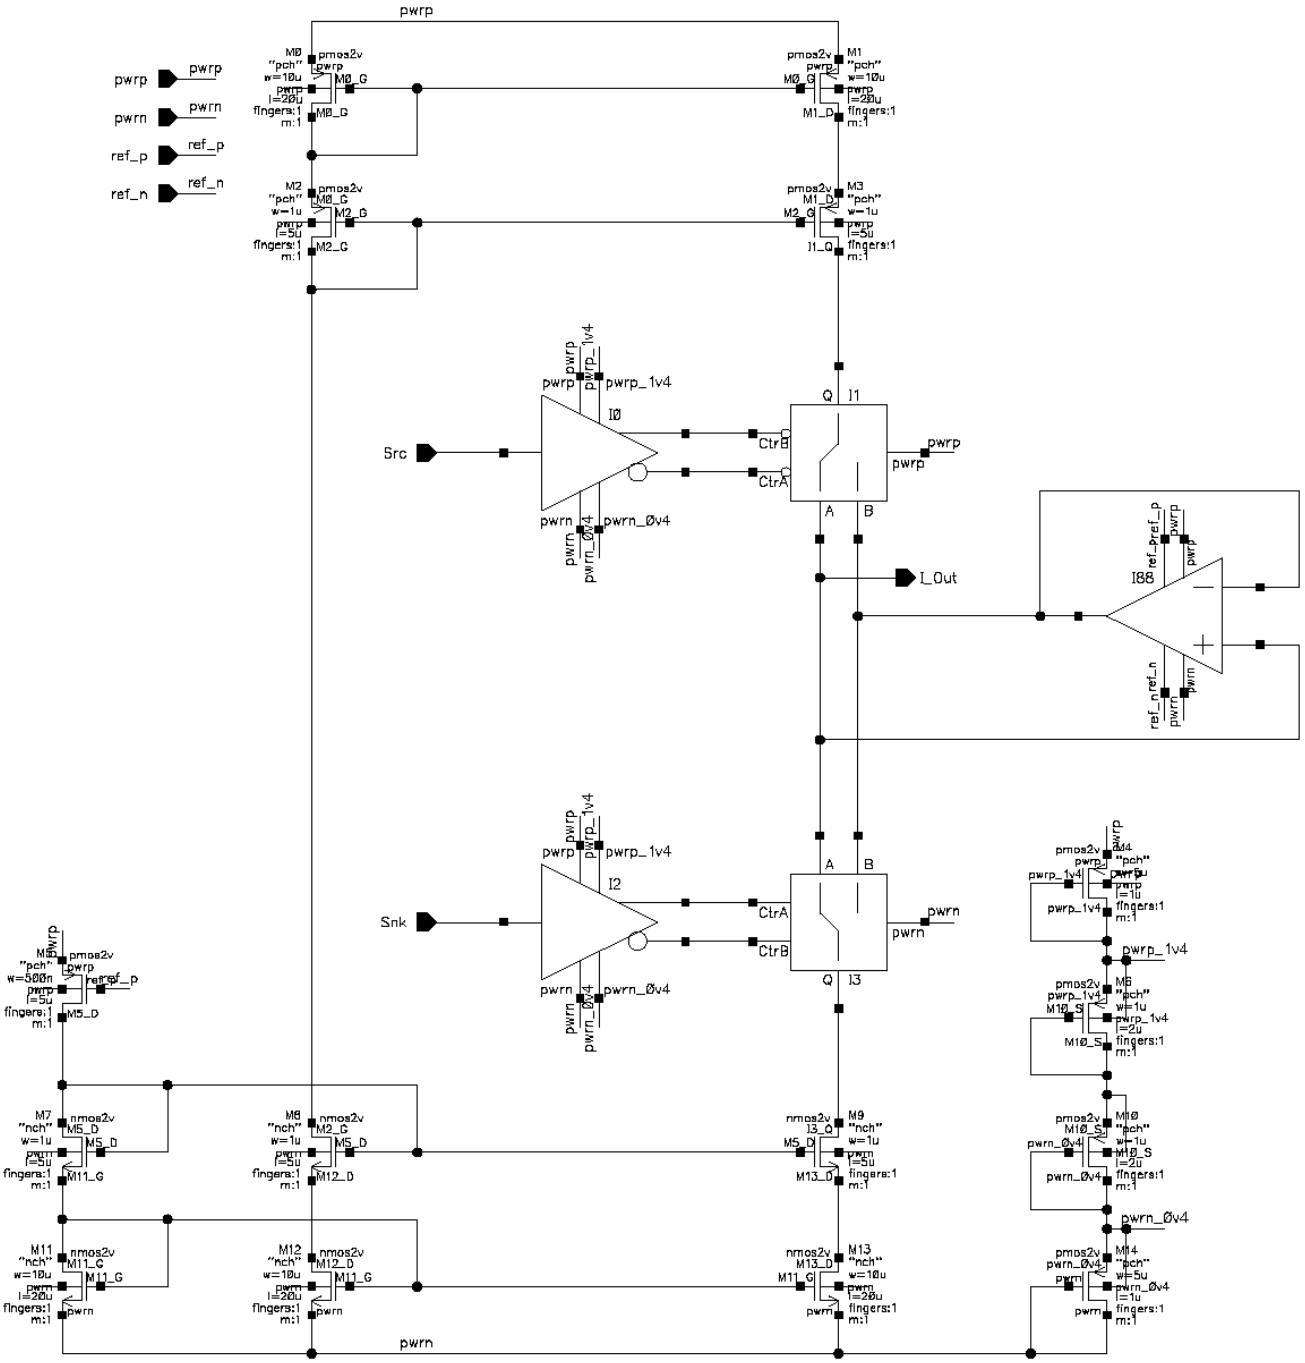

Figure 3.44: Complete charge pump

Figure 3.45 displays the result of a transient simulation, where the output is continuously switching between source and sink mode. As can be seen, the current is slightly larger than |100 nA|. This is caused by the PTAT source, whose output

voltage that slightly deviates from the value that yields 1 µA in a 1/1 device (see chapter 3.1). The peaks on the flanks are very narrow and reach magnitudes of no more than roughly 300 nA. A simulation where simple minimum sized, single transistor, deep triode switches were used showed peaks of up to 20 µA, hence the circuit has been greatly improved by adding the measures described earlier. An average current of 1.2 pA was found in this simulation, indicating a very good match between sink and source currents and transitions.

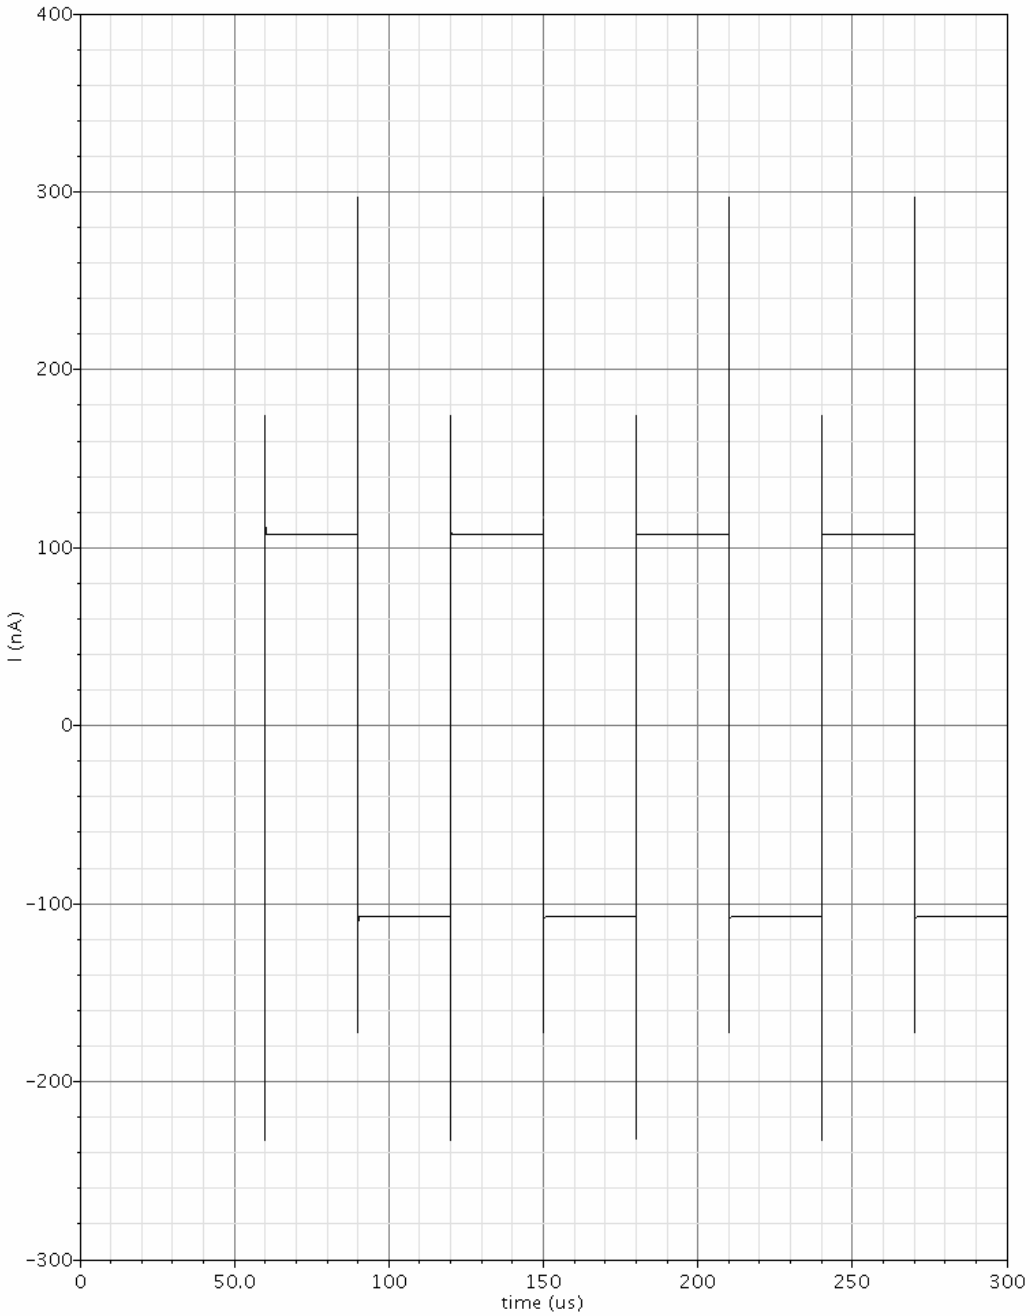

Figure 3.45: Transient of ideal charge pump

To verify the behavior of the system under mismatch and process spread, the average output current (which defines the error between the sink and source modes) has been measured in corner and Monte Carlo simulations. The results can be found in figures 3.46 and 3.47.

The corner simulations show very good behavior over all process corners. The absolute value will have to be calibrated, but the matching between sink and source current (and thus the influence on the linearity of the system) remains extremely good.

The mismatch plot shows a $1 \cdot \sigma$ error of 0.57 %. With a maximum current of 108 nA, this corresponds to 616 pA. As stated earlier, this is half of the actual mismatch current, which is thus 1.23 nA. This is actually a bit better than expected, especially since the devices are slightly shorter than originally deemed necessary. The difference is however not large enough to assume that either the prediction or the simulation result is flawed.

The final simulation verifies the behavior of the output current over temperature. As can be seen in the corner simulations, the average current remains well within limits. However, since temperature can vary for one IC, the absolute value of the current matters as well for that parameter. Unfortunately, the output current increases with temperature in the current configuration. The output current of the input stage on the other hand decreases, hence a significant error will arise. At present, the charge pump is biased by the same PTAT source as the V-I converter. The output current varies by roughly ± 16 % over the entire temperature range, and has a positive temperature coefficient. Since the PTAT reference current varies with about ± 20 % over the same temperature range (see section 3.1), much of this error is actually caused by the reference. Therefore a reference that is stable with temperature would solve most of the problem. A better possibility still would be to implement an NTAT reference that gives the charge pump current a negative temperature coefficient equal (or close to) the coefficient that is found in the input stage. Such a source has however not been implemented yet due to lack of time.

