A Highly Digital, Reconfigurable and Voltage Scalable SAR ADC

by

Marcus Yip

B.A.Sc., Engineering Science, University of Toronto

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2009

Author .....
Department of Electrical Engineering and Computer Science
May 15, 2009

Certified by .....
Anantha P. Chandrakasan
Professor of Electrical Engineering
Thesis Supervisor

Accepted by .....
Terry P. Orlando
Chair, Department Committee on Graduate Students

A Highly Digital, Reconfigurable and Voltage Scalable SAR ADC

by

Marcus Yip

Submitted to the Department of Electrical Engineering and Computer Science
on May 15, 2009, in partial fulfillment of the
requirements for the degree of
Master of Science in Electrical Engineering and Computer Science

Abstract

Micropower sensor networks have a broad range of applications which include military surveillance, environmental monitoring, chemical detection and more recently, medical monitoring systems. Each node of the sensor network requires energy efficient circuits powered off small batteries or harvested energy. In such systems, a single reconfigurable analog-to-digital converter (ADC) is needed to digitize a wide range of signals with varying bandwidth and resolution requirements. This thesis describes the design of an ADC whose power scales exponentially with resolution and linearly with frequency to maximize the system lifetime.

The proposed ADC has reconfigurable resolution from 5 to 10-bits and a scalable sample rate from 0 to 1-MS/s. The successive approximation register (SAR) architecture was chosen for its highly digital nature which enables low voltage operation. The supply voltage can be scaled from 1V down to 0.4V such that the ADC maintains a constant energy efficiency across all modes of operation when normalized with respect to sample rate and resolution. A capacitive digital-to-analog converter (DAC) in a split capacitor topology with a sub-DAC is used to minimize the DAC power and area. Top plate switches are used to decouple the MSB capacitors as resolution is scaled to avoid parasitic loading of the DAC. The DAC capacitors are laid out in a common-centroid configuration with edge effects minimized at each resolution mode to improve matching. A fully dynamic latched comparator is used to avoid static bias currents. Power gating of the digital logic is used to reduce leakage power at low sample rates. Reconfigurability between single-ended or differential modes enables a power versus performance trade-off. Lastly, programmable sampling duration and internal bootstrapping is used to maintain sampling linearity at low voltages. The ADC has been submitted for fabrication in a low power 65nm digital CMOS process and simulation results are presented.

Thesis Supervisor: Anantha P. Chandrakasan
Title: Professor of Electrical Engineering

Acknowledgments

This thesis would not have been possible without the help and support of an endless list of colleagues, friends and family. Let me attempt to express my sincere gratitude here.

First, I'd like to thank my supervisor, Professor Anantha Chandrakasan, for inspiring me to get excited about research. I am grateful for his guidance, support and insight. I would also like to thank Naveen Verma and Brian Ginsburg for answering all of my SAR-related questions, no matter how trivial.

Thanks also goes to Denis, Joyce, Yogesh, Ersin, Masood, Payam, Fred, Courtney, Patrick, Vivienne and the rest of ananthagroup for helping me with technical things somewhere along the way, and also for making the lab a fun place to be. Also, I'd like to thank Margaret for keeping everything running smoothly.

I'd like to acknowledge MOSIS for chip fabrication, and Cambridge Analog Technologies for graciously providing me with pads at the last minute, without which tape-out would not have happened.

To my friends, who all remind me that work is just that, work. There is life outside school and I am grateful to be surrounded by such wonderful friends. Edward, Alan, Anita, Cindy, Hannah and Polly, thanks for making Toronto such a great place to come home to. Bryan, Andy and George, thanks for keeping me sane with all the ridiculousness that always seems to surround us. Oh, and how can I forget to thank Young Shin, Chester and Dave Campbell for many fun-filled late night discussions. To everyone I've met at MIT, thanks for making Boston a fun place to live.

Finally, I would like to thank my family for their continued love and support.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The miniaturization of electronics has enabled the development of wireless sensor networks which consist of many distributed nodes, each equipped with sensors and low power circuits to acquire, process and transmit the signal of interest. Sensor networks have broad applications including military surveillance, environmental monitoring, chemical and biological detection and medical monitoring systems [1]. Each sensor node is typically powered by a small battery or scavenged energy, thus placing stringent requirements on the power consumption of the circuits. Furthermore, as in the case of environmental monitoring or surveillance applications, high density and ubiquity of sensor nodes requires that they be inexpensive, which limits die area. For these reasons, highly energy efficient and reconfigurable circuits are needed in order to minimize power and area.

To illustrate the structure of a typical sensor node, we will consider medical monitoring as an example. As seen in Figure 1-1, a typical medical monitoring system consists of multiple sensors, a low noise instrumentation amplifier (IA), an analog-to-digital converter (ADC) and a digital signal processor (DSP). Once digitized and processed, the data can be stored or transmitted wirelessly via a short range radio to a local relay such as a cellular phone. The cellular phone can then provide connectivity to a secure online medical server, where a physician can examine the patient's vital signs. The design of the system is primarily constrained by area and power. Small die area helps to achieve a small form factor, while low power consumption is crucial

in extending the lifetime of the system.

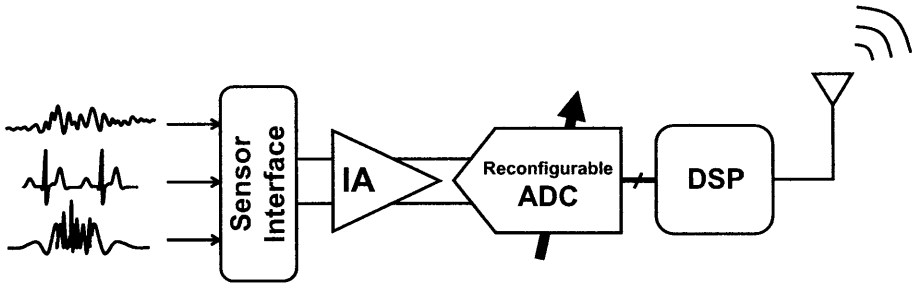

Figure 1-1: A typical medical monitoring system consisting of a sensor interface with multiple sensors, instrumentation amplifier, ADC, DSP and a short range radio.

This thesis will focus on the design and implementation of a low power, highly reconfigurable ADC. Reconfigurability in sample rate and resolution is desirable to accomodate digitization of a wide variety of signals with varying bandwidths and dynamic range requirements. Since energy is a limited resource, it is desirable to have the power consumption of the ADC scale with its performance. In this design, highly digital techniques are used to allow the ADC to maintain its energy efficiency across all operating modes.

The remainder of this chapter will be divided as follows. Section 1.1 will discuss the target applications and motivate the need for reconfigurability in performance. Section 1.2 will introduce common figure-of-merits used to compare the energy efficiency of ADCs. Section 1.3 surveys the current state-of-the-art in reconfigurable and high efficiency ADCs. Finally, Section 1.4 will discuss the architecture selection for the ADC.

1.1 Motivation for Reconfigurability

In order to motivate the need for reconfigurable circuits, we will consider the requirements of a medical monitoring system. Many medical applications require the acquisition of bio-potentials, or physiological signals, which often have different amplitudes and bandwidths [2], [3], [4]. Common bio-potentials and their characteristics are listed in Table 1.1 and plotted in the amplitude-frequency space in Figure 1-2.

In order to minimize design time and die area, as well as maximizing the number of applications, the analog front-end should be able to adjust its gain and bandwidth to accomodate a wide range of bio-potentials. In particular, the ADC should also be reconfigurable and be able to scale its performance depending on the application. The remainder of this section will outline the ADC performance requirements and motivate the need for scalability in performance.

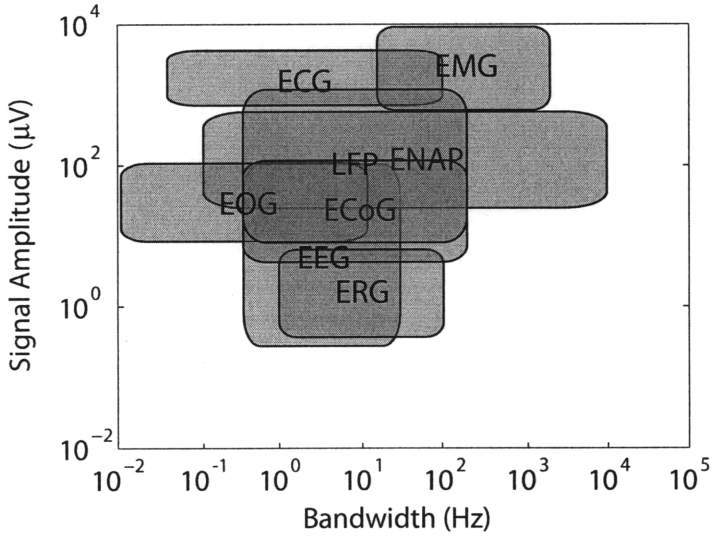

Figure 1-2: Amplitude vs. bandwidth characteristics of various bio-potentials.

1.1.1 Resolution

The dynamic range of the analog front-end in a medical monitoring system should be large enough to encompass the expected minimum and maximum signal amplitudes to be detected. Assuming that the instrumentation amplifier and front-end filters are designed for a sufficient dynamic range, the ADC will limit the dynamic range of the system by introducing quantization noise and distortion to the signal. The resolution should be chosen high enough to provide an adequate signal-to-noise and distortion ratio (SNDR). From Table 1.1, it can be seen that the signal amplitude range differs for each bio-potential, suggesting different ADC resolution requirements. Other considerations for choosing the resolution of the ADC are the algorithms to

Bio-potential	Bandwidth	Amplitude
EEG (electroencephalography)	0.5 to 40-Hz	0.5 to 100-μV
ECG (electrocardiography)	0.05 to 100-Hz	1 to 5-mV
EMG (electromyography)	20-Hz to 2-kHz	1 to 10-mV
EOG (electrooculography)	DC to 10-Hz	10 to 100-μV
ERG (electroretinography)	1 to 100-Hz	0.5 to 8-μV
ECoG (electrocortigraphy)	0.5 to 200-Hz	5 to 100-μV
LFP (local field potential)	0.5 to 200-Hz	10-μV to 1-mV
ENAP (extracellular neural action potential)	0.1-Hz to 10-kHz	50 to 500-μV

Table 1.1: Bandwidths and amplitudes of various bio-potentials.

be implemented in the digital domain. Taking the ECG as an example application, simple tasks such as heart rate extraction requires no more than 8-bits [5]. However, more sophisticated algorithms used for detecting slow changes in the ST pattern may require 12 to 16-bits of resolution [6].

One approach would be to design the ADC for the maximum resolution needed and truncate the ADC outputs as needed. However, this is very wasteful because power consumption in ADCs typically scales exponentially with the effective resolution as will be explained in Section 1.2. Thus, to allow for flexibility in the detectable signal amplitude range and processing algorithms, an ADC with reconfigurable resolution is highly desired.

In this ADC, a design choice was made to pursue a reconfigurable range of 5 to 10-bits of resolution and the primary goal was to demonstrate power scalability over a large resolution range (6-bits).

1.1.2 Sample Rate

The signal bandwidths for bio-potentials differ across applications, as shown in Table 1.1. In order for the ADC to faithfully digitize a variety of signals, the maximum sampling rate must satisfy the Nyquist requirement for the highest bandwidth application. However, since it is well known that power consumption scales with frequency in any digital or mixed-signal system, it is desirable to be able to reduce the sampling rate for lower bandwidth applications.

For medical applications, bio-potentials are typically very low bandwidth (up to

a few kHz) and the ADC sampling rate can be as low as just tens of kilo-Samples per second. In this design, a maximum sampling rate of 1-MS/s was chosen to demonstrate frequency scaling over a larger range of frequencies. This will enable characterization of the ADC in both active and leakage energy dominated regions of operation, which in turn, will allow the effects of the applied leakage reduction technique presented later in Section 3.2.5 to be quantified.

1.2 ADC Figure of Merit

In order to compare the energy efficiency of ADCs, there are two commonly used figure-of-merits (FOM). For low to moderate resolution ADCs (roughly 12-bits or less), FOM1 shown in Equation 1.1 is used, where P is the power consumption, ENOB is the effective number of bits and f_{in} is the input bandwidth at which the ENOB is calculated. The ENOB is a measure of the effective dynamic range of the ADC, and is given by Equation 1.2, where the SNDR is calculated from the FFT of a sequence of ADC outputs [7]. FOM1 is commonly referred to as the Walden empirical FOM and is based on an empirical survey of the performance of over 150 converters [8]. Essentially, the power is normalized by the input bandwidth and effective dynamic range of the converter to arrive at an energy per conversion step.

$$FOM1 = \frac{P}{2f_{in}2^{ENOB}} \tag{1.1}$$

$$ENOB = \frac{SNDR(dB) - 1.76}{6.02} \tag{1.2}$$

It is worthwhile to note that high resolution ADCs (typically 12-bits or higher) are typically at a disadvantage when FOM1 is applied. This is because high resolution ADCs are usually limited by thermal noise, which can only be reduced by consuming more power. Therefore, a modified figure-of-merit FOM2, known as the thermal FOM shown in Equation 1.3 is used. FOM2 typically applies to ADCs with an SNDR exceeding 85-dB [9].

$$FOM2 = \frac{P}{2f_{in}2^{2ENOB}} \tag{1.3}$$

It is important to emphasize that FOM1 is based only on empirical observation [8]. Ideally, it would be desirable to minimize power as much as possible (beyond the 2× per bit resolution) as resolution is scaled, but FOM1 serves as a reasonable guideline for power scaling. Furthermore, it has been observed that technology scaling along with voltage scaling has been the driving force behind improving FOM1 for low to moderate resolution ADCs [9]. Thus, in this design, a decision was made to use a low power 65-nm digital CMOS process with aggressive voltage scaling to improve the energy efficiency. For the ADC described in this thesis, FOM1 will be used as the measure of efficiency since it is not in the noise-limited regime as explained in Section 2.7.

1.3 Current State-of-the-Art

Before architecture selection, it is worthwhile to survey the current state-of-the-art in high energy-efficiency ADCs, as well as ADCs that have high degrees of reconfigurability. This may provide valuable insights into which architectures are suitable for the design space specified in the previous section. Although being energy-efficient and reconfigurable are not mutually exclusive, there is often a large overhead involved with reconfigurable architectures that limit the energy-efficiency of the ADC across all modes of operation.

1.3.1 High Energy-Efficiency ADCs

While a FOM of a few pJ/conversion-step was considered state-of-the-art a few years ago, FOMs on the order of 10's of fJ/conversion-step are now being reported [10] - [17]. With technology scaling, clever design techniques and a trend towards minimalistic design, femto-Joule energy efficient ADCs are becoming commonplace. A summary of the ADCs with the lowest reported FOMs in the last three years is given in Table

Source	Year Published	Process (nm)	Supply (V)	Power (μW)	ENOB	Sample Rate	FOM (fJ/conv)	Architecture
[10]	2007	180	0.9	2.47	7.58	200-kS/s	64.5	SAR
[11]	2007	90	1	700	7.8	50-MS/s	65	SAR
[12]	2007	180	1	25	10.55	100-kS/s	165	SAR
[13]	2008	65	1	1.9	8.74	1-MS/s	4.4	SAR
[14]	2008	90	1	820	8.56	40-MS/s	54	SAR
[15]	2008	90	1.2	2200	4.67	1.75-GS/s	50	Flash
[16]	2008	90	1	133	6.4	150-MS/s	10.4	CABS
[17]	2009	90	1.2	4500	10	50-MS/s	88	Pipelined

Table 1.2: Current state-of-the-art in femto-Joule FOM ADCs.

1.2. It can be seen that the SAR ADC architecture has been used for some of the lowest FOM ADCs in the low to moderate resolution and sample rate space [10], [11], [12], [13], [14]. In particular, a 10-bit, 1-MS/s SAR ADC described in [13] achieves the lowest reported FOM1 to date of 4.4-fJ/conversion-step. However, this design uses very small unit capacitors which compromises matching and degrades linearity. This is a trade-off that will be discussed later in Section 3.2.7.

In [15], a 5-bit, 1.75-GS/s fully dynamic folding Flash architecture was used to achieve a FOM1 of 50-fJ/conversion-step. In [16], a 7-bit, 150-MS/s comparator-based asynchronous binary-search (CABS) ADC was used to achieve a FOM1 of only 10.4-fJ/conversion-step. In this approach, a binary tree of fully dynamic comparators with static offsets was used to resolve the input. Finally, in [17], a 12-bit, 50-MS/s zero-crossing-based pipelined ADC without op-amps achieved a FOM1 of 88-fJ/conversion-step.

In order to increase the energy efficiency of ADCs, power must be reduced without sacrificing performance. Fundamentally, this can be done by either reducing the current consumption or lowering the voltage. For digital circuits, process scaling enables designers to lower the supply voltage without sacrificing performance due to reduced device dimensions and scaled threshold voltages. However, scaling leads to increased static leakage currents which places a lower limit on power consumption. Often, for leakage to become negligible, the frequency of operation must be high enough such that active power dominates. For example, Table 1.3 summarizes recently published ADCs (where a significant portion of the circuitry is digital) in 65-nm CMOS and it

Source	Year Published	Supply (V)	Power (μW)	ENOB	Sample Rate (MS/s)	FOM (fJ/conv)	Architecture
[18]	2008	1.2	4500	9.5	100	62	Pipelined
[19]	2008	0.8	1200	4.43	250	240	SAR
[20]	2008	N/A	950	12.5	150	300	$\Delta\Sigma$
[21]	2008	1.3	50000	13.2	256	340	$\Delta\Sigma$
[22]	2008	1.2	12000	5.63	800	400	Flash
[23]	2008	1.2	5510	8.73	26	499	Pipelined
[24]	2006	1.2	6000	4.5	500	755	SAR
[25]	2008	1	180000	10	200	855	Pipelined

Table 1.3: Recently published ADCs in 65-nm CMOS.

can be seen that the sample rates are all on the order of 10's to 100's of MSamples/s. Thus, for medical applications which require lower sample rates, a key design problem is leakage reduction which will be one of the focuses of this thesis.

1.3.2 Reconfigurable ADCs

While the ADCs in [10] - [17] are all extremely energy efficient, they have been designed for a fixed resolution and lack the scalability and reconfigurability needed for applications such as medical monitoring systems. Apart from these examples, there has been research on power scalable and resolution reconfigurable ADCs in the past [12], [26], [27]. In [26], a frequency scalable 10-bit pipelined ADC from 50-MS/s down to 1-kS/s was reported. However, this design did not have reconfigurable resolution. In [12], a SAR ADC with 8 and 12-bit modes was reported, however the power consumption of the ADC was reduced by only 24% when going from 12 to 8-bits. In [27], a 0-10 MS/s ADC that can reconfigure between pipeline and $\Delta\Sigma$ modes was reported. The pipeline mode was used in the range of 6 to 12-bits, while the $\Delta\Sigma$ mode was used from 13 to 16-bits. Its FOM at each resolution was competitive with custom state-of-the-art ADCs at the time, but its FOM varied by 4 orders of magnitude over the entire resolution range indicating that its energy efficiency was dependent on the resolution mode. Thus, it can be seen that it is difficult to achieve both a high degree of scalability and energy efficiency simultaneously.

1.4 Architecture Selection

The ADC architecture should be chosen based on the requirements of scalable sample rate from 0 to 1-MS/s and reconfigurable resolution from 5 to 10-bits as discussed in Section 1.1. Since energy efficiency is most critical, power consumption must also be considered in the architecture selection. Figure 1-3 illustrates the region where popular ADC architectures are most energy efficient in the bandwidth-resolution space [7], [10] - [17].

Flash ADCs can be eliminated because they are typically used for low resolution, high speed applications [15]. Flash converters occupy only the low resolution regime (7-bits or less) because the number of comparators grows exponentially with resolution, rendering it unsuitable for the desired resolution specification. Pipelined ADCs typically occupy the high resolution, high speed regime and they are typically quite power hungry due to significant analog circuitry [7]. Pipelined ADCs are most efficient at 10's to 100's of MS/s and are not suitable for the desired sample rate specification [17], [18], [25]. Oversampling ADC architectures such as $\Delta\Sigma$ modulators are quite energy efficient at low input bandwidths. They are able to achieve very high signal-to-noise ratios (SNR) by using averaging and noise shaping and are typically used for high resolution applications requiring 12-bits or more [20], [21]. It is apparent that the SAR ADC architecture, due to its simplicity and highly digital nature, is ideally suited for low speed, moderate resolution ADCs and it has been chosen for this ADC prototype.

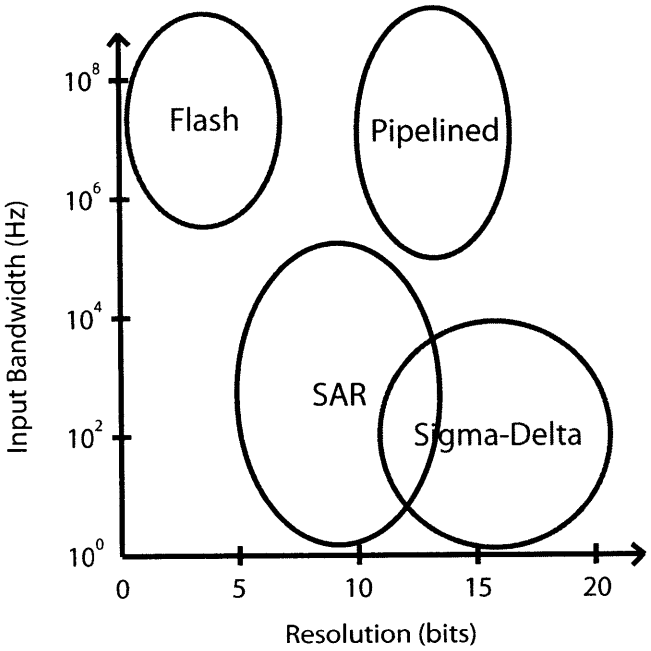

Figure 1-3: ADC architectures and their area of highest energy efficiency in the bandwidth-resolution space.

Chapter 2

Technology Limitations

This chapter will examine commonly encountered limitations and challenges with low voltage design in deep sub-micron technologies. In particular, the implications of these limitations on the design of the proposed SAR ADC will be discussed.

2.1 MOS Switch Resistance

In analog and mixed signal circuit design, ideal switches with zero resistance when ON and infinite resistance when OFF are often required for input sampling, signal multiplexing and charge redistribution as in switched-capacitor circuits. In practice, MOS transistors operating in the linear region are good approximations to ideal switches, with low ON resistance and high OFF resistance in the GΩ range. The effective ON resistance of a MOS switch is given by Equation 2.1, where μ is the carrier mobility, C_{ox} is the oxide capacitance, $\frac{W}{L}$ is the transistor aspect ratio, V_{GS} is the gate-source voltage and V_T is the transistor threshold voltage.

$$R_{ON} = \left[\frac{\partial I_{DS}}{\partial V_{DS}}\right]^{-1} = \frac{1}{\mu C_{ox}\frac{W}{L}(V_{GS} - V_T)} \tag{2.1}$$

Since a NMOS-only switch can pass signals from 0 to $(V_{DD} - V_{Tn})$ and a PMOS-only switch can pass signals from $|V_{Tp}|$ to V_{DD}, a transmission gate made up of a parallel combination of NMOS and PMOS switches must be used in order to pass

signals from rail-to-rail as shown in Figure 2-1 (a). The effective resistance of the transmission gate is $R_{ON_n}||R_{ON_p}$ as shown in Figure 2-1 (b). The simulated effective ON resistance of a transmission gate over the full scale input range from 0 to V_{DD} is plotted in Figure 2-2 for V_{DD} ranging from 0.4V to 1V.

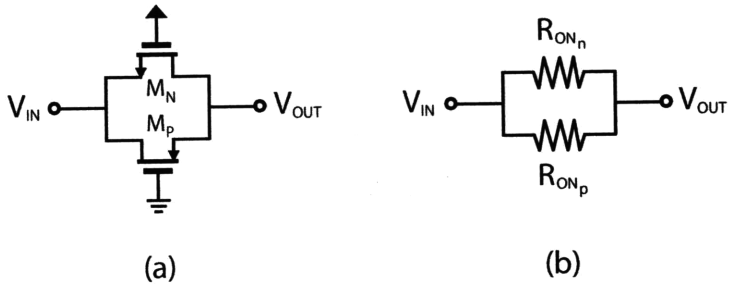

Figure 2-1: (a) CMOS transmission gate. (b) Model for effective resistance of a transmission gate.

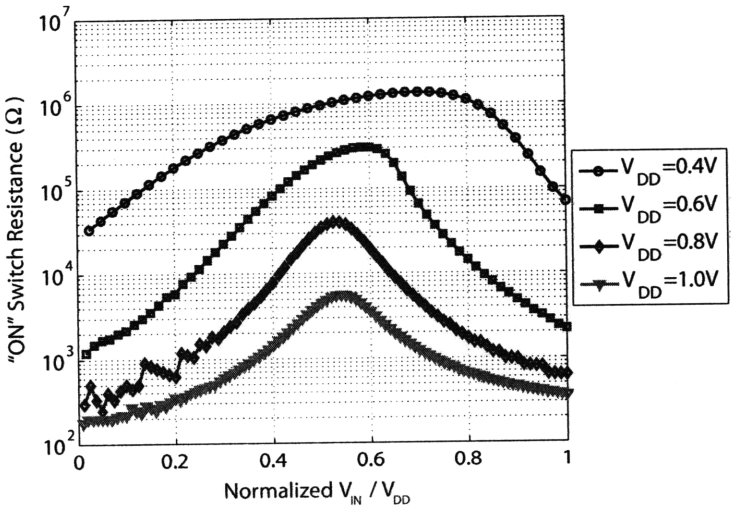

Figure 2-2: Effective large signal ON resistance of a transmission gate for an input range from 0 to V_{DD} for V_{DD} from 0.4V to 1V.

It can be seen that for low voltage designs where $V_{DD} < V_{Tn} + |V_{Tp}|$, there exists a range of inputs near mid-rail where neither the NMOS or PMOS switch conducts strongly. This results in increased switch resistance, leading to larger RC time constants which must be accounted for when designing circuits at ultra low supply voltages. Furthermore, it can be seen that the switch resistance depends on the input

voltage, resulting in non-linearity.

In this ADC where we are targeting operation down to 0.4-V, charge pumps are used to boost the overdrive voltage of sampling switches in order to reduce the switch resistance and increase linearity. This will be discussed in Section 4.2.1.

2.2 Charge Injection

Charge injection is a non-ideality of using MOS devices as switches. When a switch turns off, the channel charge q_{ch}, will leave through the source and drain terminals. Assuming that the switch is driven with V_{DD} when "on", q_{ch} is given by Equation 2.2. With regards to ADCs, charge injection from the sampling switches results in error in the sampled voltage. An example sampling circuit consisting of a simple NMOS switch and a sampling capacitor C_{SAMP} is shown in Figure 2-3. C_{OV} represents the extrinsic overlap capacitance which will be ignored for now. The actual fraction k of the channel charge that moves towards the output node V_{SH} depends on the relative impedances on either side of the switch.

$$|q_{ch}| = WLC_{ox}(V_{GS} - V_T) = WLC_{ox}(V_{DD} - V_{IN} - V_T) \tag{2.2}$$

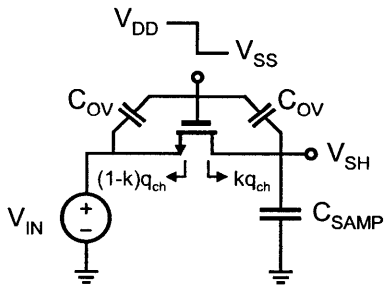

Figure 2-3: Source of charge injection and clock feedthrough errors in a sampling circuit.

The error in the sampled voltage, ΔV_{CI}, can be approximated by Equation 2.3. ΔV_{CI} can be decomposed into an offset term $V_{OS,CI}$, and a gain term $A_{CI}V_{IN}$ which are given in Equations 2.4 and 2.5 respectively. Here, we have ignored the body effect

and assumed that V_T stays constant during the sampling process. Considering the charge injection error, the final sampled voltage V_{SH} is given in Equation 2.6. The offset term $V_{OS,CI}$ can be canceled by using a differential architecture assuming good matching, however the gain term $A_{CI}V_{IN}$ is input dependent and cannot easily be canceled.

$$\Delta V_{CI} = \frac{k|q_{ch}|}{C_{SAMP}} = kWL(V_{DD} - V_{IN} - V_T)\frac{C_{ox}}{C_{SAMP}} = V_{OS,CI} + A_{CI}V_{IN} \tag{2.3}$$

$$V_{OS,CI} = kWL(V_{DD} - V_T)\frac{C_{ox}}{C_{SAMP}} \tag{2.4}$$

$$A_{CI} = -kWL\frac{C_{ox}}{C_{SAMP}} \tag{2.5}$$

$$V_{SH} = V_{IN} + \Delta V_{CI} = (1 + A_{CI})V_{IN} + V_{OS,CI} \tag{2.6}$$

In addition to charge injection error, another source of error known as clock feedthrough affects the sampled voltage in a similar manner. To see this, assume that the clock edge is very fast such that the switch turns off instantly and the only gate-source capacitance is C_{OV} (since the intrinsic C_{GS} is zero in cut-off). Thus, the voltage excursion at the gate appears at the output node through the capacitive divider formed by C_{OV} and C_{SAMP} and results in an error ΔV_{CF} as shown in Equation 2.7. Fortunately, this error is signal independent and can be made common-mode by using differential signaling.

$$\Delta V_{CF} = -(V_{DD} - V_{SS})\frac{C_{OV}}{C_{SAMP} + C_{OV}} \approx -(V_{DD} - V_{SS})\frac{C_{OV}}{C_{SAMP}} \tag{2.7}$$

The effects of charge injection must be considered when designing the input sampling network.

2.3 Capacitor Mismatch

In this ADC, a capacitive DAC is used to generate precise analog voltages in the feedback path. The DAC output is generated based on a ratio of capacitances and thus depends heavily on the matching properties of capacitors. Any capacitor mismatch will result in non-linearity in the overall ADC transfer characteristic, resulting in distortion and reducing the effective dynamic range of the converter.

In the selected process technology, available capacitors were metal-metal flux capacitors and MOS varactors. Since MOS varactors are extremely non-linear, the metal-metal flux capacitors were selected. Since capacitor matching properties were not available, a survey of existing designs in 65-nm processes was performed as a guideline for capacitor sizing for adequate linearity. Designs in older technologies (90-nm and 0.13-μm) were also surveyed and a scaling factor of $\frac{1}{\sqrt{2}}$ was applied per technology generation. For example, if a unit capacitance of 100-fF was adequate for 10-bit linearity in a 0.13-μm process, then it was assumed that a 50-fF (100-fF $\times$ $(\frac{1}{\sqrt{2}})^2$) unit capacitance would yield the same linearity in a 65-nm process. Of course, this method should be considered only as a rough guideline and adequate margin should be added to ensure good linearity.

2.4 Transistor Mismatch

The drain current of a MOSFET operating in saturation in strong and weak inversion are modeled by Equations 2.8 and 2.9 respectively [28]. Here, n is the sub-threshold slope factor and ϕ_t is the thermal voltage. The threshold voltage, V_T, is modulated by the body effect which is described by Equation 2.10, where γ is the body effect coefficient and ϕ_s is the surface potential. Note that other second order effects such as channel length modulation and drain-induced barrier lowering have been ignored for simplicity.

$$I_{D_{Strong-Inv}} = \frac{\mu C_{ox}}{2}\frac{W}{L}(V_{GS} - V_T)^2 \tag{2.8}$$

$$I_{D_{Weak-Inv}} = I_o e^{\left(\frac{V_{GS}-V_T}{n\phi_t}\right)} = \mu C_{ox} \frac{W}{L} (n-1)\, \phi_t^2 e^{\left(\frac{V_{GS}-V_T}{n\phi_t}\right)} \tag{2.9}$$

$$V_T = V_{T0} + \gamma \left(\sqrt{\phi_s + V_{SB}} - \sqrt{\phi_s} \right) \tag{2.10}$$

It can be seen that any mismatch in μ, C_{ox}, W, L, V_{T0} or γ can result in a mismatch in drain current for two nominally identical transistors. Note that μ, V_{T0} and γ are process dependent parameters that depend largely on doping. In particular at low supply voltages, V_{T0} variation is dominated by random dopant fluctuation (RDF) and is typically modeled with a Gaussian distribution [29]. The standard deviation of the threshold voltage mismatch, ΔV_T, is given by Equation 2.11 [30]. Thus, to reduce V_T mismatch, which is particularly important for sub-threshold operation, devices should be sized relatively large. It is interesting to note that the proportionality constant, A_{VT} scales down with gate oxide thickness [31], which implies that V_T mismatch improves with process scaling for the same device area.

$$\sigma_{\Delta V_T} = \frac{A_{VT}}{\sqrt{WL}} \tag{2.11}$$

In circuits with differential pairs such as amplifiers or comparators, transistor mismatch results in input referred DC offsets, increased even-order distortion and reduced common-mode rejection [32]. In architectures where many comparators operate in parallel such as a Flash ADC, offsets can limit the linearity of the converter. For a SAR ADC where there is only one comparator, the comparator offset manifests itself as offset in the overall ADC transfer characterisitic and can limit the full scale range of the converter.

Good layout practices can often eliminate many first order effects of geometry mismatch. For example, using common-centroid arrangements and keeping matched transistors in close proximity can null out mismatch due to a gradient in oxide thickness across the wafer. Using dummy transistors to create uniform etch environments can also improve matching.

2.5 Leakage in Advanced Digital CMOS Processes

Process scaling is driven by the need to operate digital circuits at ever increasing speeds. This is enabled by shrinking device geometries leading to reduced parasitic capacitances and higher transistor f_t. However, the required power density increases dramatically due to the increased frequencies and transistor density. It is well known that the active switching power of a digital circuit is given by Equation 2.12, where α is the activity factor, f is the clock frequency, C_L is the total load capacitance and V_{DD} is the supply voltage.

$$P_{active} = \alpha f C_L V_{DD}^2 \tag{2.12}$$

An obvious way to reduce active power consumption is to reduce V_{DD}, however, this reduces the transistor drive strength and limits performance. In order to maintain performance at reduced V_{DD}, transistor threshold voltages are also being scaled. The consequence, as seen from Equation 2.9, is that reducing V_T leads to exponentially increasing sub-threshold current which is static leakage that is present even if the circuit is idle. Approximately, every 100-mV reduction in V_T increases leakage by an order of magnitude. The total power consumption of a digital circuit considering active and leakage components is given by Equation 2.13, where I_{LEAK} is the total sub-threshold and gate leakage current. In deep-submicron technologies, leakage power can become a significant portion of the total power. Thus, leakage reduction techniques must be applied to reduce power consumption and this will be discussed in Section 3.2.5.

$$P_{tot} = \alpha f C_L V_{DD}^2 + V_{DD} I_{LEAK} \tag{2.13}$$

In the context of ADCs which are inherently mixed-signal systems, the problem of leakage is still important due to the trend towards digitally assisted architectures [9]. Modern ADCs designed in advanced digital CMOS processes often have a large digital component to correct for errors made in the analog domain. In particular, the

ADC described in this thesis is fully dynamic and uses no static biases. Thus, leakage reduction techniques will be important in order to maintain energy-efficiency at very low sample rates.

2.6 Sub-Threshold Circuit Operation

As mentioned in the previous section, it is beneficial to lower V_{DD} in order to reduce the dynamic power dissipation of digital circuits. In particular, for energy constrained systems where performance is not an issue, sub-threshold operation where $V_{DD} < V_T$ enables siginifcant energy savings. The main difference between above and sub-threshold logic is the value of the "on" current, I_{ON}, used to switch the output of the gate. Above threshold, I_{ON} is a quadratic function of V_{GS} (linear, if velocity saturated) and is typically a few orders of magnitude larger than I_{ON} in sub-threshold, which depends exponentially on V_{GS}. A consequence of this is that the propagation delay, and accordingly the clock period T_{cycle}, of sub-threshold logic increases exponentially as V_{DD} is reduced. Thus, the total energy per clock cycle is given by Equation 2.14 and is plotted in Figure 2-4 for a representative 32-bit adder [29]. It is evident that the opposing trends of E_{DYN} and E_{LEAK} with V_{DD} results in an optimal supply voltage V_{DDopt}, for minimizing the energy per cycle [29], [33]. Therefore, if the application is constrained by energy rather than performance, it is beneficial to operate at this minimum energy point.

$$E_T = E_{DYN} + E_{LEAK} = C_L V_{DD}^2 + V_{DD} I_{LEAK} T_{cycle} \tag{2.14}$$

Note that operating at the minimum energy point does not come without challenges. The effects of process variation are more profound in sub-threshold since sub-threshold current depends exponentially on V_T and the ratio I_{ON}/I_{OFF} is much lower than in above-threshold operation. This leads to reduced performance and increased functional failures due to reduced logic noise margins [29].

With respect to the ADC described in this thesis which is highly digital in nature, a theoretical V_{DDopt} exists at low sample rates where the digital logic can operate in

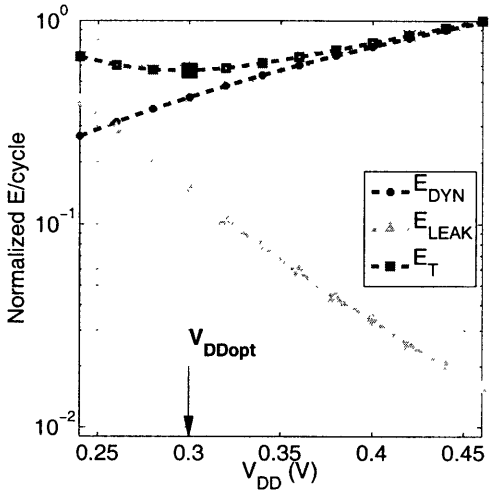

Figure 2-4: Dynamic (E_{DYN}), leakage (E_{LEAK}) and total (E_T) energy per cycle in a 32-bit adder (Data courtesy of J. Kwong, MIT).

sub-threshold.

2.7 Device Noise

For ADCs, a major source of noise is the integrated thermal noise of MOS switches during the sampling process. The sampled noise power is given by Equation 2.15, where C_{DAC} is the sampling capacitance in a SAR ADC. To determine whether this noise is significant, it should be compared to the ADC quantization noise power, v_q^2, given by Equation 2.16, where V_{LSB} is the LSB voltage equal to the ADC full scale range V_{FS} divided by the total number of quantization levels, 2^N.

$$v_{n,samp}^2 = \frac{kT}{C_{DAC}} \tag{2.15}$$

$$v_q^2 = \frac{V_{LSB}^2}{12} = \frac{(V_{FS}/2^N)^2}{12} \tag{2.16}$$

For a 10-bit ADC with a full scale range of 1-V, $v_q = 282\text{-}\mu V_{rms}$. If we conservatively assume C_{DAC} = 1-pF (in practice, 10-bit ADCs will have C_{DAC} on the order of 5 to 10-pF), then $v_{n,samp} = 64.3\mu V_{rms}$ which is below the quantization noise

floor. Thus, the dynamic range of the ADC described in this thesis is limited by quantization noise rather than sampled $\frac{kT}{C}$ noise.

2.8 Substrate Noise

In modern process technologies, the substrate is usually heavily doped to reduce substrate resistance in order to minimize the occurences of latch-up. However, these low resistance paths can lead to unwanted coupling between certain devices in the circuit. For example, any digital switching noise injected into the substrate can change the V_T of a nearby device through the body effect. If that transistor is part of a sensitive analog circuit, then the substrate noise essentially acts as an input noise source. In a mixed-signal system such as an ADC, care must be taken to ensure that any digital noise (especially clocks and I/O) does not corrupt any crucial analog signals.

In this design, several techniques are used to minimize the effects of substrate noise. First, a differential architecture is used to reject any common-mode noise. Secondly, all sensitive analog blocks are placed in a p-well, separated from the substrate by a guard-ringed n-well. Note that this option is only possible in triple well processes. Furthermore, while the input is being sampled by the ADC, the system clocks are gated and digital switching is suspended. Lastly, the ability to selectively turn off the I/O circuits can greatly reduce substrate noise.

2.9 Summary

Having examined the challenges imposed on the design of the ADC due to limitations in process technology, the next two chapters will discuss in detail how each limitation is overcome. High switch resistance at low supply voltages is solved by applying bootstrapping with charge pumps. Charge injection is minimized by using a differential architecture and by directing charge in a controlled manner. Mismatch is reduced by appropriate sizing. Leakage is managed by applying a technique known as power

gating. Process variation in sub-threshold operation is mitigated through sizing and allowing for sufficient delay. Finally, substrate noise coupling is minimized through good analog layout techniques.

Chapter 3

Architecture Design Considerations

This chapter will describe the global and block-level architecture of the proposed ADC. In Section 3.1, the basic operation of a SAR ADC will be presented. In Section 3.2, design issues affecting the global architecture will be presented. This includes the SAR conversion process, scalability and reconfigurability, leakage reduction techniques and power versus performance trade-offs. Finally, a high-level description of each of the constituent blocks of the ADC will be presented in Section 3.3.

3.1 Successive Approximation Conversion Basics

This section will describe the architecture and operation of a conventional SAR ADC [34]. A block diagram of a 4-bit SAR ADC is shown in Figure 3-1. It consists of a sample and hold, a comparator, a binary weighted capacitive DAC and control logic implementing the successive approximation algorithm, which is essentially a binary search algorithm. In practice, the sample and hold function is often combined with the DAC where the input voltage is sampled across the DAC capacitors. Note that the DAC function can be implemented using any DAC topology and is not limited to a capacitive DAC.

Each conversion consists of two phases; the sampling phase where the input is

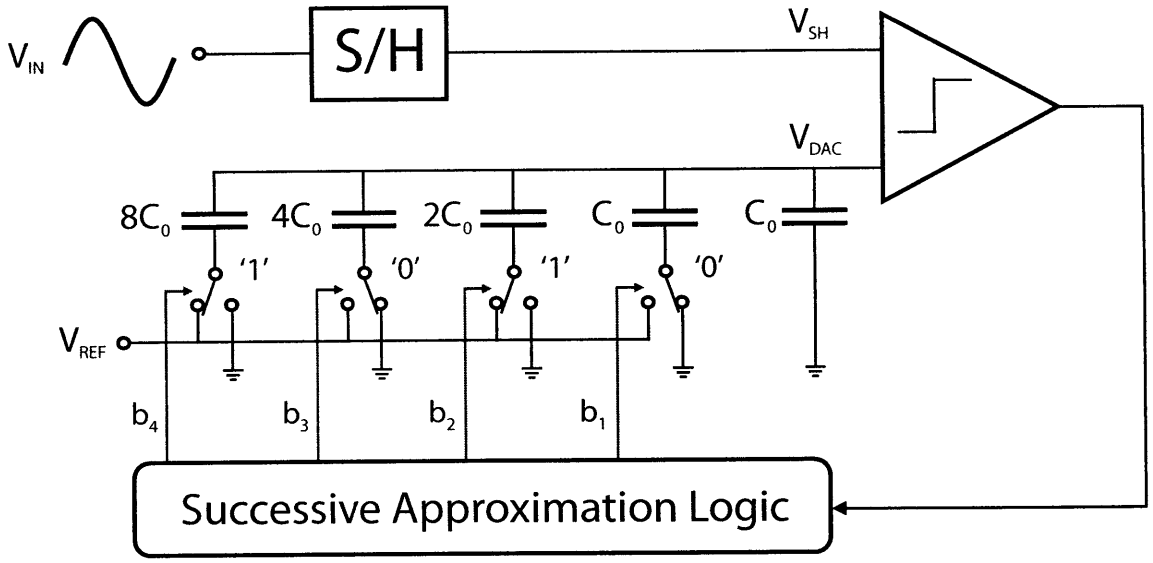

Figure 3-1: Block diagram of a basic 4-bit SAR ADC with a feedback capacitive DAC.

sampled and stored, and the bit cycling phase where the digital output bits are resolved. An example of the bit cycling phase waveforms are shown in Figure 3-2 for an input of V_{IN} = 0.67-V, where the ADC reference voltage is V_{REF} = 1-V.