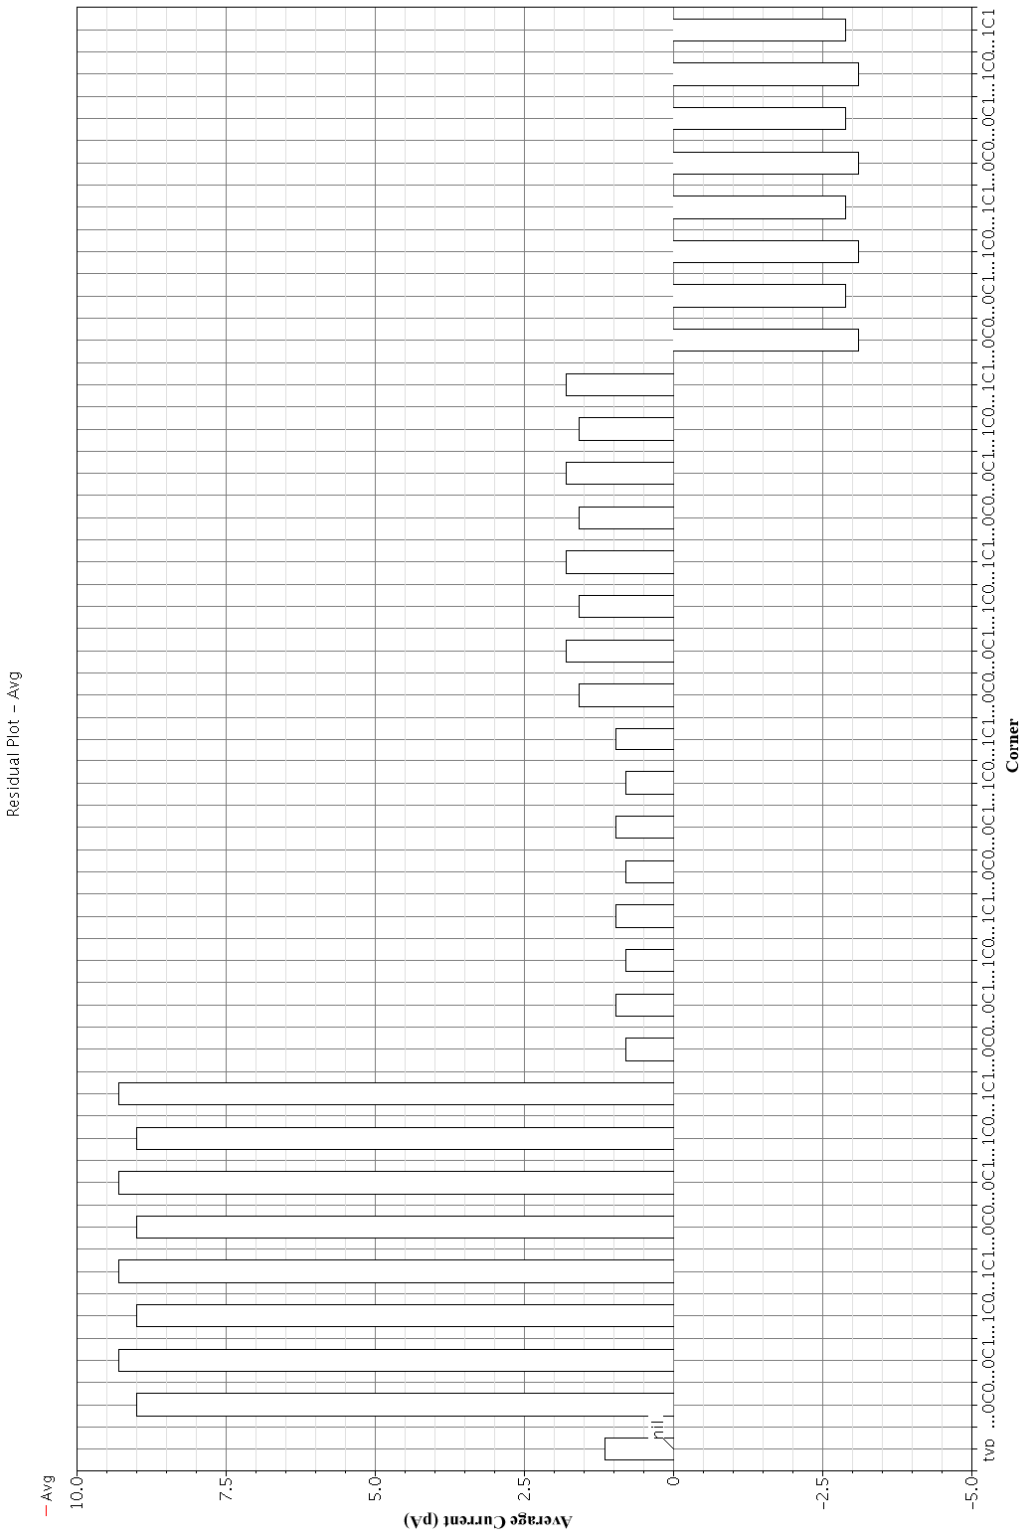

Figure 3.46: Error in average due process spread

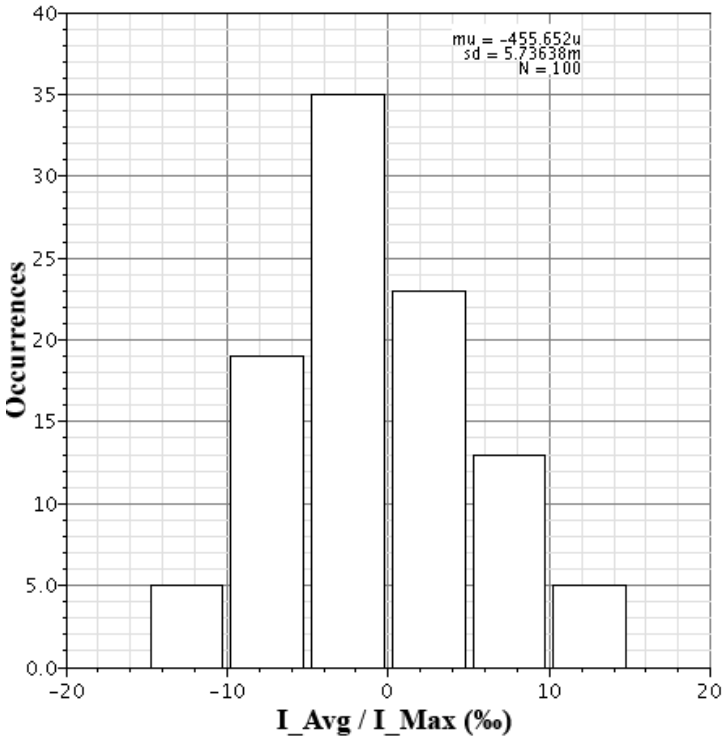

Figure 3.47: Error in average due to mismatch

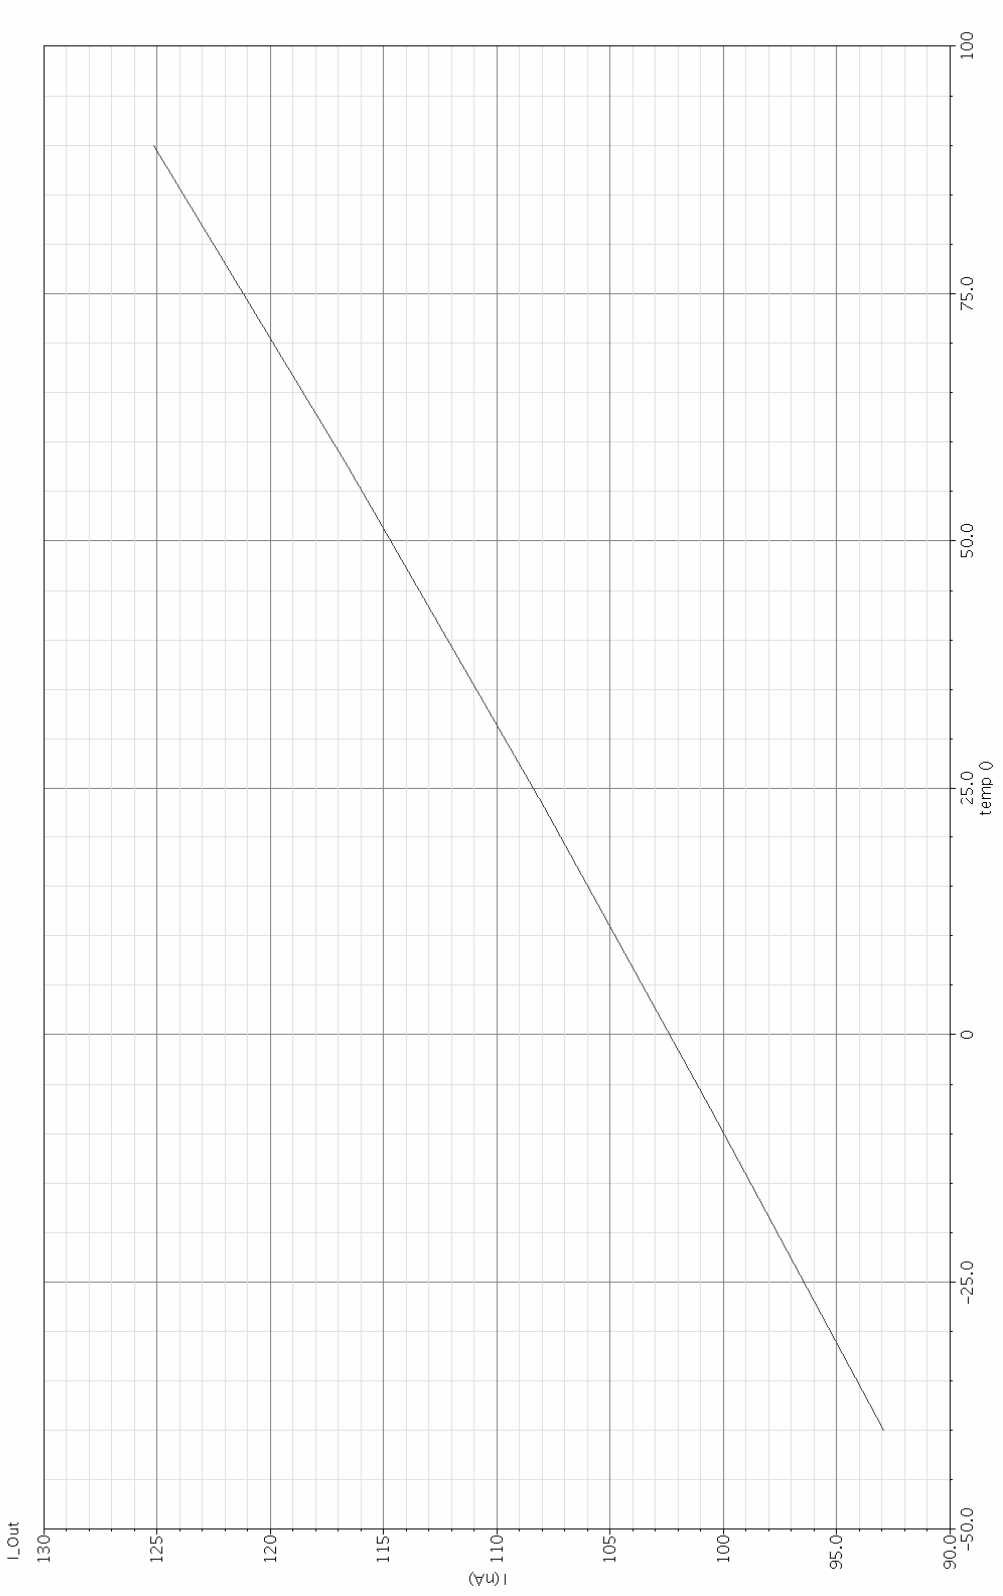

Figure 3.48: Current magnitude over temperature

4. Chopping

As mentioned in the foregoing chapters, a chopping mechanism has to be added to the system in order to remove the offset from the input stage and any mismatch in the charge pump from the measurement result. This chapter explains the simple mechanism that is used in the current system, and also describes two alternative methods that were considered but not implemented.

4.1 Applied Strategy

4.1.1 Theory

In order to accurately remove errors from the measurement result by means of chopping, it is important that the system spends an equal amount of time in both possible modes (normal and inverted). An easy way to do this is by always swapping the inputs after a fixed amount of time. This is a very simple strategy, but it has one major drawback. Figure 4.1 displays this drawback graphically. The solid line shows the integrator output voltage in time, and the dotted line shows the same output if no chopping had occurred. As can be seen, an error will be introduced that is dependant on the output voltage of the integrator at the moment of chopping.

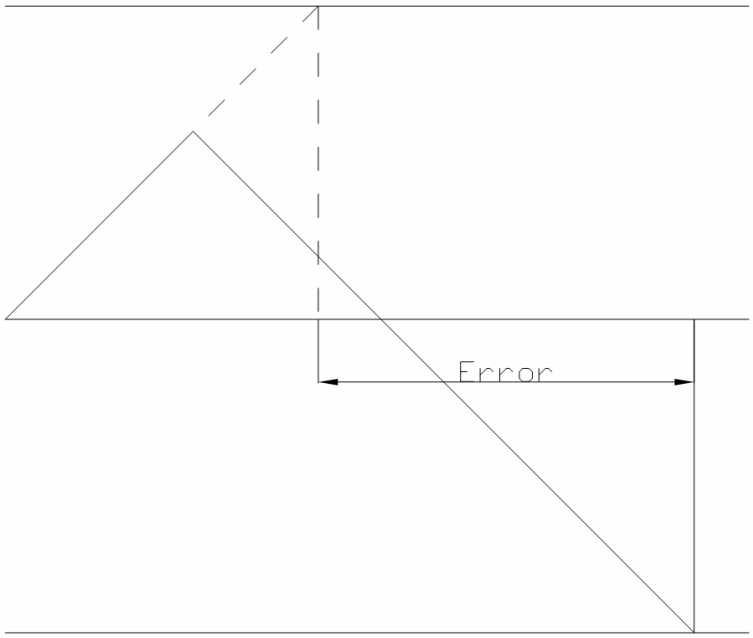

Figure 4.1: Error introduced by chopping

This error can be kept within limits by increasing the time between chopping events. The worst case error occurs if the inputs are switched around an infinitesimal amount of time before either comparator would trigger. In that event the entire range between the positive and negative comparator thresholds needs to be covered before the next charge dump occurs. Since ideally the amount of time this takes is proportional to V_{In} in the same way as the normal ‘period’, the error is a constant factor of the input. This means that if the error caused by this chopping mechanism can be found for one input value, it is known for all inputs.