The successive approximation algorithm works as follows. In the first bit cycle, the MSB b_4 is switched to a '1', thus connecting the MSB capacitor to V_{REF} while all other capacitors remain grounded. This generates a DAC output voltage of V_{DAC} = $0.5V_{REF}$ = 0.5-V, which is compared with the sample and held value of V_{SH} = 0.67-V. Since the input is greater than V_{DAC}, b_4 is set as '1' and the MSB capacitor remains connected to V_{REF} for the next bit cycle. Note that if the input happened to be less than V_{DAC} for this first bit cycle, b_4 would be returned to '0' and the MSB capacitor would be returned to ground.

After b_4 has been resolved, the MSB-1 bit b_3 is then switched to a '1', thus generating a V_{DAC} = $0.75V_{REF}$ = 0.75-V. Since $V_{IN} < V_{DAC}$ for this cycle, b_3 is returned to '0'. Bit cycling continues this way until all bits have been resolved and V_{DAC} has converged to V_{SH}. In the 4-bit example shown in Figure 3-1, the final digital output is $B_{OUT} = b_4b_3b_2b_1$ = 1010 corresponding to the input of 0.67-V.

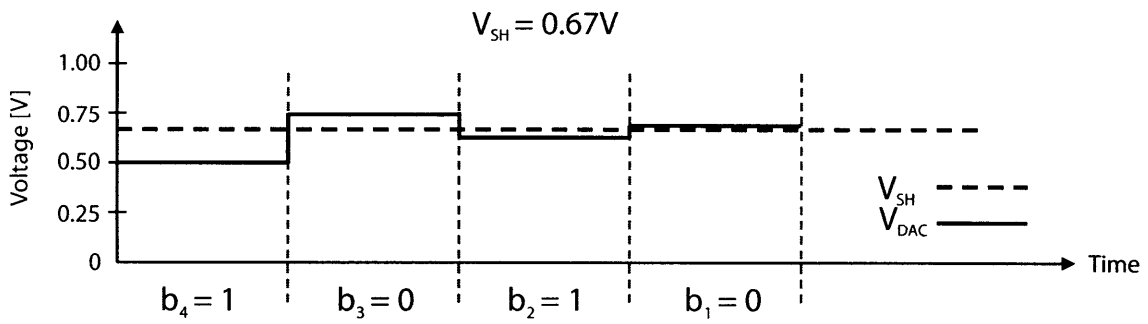

Figure 3-2: Waveforms of the DAC output voltage and the sampled and held input for a complete bit cycling phase.

3.2 Global Architecture Considerations

This section will present the strategies used on a global level to achieve a high degree of scalability and reconfigurability. The possible power versus performance trade-offs enabled by these global approaches will also be analyzed.

3.2.1 SAR Conversion Cycle

In this ADC, each input sample conversion consists of three active phases and an optional *SLEEP* mode. In the first phase, one clock cycle is used to zero the capacitor array and purge it of any charge from the previous cycle. The second phase is the input sampling phase, which can be made programmable between one or two clock cycles. The last phase is the bit cycling phase which requires N_{BITS} clock cycles, where N_{BITS} is the resolution mode of the ADC from 5 to 10-bits. Figure 3-3 illustrates the basic operation of the ADC with the *SLEEP* mode disabled.

Figure 3-4 shows the conversion cycle when *SLEEP* mode is enabled. The purpose and benefit of the *SLEEP* mode is discussed in Sections 3.2.2 and 3.2.5.

3.2.2 Sample Rate Scaling

Since the ADC is designed to be fully dynamic, sample rate scaling is easily achieved by scaling the clock frequency f_{clk}, or by operating at the maximum clock frequency and then duty cycling. The sampling frequency f_{samp} is given by Equation 3.1, where

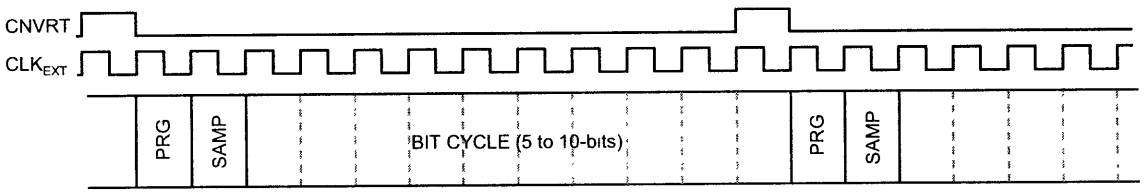

Figure 3-3: SAR ADC conversion plan without SLEEP mode.

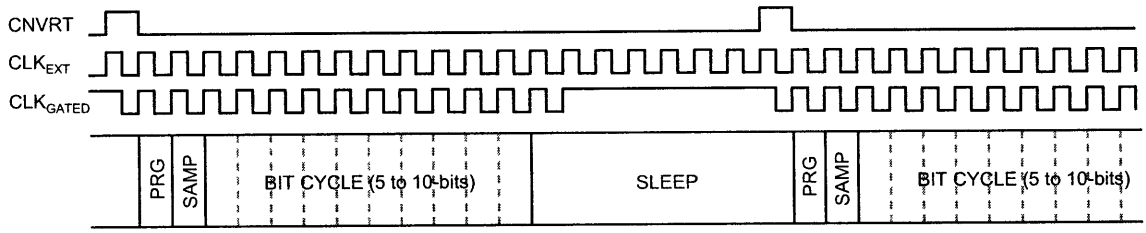

Figure 3-4: SAR ADC conversion plan with SLEEP mode.

N_{SAMP} is the number of cycles used for sampling (1 or 2) and N_{SLEEP} is the number of cycles in $SLEEP$ mode. The denominator is basically the number of clock cycles used per conversion. The additional cycle is due to the capacitor purge phase. For a given f_{clk}, the sample rate can be scaled by adjusting N_{SLEEP}.

$$f_{samp} = \frac{f_{clk}}{N_{SAMP} + N_{BITS} + N_{SLEEP} + 1} \tag{3.1}$$

To achieve a targeted f_{samp}, there are two approaches. The first, where $SLEEP$ mode is disabled ($N_{SLEEP} = 0$), f_{clk} is chosen such that all required phases are completed within each conversion period. The second option is to use a higher f_{clk} to complete all required phases in a shorter amount of time and then put the ADC into $SLEEP$ mode where the external clock CLK_{EXT} is gated to produce CLK_{GATED} as shown in Figure 3-4. In the latter approach, only the required number of clock cycles per conversion are executed, therefore the CV^2 switching energy per conversion is the same as the first approach. However, the benefit of of using a using a higher f_{clk} (and hence a larger N_{SLEEP} for a given f_{samp}), is that leakage reduction techniques can be applied while the ADC is idle (see Section 3.2.5). This is illustrated in Figure 3-5, where the dynamic and leakage components of power during one conversion period

T_{conv} is shown. When $SLEEP$ mode is disabled, the total energy per conversion $E_{conv,nosleep}$ is given by Equation 3.2, where P_{dyn} and $P_{leak-active}$ are the dynamic and active mode leakage powers respectively. When $SLEEP$ mode is used, the total energy per conversion becomes $E_{conv,sleep}$ given in Equation 3.3, where $P_{leak-sleep}$ is the $SLEEP$ mode leakage power, t_{active} is the duration of the active phase and t_{sleep} is the amount of time in $SLEEP$ mode ($T_{conv} = t_{active} + t_{sleep}$). It can be seen that for non-zero t_{sleep}, $E_{conv,sleep} < E_{conv,nosleep}$, assuming the same supply voltage.

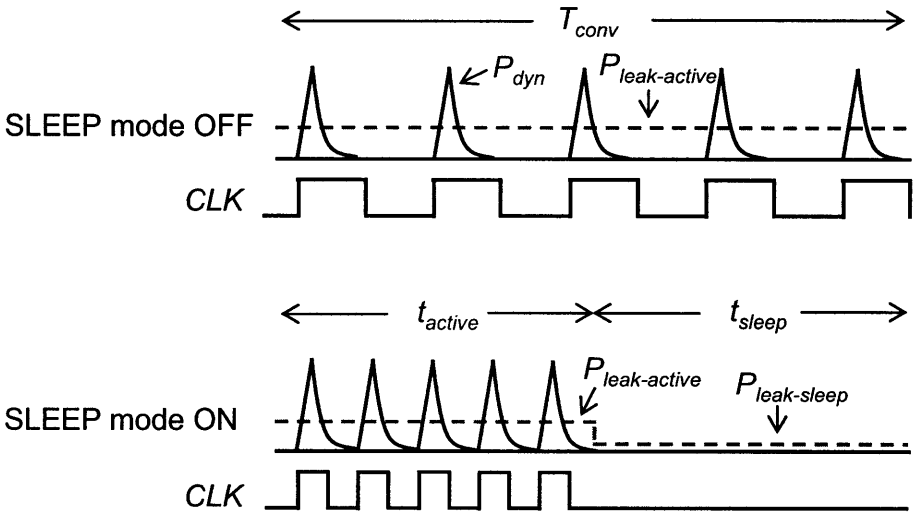

Figure 3-5: Active and leakage components of power during a conversion period with and without $SLEEP$ mode.

$$E_{conv,nosleep} = \int_{<T_{conv}>} P_{dyn}\, dt + \int_{<T_{conv}>} P_{leak-active}\, dt \tag{3.2}$$

$$E_{conv,sleep} = \int_{<T_{conv}>} P_{dyn}\, dt + \left(\int_{<t_{active}>} P_{leak-active}\, dt + \int_{<t_{sleep}>} P_{leak-sleep}\, dt \right) \tag{3.3}$$

3.2.3 Resolution Scaling

For a SAR ADC, the resolution can easily be scaled by cycling just the desired number of bits in a conventional binary weighted N-bit DAC shown in Figure 3-6, where N is the maximum desired resolution. This can be done in two ways. First, it is possible to

just bit cycle starting at the MSB capacitor and stopping when the desired resolution is achieved. The alternative is to start somewhere in the middle of the binary weighted DAC and bit cycle to the LSB capacitor. However, both methods have their own disadvantages. The former method is very energy inefficient because most of the power in the DAC is consumed in bit cycling the largest capacitors [12]. The latter method is also undesirable because the largest MSB capacitors which are not bit cycled become parasitic to the DAC and can greatly attenuate the DAC output. Attenuating the DAC output unnecessarily increases the resolution requirements on the comparator.

The proposed solution is to insert switches between the top plates of the capacitors as shown in Figure 3-7. With these switches in place, the resolution can be scaled by starting bit cycling in the middle of the DAC which saves power, and the MSB capacitors can be decoupled from the DAC by turning off the appropriate switches, thereby not attenuating the DAC output.

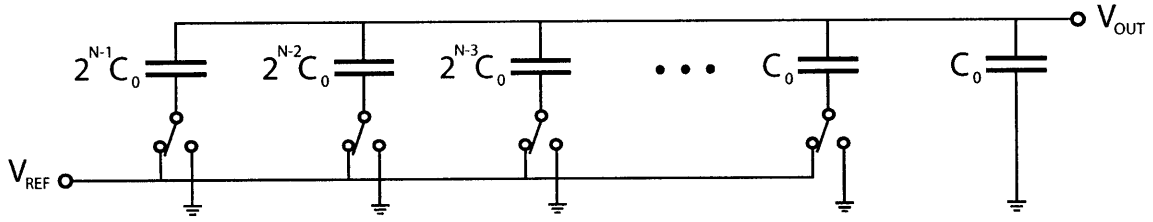

Figure 3-6: A conventional N-bit binary weighted capacitor DAC.

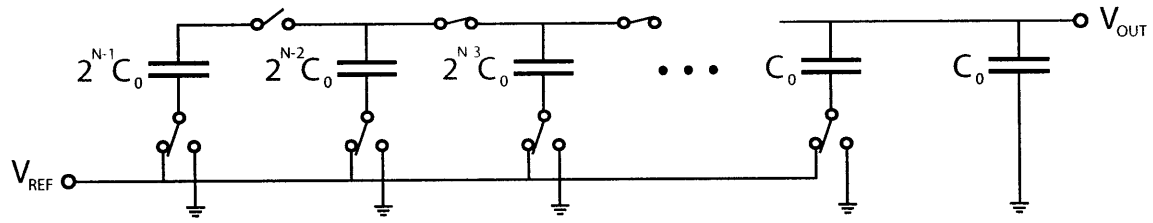

Figure 3-7: A conventional N-bit binary weighted capacitor DAC with switches inserted between top plates to decouple the MSB capacitors as resolution is scaled.

However, since this approach uses analog switches at the DAC output node which is critical for generating precise voltages, care must be taken to ensure that the finite

switch resistance and switch parasitic capacitances do not affect the operation of the DAC.

3.2.4 Voltage Scaling

The previous section described a technique to exponentially scale the DAC power as resolution is reduced. However, due to the binary nature of the SAR algorithm, the power consumed by the rest of the ADC scales only linearly with resolution (since 1-bit of resolution corresponds to one bit cycle). As discussed in Section 1.2, in order to maintain a constant FOM, the ADC power must scale exponentially with resolution. Thus, voltage scaling is used to bridge the gap between linear and exponential power savings.

With respect to the ADC power, let d_L and d_E be the fractions of the total power with linear and exponential power relationships with resolution respectively. Also, let b_H and b_L be arbitrary high and low resolution modes respectively. If $V_{DD,H}$ is the supply voltage used at a resolution of b_H, then the voltage when operating at b_L must be scaled to $V_{DD,L}$ which is given by Equation 3.4 in order to maintain a constant FOM. This relationship is derived in Appendix A.

$$V_{DD,L} = \frac{V_{DD,H}}{\sqrt{d_E + d_L \cdot 2^{(b_H - b_L)} \cdot \frac{b_L}{b_H}}} \tag{3.4}$$

Note that three main assumptions must be satisfied in order for Equation 3.4 to be valid. First, the power that is fixed over resolution (i.e. leakage) represents only a small portion of the total power. This assumption is generally valid at higher sample rates where active power dominates. Secondly, all blocks are assumed to scale quadratically with the supply voltage. This assumption is generally valid for a highly digital architecture such as the one proposed in this thesis. The last assumption is that the fractions d_E and d_L stay relatively fixed over resolution, which can be achieved by turning off digital logic that is not required as resolution is reduced. Practically, these fractions can vary by 10% to 20% over 5 to 6-bits of scalability for reasonably sized DACs. However, Equation 3.4 can still serve as a general guideline

for understanding the factors affecting the effectiveness of voltage scaling.

Figure 3-8 shows simulated results of the ADC power versus resolution with and without voltage scaling and it can be seen that if the voltage is scaled properly, it is possible to maintain constant energy efficiency across the desired resolution range.

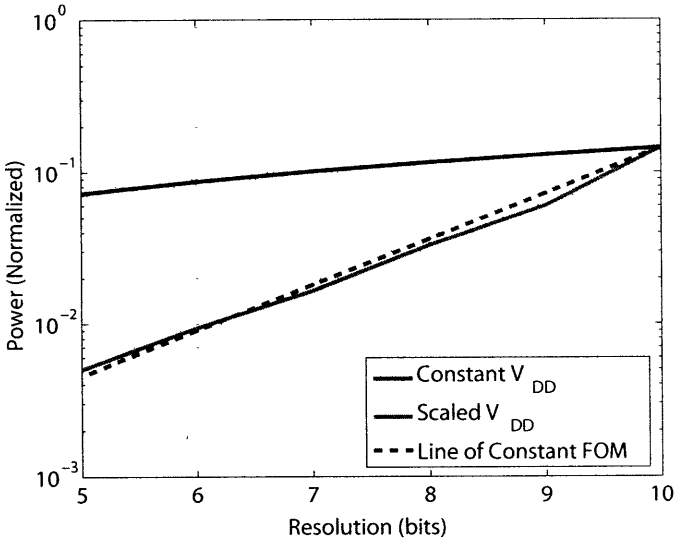

Figure 3-8: ADC power vs. resolution with and without voltage scaling.

It is worth noting that voltage scaling to achieve exponential power savings is only effective for a limited range of resolution, and this is highly dependent on the value of d_E and d_L. Practical values of $V_{DD,L}$ are limited to around 0.2-V to 0.3-V before propagation delay becomes excessive. When $V_{DD,L}$ is around V_T or less, the sample rate is drastically reduced, the first assumption breaks down and leakage will start to degrade the FOM. Thus, leakage reduction techniques will have to be used to manage the leakage component of the overall power and this will be discussed in Section 3.2.5.

Effect of Voltage Scaling on FOM

Up until now, we have not yet considered the effect of reduced voltages on sampling linearity and SNDR. As V_{DD} is reduced, the input range is reduced, while the harmonic distortion increases (due to reduced linearity in the sampling switches). The net effect is that the SNDR, and hence the ENOB, is degraded. Even though the sampled $\frac{kT}{C}$ noise remains constant as voltage is reduced, this does not appreciably

degrade the SNDR since the ADC is not thermal noise limited. Lastly, note that the signal-to-quanization noise ratio (SQNR) remains the same since the quantization noise is directly proportional to the full scale input range as seen by Equation 2.16. Nonetheless, the degradation in ENOB due to added distortion will limit the FOM even though power is being scaled exponentially.

3.2.5 Power Gating for Leakage Management

Power gating is a common technique used in digital circuit design to minimize leakage during periods of inactivity [35], [36]. This technique is also known as multi-threshold CMOS (MTCMOS) because a high-V_T (HVT) sleep transistor is placed between actual ground and the virtual ground as shown in Figure 3-9. During normal operation, the HVT sleep transistor is turned on and the virtual ground node is pulled down to ground and the circuits (which are implemented using low-V_T (LVT) devices for increased performance) operate as usual. However, when the ADC is put in $SLEEP$ mode, the HVT sleep transistor is turned off and this helps suppress leakage current in several ways. First, the HVT device can reduce the leakage by several orders of magnitude due to the exponential dependence of leakage current on V_T in sub-threshold operation. Secondly, by placing the HVT device in series with the rest of the digital circuits, the stack effect further reduces leakage [37]. Lastly, since the virtual ground will tend to drift up to V_{DD}, the effective V_T of the LVT devices increases due to the body effect.

However, this technique does not come without penalty and overhead. First, there is overhead power associated with switching the parasitic gate capacitance C_{GATE} of the HVT sleep transistor which is typically sized very large to ensure that the virtual ground node is as close to ground as possible. Secondly, there is a recovery energy associated with discharging the virtual ground node back to ground after each idle period. This implies that there is a break even time in which power gating becomes beneficial. A thorough analysis of the overhead associated with power gating and the break even time can be found in [38]. Other issues to be aware of include the wake up time associated with recovering the virtual ground node and the performance degra-

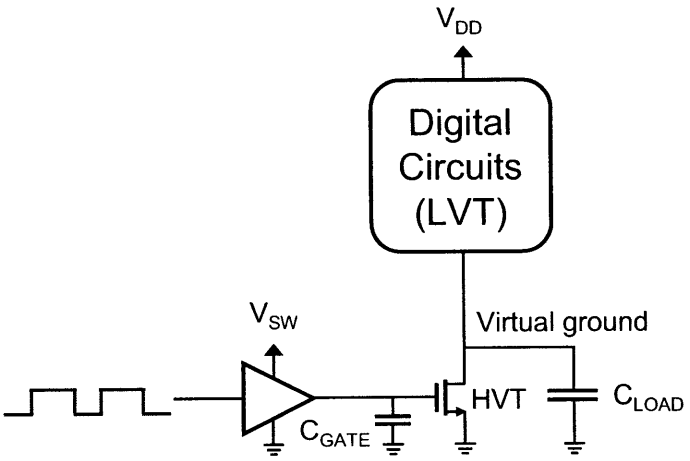

Figure 3-9: Schematic showing the digital circuits being power gated with a HVT footer and the associated parasitic capacitance which affects the break even time.

dation due to any virtual ground bounce. The design of the power gating circuitry will be presented in Section 4.4.1.

3.2.6 Sampling Duration

The sampling network was designed to settle to the desired accuracy within one clock cycle. However, since aggressive voltage scaling is applied in this design, the sampling phase was made programmable between one or two clock cycles to accomodate increased device variation at low voltages.

3.2.7 Power, Area and Performance Trade-offs

The power and area consumption of the ADC is largely determined by linearity, bandwidth, robustness and scalability requirements. This section will briefly discuss how power and area can be traded off.

Power versus Linearity

The degradation in linearity of the ADC due to capacitor mismatch can virtually be eliminated by sizing the capacitors very large. However, the DAC power and area scales directly with capacitance. Also, the power of the ADC driver also scales with

this capacitance. Therefore, it is important to size the unit capacitor in the DAC to just achieve the desired linearity. This will be discussed in detail in Section 3.3.1.

Furthermore, since aggressive voltage scaling is used in this design, it is important to understand what limits the degree of supply voltage scaling. Generally, the minimum supply voltage must satisfy both the bandwidth and linearity requirements. However, in higher bandwidth applications, the minimum supply voltage is typically set by bandwidth, while in ultra low bandwidth applications, it is set by the linearity requirement. In this manner, the linearity requirement sets the power efficiency of the ADC.

Single-Ended versus Differential Conversion

A fully differential architecture provides benefits such as better power supply and common mode rejection, as well as elimination of second order harmonics (assuming good layout practices are followed). However, this comes at the expense of doubling the power consumed by the DAC. In this ADC, with the goal of being as reconfigurable as possible, both single-ended and differential operation are supported. When power is the primary constraint and the aforementioned issues can be tolerated, then single-ended operation can provide a means to achieve lower power operation. However, note that in this implementation, an extra common mode reference voltage V_{CM} is required for single-ended operation. When the ADC is configured in differential mode, no common mode reference is required since it is established passively as will be discussed in Section 4.1.3.

3.3 Block Architecture

In this section, the high level considerations for each block of the ADC is presented.

3.3.1 Resolution Scalable DAC

The primary function of the DAC is to generate analog voltages at the transition voltages of the ADC. Therefore, the DAC must provide the linearity required by the

ADC. As described in Section 3.2.3, it must also be resolution scalable. In addition to these requirements, it must be area and energy efficient, and also be able to support both single-ended and differential modes.

In this design, a binary-weighted capacitive DAC architecture was chosen because it is generally area efficient and does not consume any static current. In particular, the concept of a split-capacitor array [39] is combined with a sub-DAC [40] to reduce the power and area consumption of the DAC respectively. Boosted switches are used to implement the resolution scalability and to reconfigure between differential and single-ended modes. Circuit details are provided in Section 4.1.

Lastly, with the capacitors in the given process technology, sizing for 3-σ matching at the 10-bit level would result in unit capacitors that are prohibitively large. Therefore, a design decision was made to size the unit capacitors of the DAC for 3-σ matching at the 8-bit level. With the same sized capacitors, this approximately results in 2-σ and 1-σ matching at the 9 and 10-bit levels respectively. At resolutions below 8-bits, capacitor mismatch should not be an issue. More details on capacitor matching can be found in Section 4.1.

3.3.2 Latched Comparator

In a SAR ADC, the comparator is required to make a decision based on the DAC output which controls the binary search algorithm. In order to interface with the digital state machine, a full swing digital latch is used. Usually, a linear preamplifier is used to amplify the analog inputs above the offset floor of the latch, as well as reduce the effects of latch kickback. However, in this design, the preamplifier was eliminated to allow the use of aggressive voltage scaling. Since dynamic latched comparators typically have large input referred offsets on the order of 10's to 100's of mV, appropriate input device sizing and offset compensation techniques must be employed to avoid losing a significant portion of the ADC output range.

3.3.3 SAR Control Logic

The SAR digital state machine performs several functions in this ADC. First, based on an external signal to begin each conversion, it generates the appropriate purge and sample control signals for the DAC. Next, based on the decision from the comparator, the digital state machine generates the control signals to bit cycle the DAC. Aside from these core functions, the digital logic also implements control for power gating, clock gating and resolution scaling. A custom digital library was designed and the logic was designed and routed manually.

3.4 Summary

In this chapter, the basic operation of a SAR ADC was described. Global architecture considerations such as the SAR conversion plan and scaling of the sample rate, resolution and voltage were also discussed. The concept of power gating was introduced, and performance tradeoffs were considered with respect to power and area. Lastly, the functions of the constituent blocks of the ADC were presented.

Chapter 4

Circuit Design Details

This chapter will describe the implemented analog and digital circuits at the transistor and logic level respectively. The main design objectives for each circuit will be presented and the strategies used to achieve the objectives will be described. Section 4.1 will present the structure of the DAC which is crucial in determining the static and dynamic performance of the overall ADC. Section 4.2 will describe the design of the switch network used to switch the DAC capacitors during the different phases of operation. Section 4.3 will present the comparator and finally, Section 4.4 will describe the digital state machine and discuss the overhead that reconfigurability requires.

4.1 DAC Circuit Design

As discussed in Section 3.3.1, the primary design objectives for the DAC is energy efficiency, adequate linearity for the targeted resolutions, resolution scalability and differential and single-ended support.

In order to size the DAC capacitors for adequate matching, the scenario for the worst case differential non-linearity (DNL) is considered. The DNL (in units of LSB) for each code transition is defined as the deviation from the an ideal LSB, $DNL_i = \Delta V_i - 1$, where ΔV_i is the transition code width for the i-th code (in units of LSB). For an N-bit binary weighted DAC, this occurs at the $i = 2^{N-1}$ transition (i.e. the

MSB transition) from code $2^{N-1}-1$ to 2^{N-1}, where the MSB capacitor is skewed by a fraction $\pm\delta$, while the rest of the capacitors are skewed by $\mp\delta$. This translates to a worst case DNL given by Equation 4.1.

$$DNL_{max} = \Delta V_{2^{N-1}} - 1 = (1 \pm \delta)\, 2^{N-1} - (1 \mp \delta)\left(2^{N-1} - 1\right) - 1 \tag{4.1}$$

Rearranging Equation 4.1, we arrive at Equation 4.2 for δ, which was used to appropriately size the DAC unit capacitor based on the required maximum DNL, resolution N and matching data from the vendor. In this ADC, a unit capacitance of $C_0 = 85.4$-fF was adequate for the methodology outlined in Section 3.3.1.

$$\delta = \frac{DNL_{max}}{2^N - 1} \tag{4.2}$$

The remainder of this section will describe the circuit techniques used to achieve the other design goals of the DAC.

4.1.1 Split Capacitor Array

For a binary weighted capacitive DAC, there exists many ways to switch the capacitors between the reference voltage V_{REF} and ground. Common methods discussed in [39] include 1-step or 2-step switching, charge sharing or capacitor splitting. For all four methods, the energy required to perform an UP transition (i.e. switching a capacitor from ground to V_{REF}) is identical. For an N-bit DAC, generating the last output code $2^N - 1$ requires that all the capacitors be switched in an UP transition to V_{REF}. Following the analysis shown in [39], it can be shown that the total energy required to complete the N UP transitions is given by $E_{tot,UP}$ in Equation 4.3 which applies to all four methods.

$$E_{tot,UP} = \frac{2^N}{3}\left[1 - \left(\frac{1}{4}\right)^N\right] C_0 V_{REF}^2 \approx \frac{2^N}{3} C_0 V_{REF}^2 \tag{4.3}$$

However, the energy for a DOWN transition differs for each method. For this reason, the capacitor splitting approach is chosen since it requires the least amount

of energy for a DOWN transition. Details on the implementation of a split capacitor array can be found in [39], but a brief overview is provided here. To create a split capacitor array, the MSB capacitor of a conventional binary weighted array is *split* into a separate binary weighted array (herein called the MSB array), which is identical in structure to the rest of the main capacitor array (herein called the LSB array). A 3-bit example is shown in Figure 4-1(b).

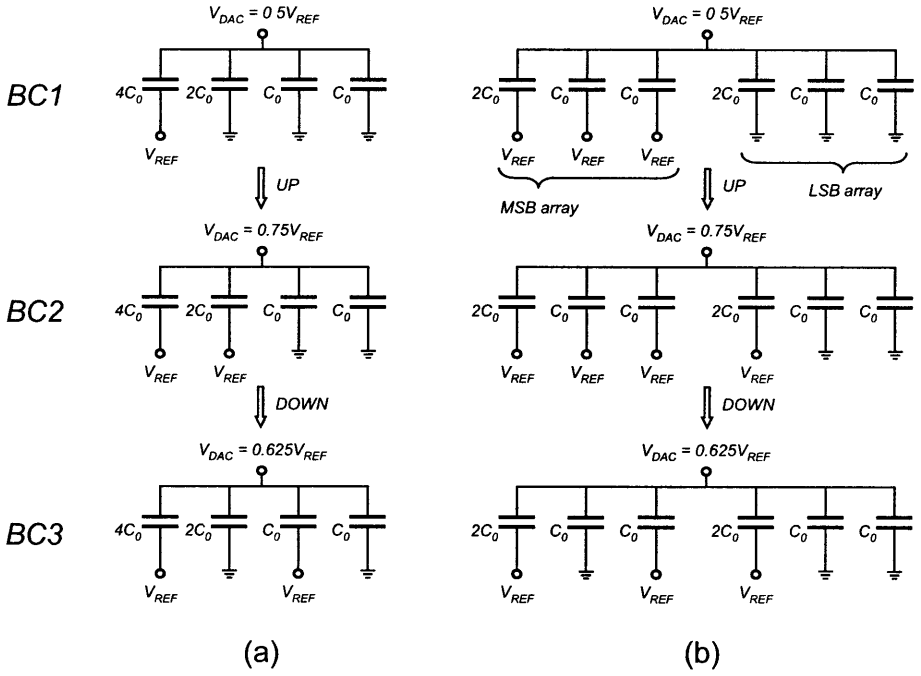

Figure 4-1: Comparison of bit cycling with (a) a conventional binary weighted array (b) a split capacitor array.

For the first bit cycle, the entire MSB array is charged to V_{REF}, identical to the conventional array. For subsequent bit cycles, any UP transition is accomplished by charging a capacitor to V_{REF} in the LSB array, while any DOWN transition is accomplished by switching a capacitor in the MSB array to ground. The operation of a split capacitor array vs. 1-step switching of a conventional array is illustrated in Figure 4-1. The energy required for the first two bit cycles $BC1$ and $BC2$ for both methods is the same since they are both UP transitions. However, the difference can be seen in the third bit cycle $BC3$ which is a DOWN transition. In the case of the conventional array, energy is consumed in driving the change in V_{DAC} as well

as charging the bit cycled capacitor to V_{REF}. In contrast, with the split capacitor array, energy is only required in driving the change in V_{DAC}, and no energy is spent on charging any capacitors to V_{REF} during a DOWN transition.

The average energy savings of the split capacitor array when compared to the conventional array is 37% [39]. Another important benefit of the split capacitor array topology is that the switching energy is relatively constant over all output codes, whereas the switching energy for the conventional array increases as the output code decreases. For the split capacitor array, this means that the overall ADC power is roughly independent of the input signal which may be desirable in many applications.

4.1.2 Sub DAC Interpolation

A major limitation of an N-bit binary weighted capacitor array is that the ratio of the MSB capacitor to the unit capacitor is 2^{N-1}. For moderate to large N, this results in excessively large area consumption. In order to reduce the area of the capacitor array, a separate capacitor array known as a sub-DAC can be used [40].

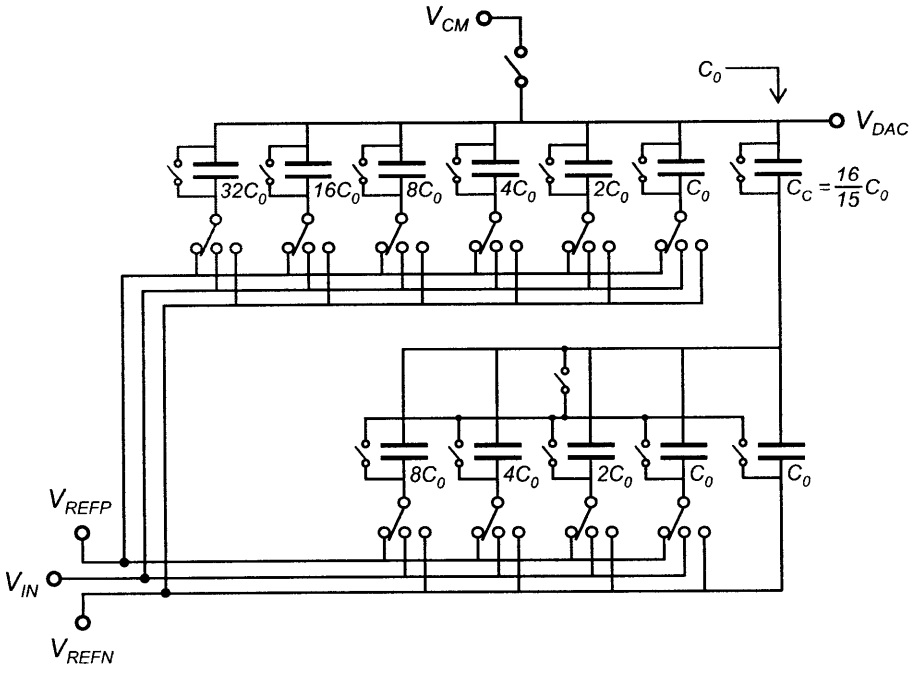

Figure 4-2: Schematic of a 10-bit DAC split into a 6-bit main-DAC and a 4-bit sub-DAC.

Resolution	Conventional DAC	Main DAC with 4-bit Sub-DAC	Area Savings
10	$1024C_0$	$80\frac{1}{15}C_0$	12.8×
9	$512C_0$	$48\frac{1}{15}C_0$	10.7×
8	$256C_0$	$32\frac{1}{15}C_0$	8.0×
7	$128C_0$	$24\frac{1}{15}C_0$	5.3×
6	$64C_0$	$20\frac{1}{15}C_0$	3.2×
5	$32C_0$	$18\frac{1}{15}C_0$	1.8×

Table 4.1: Total DAC capacitance versus resolution for a conventional binary weighted DAC and a main DAC with a 4-bit sub-DAC.

An example of a 10-bit DAC employing a 4-bit sub-DAC is shown in Figure 4-2. Essentially, the main DAC consists of the first 6 MSB capacitors and the series combination of the coupling capacitance C_C and the 4-bit sub-DAC. Through appropriate sizing of C_C ($C_C = \frac{2^4}{2^4-1}C_0 = \frac{16}{15}C_0$ for a 4-bit sub-DAC), the series combination looks like C_0. Therefore, the first 6 MSBs generate the 64 main DAC transitions, which must be 10-bit linear. To achieve 10-bit resolution, the sub-DAC interpolates between each main DAC transition, resulting in $1024 = 64 \times 16$ codes. The area savings from using a 4-bit sub-DAC can be seen in Table 4.1. Here, as resolution is scaled from 10-bits, it is assumed that the MSB capacitor is removed. Rationale behind using 4-bits in the sub-DAC is provided in Section 4.1.3.

Note that the sub-DAC can be implemented with a variety of DAC topologies, but the sub-DAC in this prototype is fully passive which eliminates the need for any active amplifiers. However, the sub-DAC interpolation is sensitive to the top plate parasitics C_P of the sub-DAC as shown in Figure 4-3. This has the effect of compressing the transition voltages of the sub-DAC, thus creating errors in the interpolation leading to integral non-linearity (INL) and DNL errors with a period equal to the sub-DAC interpolation range.

Fortunately, this error can be corrected by compensating the sub-DAC transmission gain as outlined in [41]. In effect, the value of C_C should be increased by the factor R_{TP} given in Equation 4.4, where C_{subDAC} is the total sub-DAC capacitance.

$$R_{TP} = 1 + \frac{C_P}{C_{subDAC}} \tag{4.4}$$

By increasing C_C, the capacitance looking into C_C is no longer C_0, which leads to gain error in the ADC. However, the sub-DAC interpolation is corrected and INL and DNL errors are minimized. Note that this technique requires accurate extraction of the top plate parasitics from layout. More discussion of linearity errors due to top plate parasitics is provided in Section 5.4.1. For the 4-bit sub-DAC implemented in this design, $C_{subDAC} = 2^4 C_0 = 1366.4$-fF, while C_P was extracted to be 396-fF.

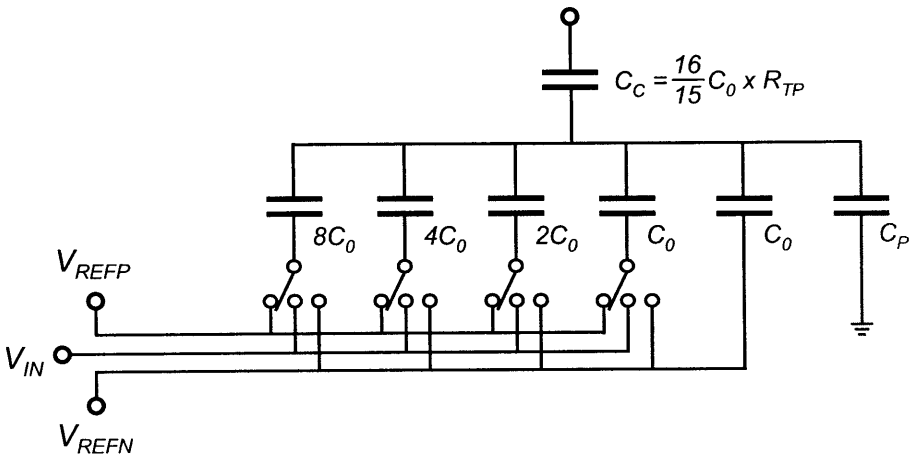

Figure 4-3: Schematic of the 4-bit sub-DAC with top plate parasitic and the required coupling capacitance adjustment.

4.1.3 DAC Schematic

In order to control many of the switches to reconfigure the DAC, three configuration bits $RM[2:0]$ were used to set the resolution. Table 4.2 shows the derivatives of the $RM[2:0]$ bits in thermometer and one-hot encoding which were used throughout the ADC.

The schematic of the fully differential, resolution scalable DAC used in the ADC is shown in Figure 4-4. The core DAC consists of a split capacitor array combined with a sub-DAC. The MSB capacitor (highlighted in gray) of the DAC shown in Figure 4-2 is split into a 9-bit MSB array (also highlighted in gray), of which 4-bits are part of a sub-DAC. This arrangement of the MSB array matches the structure of the LSB array of the DAC. In Figure 4-4, the i-th bit capacitor of the MSB and LSB arrays

Resolution	$RM[2:0]$	$RM_{TH}[4:0]$	$RM_{OH}[5:0]$
5	000	00000	000001
6	001	00001	000010
7	010	00011	000100
8	011	00111	001000
9	100	01111	010000
10	101	11111	100000

Table 4.2: Resolution scaling logic used to control the resolution mode in the SAR state machine as well as the DAC switches.

for i from 5 to 9 are drawn in pairs with top plate switches in between each pair. This group of capacitors constitutes the main DAC. The other bit capacitors for i from 1 to 4 for both MSB and LSB arrays make up the sub-DACs.

The top plate switches controlled by $RM_{TH}[4:0]$ are CMOS transmission gates, with boosted NMOS devices. These switches are used to decouple the MSB capacitors as resolution is scaled. For example, in 7-bit mode, $RM_{TH}[4:0] = 00011$, turning off the three largest bit cycling capacitors as shown in Figure 4-5. Note that as resolution is scaled, capacitors from both the MSB and LSB array are removed, reducing the size of the main DAC. However, the sub-DACs remain at 4-bits in resolution which sets the minimum resolution of the overall DAC at 5-bits as desired. In 5-bit mode, all capacitors in the main DAC are turned off and only the sub-DACs are used in the conversion. Note that the top plate switches were sized to have very small ON resistance so as to not introduce any additional RC delay in the DAC. Sufficient time was allotted between bit cycles to allow the charge to redistribute through those switches.

During the sampling phase, the top plates of both the positive and negative DACs are shorted together with boosted NMOS switches controlled by $TPSAMP[5:0]$. This allows the top plates to passively settle to the input common mode, provided that both arrays are initially purged of all charge. This provides good common mode rejection and eliminates the need for a mid rail reference voltage. The DAC is purged with switches across each capacitor before the sampling phase. In order to support single-ended conversion, the negative DAC can be turned off and the top plate of the positive DAC can be connected to a mid rail reference voltage V_{CM} during the

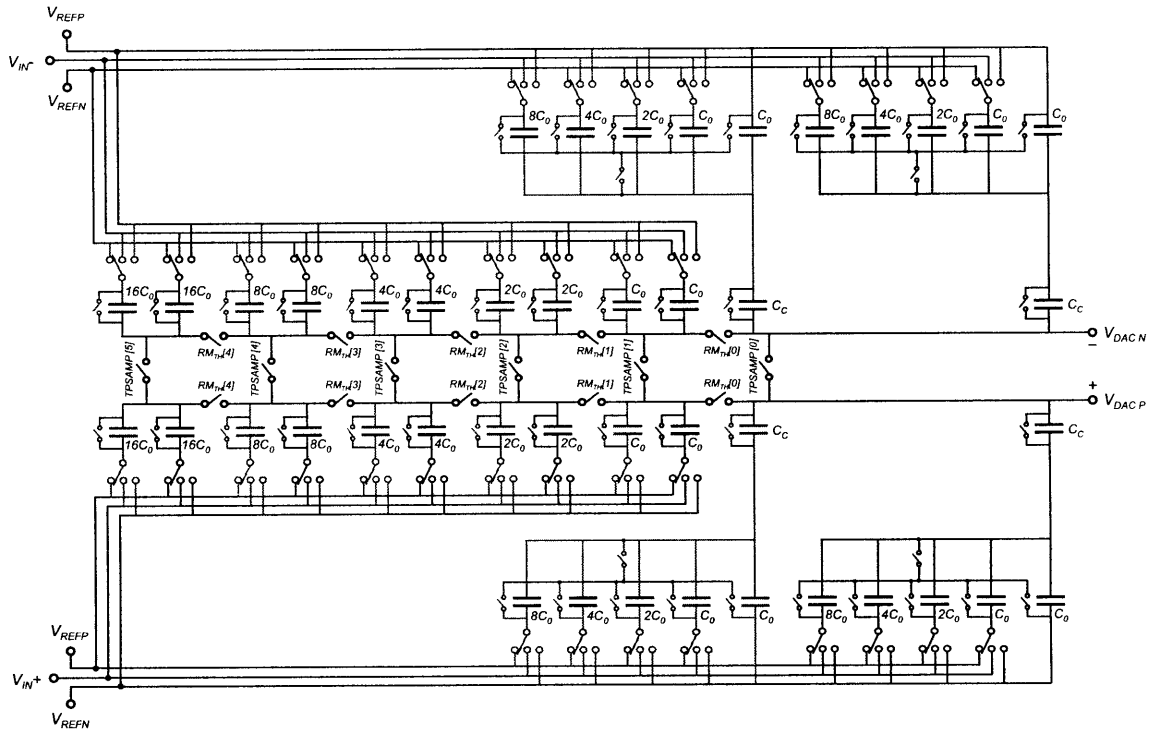

Figure 4-4: Schematic of implemented 5 to 10-bit reconfigurable, fully differential DAC.