In order to find the actual error, a signal of 100 µV is assumed to be present at the input of the system. With that signal, the output current of the V-I converter is approximately 28 pA (see chapter 3.1), which is significantly smaller than the charge pump output current (108 nA, chapter 3.4) and the net current during a charge dump can be assumed to be simply 108 nA. Also, the integration period with such a small input signal is much longer than the duration of the charge dump phase, meaning said duration can be considered infinitesimal. Assuming no error in the size of the integration capacitor, the voltage difference built up during a charge dump phase is then:

$$\Delta V = \frac{\Delta Q}{C}$$

$$\Delta V = \frac{108\ nA \cdot \frac{1}{1.152\ MHz} \cdot 32}{4\ pF} = 750 mV \qquad \textbf{4.1}$$

Since the output current of the V-I converter is 28 pA, the integration period (or time between charge dumps) is:

$$\Delta t = \frac{C \cdot \Delta V}{I}$$

$$\Delta t = \frac{4\ pF \cdot 750\ mV}{28\ pA} \approx 107\ ms \qquad \textbf{4.1}$$

The voltage difference between the thresholds of the upper and lower comparators is normally 1.0 V, which is 1.33 times the voltage step caused by a charge dump. This means that in the worst case scenario, the chopper introduces an extra 143 ms between two charge dumps. This in turn means that the maximum error caused by the chopper is equal to 1.33 as well. To keep the error within 1% the inputs must be swapped no more than once every 133 periods, or once every 14.2 seconds. Waiting even longer further reduces the error. That being said, the error found here is the worst case scenario. Also, for higher inputs, the relative error may be the same but the time between dumps is much shorter, meaning the error is averaged out better. On the downside, the errors may be larger due to process spread (mostly in the capacitor). Bearing this in mind, chopping every 30 seconds should be good enough for overall accuracy.

The biggest drawback of this chopping scheme is the long testing time. A good idea would be to functionally test the system using a much shorter interval and then setting the chopping interval to 30 seconds for normal operation.

Finally, a verification if the result is accurate enough under mismatch is in order. For this verification, an offset of 6 mV at the input of the system is assumed, and the smallest specified input signal of 100 µV is applied. Since the offset is 60 times the input signal, the time between dump phases is roughly 60 times shorter than without offset, or 1.78 ms. This means that almost 17,000 charge dumps will occur between two chopping events. Since the input is 60 times smaller than the offset, at least 60

charge dumps must occur to be able to read a significant difference between the offset and the input signal. As can be seen, this is not an issue.

4.1.2 Incorporation Into the System

Incorporating the aforementioned chopping strategy into the system is quite trivial. Four switches are required at the input. These can be simple minimum sized switches, since they connect to a very high impedance. Since chopping events are relatively rare, errors due to charge injection and clock feedthrough are not considered either.

All that needs to be done now is to create a timer that has an output that toggles once every 30 seconds. Using a clock of 1.152 MHz, a 26 bit counter will overflow after 58.25 seconds. This means that the MSB of such a counter will toggle once every 29.13 seconds, making it very suitable for this system. Code fragment 4.1 shows the Verilog implementation of such a counter. Note that a test mode can be activated by commenting out the lines for the 60 second timer and removing the comment from the lines for the 1 second timer.

```
//Verilog HDL for "thesis", "chopper_control" "functional"


module chopper_control ( Q, Qn, clk, pwrp, pwrn );

  //Inputs & Outputs
  input clk;
  input pwrp;
  input pwrn;
  output Q;
  output Qn;

  //Internal Registers
  reg[25:0] counter;

  initial
  begin
    counter <= 0;
  end

  always @(posedge clk)
  begin
    counter <= counter + 1;
  end

  //60 second timer
  assign Q = counter[25];
  assign Qn = !counter[25];

  //1 second timer for testing
  //assign Q = counter[19];
  //assign Qn = !counter[19];

endmodule
```

Code Fragment 4.1: Verilog code for chopper control

4.2 Alternate Strategies

4.2.1 Chopping at Mid-Scale

A different chopping strategy has been devised that waits for the signal to reach mid-scale before chopping. Figure 4.2 displays this graphically. As can be seen, the time error is eliminated with this strategy. The only error that remains is apparently the mismatch between the two voltage ranges (lower comparator threshold to mid-scale and mid-scale to the upper comparator threshold).

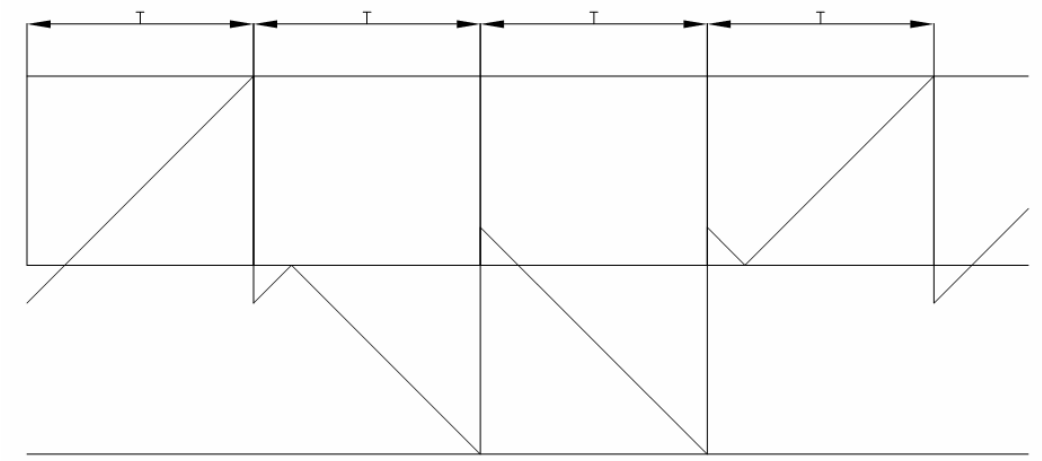

Figure 4.2: Removal of error by chopping at mid-scale

Unfortunately, this chopping scheme suffers from two major drawbacks. First off there is the added complexity. A comparator must be added to detect mid-scale crossings and quite an elaborate controller needs to be added as well, that does the following:

- Green-light a chopping event
- Wait for the mid-scale crossing detector to trigger
- Verify that no charge dump phase is in process

For the implementation of the mid scale crossing detector a single comparator is desired. Using two comparators and triggering when the signal is between the thresholds of these is an easier option, but would require two thresholds that are very close to each other. This introduces the risk that the comparators 'swap around' due to offsets. If one comparator is to be used, a memory has to be added that remembers the polarity of the signal. This can be done by detecting whether the last charge dump was a sink or a source operation. All in all a lot of extra control circuitry has to be added.

The second downfall of this chopping scheme is that the chopping events are no longer equally spaced in time. Figure 4.2 does not show this, but if a different input signal is applied, the tangent of the integrator output is different, and thus the amount of time it takes for the integrator output to reach mid-scale varies as well. Even with a constant input signal this effect occurs due to mismatch in the V-I converter. This means that there is an imbalance between the amount of time the system spends in both chopping modes, which introduces an error in time, not unlike the error caused by the strategy mentioned in chapter 4.1.

For these reasons, the chopping mechanism described here has not been implemented.

4.2.2 Swapping Integrator Capacitor Polarity

The last option that was considered was to swap the polarity of the integration capacitor when the inputs are toggled. Theoretically, this would void all errors, since the output voltage of the integrator would be mirrored around mid-scale and at the same time the direction of the integration would toggle. In practice however, a significant part of the capacitance consists of parasitics to the bottom plate (in fact

about 730 fF for a 4 pF capacitor as will be shown later). Part of this capacitance is connected to the input of the integrator, which is fixed at 900 mV and part is connected to the variable voltage at the output. Swapping the capacitor around will cause the part that used to be connected to the output to be charged to 900 mV. This charge has to come from (or flow on to) the integration capacitor, introducing an error. For this reason this scheme has not been applied. It may still be implemented at a later stage, since the error is smaller than when the inputs are simply swapped around, especially if the bottom plate parasitic can be reduced in some way. For this work however, the error reduction was not deemed to warrant the extra complexity.

5. System Level

5.1 Implementation

With all the building blocks complete, the top level of the system can be implemented and tested. Figure 5.1 shows the implementation of the whole system.

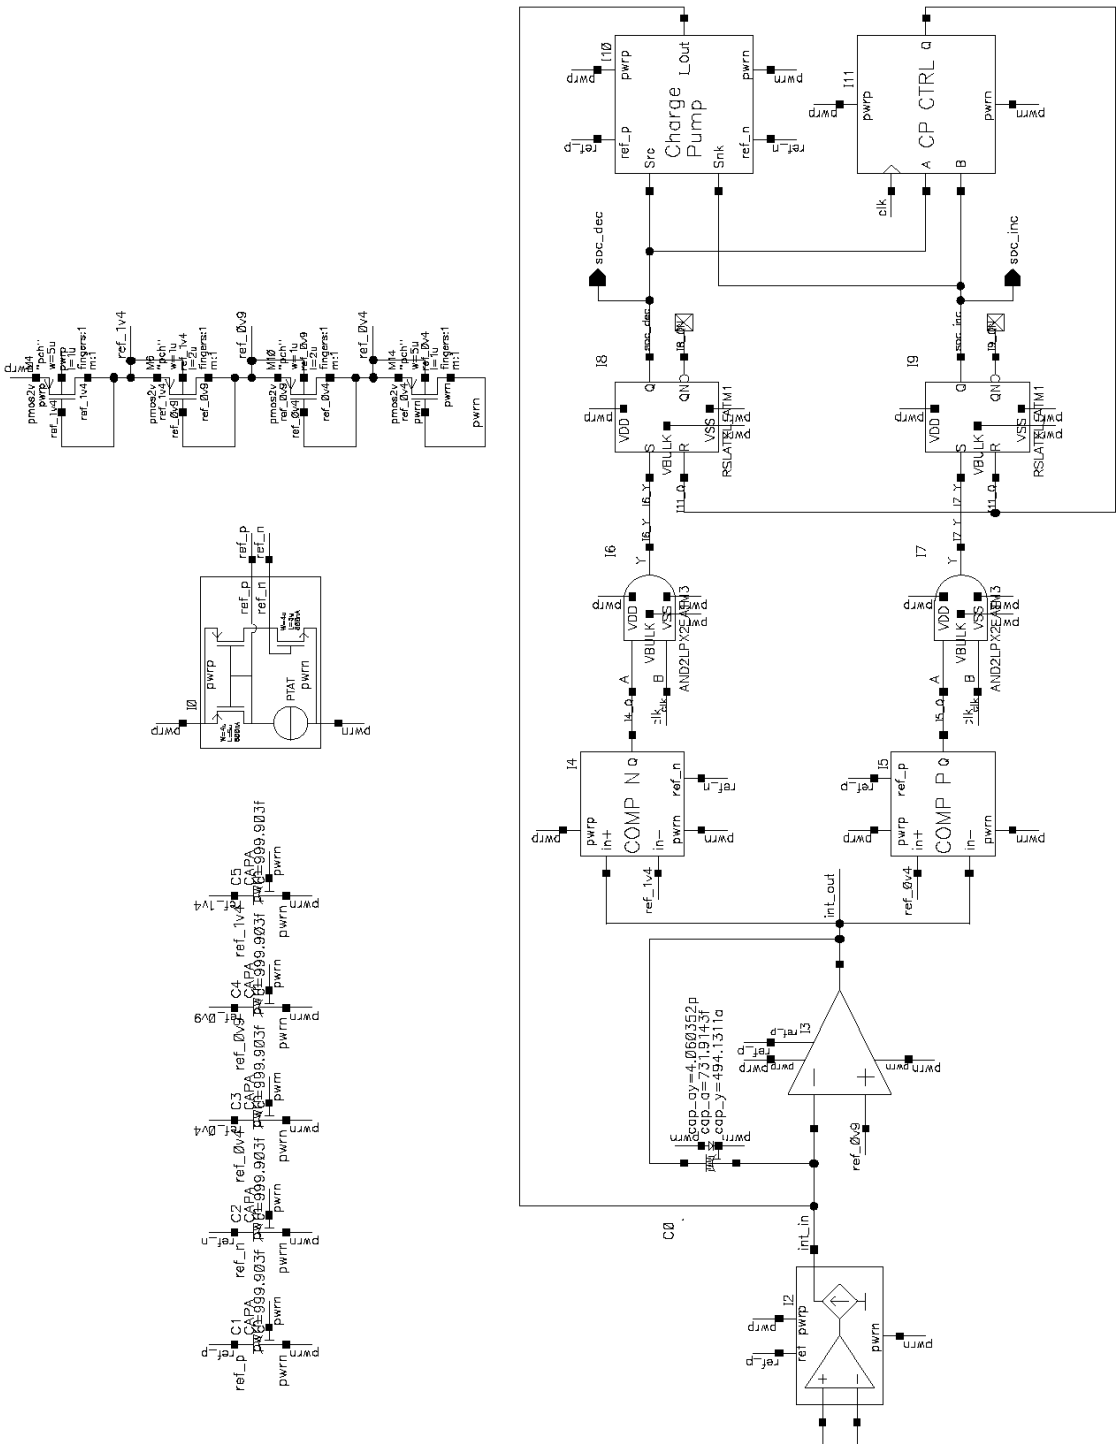

Figure 5.1: System top-level

A number of buffer capacitors can be seen, as well as a divider string to create the comparator reference voltages. These buffer capacitors were added when the switching comparators were still in place and were not removed after said comparators were replaced. Removal should be safe, but simulations will have to confirm this. Also visible is the PTAT reference generator that creates the reference voltages for the current sources in the system. The rest of the blocks are, from left to right, top to bottom:

- The V-I converter
- The integrator
- The NMOS input and PMOS input comparators
- The AND gates to synchronize the comparators with the clock
- The SR latches that control the charge pump
- The charge pump
- The charge pump control block

The soc_inc and soc_dec outputs can be used to update the counter that keeps track of the actual state of charge. The effect of these outputs depends on the chopper. If the input signal is connected normally, the soc_inc signal increases the counter and the soc_dec signal decreases it. If the input signal is inverted, the meaning of these signals toggles as well.
The chopping mechanism is depicted in figure 5.2. It was left out of figure 5.1 for readability reasons. The chop_act output can be used to determine the meaning of the aforementioned soc_inc and soc_dec ports.

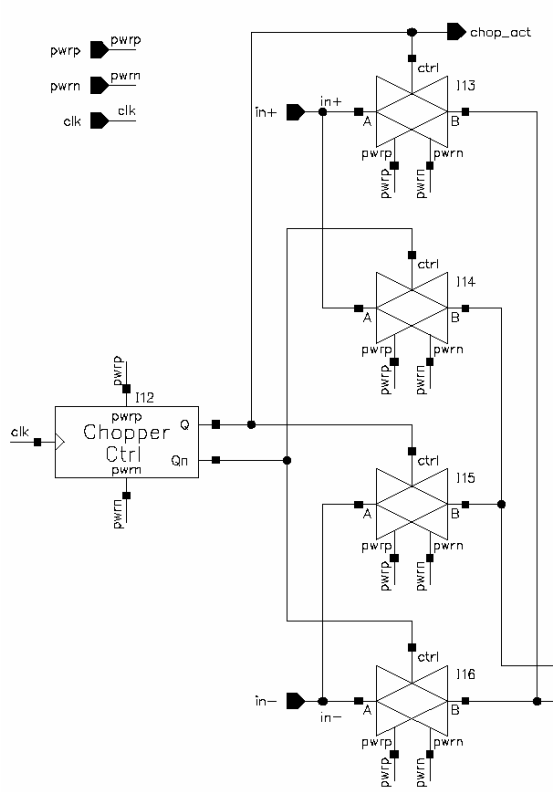

Figure 5.2: Chopper

5.2 Simulations

5.2.1 System Level Simulations Without Chopping

The first tests verify the proper functionality of the system with the chopper disabled. To do this, the tests are performed with the chopper controller set to a 30 second interval and the transients are performed over less than 30 seconds. Figures 5.3 through 5.6 show the integrator output voltage and the output ports of the system for inputs of -100 mV, -100 µV, 100 µV and 100 mV respectively. As can be seen, the system is functioning as expected. A certain amount of charge is slowly built up on the integrator until a threshold is reached. At that point a charge package is removed and the process repeats itself.

In order to quantitatively evaluate the system, a number of simulation runs at different input voltages have been performed. In these runs, the duration of 10 complete cycles has been measured and the number of pulses per second calculated from that duration. The same procedure was then repeated, with the input signals swapped around. The latter result was interpreted as if the chop_act output was high (hence the meaning of soc_inc and soc_dec was swapped) and once again the number of pulses per second was found. The two results were added, yielding the number of pulses in two seconds, assuming an error-less chopping mechanism.

These runs have been performed at input offsets of 0 mV, -6 mV and +6 mV. Tests have also been performed under maximum input offset with a charge pump mismatch of +1.5 % and –1.5 %. A test covering a set of temperatures has also been carried out. Finally tests have been performed to verify the voltage swing at different capacitor corners. As mentioned before, the initial tests were carried out using a longer dump phase. The last test is a zero-offset test, that has been performed to verify that the system still functions properly after redefining the dump phase duration.

The result of the first test, which reflects the response of the ideal system to a range of input voltages can be found on page 94 in table 5.1. Note: in this table, PPS means pulses per second.