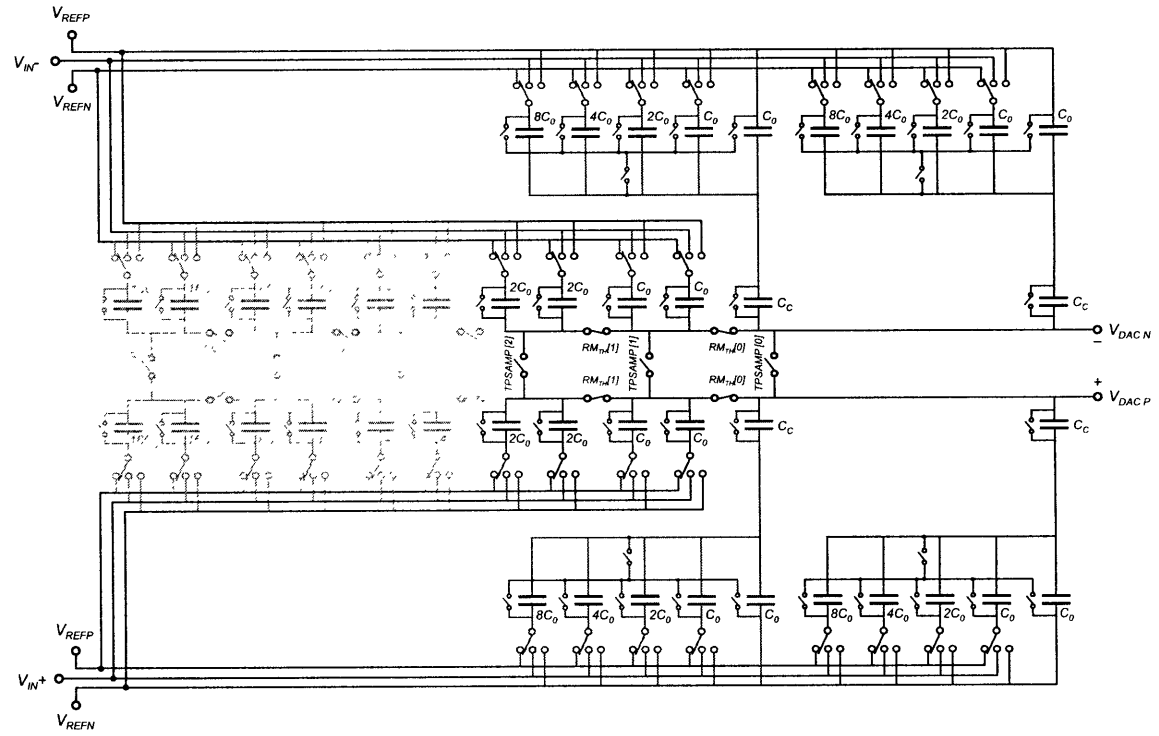

Figure 4-5: Schematic of fully differential DAC configured in 7-bit mode.

sampling phase. This is shown in Figure 4-6.

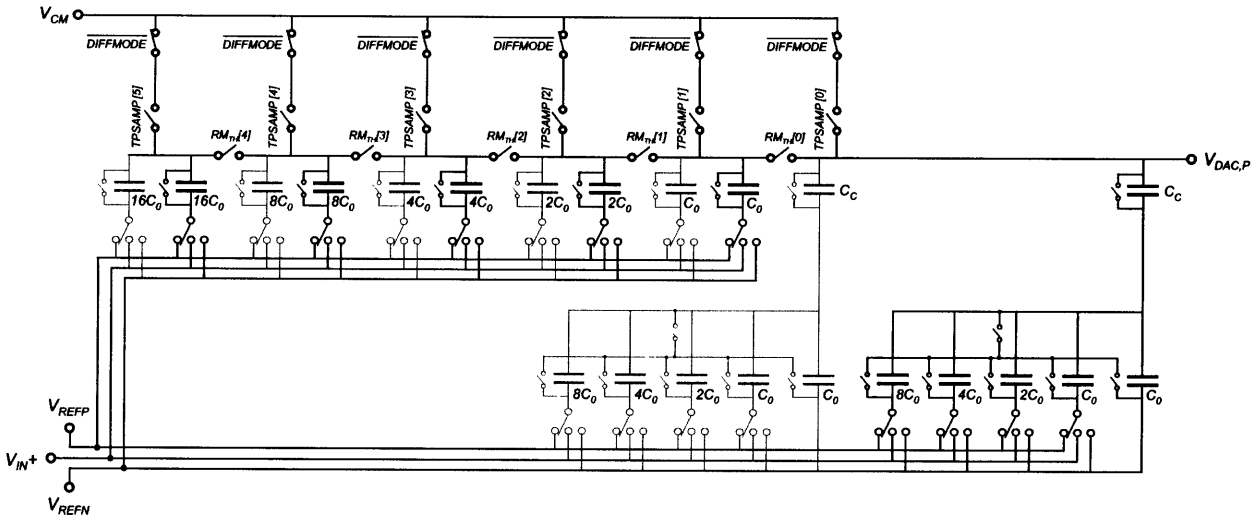

Figure 4-6: Schematic of positive DAC reconfigured in single-ended mode with the top plate connected to V_{CM} during the sampling phase.

4.1.4 DAC Layout

Good layout of the capacitor array is crucial to ensure good matching and to minimize edge effects and non-linearities resulting from coupling between top and bottom plates. Common centroid layout was used to mitigate against first order gradient effects. To ensure a uniform processing environment for the capacitors at the edge of the array, a border of grounded dummy capacitors was used. However, this may introduce some mismatch due to edge effects, since capacitors in the center of the array couple to neighboring capacitors which are not at ground, while those on the edge couple to the grounded dummy capacitors.

To minimize this source of mismatch, equal-edge ratio common centroid layout was used [41]. This technique attempts to match the ratio of edges to area for the largest capacitors in the array, since they introduce the largest linearity errors. The layout of the main DAC is shown in Figure 4-7, where the active capacitors are highlighted for each resolution mode. The edge to area ratios for each resolution setting is listed in Table 4.3. The layout of the sub-DAC following the equal-edge

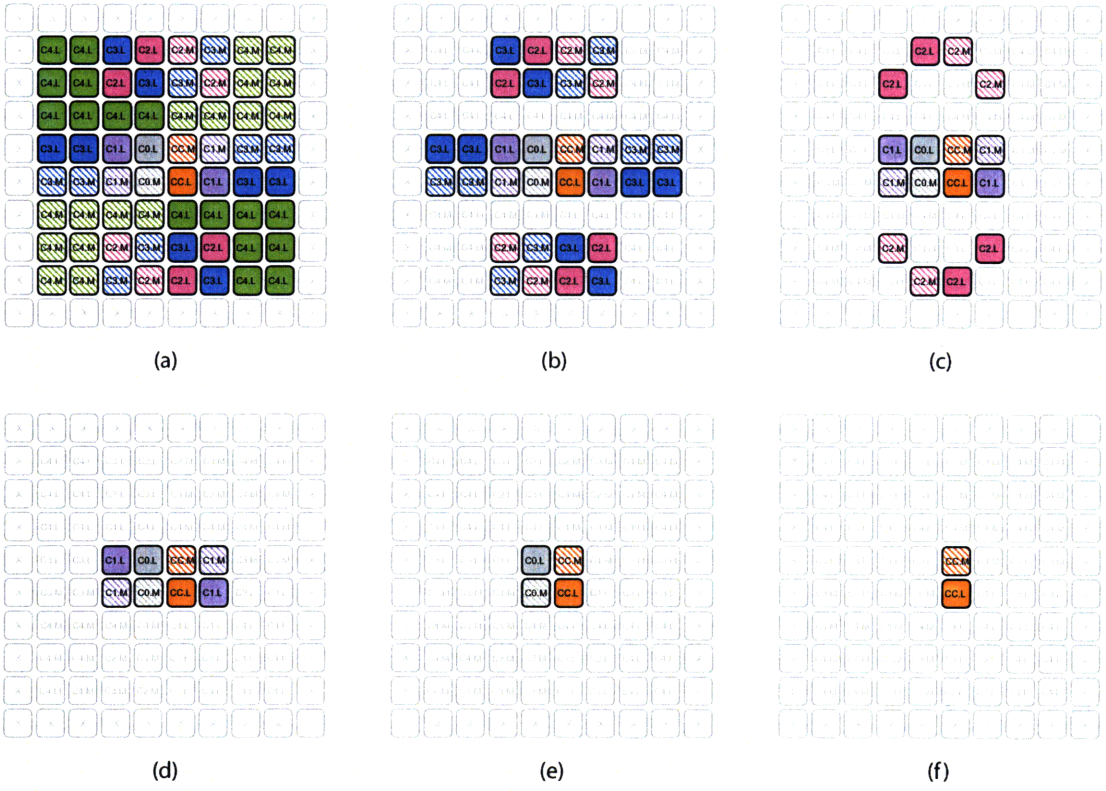

Figure 4-7: Layout of the main DAC in common centroid arrangement. 10-bit to 5-bit configuration is shown in (a) to (f) respectively.

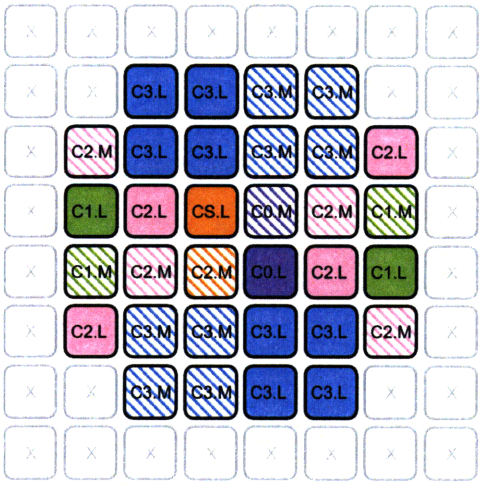

Figure 4-8: Layout of the sub-DAC in common centroid arrangement with minimized edge effects.

Resolution	$16C_0$	$8C_0$	$4C_0$	$2C_0$	C_0	C_C
10	0.625	0.5	0.5	0	0	0
9	N/A	1.5	1.5	1	1	1
8	N/A	N/A	3.5	2	1	1
7	N/A	N/A	N/A	2	1	1
6	N/A	N/A	N/A	N/A	2	2
5	N/A	N/A	N/A	N/A	N/A	3

Table 4.3: Edge to area ratio for the main DAC capacitors as resolution is scaled.

ratio common centroid technique is also provided in Figure 4-8. Finally, to minimize coupling between top and bottom plate routing, an electrostatic shield was placed in between.

4.2 Switch Design

The bottom plate switch network used to control the purge, sampling and bit cycling phases of the positive DAC are shown in Figure 4-9. The switches for the negative DAC are identical except that the control signals are opposite in polarity.

Due to the reconfigurability of the DAC and the use of a split capacitor array, the control signals for the main and sub-DACs in the MSB and LSB array are all uniquely derived. Details on the control logic generating these signals can be found in Section 4.4.

4.2.1 Sampling Switches

The equivalent input sampling network showing the sampling and top plate switch resistances R_{SAMP} and R_{TP} respectively, is shown in Figure 4-10. The THD of the $V_{IN,SAMP}$ signal was designed to be at least 10-dB better than the required SNDR at a specified resolution. From the RC circuit shown in Figure 4-10, the time constant τ was calculated and the time required for sufficient settling t_{settle} for a specified resolution is given by Equation 4.5, where x is the percentage settling required. For 10-bit settling (better than 1 part out of 1024), x is approximately 99.9% and $t_{settle} = 6.9\tau$.

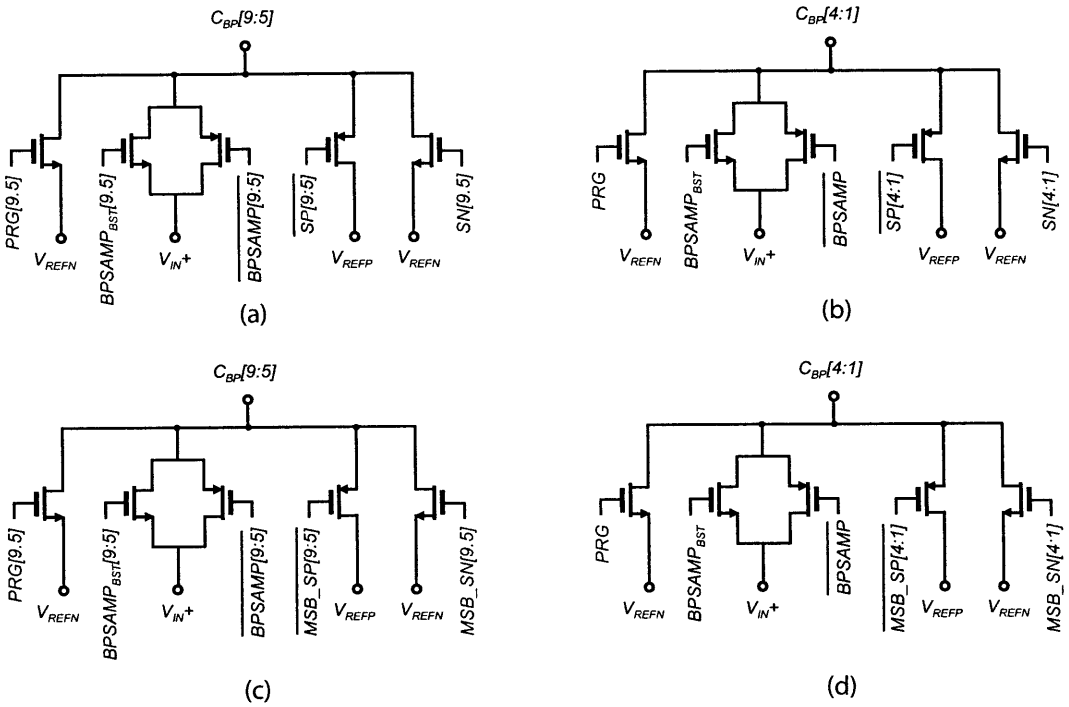

Figure 4-9: Schematic of DAC switches for the positive DAC. (a) and (b) show the switches for the main and sub-DAC respectively in the LSB array, while (c) and (d) show the switches for the main and sub-DAC respectively in the MSB array.

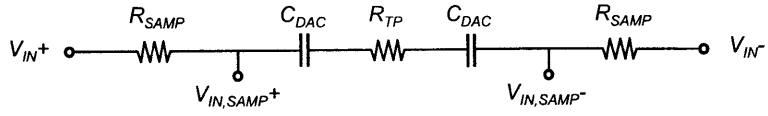

Figure 4-10: Equivalent input network showing the non-zero switch resistances of the sampling switches.

$$t_{settle} = \tau \ln\left(\frac{1}{1-x}\right) \tag{4.5}$$

Charge Pump Circuit

To boost the input sampling switches, the charge pump circuit shown in Figure 4-11 from [42] was used. During normal operation ($BYPASS = 0$), the voltage across the boosting capacitor C_{BST} is charged to V_{DD} when the input is 0. When the input switches high, $MN1$ and $MP1$ turn off while $MP2$ turns on, and node V_X is boosted above V_{DD} since C_{BST} has a voltage of V_{DD} across it. The output logic high is now the boosted V_X. In practice, due to charge sharing between C_{BST} and parasitic capacitance, the boosted output is approximately $1.5V_{DD}$.

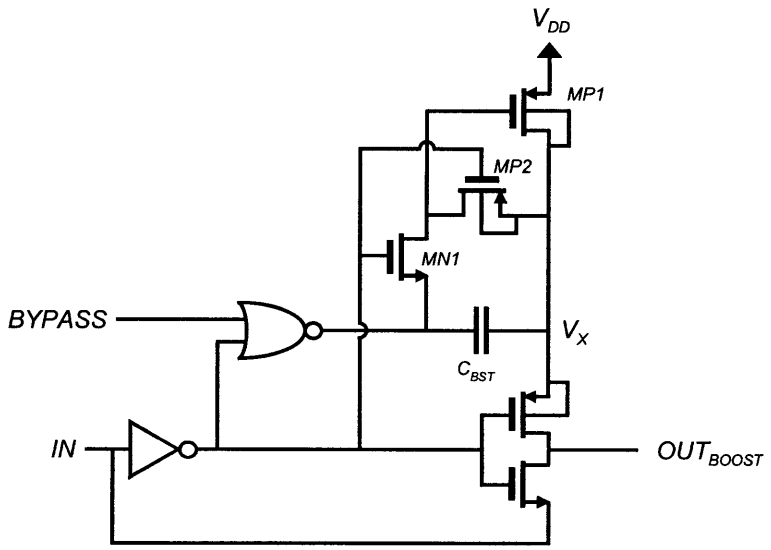

Figure 4-11: Schematic of the charge pump used to boost the input sampling switches.

To disable the charge pump, the $BYPASS$ signal is set to 1, keeping one plate of C_{BST} always at ground, so V_X never rises above V_{DD}. Note that the n-well of all PMOS devices should be connected to V_X to prevent forward biasing any junctions.

4.2.2 Reference Switches

The reference switches used to bit cycle the DAC capacitors were sized such that the required settling time given by Equation 4.5 is less than the shortest intended clock period.

4.3 Comparator Circuit Design

In this section, the comparator circuit design details are presented.

4.3.1 Comparator Description

The schematic of the latched comparator used in this ADC is shown in Figure 4-12 and waveforms from a transient simulation is provided in Figure 4-13. The comparator resets when CLK is high and all internal nodes are grounded, eliminating any memory effects from the previous cycle. When CLK switches low turning on $MP0$, the input devices $MP3$ and $MP4$ will start to pull the two branches up at different rates determined by the inputs $V_{IN}+$ and $V_{IN}-$. The cross coupled inverters made up of $MN1$, $MN2$, $MP1$ and $MP2$ will start to regenerate due to positive feedback, and latch when V_M+ and V_M- reach the inverter threshold. Output buffers with a high threshold (strong PMOS) are used to prevent output glitching while the comparator is latching.

Regenerative amplifiers such as this comparator are very power efficient and are functional at low voltages, however they are prone to large input referred offsets. To accomodate large offsets, 4-bit capacitor banks C_{LoadP} and C_{LoadN} are used to intentionally compensate for any random device mismatch due to process variation.

Another issue with latched comparators is the kickback noise at the input due to capacitive coupling from the outputs when they start to regenerate. However, kickback has little effect on the operation of the ADC since the noise is roughly common mode when the output nodes begin rising. Once the outputs latch to opposite rails, the comparator has already made its decision.

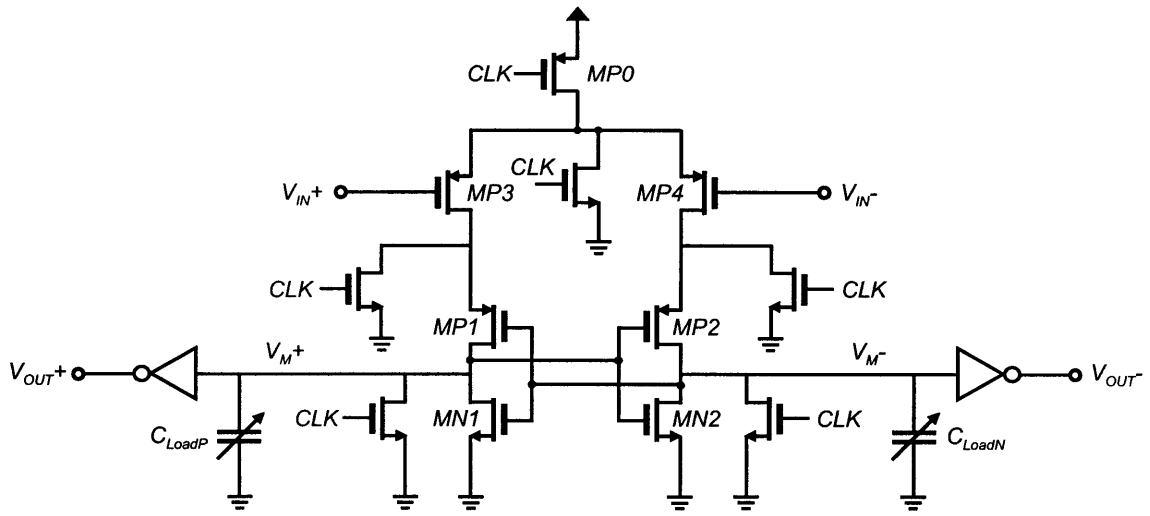

Figure 4-12: Schematic of dynamic latched comparator used in the ADC.

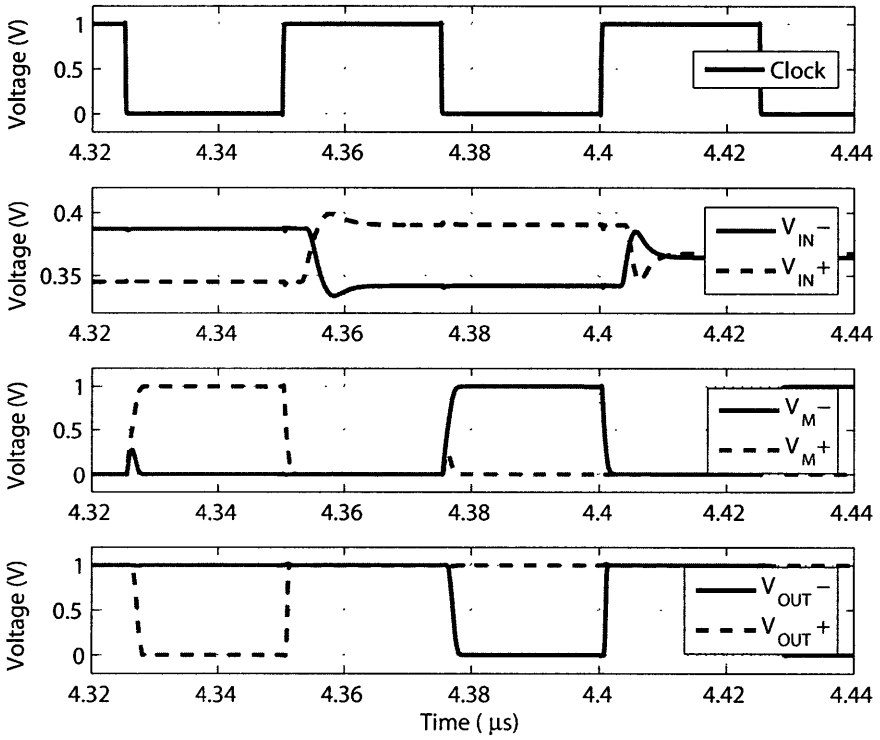

Figure 4-13: Latched comparator transient waveforms.

Finally, since this ADC supports both single-ended and differential modes, two separate comparator paths shown in Figure 4-14 are used to avoid using additional switches in the signal path. Since the positive DAC is used for single-ended conversion, a dummy load is applied to $V_{DAC,N}$ to maintain good differential matching.

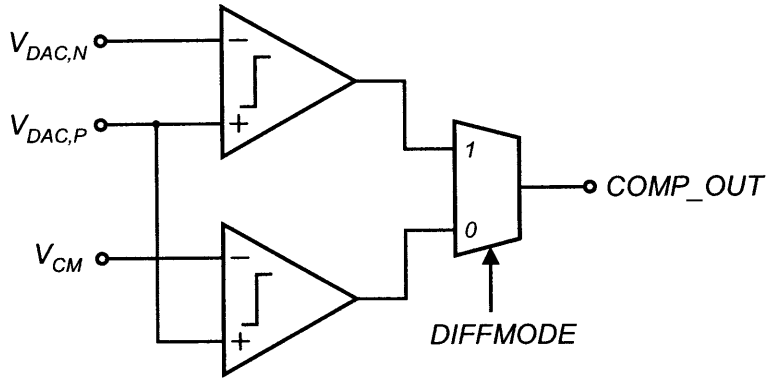

Figure 4-14: Comparator selection schematic for differential or single-ended mode.

4.3.2 Comparator Performance Optimization

The size of the input devices, cross coupled inverters and tail transistor $MP0$ were all sized for optimal performance in terms of power, delay and offset. Figures 4-15 and 4-16 plots the comparator power and delay vs. transistor sizing respectively. Furthermore, the comparator offset is mostly dependent on the input transistor sizing. Note that the offset reduction achieved by increasing the size of the input devices eventually tails off which is consistent with Equation 2.11, while the power consumption increases. Based on these trends, the following sizing methodology was followed: the input transistors were sized relatively large (8× the minimum size) for low offset and delay, the tail transistor was sized large (8× the minimum size) to reduce delay without much expense in power, and the cross coupled inverters were minimum sized to achieve both low power and delay simultaneously. Furthermore, all devices in this comparator are LVT devices to help with low voltage operation.

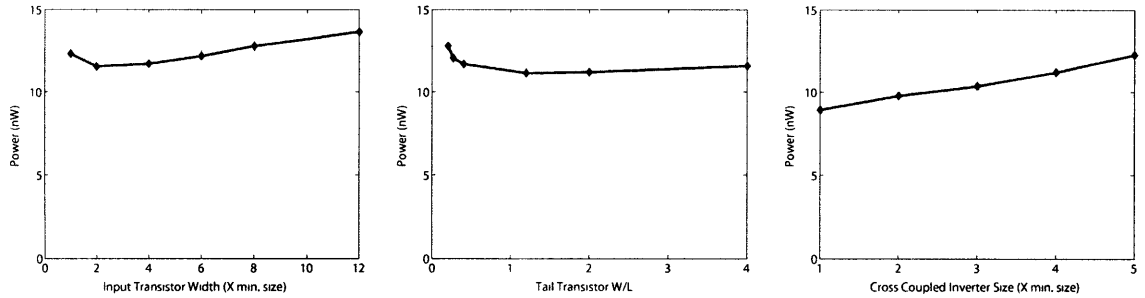

Figure 4-15: Comparator power vs. transistor sizing at V_{DD} = 1-V and f_{clk} = 1-MHz.

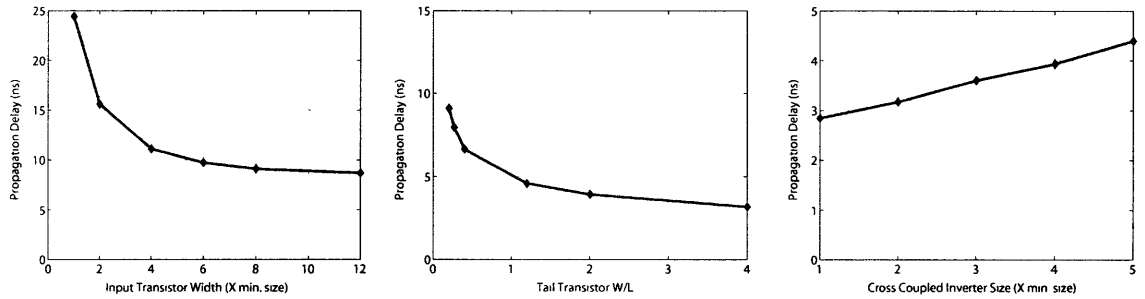

Figure 4-16: Comparator propagation delay vs. transistor sizing at V_{DD} = 1-V.

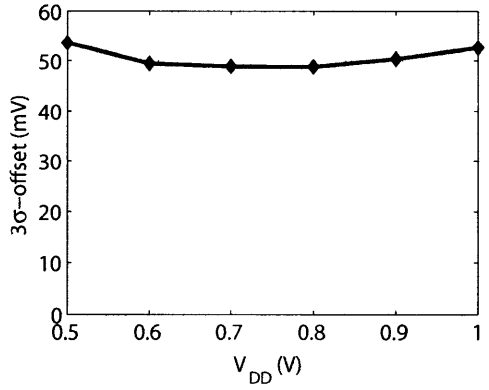

Figure 4-17: 3-σ-offset vs. V_{DD} for the comparator shown in Figure 4-12.

4.3.3 Comparator Offset

The 3-σ offset of the comparator was simulated to be relatively constant over voltage as plotted in Figure 4-17. The offset of the comparator manifests itself as offset in the overall ADC transfer characteristic. As mentioned earlier in Section 4.3.1, a 4-bit capacitor bank shown in Figure 4-18 is used to cancel the offset so as to not limit the available input range of the ADC.

In order to cover the entire 3-σ offset of approximately 50-mV, it was determined that a total load capacitance range of 32-fF was required. For a 4-bit binary weighted capacitor bank, this translates to a unit capacitance of 2-fF which is on the order of parasitics. Therefore, in order to make the capacitances manufacturable, a coupling capacitance C_C of 50.3-fF was added in series such that the unit capacitance became $C_U = 5.4$-fF. Note however, this technique of capacitive compensation increases the power consumption of the comparator. In this ADC prototype, the comparator consumes approximately 5% of the total power so this increase was acceptable.

Simulations were performed with a 50-fF load capacitance at the outputs to ensure that the 4-bit capacitor bank did not limit the speed of the comparator. The maximum clock frequency vs. V_{DD} is plotted in Figure 4-19. This maximum frequency determines the maximum sample rate according to Equation 3.1.

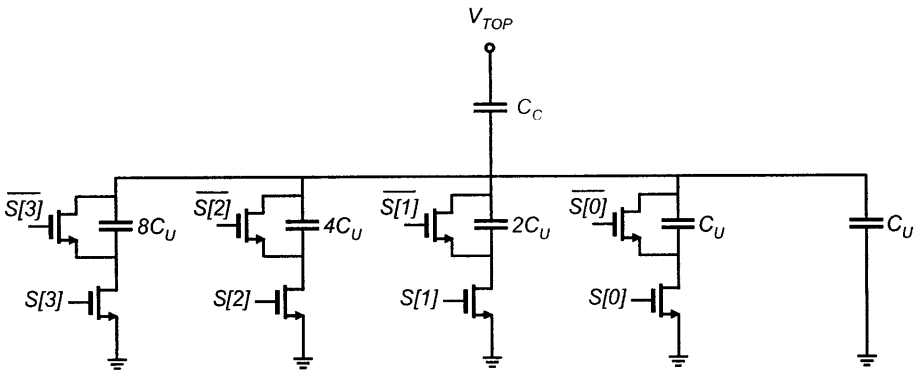

Figure 4-18: Schematic of 4-bit programmable load capacitor used for offset compensation.

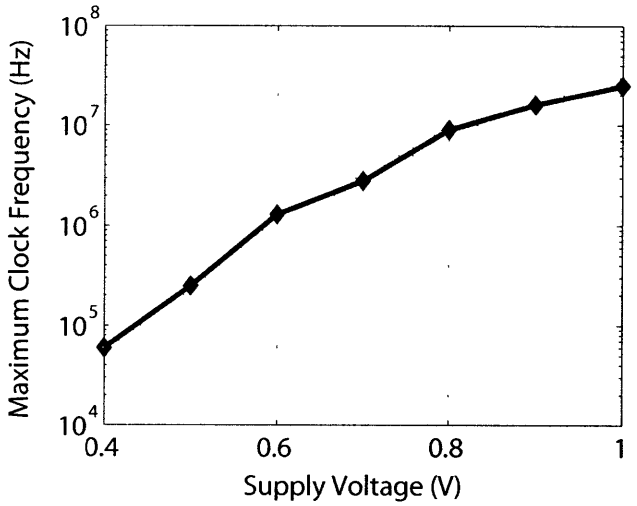

Figure 4-19: Maximum comparator strobe frequency vs. V_{DD} assuming a 50-fF load capacitance and correct resolution of an input of one LSB (10-bit mode) within $1/5^{th}$ of a clock period.

4.4 SAR Control Logic Design

As with the resolution scalable DAC, the resolution setting for the SAR control logic is digitally set by $RM[2:0]$ and the associated thermometer and one-hot representations shown in Table 4.2. The decoders used to decode $RM[2:0]$ are shown in Figure 4-20, and are implemented with HVT logic to minimize leakage since these signals are fixed for a given resolution setting.

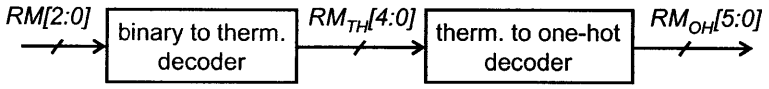

Figure 4-20: Decoder logic used to generate resolution scaling control signals. The truth table is given in Table 4.2.

Each conversion cycle is initiated from an external signal $CNVRT\_EXT$ which propagates through the shift register shown in Figure 4-21. The ADC is designed to accomodate any duty cycle from $CNVRT\_EXT$ to allow for flexibility when interfacing with a DSP. Internally, $CNVRT\_PLS$ signifies the start of a conversion cycle. In order to allow the digital circuits to recover from $SLEEP$ mode where power gating is applied, $CNVRT\_PLS$ is delayed by one clock cycle. Following this, PRG

and $SAMP$ are asserted to purge the DAC and sample the input signal respectively. The duration of $SAMP$ can be made programmable between one or two clock cycles, although this is not shown in the schematic. Both PRG and $SAMP$ are gated with $RM_{TH}[4:0]$ as shown in Figure 4-22 to derive the signals to drive the resolution scalable DAC. Finally, $\overline{SAR\_RST}$ initiates the start of the bit cycling phase.

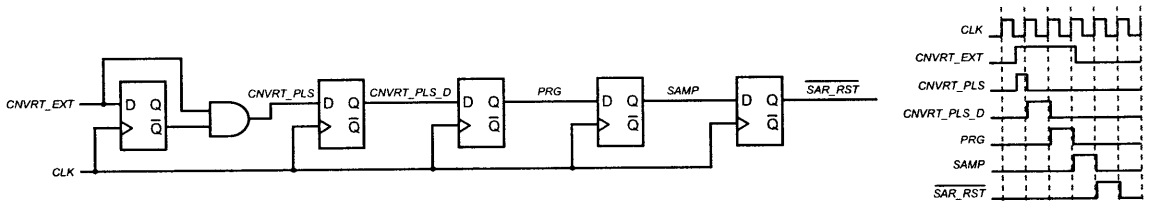

Figure 4-21: Simplified schematic of shift register generating the control signals for the purge, sample and bit cycling phases.

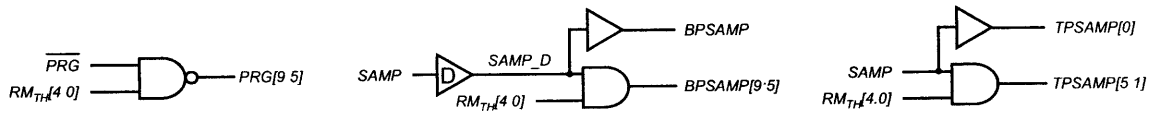

Figure 4-22: Logic used to generate control for the purge and sampling switches shown in Figure 4-9.

The synchronous digital state machine which controls bit cycling is shown in Figure 4-23. The top shift register is reset from the previous cycle using the $BITCYC$ signal which is high only during bit cycling. $\overline{SAR\_RST}$ is gated with $RM_{OH}[5:0]$ to set the appropriate register in order to start bit cycling with the desired DAC capacitor for the required resolution setting. For example, in 6-bit mode, the $L[6]$ register is set and the signal propagates through the top shift register, successively asserting $L[i]$, which clocks a corresponding register that determines $BP[i]$, which are intermediate signals used to derive the control for the DAC reference switches.

Since the split capacitor array was used in this design, the logic shown in Figure 4-24 is required to generate the intermediate signals $MSBP[9:1]$ for the MSB array. A delay cell was used to prevent glitches in $MSBP[9:1]$. Finally, to generate the control signals for the DAC reference switches shown in Figure 4-9, $BP[9:1]$ and $MSBP[9:1]$ are gated with $BITCYC$ and $RM_{TH}[4:0]$ as shown in Figure 4-25.

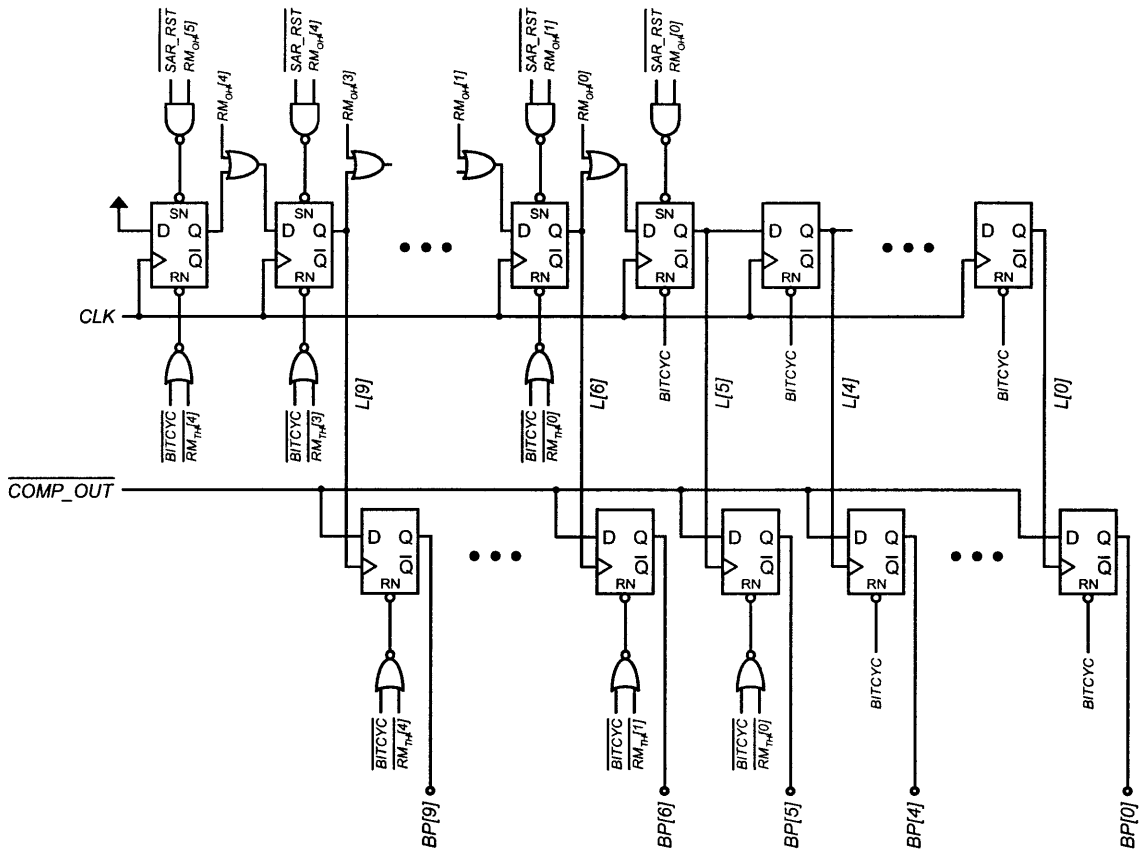

Figure 4-23: Schematic of the shift register based SAR control logic with peripheral logic to set the resolution mode.

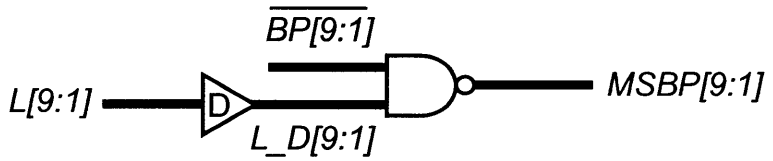

Figure 4-24: Logic to generate bit cycling control for the MSB capacitor of the split capacitor array.

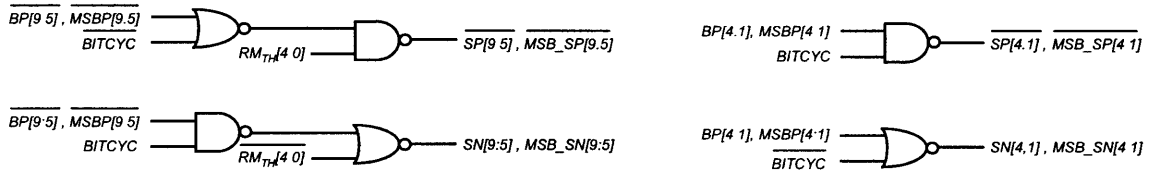

Figure 4-25: Logic used to generate control for the bit cycling switches shown in Figure 4-9.

Lastly, in order to wake up the ADC from $SLEEP$ mode, the circuit shown in Figure 4-26 is used. The active low signal $BCDONE\_PLS$ is asserted when bit cycling is finished, setting $SLEEP$ to 1. $SLEEP$ remains at 1 until the next positive edge of $CNVRT\_EXT$. Note that the control signals for the DAC switches are all gated with a power gating interface to ensure that the DAC is not being switched inadvertently during power gating. The purge switches are used to keep the decoupled MSB capacitors grounded, and to zero the DAC during power gating.

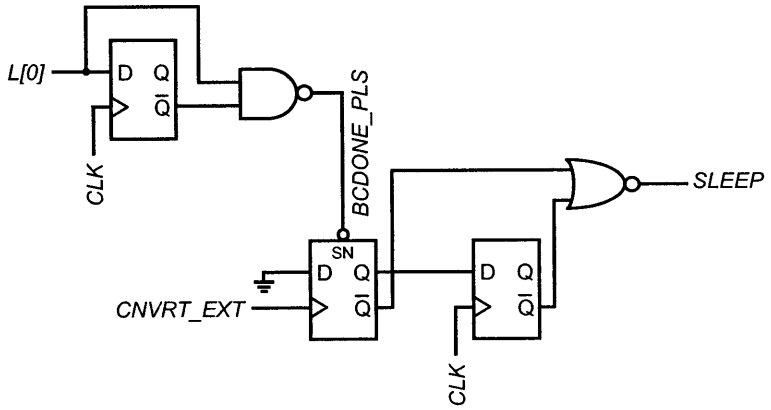

Figure 4-26: Logic to generate the $SLEEP$ signal for power gating and clock gating.

4.4.1 Power Gating Circuit Design

The power gating footer switch is controlled by the $SLEEP$ signal and power gating can be disabled by clearing $PGATEMODE$ as shown in Figure 4-27. The HVT footer was sized to be 48-μm wide such that the maximum spike on the virtual V_{SS} node was less than 10-mV over all modes of operation.

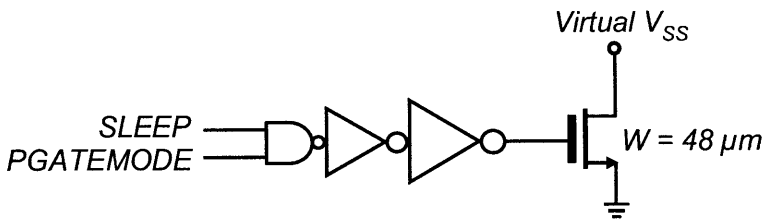

Figure 4-27: Logic used to drive the HVT power gating footer.