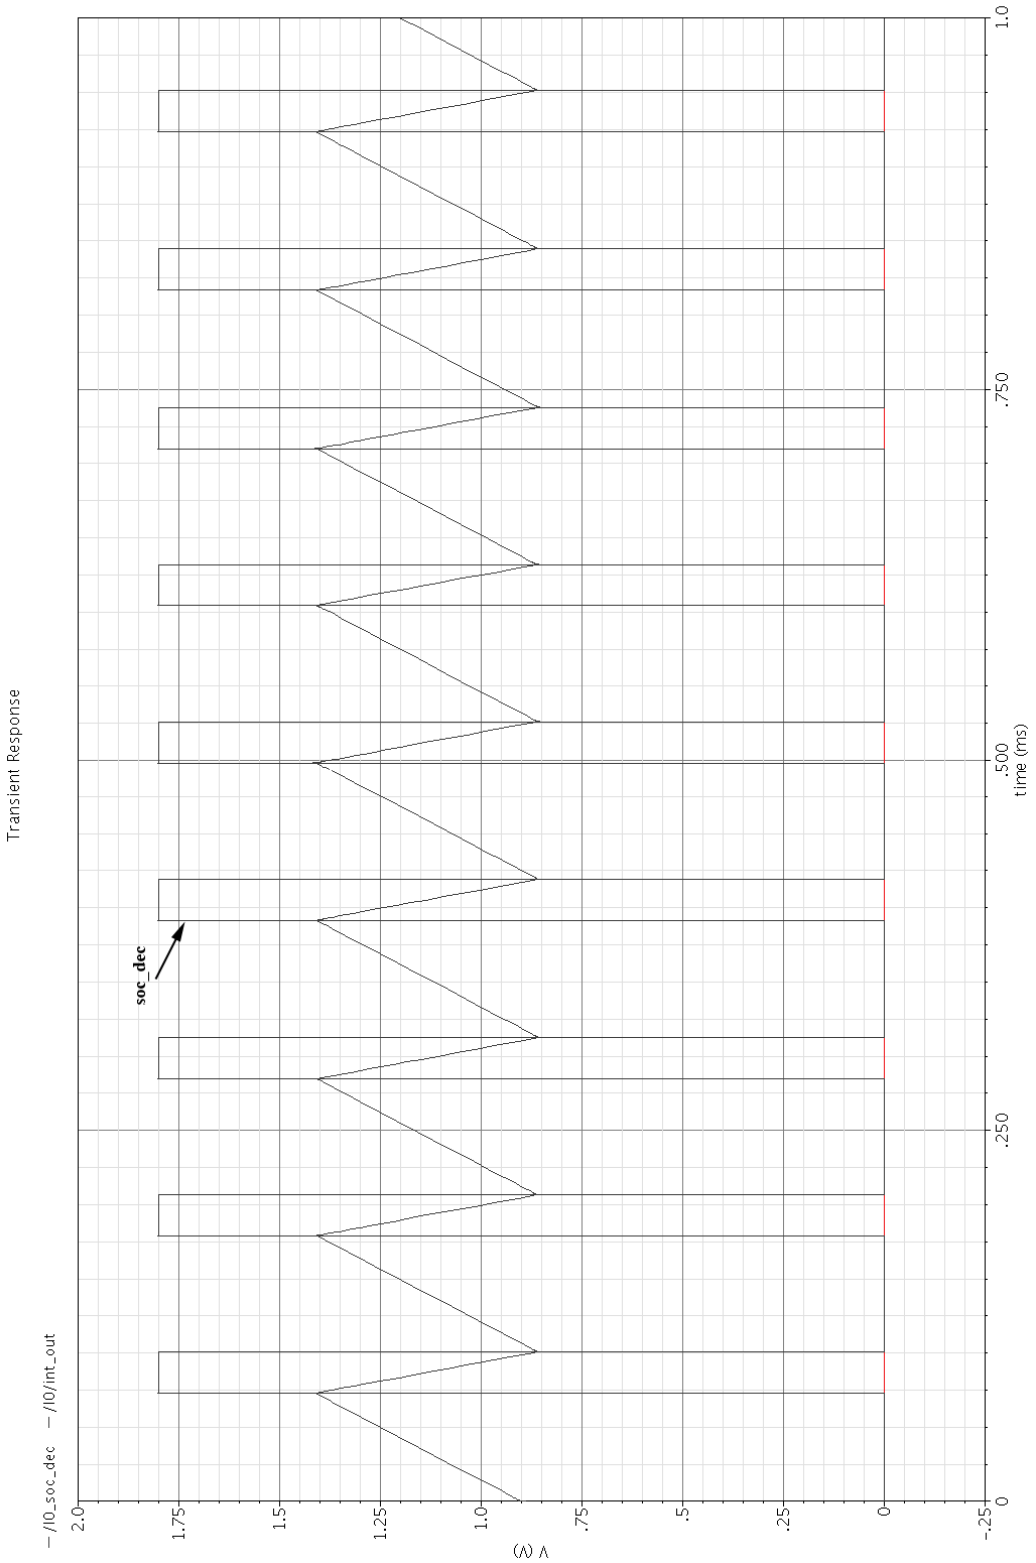

Figure 5.3: Transient simulation; -100 mV input

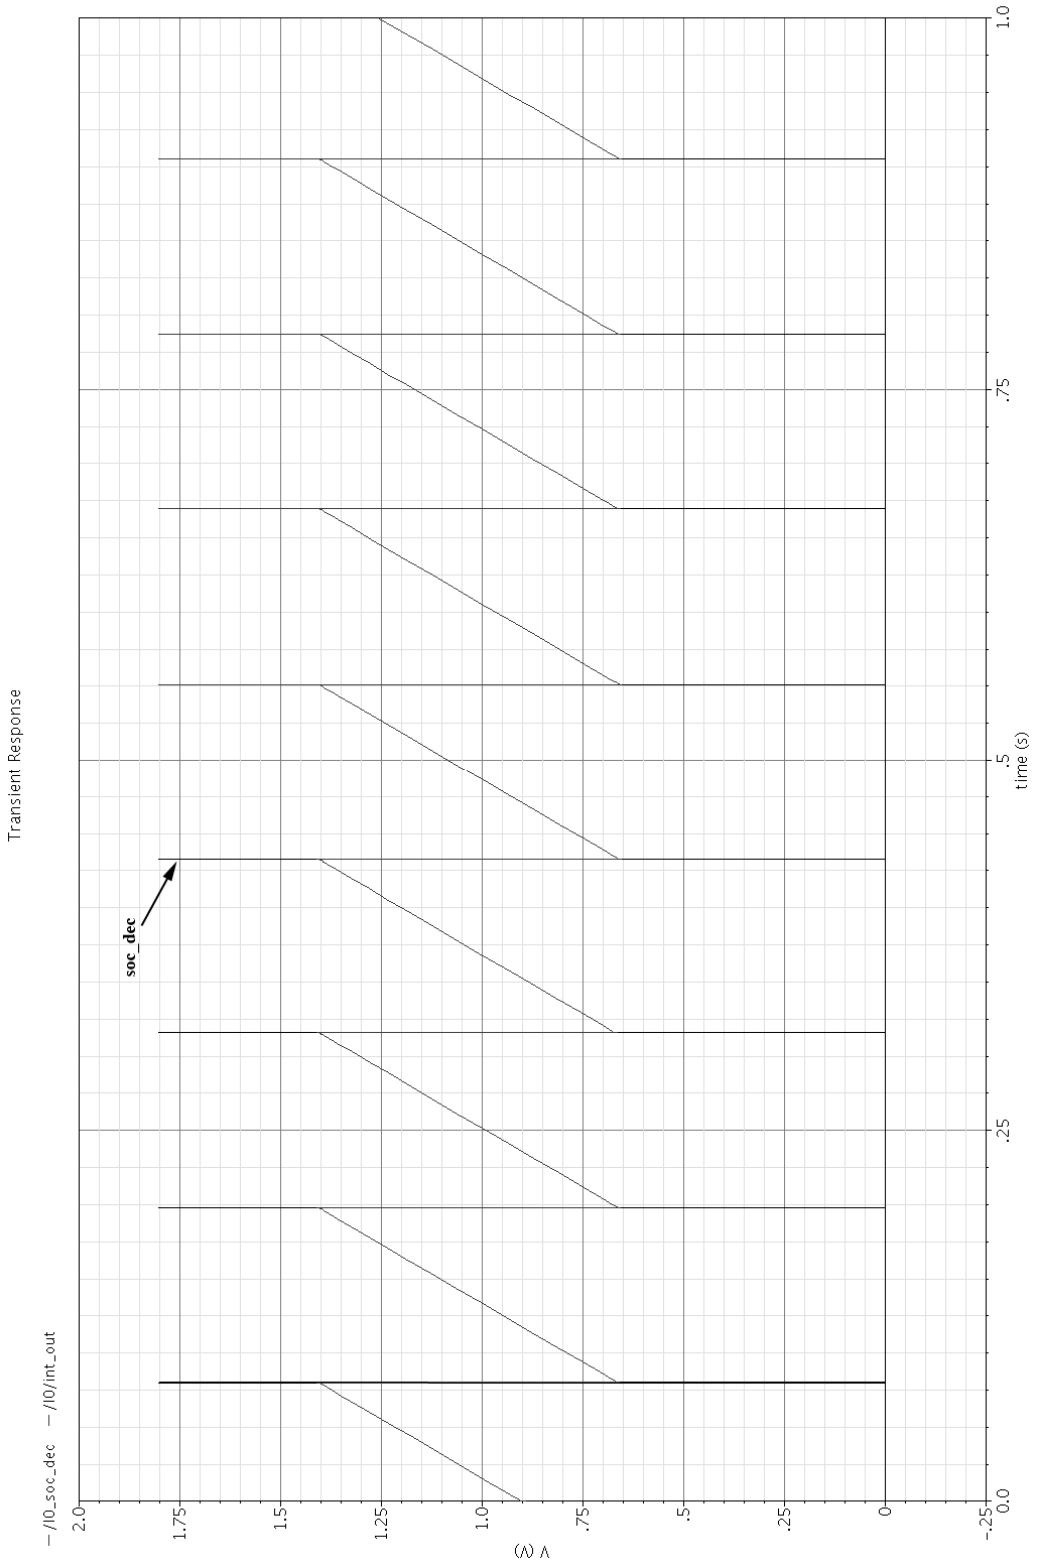

Figure 5.4: Transient simulation; -100 µV input

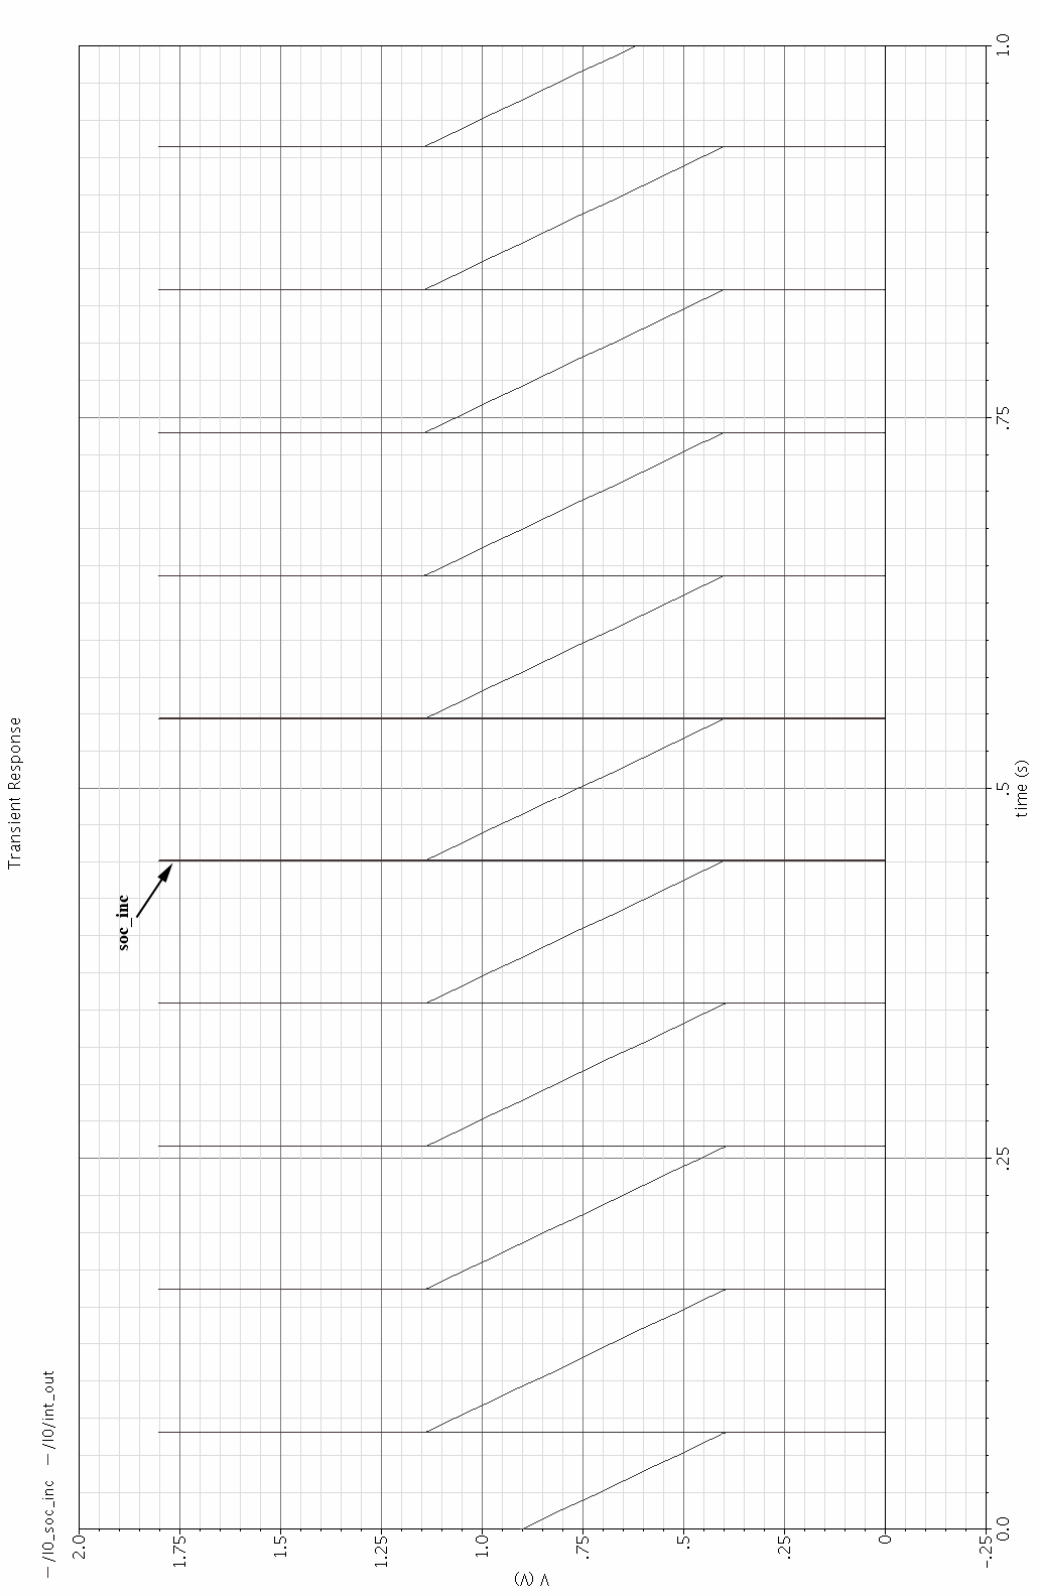

Figure 5.5: Transient simulation; +100 µV input

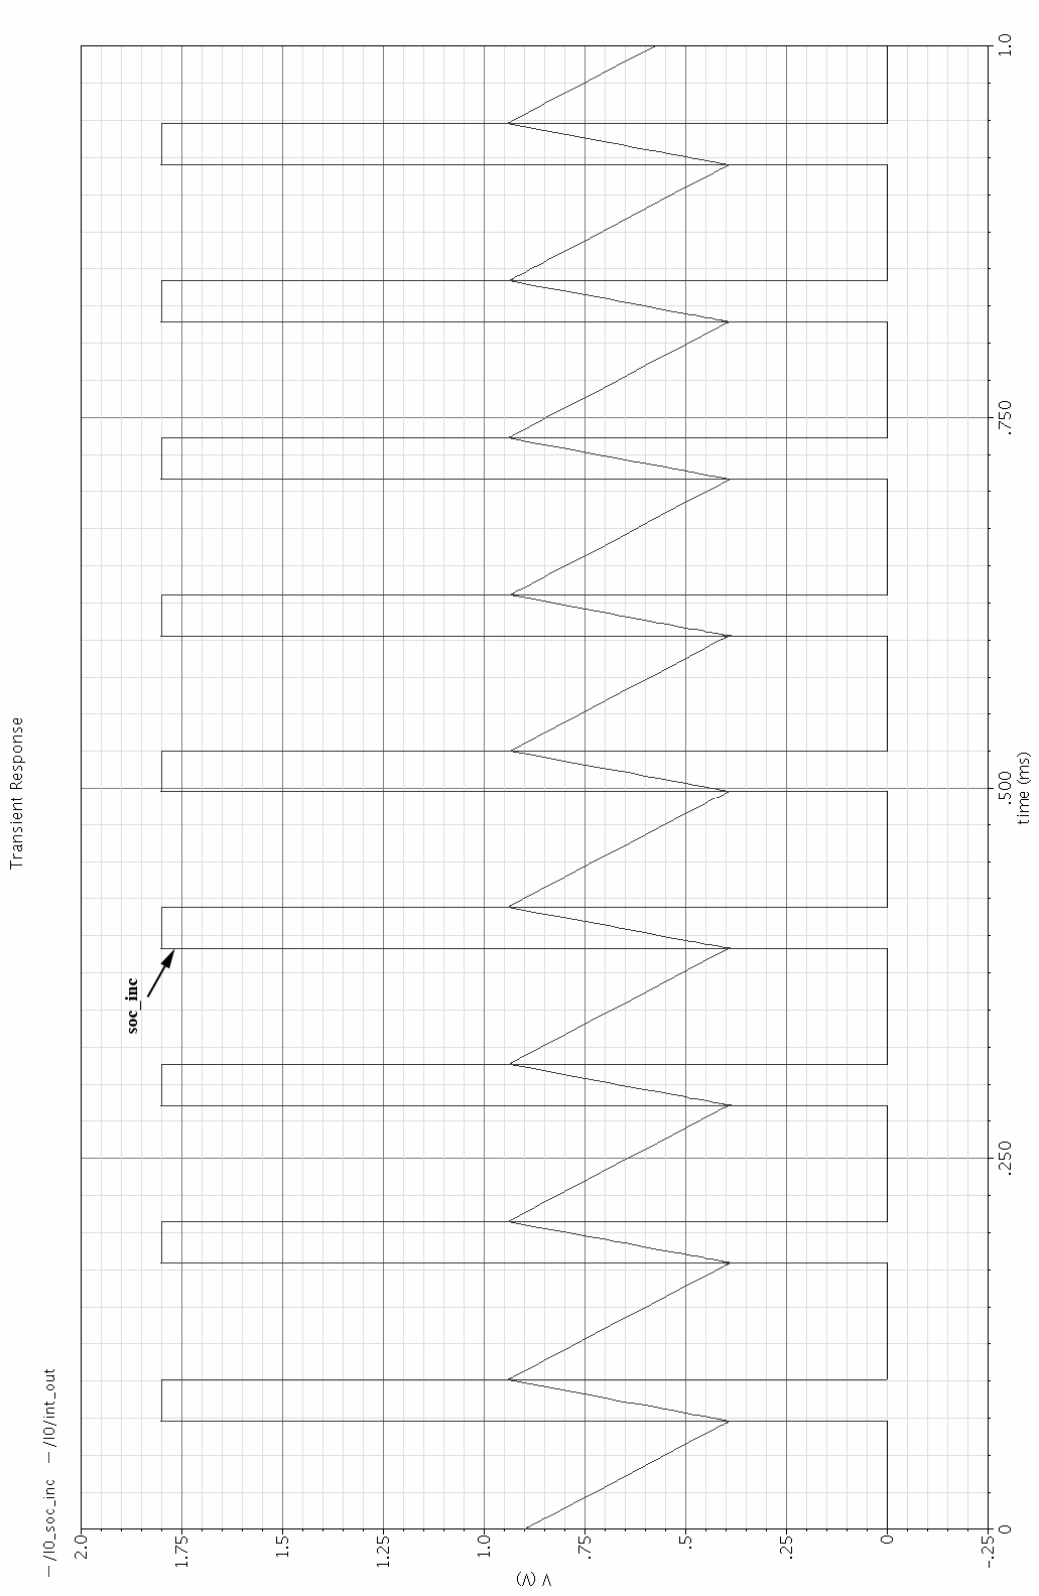

Figure 5.6: Transient simulation; +100 mV input

V_{In}:	PPS (chop_act=0):	PPS (chop_act=1):	Total (2 seconds):
-100 mV	-8076.1	-8040.0	-16116.1
-10 mV	-806.76	-809.02	-1615.8
-1 mV	-79.497	-81.643	-161.14
-100 µV	-7.2138	-8.9138	-16.128
100 µV	8.9168	7.2156	16.132
1 mV	81.549	79.493	161.04
10 mV	808.64	805.90	1614.5
100 mV	8046.3	8046.3	16092.6