Also, the leakage of the HVT footer was verified to be a small percentage of the leakage of the power gated logic. A plot of the HVT footer leakage over voltage is shown in Figure 4-28, showing that it is less than 3% of the total leakage from the power gated logic. This implies that most of the overhead associated with power gating is due to switching and recovery power.

Note that increasing the size of the footer beyond what is necessary will increase the parasitic gate capacitance as well as the residual leakage, thus lowering the break even frequency for power gating. To help with this tradeoff, voltage boosting was implemented (not shown in the schematic) and may be used to further reduce the virtual V_{SS} spike if digital performance is an issue at low voltages.

Since the virtual V_{SS} node can drift during power gating, it is necessary to interface the power gated logic with non-power gated logic to prevent any static currents. The interface shown in Figure 4-29 is used, where X and Y are arbitrary signals that drive non-power gated cells.

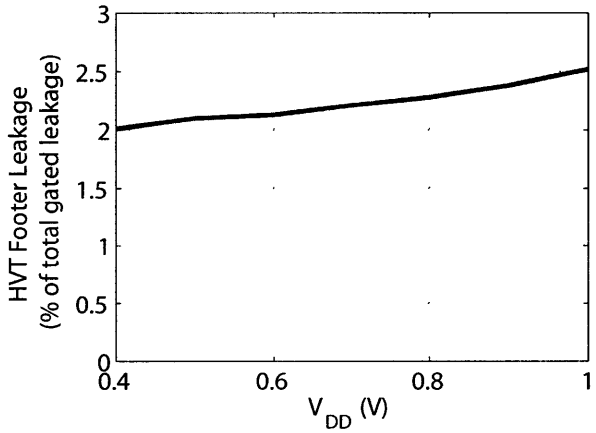

Figure 4-28: The HVT footer leakage as a percentage of the total leakage from the power gated logic.

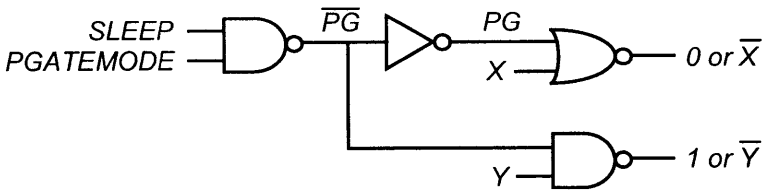

Figure 4-29: Logic to interface between power-gated and non-power-gated logic.

4.4.2 ADC System Clock

The clock provided off-chip, CLK_{EXT}, is gated with the $SLEEP$ signal as shown in Figure 4-30 in order to save power between conversions. Clock gating can be disabled by clearing $CGATEMODE$.

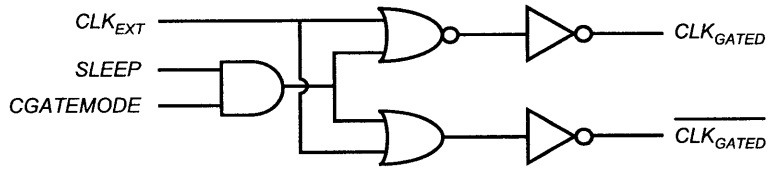

Figure 4-30: Clock gating logic

4.4.3 Overhead of Reconfigurability

The overhead as a result of reconfigurability can be quantified in terms of power and area. However, it is difficult to qualify exactly which circuits should be considered as overhead for reconfigurability. In this design, the power overhead results from the extra digital logic required to implement resolution scaling in the DAC. From simulation, it was determined that the extra resolution scaling logic increased the total power by less than 10%.

In terms of area, the overhead consists of the area of the extra comparator and the area of the extra resolution scaling logic. When compared to the area of the DAC which dominates the chip area, this overhead is negligible.

Due to a high degree of reuse of existing circuit blocks, especially in the DAC, the reconfigurability overhead in this ADC was minimized.

4.5 Summary

In this chapter, transistor and logic level circuit details for each block of the ADC was presented. The resolution scalable DAC employing both a split capacitor array and a sub-DAC was presented. The design of the input sampling and bit cycling switches was discussed. The sizing methodology for the latched comparator was presented, optimizing it for power, delay and offset. Finally, the SAR control logic design was described, and consideration was given to the overhead of reconfigurability.

Chapter 5

Simulation Results

The ADC was designed in a low power 65-nm digital CMOS process and has been submitted for fabrication. The chip layout can be seen in Figure 5-1. The total die area including pads is 2-mm $\times$ 1-mm, while the size of the ADC core measures 750-μm $\times$ 470-μm. This chapter will present simulation results demonstrating functionality, performance and the effectiveness of the applied techniques on maintaining energy efficiency across all modes of operation. All presented simulation results are from post-layout extraction HSPICE simulation unless otherwise indicated.

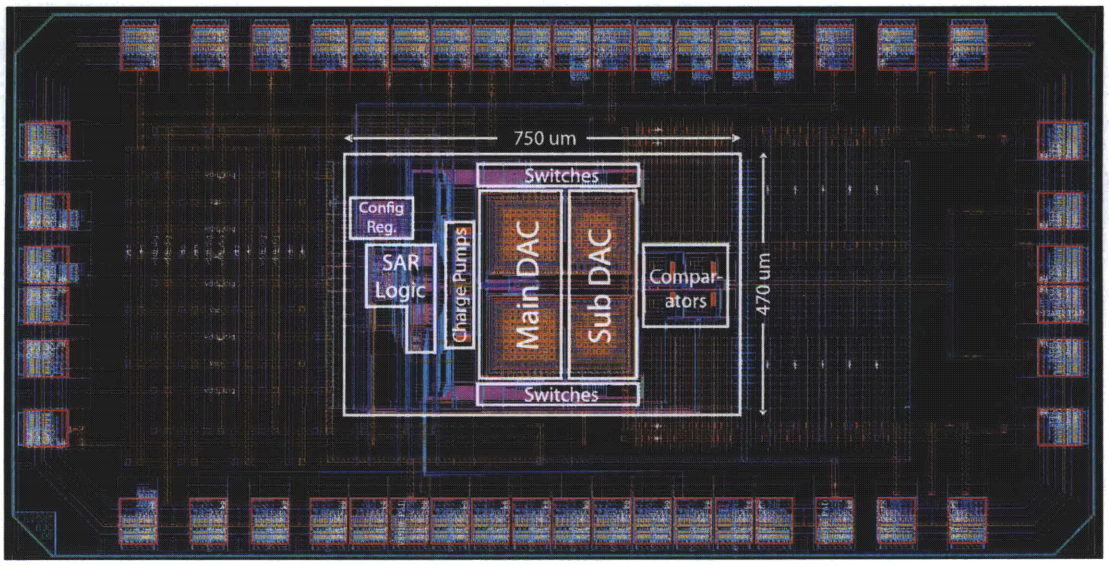

Figure 5-1: Layout of the ADC in a 2-mm $\times$ 1-mm die area including pads.

5.1 Input Sampling Network

At the highest sampling rate, the input sampling network must be able to accomodate signals up to 500-kHz with a peak-to-peak input of $\pm V_{DD}$ differentially, or V_{DD} single-endedly. The FFT of the sampled input in 10-bit mode with a full scale 472.65625-kHz input signal without and with voltage boosting is shown in Figure 5-2. Without boosting, the THD is only -39.57-dB corresponding to a linearity of 6.3 effective bits. By enabling the charge pumps during sampling, the THD improves to -74.72-dB and the linearity increases to 12.1-bits, sufficient for the our desired maximum resolution of 10-bits.

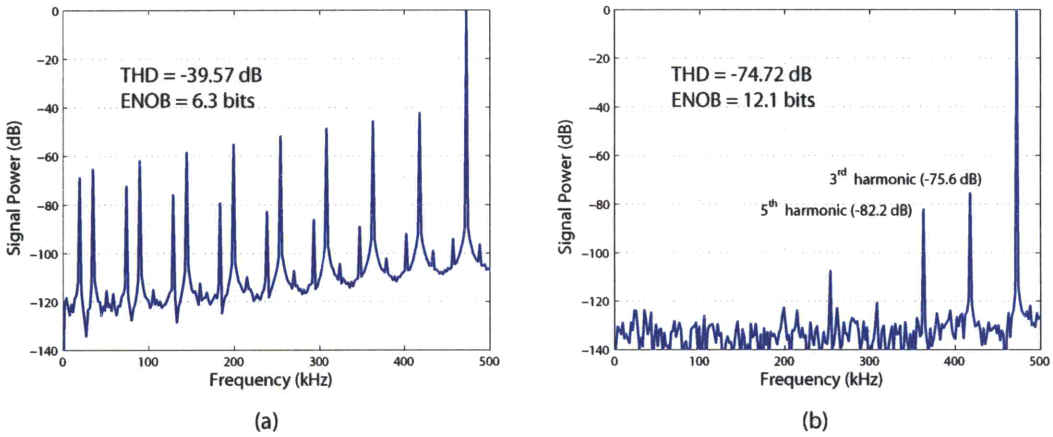

Figure 5-2: Simulated FFT of input sampling network in 10-bit mode with a 472.65625-kHz input tone at V_{DD} = 1-V (a) without voltage boosting and (b) with voltage boosting.

Since the sampling switch for each DAC capacitor was sized for a constant RC time constant, at a constant V_{DD}, the THD of the sampling network remains constant as the resolution is reduced. However, since aggressive voltage scaling is used in this design, the THD (with voltage boosting) versus input frequency was simulated over a range of supply voltages and the results are plotted in Figure 5-3. Note that the boosted sampling clock is approximately $1.5V_{DD}$. Simulations confirm that the sampling network provides more than enough linearity for the desired input frequency range at each supply voltage.

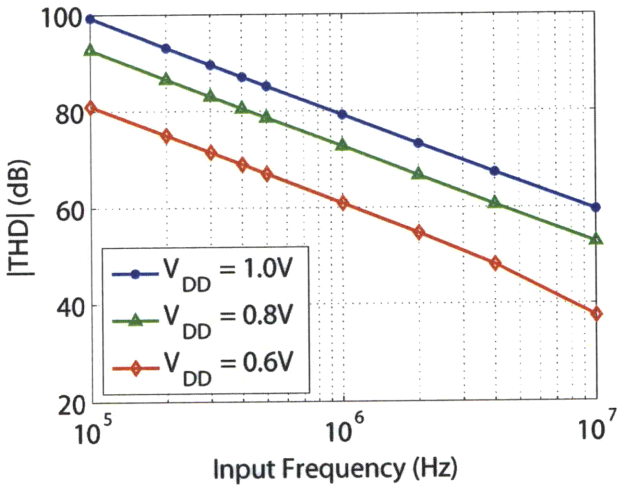

Figure 5-3: Simulated THD vs. input frequency over various supply voltages for the input sampling network.

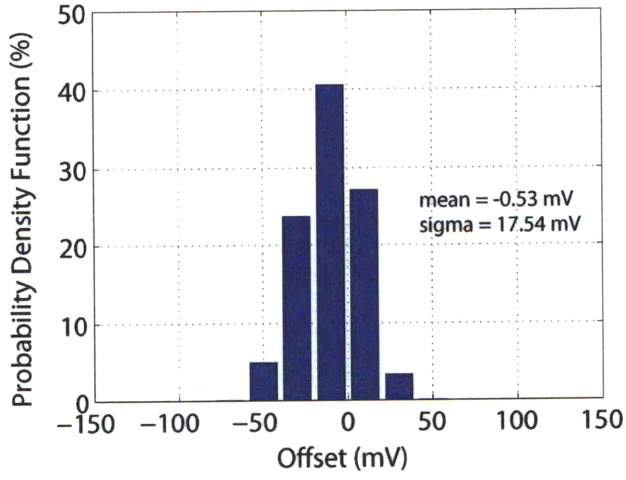

Figure 5-4: Simulated distribution of comparator offset from 1000 monte carlo iterations.

5.2 Comparator Offset Compensation

The distribution of the comparator offset was simulated with monte carlo analysis and is plotted in Figure 5-4. The mean was -0.53-mV while the standard deviation was 17.54-mV. Simulations were performed to ensure that the implemented offset compensation capacitors were able to cover the 3-σ offset of the comparator. The simulated input referred offset versus the digital calibration code is shown in Figure 5-5, showing that offsets in the range of $\pm$50-mV can be compensated.

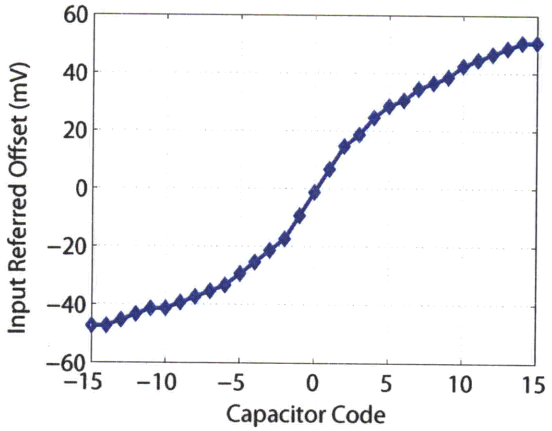

Figure 5-5: Digital code vs. systematic comparator input referred offset at V_{DD} = 1-V. The negative and positive code range corresponds to C_{LoadN} and C_{LoadP} respectively.

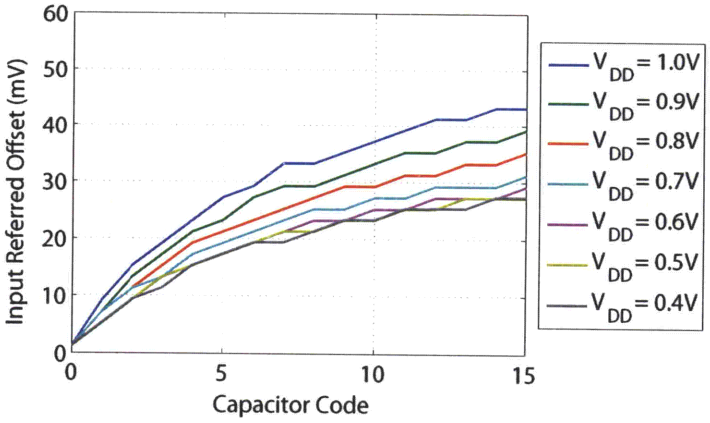

Figure 5-6: Digital code vs. systematic comparator input referred offset for V_{DD} from 0.4-V to 1-V.

Figure 5-6 shows the effect of capacitive offset compensation across supply voltages. As the supply voltage is lowered, the tolerable offset range also decreases. However, since the ADC operates at lower resolutions for reduced supplies, any offset that cannot be cancelled will have a less severe impact on the ADC output characteristic since the LSB voltage is larger at lower resolutions.

5.3 Power Gating Simulation Results

The effect of power gating in reducing the overall power consumption of the ADC was simulated in HSPICE. As discussed in Section 3.2.2, by completing each conversion at the maximum clock frequency and then duty cycling, the ADC can be put into *SLEEP* mode. Figure 5-7 shows a plot of the ADC power versus the *SLEEP* mode duration (as a percentage of the total conversion period) with and without power gating applied.

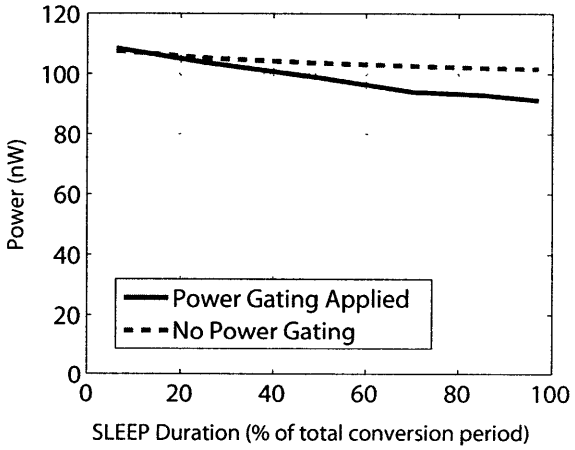

Figure 5-7: Simulated power versus amount of *SLEEP* time showing the effectiveness of power gating on the total power at 10-bits, V_{DD} = 1-V, f_{samp} = 5-kS/s.

It can be seen that without power gating, the ADC power remains relatively constant irrespective of the amount of duty cycling because the CV^2 switching energy is constant per conversion cycle, and the active mode leakage is not reduced as described in Section 3.2.2. However, with power gating applied, as the *SLEEP* mode

duration increases, the active mode leakage is suppressed for a longer period of time, thus decreasing the overall power. At 10-bits, f_{samp} = 5-kS/s and a 1-V supply, it can be seen that the total power is reduced from 110-nW with virtually no power gating, to just 90-nW when the leakage is power gated for 97% of the time, corresponding to a power savings of 18%. As the sample rate is reduced further, even greater power savings can be achieved from duty cycling and power gating since the amount of available idle time is inversely proportional to the sampling rate. In the extreme case of a zero sampling rate, the leakage power of the digital circuits is limited by the leakage of the HVT footer device, plus any digital logic that is not power gated.

5.3.1 Simulated Break Even Frequency

As mentioned in Section 3.2.5, it is worth looking at the break even point for power gating. The break even frequency f_{BE}, defined as the sample rate at which power gating becomes beneficial, was simulated over a range of voltages and the results are plotted in Figure 5-8. At supply voltages of 0.4-V, 0.6-V, 0.8-V and 1.0-V, the ADC was configured in 5, 7, 8 and 10-bit modes respectively. The overhead power consists of the power required to switch the HVT footer, the power required to recover the virtual ground node, and the leakage of the HVT footer device (which was shown to be almost negligible in Figure 4-28).

From Figure 5-8, it can be seen that f_{BE} decreases as the supply is reduced. At 1-V, f_{BE} is well over 100-kS/s, while at 0.4-V, f_{BE} = 1-kS/s. This can be explained by the fact that the clock was operated at the maximum frequency for the given supply such that the purge, sample and bit cycling phases complete in the shortest time possible per sampling period, thus maximizing the amount of time that the circuits are power gated. A higher supply voltage allows the $SLEEP$ mode duration to be increased per conversion period.

Also, note that the worst case leakage from the digital circuits occurs when the ADC is configured in 10-bit mode. This translates to a vertical shift in the leakage power curves, increasing f_{BE} slightly. However, since these are pre-extraction simulations, the actual f_{BE} will be lower since the overhead power will increase due to

parasitics.

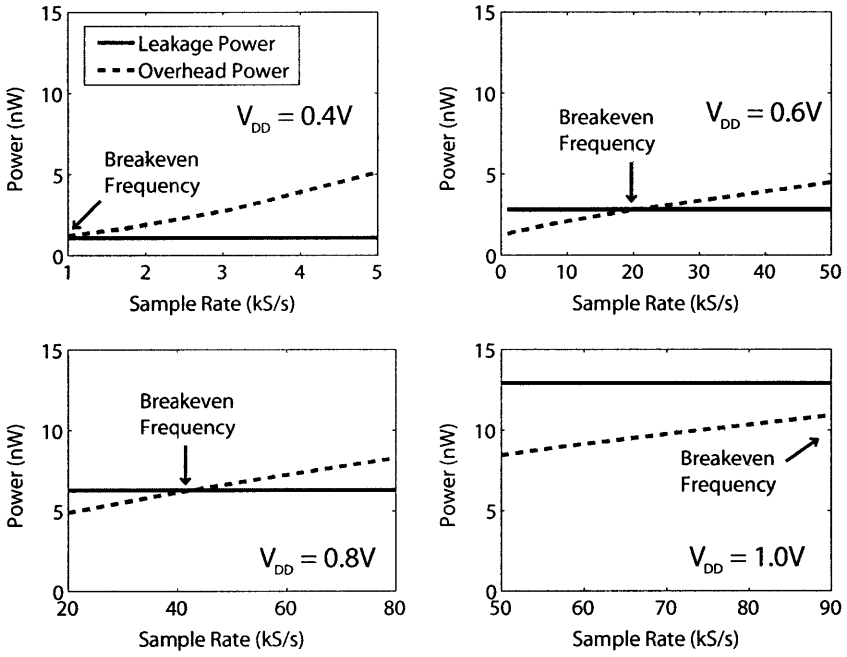

Figure 5-8: Simulated break even frequency over supply voltage (pre-extraction simulation).

5.4 Static Linearity

As discussed in Section 2.3, capacitor mismatch leads to errors in the ADC transition voltages degrading the linearity of the ADC. Apart from capacitor mismatch due to processing, systematic errors such as mismatched parasitic wiring capacitance due to layout, coupling between top and bottom plates and error in the sub-DAC coupling capacitance can cause linearity errors as well. The primary purpose of simulating the static linearity of the ADC is to quantify the linearity errors due to the systematic errors, since process induced capacitor mismatch is not accounted for.

Experimentally, the static transfer characteristic of an ADC can be derived through code density testing [43]. However, in simulation, we are limited to stepping the input voltage and recording the corresponding output digital code. The simulated INL

and DNL are only as accurate as the resolution of the input steps. For 0.25-LSB accuracy (4 steps per LSB voltage) in a 10-bit ADC, $4 \times 2^{10} = 4096$ samples must be simulated, which requires several weeks to simulate. Therefore, at 9 and 10-bit modes, only the first few MSB transitions of the ADC transfer characteristic were simulated. At 8-bits or less, the full transfer characteristic could be simulated in a reasonable time at a resolution of 4 steps per LSB.

Simulations confirmed that the parasitic wiring capacitance in the DAC was well matched (ratiometrically) and coupling between top and bottom plates was negligible. Therefore, the rest of this section will focus on issues with the sub-DAC.