Table 5.1: Pulses per second under ideal conditions

As can be seen, the system performs well. The largest non-linearity, between 100 mV and 100 µV, is 0.24%. These tests have been repeated using an input offset of +6 mV (table 5.2) and -6 mV (table 5.3)

V_{In}:	PPS (chop_act=0):	PPS (chop_act=1):	Total (2 seconds):
-100 mV	-7574.0	-8539.7	-16113.7
-10 mV	-321.98	-1291.2	-1613.2
-1 mV	402.74	-565.04	-162.30
-100 µV	476.32	-492.38	-16.06
100 µV	492.38	-476.39	15.99
1 mV	565.90	-404.47	161.43
10 mV	1291.2	321.15	1612.4
100 mV	8530.2	7590.7	16120.9

Table 5.2: Pulses per second, +6 mV offset

V_{In}:	PPS (chop_act=0):	PPS (chop_act=1):	Total (2 seconds):
-100 mV	-8539.0	-7577.0	-16116.0
-10 mV	-1287.0	-322.90	-1609.90
-1 mV	-562.80	401.60	-161.20
-100 µV	-490.49	474.23	-16.26
100 µV	474.58	490.70	16.12
1 mV	-401.60	563.50	161.90
10 mV	322.80	1287.0	1609.80
100 mV	7578.0	8555.0	16133.0

Table 5.3: Pulses per second, -6 mV offset

As can be seen, the largest error is slightly worse now. The worst outcome is observed in table 5.2 at an input of 100 µV. This value shows a non-linearity of 1.5% compared to the value obtained at an input of -1 mV. However, this is the only value that falls outside the tolerance and obtaining correct values for the smallest voltages already required high timing accuracy and very strict absolute and relative tolerance levels in the simulator. The following tests have been performed to verify whether this one extreme value was caused by the system, or by insufficient simulator precision:

- The output current of the V-I converter was measured under offset (chapter 3.1). The average current was found the be highly linear.

- The current from the V-I converter completely flows on to the integration capacitor. No leakage into the charge pump was found.
- The results do not change if the integration capacitor is replaced by an ideal capacitor.
- No change is observed after also replacing the integrator OPAMP with an ideal version with a gain of 120 dB.
- With these measures in place, the charge pump has been replaced by a set of voltage controlled current sources and the chopper was removed. Still no change in the behavior was observed.

From these steps, along with the fact that only one measurement is outside of the tolerance, the conclusion was drawn that the accuracy of the simulator is the limiting factor and that the accuracy of the system is sufficient. Note that, when ignoring these smallest values, the worst accuracy is expected between the |1 mV| and |100 mV| inputs. The worst accuracy observed between those is 0.7% (6 mV offset, linearity error between -100 µV and -100 mV)

The next step is to add an imbalance in the charge pump. To do this, the width of one of the PMOSTS was first increases by 1.5% and then decreased by 1.5%. Both conditions were simulated with input voltages of +10 mV and -10 mV, with input offsets of +6 mV and -6 mV. The results can be found in table 5.4.

V_{In}:	**PPS (chop_act=0):**	**PPS (chop_act=1):**	**Total (2 seconds):**
I_{Src} = 1.015 · I_{Snk}, Input offset = +6 mV			
-10 mV	-317.0	-1293	-1610.0
10 mV	1291	316.9	1607.9
I_{Src} = 1.015 · I_{Snk}, Input offset = -6 mV			
-10 mV	-1272	-322.7	-1594.7
10 mV	322.5	1271	1593.5
I_{Src} =0.985 · I_{Snk}, Input offset = +6 mV			
-10 mV	-326.8	-1293	-1619.8
10 mV	1289	326.3	1615.3
I_{Src} =0.985 · I_{Snk}, Input offset = -6 mV			
-10 mV	-1308	-322.5	-1630.5
10 mV	323.1	1311	1634.1

Table 5.4: Pulses per second, charge pump mismatch

The errors between the positive and negative inputs in these measurements are:
I_{Src} = 1.015 · I_{Snk}, Input offset = +6 mV: 0.13%
I_{Src} = 1.015 · I_{Snk}, Input offset = -6 mV: 0.08%
I_{Src} = 0.985 · I_{Snk}, Input offset = +6 mV: 0.28%
I_{Src} = 0.985 · I_{Snk}, Input offset = -6 mV: 0.22%

These are all well within tolerance. Mismatch in the charge pump, even combined with mismatch in the input stage, will thus be properly removed by the chopper.

The temperature test is the next test to be performed. Once again an input of +10 mV and -10 mV was used. No chopping was applied, since no offset was added and the only quantity of interest is the difference between the results at various temperatures. Six different temperatures were simulated, ranging from -40 °C to +90°C. The results,

i.e. pulses per second, can be found in table 5.5. As expected, the differences are huge at this point (roughly ± 16 %) and are mostly caused by the charge pump. Implementing an NTAT reference for this block should solve the issue as mentioned in section 3.4. The difference between positive and negative voltages is due to a shift in the (small) offset current of the input stage. This effect will be canceled by chopping.

V_{In}:	-40 °C	-20 °C	0 °C	25 °C	60 °C	90 °C
-10 mV	-940.0	-895.9	-849.6	-795.2	-733.1	-682.7
10 mV	956.9	902.3	855.6	798.7	729.1	678.3

Table 5.5: Pulses per second, temperature sweep

The final test has been performed after re-dimensioning the integration capacitor. Tests have been performed to verify the voltage step caused by the charge dump with minimum capacitor size and input voltage, as well as maximum capacitor size and input voltage. These tests cover the extreme values that can occur here.
With an input voltage of 100 µV and a capacitance of 3.4 pF, a voltage difference of 882 mV is expected. At 100 mV and 4.6 pF the expected value is 483 mV. The results in table 5.6 are both within 1% of the theoretical values. This means that a margin of 125 mV will always exist between the voltage step and the difference between the comparator thresholds. The final step is to recreate table 5.1 to verify that the system is still linear with this new charge dump duration. The results can be found in table 5.7 and they do not give any cause for concern on this front.

V_{In}:	C (pF):	ΔV_{Int} (mV):
100 µV	3.4	875
100 mV	4.6	479

Table 5.6: Voltage step over integrator capacitor

V_{In}:	PPS (chop_act=0):	PPS (chop_act=1):	Total (2 seconds):
-100 mV	-9428	-9428	-18856
-10 mV	-941.9	-942.5	-1884.4
-1 mV	-93.47	-95.71	-189.18
-100 µV	-8.466	-10.39	18.856
100 µV	10.38	8.464	18.844
1 mV	95.55	93.39	188.94
10 mV	944.5	942.8	1887.3
100 mV	9454	9439	18893

Table 5.7: Pulses per second under ideal conditions, new charge dump duration

5.2.2 System Level Simulations With Chopping

As mentioned earlier, the testing the chopper with its normal interval will take a long time. This applied for an actual circuit, but also for the simulations. In fact, the amount of data gathered during a simulation run on a transient of one minute (to include both chopper states) with the required accuracy would generate too much data to be processed in a reasonable amount of time. Therefore, the only thing that was tested was the functionality of the chopper, by setting the interval to 455 ms.

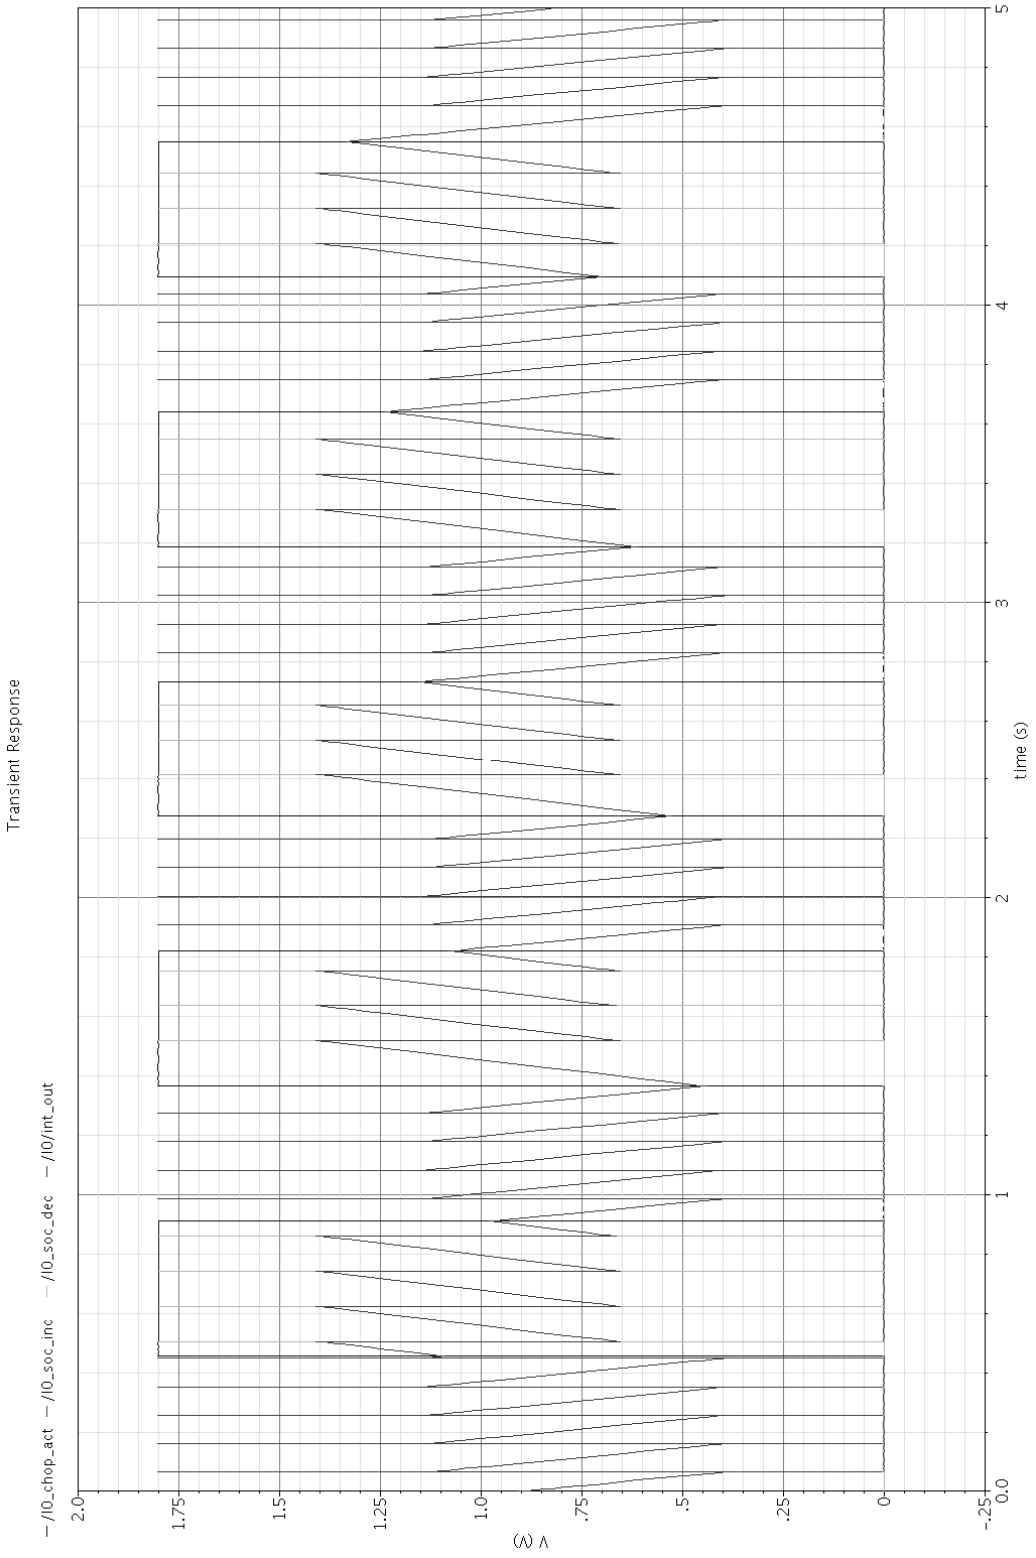

Figure 5.7: Transient simulation; +100 µV input, chopper active

Figure 5.7 confirms that the chopper works as intended. When the chop_act signal goes high, the inputs toggle and as a result the direction of the integrator output voltages toggles as well.

With the system now complete, the total current consumption can be measured. This does not include the digital blocks since they are simulated using behavioral descriptions, but due to the low clock speed no high current consumption is expected there. The result can be found in figure 5.8. An average supply current of 8.0 µA is found, which is well within limits.

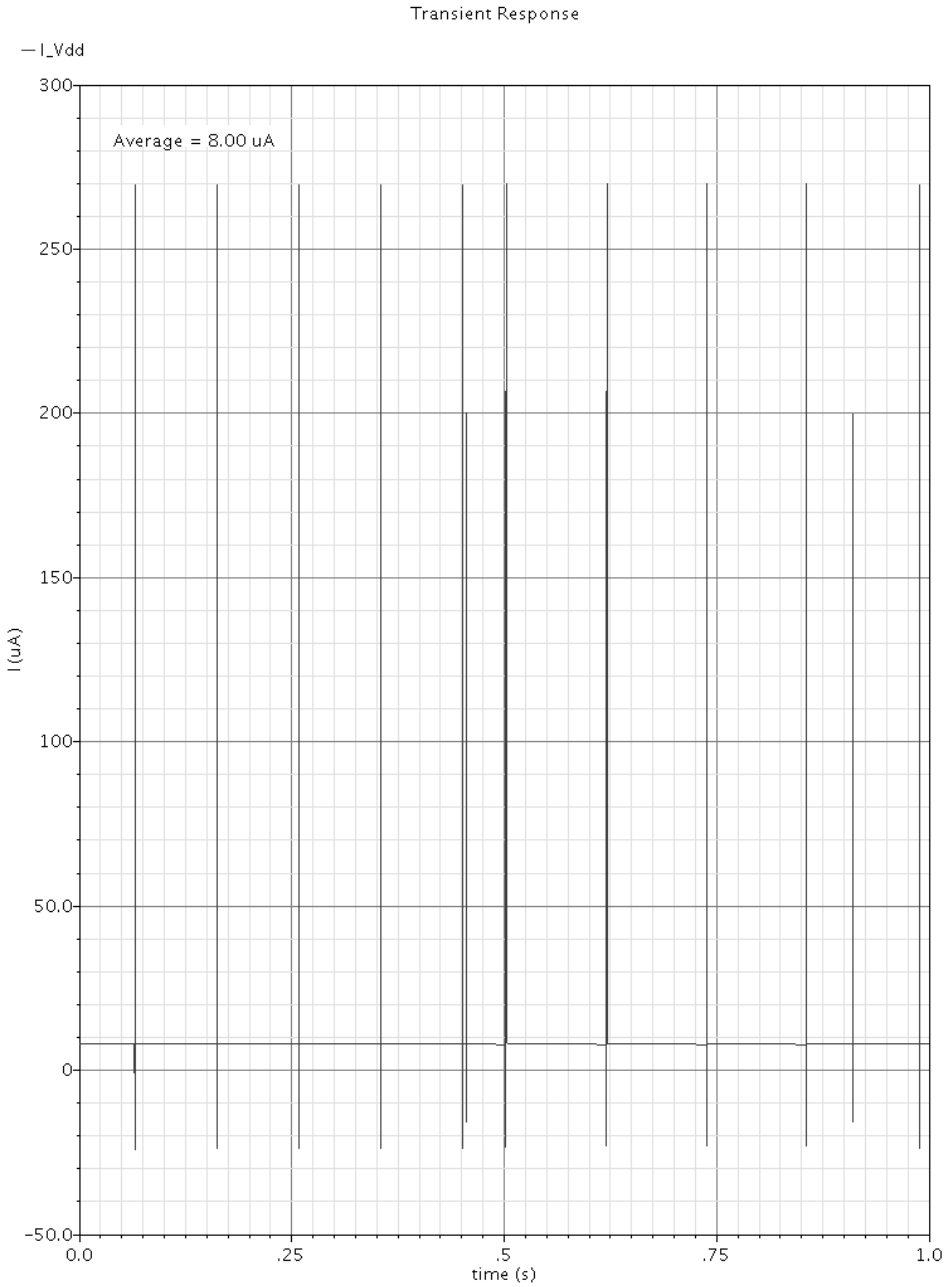

Figure 5.8: Supply current consumption

6. Conclusions and Recommendations

6.1 Conclusions

The functionality of the system is in confirmation with the theories upon which it is based. The figures in chapter 5 show that the integrator voltage varies linearly in time, with a dV/dt that varies with the input voltage, until either threshold is reached. At this point an amount of charge is quickly added to or removed from the integrator and the process repeats itself.

The linearity is well within limits if the components exhibit perfect matching (a non-linearity error of 0.24% was found). When mismatch in the V-I converter is also simulated, the accuracy drops somewhat. The worst case mismatch found in this case that was not caused by accuracy limitations in the simulator was 0.7%, which is still well within the 1% limit.
The accuracy is maintained if mismatch is added to the charge pump and when the maximum voltage swing at the integrator output is reduced by altering the duration of the charge dump phase. The voltage swing remains within the boundaries set by the comparator thresholds if the size of the integration capacitor changes due to process spread, guaranteeing proper functionality under all circumstances.
The only linearity issue at present is the accuracy of the system under temperature variations. A total error of approximately ± 16 % is achieved over a temperature range of -40 °C … 90 °C. This error is mostly caused by the temperature dependence of the charge pump (and to a lesser degree by the V-I converter). Finally, the current consumption remains well within the 100 µA limit, averaging at 8.0 µA.

6.2 Recommendations

As mentioned before, the temperature stability of the system is not sufficient as of yet, and this is mostly due to the fact that the charge pump has not been properly compensated for temperature variations. In fact, by using the same PTAT reference that is used by the V-I converter, a PTAT behavior is found in the charge pump. The recommended solution would be to add an NTAT source to bias the charge pump that would give it the same relative current change over temperature that is found in the V-I converter.

A second recommendation is to reconsider the chopping mechanism. Several possible implementations have been treated, and especially the mechanism where the polarity of the integration capacitor is swapped is promising. If a measure can be found that alleviates the influence of the parasitic capacitance to ground, or a capacitor structure is found that has a lower parasitic, this measure can serve to significantly reduce the time between chopping events.

Finally, a layout of the system needs to be made to verify that functionality is maintained when layout parasitics are added. This would also pave the road to an actual tape out and the possibility to test the physical system in practice.

7. Literature

[1] H.J. Bergveld: *Battery Management Systems, Design by Modeling*
Royal Philips Electronics, 2001
ISBN 90-74445-51-9

[2] MAXIM: *Frequency Undersampling in Coulomb Counting*
December 2002
http://pdfserv.maxim-ic.com/en/an/AN1798.pdf

[3] Behzad Razavi: *Design of Analog CMOS Integrated Circuits*
McGraw Hill, 2001
ISBN 0-07-118839-8

[4] Afshin Haft-Baradaran: *Basic and Advanced Current References*
University of Toronto
www.eecg.toronto.edu/~kphang/papers/2001/haft_ref.pdf

[5] M.J.M. Pelgrom: *Matching Properties of MOS Transistors*
IEEE Journal of Solid-State Circuits, vol. 24, no. 5, October 1989

Appendix A: Assignment Description

Integrated 'State of Charge' circuit for rechargeable batteries.

Most wireless 'Portable Part' applications make use of rechargeable batteries as a source of power. To be able to develop an accurate battery management algorithm it is essential to 'know' or keep track of the amount of charge present in the battery cell at any point in time.

The graduation project will consist of selecting or inventing the best principle to keep track of the amount charge that is flowing in and out of the battery and implementing this principle into a practical CMOS circuit. History has shown that this quite a challenging task to perform. A software algorithm must be able to readout the 'State of Charge' from a register (e.g. in the form of a counter value).

Characteristics:	- standard CMOS mixed-signal design.	
	- evaluation of silicon design.	
Specifications:	- current range	-1…+1A.
	- accuracy	±1% over 60 dB dynamic range (1mA – 1A)
	- current consumption	<100uA

Appendix B: Component Parameter Estimations

A number of simulations have been performed to estimate the model parameters of several devices available in the tsmc018rf library, which was used for this project.

B.1 MOSFET Threshold Voltages and K-factors

The most important parameters are the threshold voltages and K-factors of the N-MOST and P-MOST devices. These have been estimated by sweeping the gate-source voltage of both device types while keeping the drain-source voltage constant. Long devices were used to avoid errors caused by channel length modulation. The estimations have been done by plotting the square root of the drain current, yielding the relationship depicted in formula B.1.

$$I_D = \frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L} \cdot (V_{GS} - V_{TH})^2$$
$$\sqrt{I_D} = \sqrt{\frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L}} \cdot (V_{GS} - V_{TH}) \qquad \textbf{B.1}$$

This means that, as long as the quadratic relationship holds, the plot will be linear. By adding a tangent line to the part of the curve belonging to strong inversion, and extrapolating this to 0 A, one can thus find the threshold voltage. The K-factor can be found by differentiating B.1.

$$\frac{d\sqrt{I_D}}{dV_{GS}} = \sqrt{\frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L}} \qquad \textbf{B.2}$$

Again, this is only true where B.1 is valid, i.e. the strong inversion region. It is thus possible to find the K-factor by differentiation of the tangent line that was used to find the threshold voltage.

The following parameters were used for both N-MOST and P-MOST devices:

- $|V_{DS}|$ 1.8 V
- $|V_{GS}|$ 0 V … 1.0 V
- L 10 µm
- W 10 µm
- W/L 1

Figure B.1 shows the result for the N-MOST device and figure B.2 shows the result for the P-MOST device. Curves created by using the parameters found in the estimation where also added for verification.

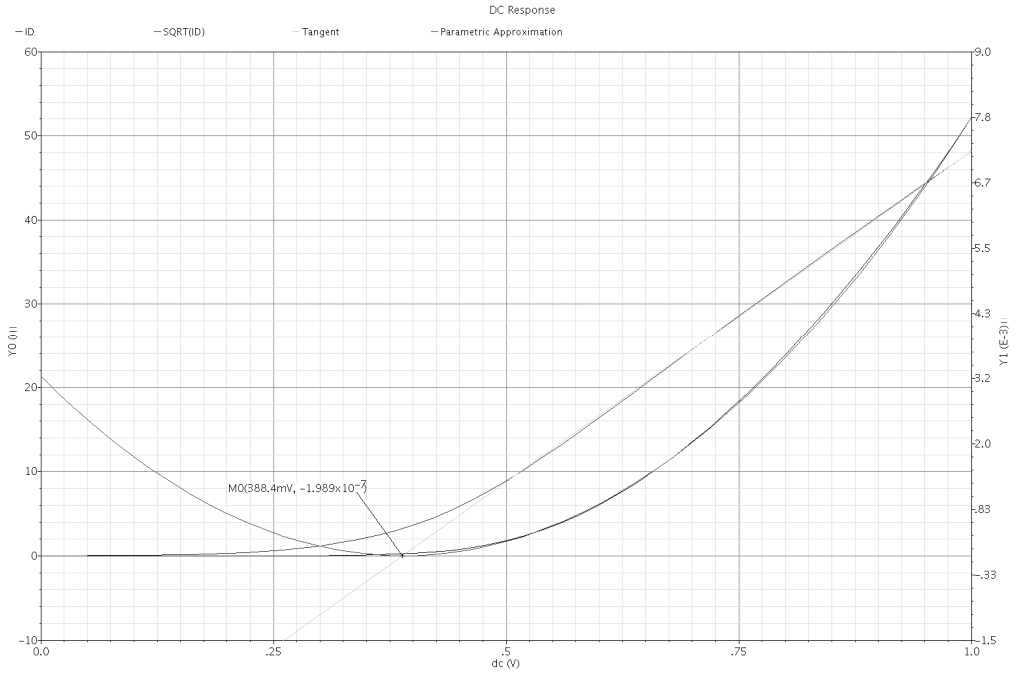

Figure B.1: N-MOST parameter estimation

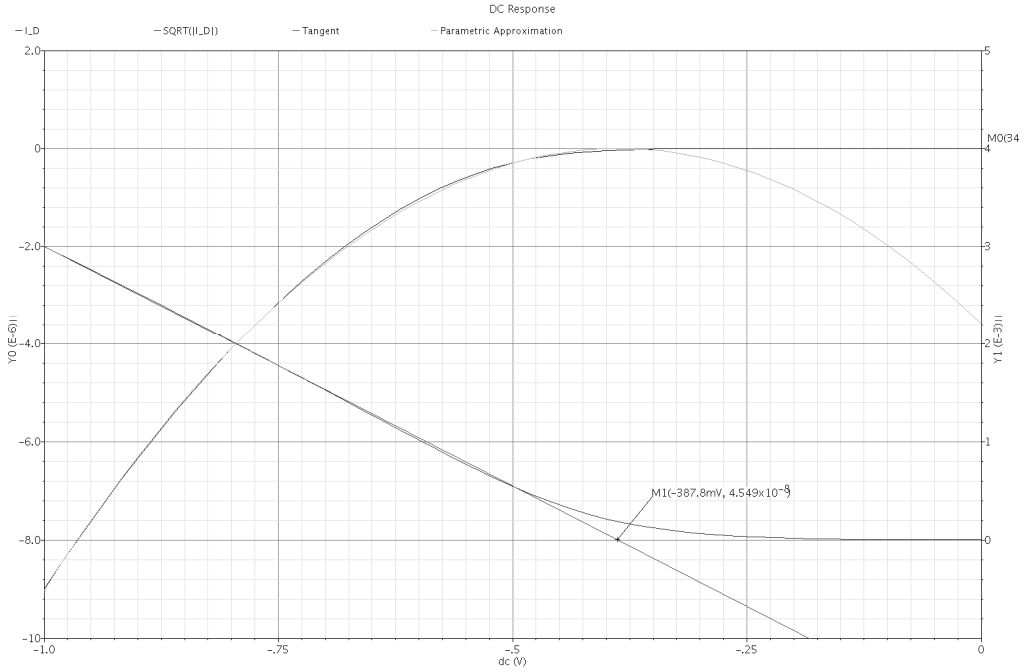

Figure B.2: P-MOST parameter estimation

For the N-MOST:

$$\frac{d\sqrt{I_D}}{dV_{GS}} \approx 11.8 \cdot 10^{-3} \frac{\sqrt{A}}{V} \Rightarrow \mu_n \cdot C_{OX} \approx 0.28 \frac{mA}{V^2} \qquad \textbf{B.3}$$

$$V_{GS} \approx 390\ mV$$

For the P-MOST:

$$\frac{d\sqrt{I_D}}{dV_{GS}} \approx 4.90 \cdot 10^{-3} \frac{\sqrt{A}}{V} \Rightarrow \mu_n \cdot C_{OX} \approx 0.048 \frac{mA}{V^2} \qquad \textbf{B.4}$$
$$V_{GS} \approx -390\ mV$$

As can be seen in the figures, a good match between the square root model and the simulated curves is achieved.

B.2 Lambda Estimation

An estimate of how the MOSFET devices behave when their drain source voltage changes has also been made. This has been done by sweeping the drain-source voltage of both N- and P-MOSTS of different lengths, while keeping their gate-source voltage constant, and maintaining saturation. By taking two points on the resulting curves, an estimation for λ can be made. Note that since the square root model was made under the assumption that λ is small enough to be neglected for 10 µm devices, the value that will be found here is not the true value. It will however give an impression of the shape of the curve, which can be useful to choose device lengths where large output resistances are desirable. In B.5, a formula is derived that is used to find λ from two points along the output curve.

$$I_{D,A} = \frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L} \cdot (V_{GS} - V_{TH})^2 \cdot (1 + \lambda \cdot V_{DS,A})$$
$$I_{D,B} = \frac{1}{2} \cdot \mu \cdot C_{OX} \cdot \frac{W}{L} \cdot (V_{GS} - V_{TH})^2 \cdot (1 + \lambda \cdot V_{DS,B})$$
$$\frac{I_{D,A}}{I_{D,B}} = \frac{(1 + \lambda \cdot V_{DS,A})}{(1 + \lambda \cdot V_{DS,B})} \qquad \textbf{B.5}$$
$$\lambda = \frac{I_{D,B} - I_{D,A}}{I_{D,A} \cdot V_{DS,B} - I_{D,B} \cdot V_{DS,A}}$$

For the estimation, the following parameters were used:

- $|V_{GS}|$ 1.0 V
- $|V_{DS,A}|$ 750 mV
- $|V_{DS,B}|$ 1.8 V
- W 1 µm

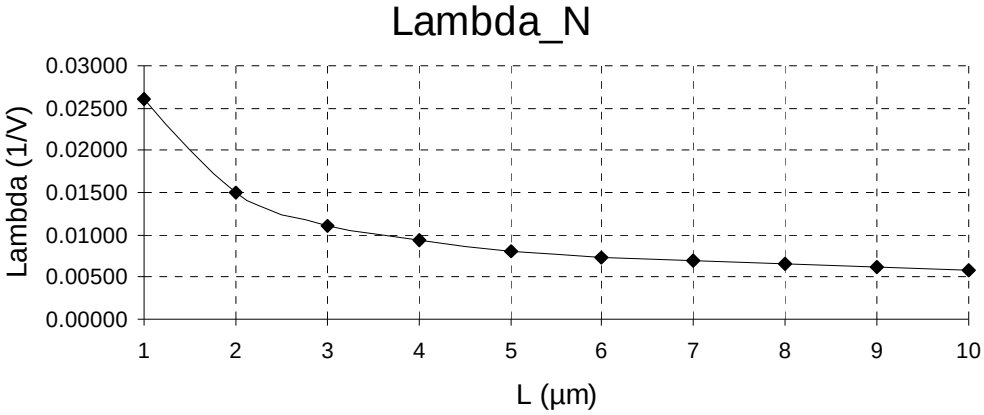

Figure B.3: λ estimation for N-MOST devices

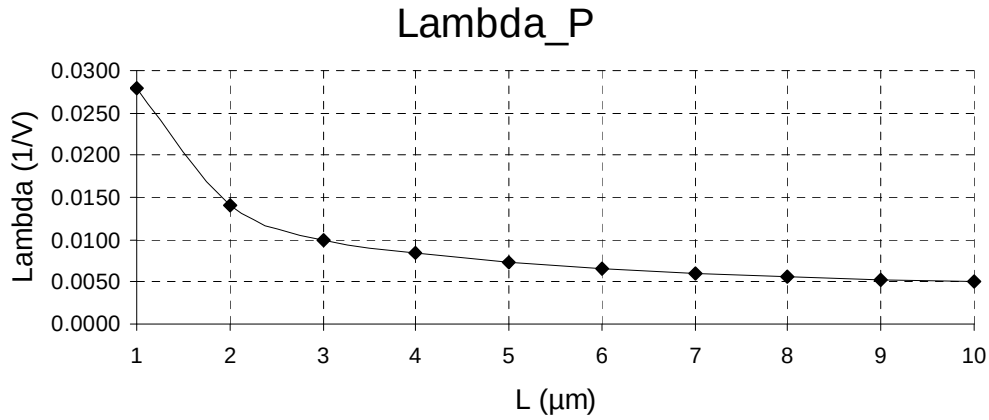

Figure B.4: λ estimation for P-MOST devices

Both figures show a very low λ for 10 µm devices as expected. Theoretically these values should even be zero, but some deviation between the quadratic model and the very high order models used in Cadence is not surprising. What is interesting to see is that the curve more or less flattens out for lengths above 5 µm.

B.3 Matching Parameters

Two important MOSFET matching parameters have been determined through Monte Carlo simulations (100 runs), namely the threshold mismatch (A_{VTH}) and K-factor mismatch (A_K). These factors have been found by performing operating point simulations on MOSFET diodes biased at 200 nA and looking at the standard deviation (σ) for the VTH0 and BETAEFF parameters respectively. The following formulas have been used to find the factors from these parameters [5]:

$$\sigma_{VTH} = \frac{A_{VTH}}{\sqrt{W \cdot L}}$$
$$\frac{\sigma_K}{K} \approx \frac{A_K}{\sqrt{W \cdot L}} \qquad \textbf{B.5}$$

Note that the formula for K-factor mismatch is given incorrectly in [3]. It was copied from [5] but Pelgrom uses K as a factor in the substrate dependant part of the threshold voltage, and uses β for the transconductance parameters.

Threshold mismatch:

W = 2 µm	**NMOST**		**PMOST**	
L (µm):	**σ (mV):**	**A_{VTH}:**	**σ (mV):**	**A_{VTH}:**
0.18	6.08	$3.65 \cdot 10^{-3}$	7.82	$4.69 \cdot 10^{-3}$
1	2.44	$3.45 \cdot 10^{-3}$	3.15	$4.45 \cdot 10^{-3}$
2	1.71	$3.42 \cdot 10^{-3}$	2.22	$4.44 \cdot 10^{-3}$
5	1.08	$3.42 \cdot 10^{-3}$	1.40	$4.43 \cdot 10^{-3}$
10	0.764	$3.42 \cdot 10^{-3}$	0.987	$4.41 \cdot 10^{-3}$

Table B.1: Threshold mismatch

The following averages can be found from table B.1:

$A_{VTH,N} \approx 3.4 \cdot 10^{-3}$
$A_{VTH,P} \approx 4.4 \cdot 10^{-3}$

The deviation at minimum size comes from the significant difference between drawn length and effective length at these sizes. These values are therefore not used in the calculation of the average.

K-factor mismatch:

N-MOST	**W = 2 µm**		**W = 5 µm**	
L (µm):	**σ/K (10^{-3}):**	**A_K (10^{-3}):**	**σ/K (10^{-3}):**	**A_K (10^{-3}):**
1	4.72	6.68	2.95	6.59
2	3.42	6.84	2.13	6.74
5	2.11	6.67	1.31	6.56
10	1.48	6.61	0.917	6.49

Table B.2: N-MOST K-factor mismatch

P-MOST	**W = 2 µm**		**W = 5 µm**	
L (µm):	**σ/K (10^{-3}):**	**A_K (10^{-3}):**	**σ/K (10^{-3}):**	**A_K (10^{-3}):**
1	3.96	5.60	2.46	5.50
2	2.79	5.58	1.73	5.48
5	1.80	5.68	1.12	5.60
10	1.29	5.77	0.799	5.65

Table B.3: P-MOST K-factor mismatch

$A_{K,N} \approx 6.6 \cdot 10^{-3}$
$A_{K,P} \approx 5.6 \cdot 10^{-3}$

With these values, an quick, first order estimation of mismatch can be made where needed.

B.4 CAPA Capacitance versus Voltage

A sweep has been performed to find the capacitance of a CAPA. This has been done by measuring the impedance of a CAPA at a range of DC bias voltages. The result can be found in figure B.5.

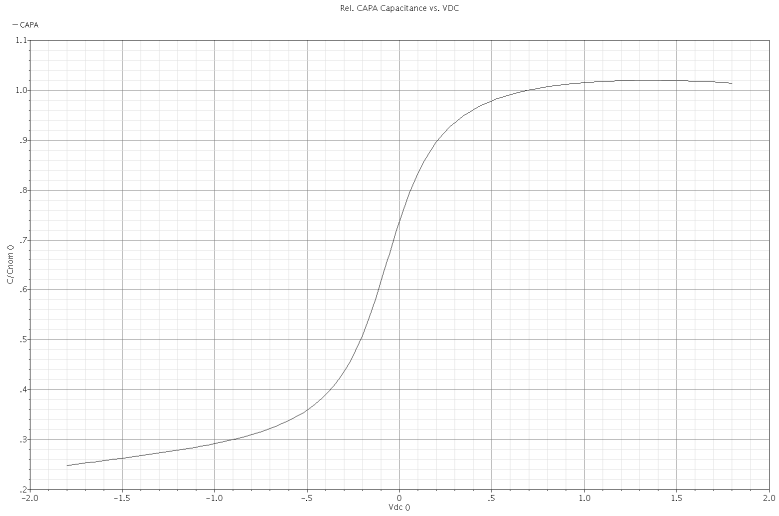

Figure B.5: Relative CAPA capacitance versus DC voltage

B.5 Summary and Some Other Constants

A summary of the results from this appendix, as well as a few other constants that are convenient to have at hand are listed here.

Name:	**Symbol:**	**Dimension:**	**Unit**:
Boltzmann's constant	K	$1.38 \cdot 10^{-23}$	J/K
Magnetic constant	μ_0	$4 \cdot \pi \cdot 10^{-7}$	H/m
Electric constant	ε_0	$8.85 \cdot 10^{-12}$	F/m
Electron charge	e	$1.6 \cdot 10^{-19}$	C
Room temperature	T_0	300	K
Thermal voltage	V_T	$25.9 \cdot 10^{-3}$	V
Inverse thermal voltage	V_T^{-1}	38.6	V^{-1}
PNP saturation current*	I_S	$2.38 \cdot 10^{-17}$	A
NMOST threshold voltage*	V_{TH}	$390 \cdot 10^{-3}$	V
NMOST K-factor*	$\mu_n \cdot C_{OX}$	$0.28 \cdot 10^{-3}$	A/V^2
NMOST threshold mismatch*	A_{VTH}	3.4	-
NMOST K-factor mismatch*	A_K	$6.6 \cdot 10^{-3}$	-
PMOST threshold voltage*	V_{TH}	$-390 \cdot 10^{-3}$	V
PMOST K-factor*	$\mu_n \cdot C_{OX}$	$0.048 \cdot 10^{-3}$	A/V^2
PMOST threshold mismatch*	A_{VTH}	4.4	-
PMOST K-factor mismatch*	A_K	$5.6 \cdot 10^{-3}$	-
Minimum device length	L_{Min}	180	nm
Maximum device length	L_{Max}	20	µm

* Found by simulation