5.4.1 Linearity Errors Due to the Sub-DAC

With regards to the maximum INL and DNL (in units of LSB) resulting from error in the sub-DAC coupling capacitance C_C, it is only necessary to simulate a portion of the transfer characteristic equal to the sub-DAC interpolation period. For this implementation, due to the presence of two 4-bit sub-DACs (one in each of the MSB and LSB arrays), the interpolation period is $2^5 = 32$ output codes. Note that since the size of the sub-DAC is constant for all resolution modes of the ADC from 6 to 10-bits, the maximum INL and DNL (in units of LSB) due to error in C_C is the same for all resolution modes. To illustrate this, an example is shown in Figure 5-9, showing the maximum INL due to errors in the sub-DAC transitions in N-bit mode and $(N+1)$-bit mode for a 2-bit sub-DAC. Since the sub-DAC interpolates the main-DAC transitions, the voltage linearity error is halved for each extra bit of resolution, but since the LSB voltage is also halved, the maximum INL and DNL (in units of LSB) stays constant over all resolution modes.

As discussed in Section 4.1.2, the top plate parasitic capacitance of the sub-DAC C_P has the effect of compressing the sub-DAC transitions through capacitive attenuation, leading to INL errors as shown in Figure 5-9. As described in [41], a way to mitigate this error is to increase the size of C_C. However, the amount by which to adjust C_C depends strongly on the ratio between C_P and the total sub-DAC capacitance C_{subDAC}, which must be obtained through post-layout extraction.

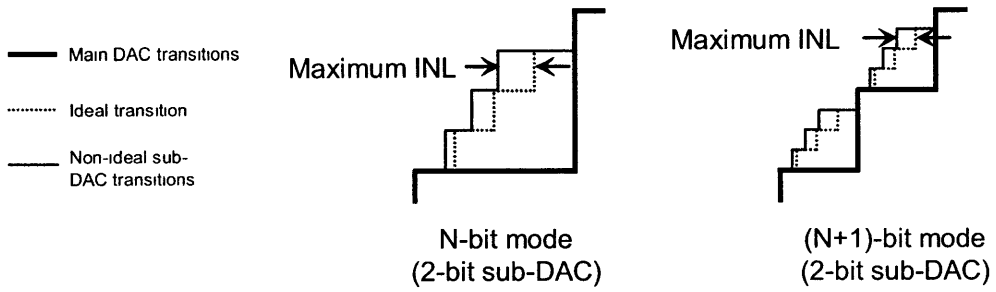

Figure 5-9: Effect of error in the sub-DAC coupling capacitance over different resolution modes.

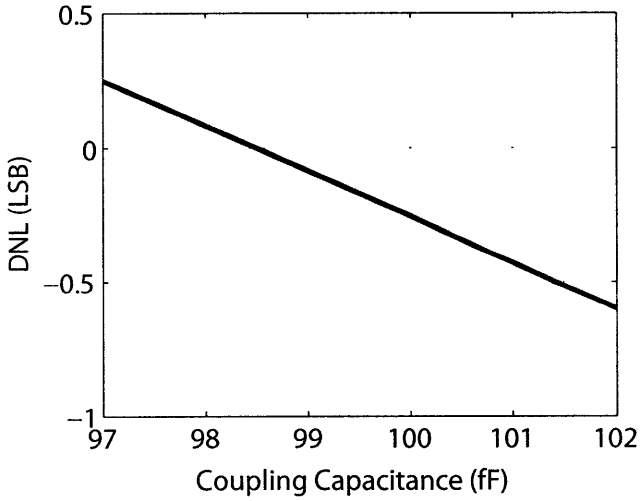

Figure 5-10: Simulated DNL vs. sub-DAC coupling capacitance.

Simulations were performed to quantify the sensitivity of the DNL on the sizing of C_C and the results are shown in Figure 5-10. This suggests that C_C should be set to 98.5-fF to minimize the DNL. Note that if somehow the fabricated C_C is too large (due to variation in parasitics), then it is possible that the ADC will suffer from negative DNL spikes. In the extreme case where the DNL < -1 LSB, then the ADC would have missing codes. Therefore, in order to be conservative, C_C was chosen to be slightly smaller.

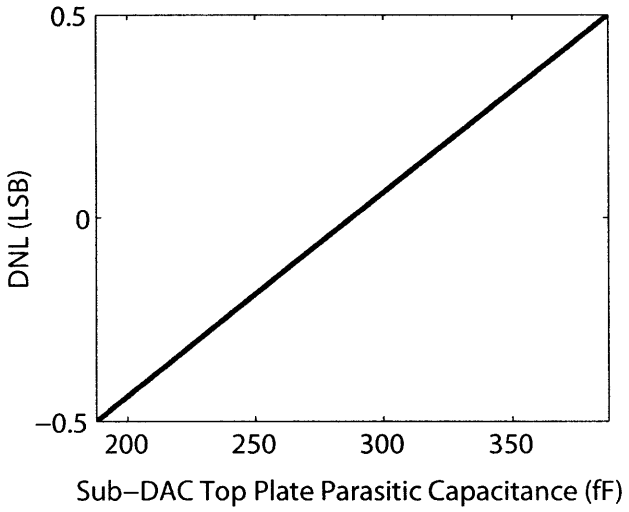

Figure 5-11: Simulated DNL vs. sub-DAC top plate parasitic capacitance.

Simulations were also performed to see the sensitivity of the DNL error to C_P for a fixed C_C and the results are shown in Figure 5-11. It can be seen that a large range of C_P can be tolerated. Figure 5-12 shows the simulated INL and DNL of the ADC in 6-bit mode (accurate to 0.25-LSB), with the DNL plot showing a slight positive spike at the MSB transition due to intentionally under-sizing C_C. A similar sawtooth pattern with a period of 32 codes in the INL is observed at higher resolution modes. Note that in 5-bit mode, there is no error due to the sizing of C_C since the sub-DAC acts as the entire main-DAC in that case.

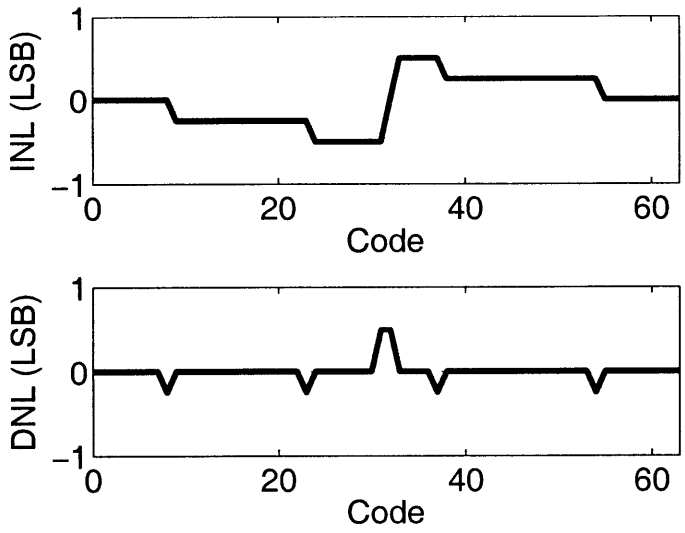

Figure 5-12: Simulated INL and DNL of the ADC in differential mode at 6-bits.

5.5 Dynamic Performance

The dynamic performance of the ADC was simulated by applying an input tone from DC to the Nyquist rate and computing the FFT (using 512 samples). Figure 5-13 shows the FFT of the ADC output in 10-bit mode at 1-MS/s and 1-V, with an input tone of 472.65625-kHz near the Nyquist frequency for both (a) differential and (b) single-ended operation.

When configured differentially, the ADC achieves an ENOB of 9.28 (SNDR = 57.6-dB). In single-ended mode, the ENOB degrades slightly due to the presence of the 2nd harmonic distortion term which is 65-dB below the fundamental. The ENOB versus input frequency in 10-bit differential mode is shown in Figure 5-14.

Figure 5-15 shows the ENOB at the Nyquist rate at each resolution mode, at a sampling rate of 1-MS/s with a 1-V supply. Note that simulations at scaled voltages around V_T were not performed. This was due to the fact that at low voltages, the low sample rates result in excessively long simulations. In order to acquire enough samples to compute the ENOB, simulations quickly become infeasible.

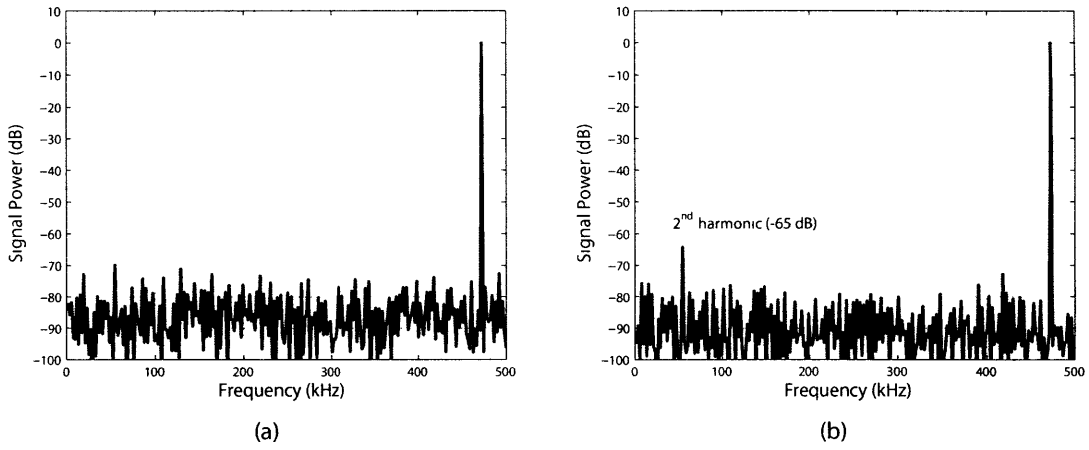

Figure 5-13: Simulated 512-point FFT of ADC in (a) differential mode and (b) single-ended mode, at 10-bits, with a 472.65625-kHz input tone at V_{DD} = 1-V, f_{samp} = 1-MS/s.

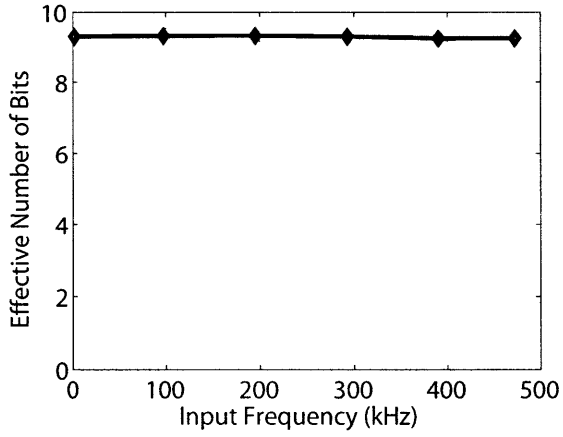

Figure 5-14: Simulated ENOB vs. input frequency for the ADC in differential, 10-bit mode at V_{DD} = 1-V.

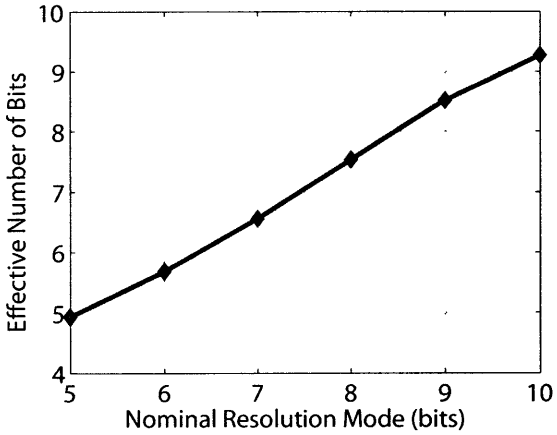

Figure 5-15: Simulated ENOB vs. the nominal resolution mode for the ADC in differential configuration with a 472.65625-kHz input tone at V_{DD} = 1-V, f_{samp} = 1-MS/s.

Resolution (bit)	10*	9*	8*	7**	6**	5**
Supply (V)	1	0.85	0.7	0.6	0.45	0.4
Sample Rate	1 MS/s	1 MS/s	400 kS/s	150 kS/s	25 kS/s	13 kS/s
DAC Power (differential)	8.93 μW	4.39 μW	0.896 μW	0.106 μW	10.1 nW	3.15 nW
SAR Logic Power	7.2 μW	4.88 μW	1.23 μW	0.233 μW	24.4 nW	15.0 nW
Clock Power	4.98 μW	3.34 μW	0.837 μW	0.210 μW	18.9 nW	7.82 nW
Comparator Power	1.19 μW	0.758 μW	0.18 μW	50.1 nW	4.32 nW	1.80 nW
Charge Pump Power	1.87 μW	1.26 μW	0.292 μW	42.3 nW	4.24 nW	1.44 nW
Total Power	24.2 μW	14.6 μW	3.4 μW	0.64 μW	62.0 nW	29.2 nW

Table 5.1: Simulated power consumption of ADC blocks at each resolution mode with scaled voltages (no power gating applied). (*differential mode **single-ended mode)

5.6 ADC Power Consumption

The simulated power consumption of the ADC at each resolution mode is provided in Table 5.1. In 10-bit mode at the maximum sample rate of 1-MS/s, the entire ADC consumes 24.2-μW from a 1-V supply. In 5-bit mode at a sample rate of 13-kS/s, the ADC consumes 29.2-nW from a 0.4-V supply, of which 15.2-nW is leakage power. The power consumption versus sample rate of the ADC in 10-bit, differential mode without power gating at 1-V is shown in Figure 5-16, showing a linear relationship as expected for a fully dynamic ADC.

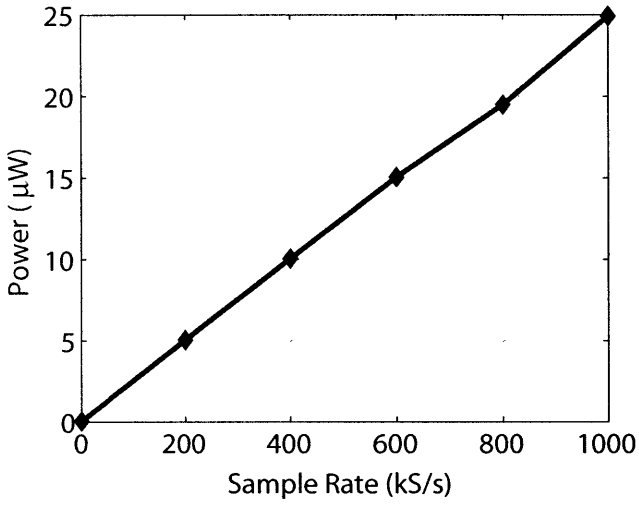

Figure 5-16: Simulated ADC power vs. sample rate in differential, 10-bit mode without power gating at V_{DD} = 1-V.

5.6.1 Simulated FOM

The effectiveness of voltage scaling in maintaining a constant energy efficiency over resolution can be seen by looking at the FOM of the ADC. Figure 5-17 shows the simulated FOM versus resolution with supply scaling from 1-V at 10-bits down to 0.4-V at 5-bits. The data points from 7 to 10-bits are from full chip simulations, however, the ENOB at 5 and 6-bits was not simulated due to simulation limitations as discussed in Section 5.5. Therefore, as a reasonable estimate, the FOM at 5 and 6-bits is plotted assuming that the ENOB is 0.75-bits below the nominal resolution (i.e. an ENOB of 4.25-bits in 5-bit mode). Note that from 5 to 7-bits, the ADC is operated in single-ended mode to reduce the DAC power. This degrades the ENOB slightly due to increased 2nd order harmonics.

The simulated FOM in 10-bit mode at a 1-V supply was 39-fJ/conversion step. It can be seen that the FOM degrades as resolution is reduced. This is because the DAC power (which is the only block with an exponential dependence on resolution) becomes a smaller percentage of the overall power and voltage scaling becomes less effective. Also, leakage becomes more significant at scaled voltages, further degrading the FOM. Overall, there is only a 3× increase in the FOM over the entire resolution

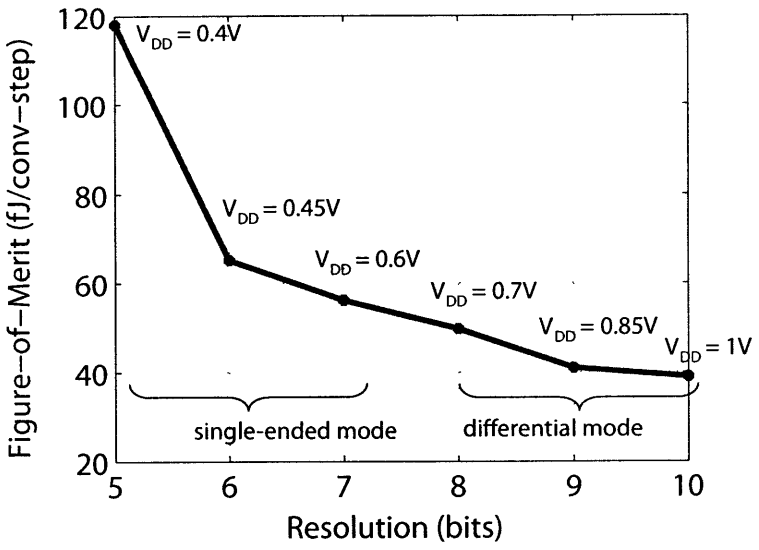

Figure 5-17: Simulated ADC FOM at each resolution mode with voltage scaling applied according to Table 5.1. The ENOB at 5 and 6-bits is estimated to be 0.75-bits below the nominal resolution.

range of 5 to 10-bits. However, the degradation of the FOM at low resolutions can partially be recovered by using power gating. As discussed in Section 5.3, this can reduce the overall power and accordingly, the FOM, at sample rates less than the break even frequency.

5.7 Summary

The effectiveness of the applied techniques such as voltage scaling, power gating and DAC reconfigurability in achieving the desired performance and scalability has been presented through HSPICE simulation results. The input sampling network was designed to accomodate the targeted range of voltages and resolution modes. The 3-σ offset of the comparator can be cancelled with off-line calibration. Power gating was simulated to be effective at sample rates in the range of 10's of kS/s or less. The power gating break even frequency decreases with decreasing supply voltage. In 10-bit mode at 5-kS/s and 1-V, it was shown that the overall power can be reduced by as much as 18% through the use of power gating. More power savings can be

expected as the sample rate is reduced even further since the amount of available idle time increases. The static linearity and dynamic performance were simulated and verified to be adequate for the targeted resolution modes. Finally, through the use of voltage scaling, the FOM of the ADC was kept reasonably constant over the range of 5 to 10-bits. The simulated ADC achieves a FOM of 39-fJ/conversion step in 10-bit mode, and increases by only 3× to 118-fJ/conversion step in 5-bit mode.

Chapter 6

Conclusion

A highly digital, reconfigurable and voltage scalable ADC has been presented. The ADC power scales linearly with sample rate and near exponentially with resolution. The ability to scale the power with performance results in a very energy efficient implementation over a wide frequency and resolution space, suitable for applications such as medical monitoring in which bio-potentials often have varying bandwidth and resolution requirements. This chapter summarizes the key results of this research and suggests opportunities for future work.

6.1 Discussion and Summary of Contributions

The ADC presented in this thesis has a scalable sample rate from 1-MS/s down to 0, reconfigurable resolution from 5 to 10-bits, and operates from a supply of 1-V down to 0.4-V. For the targeted resolution range, thermal noise is not a limitation and therefore the energy efficiency of the ADC can benefit from aggressive voltage scaling. Since conventional analog circuits such as linear amplifiers cannot easily be operated at supplies close to the transistor threshold voltage, the SAR ADC architecture was selected for its highly digital nature. In this implementation, static bias currents were avoided resulting in a fully dynamic design where the power scales quadratically with the supply, and linearly with frequency as desired.

In order to achieve exponential power scaling with resolution, two techniques were

employed. First, since the DAC consumes a significant portion of of the overall power, a resolution scalable DAC was designed such that the DAC capacitance, and hence power, scaled exponentially with resolution. Secondly, as desribed in Section 3.2.4, since the binary algorithm of a SAR ADC only results in linear power reduction with resolution, voltage scaling is used to achieve exponential savings in the rest of the ADC.

As with any circuit technique, the tradeoffs of voltage scaling must be considered. At low voltages, the performance of the input sampling switches are degraded due to increased ON resistance. Boosted switches are used to achieve the desired sampling linearity. Also, digital logic delays and DAC settling times increase exponentially at low supplies near the threshold, resulting in a reduction of the ADC bandwidth. Furthermore, increased process variation at low voltages requires additional safeguard when designing for a targeted bandwidth. However, since a target application could be the digitization of low bandwidth bio-potentials, the maximum sample rate of 13-kS/s at a supply of 0.4-V is more than adequate.

In Section 3.2.4, the effect of voltage scaling on the ADC FOM was analyzed. At a given resolution, degradation in the SNDR is mainly due to increased harmonic distortion. Although the input signal amplitude decreases, so does the quantization noise. Moreover, even though the sampled $\frac{kT}{C}$ remains constant over voltage, it does not appreciably degrade the SNDR since the ADC is not thermal noise limited.

Another consequence of scaled supply voltages is that leakage becomes a larger fraction of the overall power due to reduced sample rates. To overcome this problem, power gating, a technique commonly used in digital circuit design, is applied to minimize leakage when the ADC is in $SLEEP$ mode. Through simulation, it was determined that power gating becomes beneficial at 10's of kS/s or less. At 1-V and a sample rate of 5-kS/s, power gating reduces the overall power by 18%. However, the power gating break even frequency decreases as the supply voltage is lowered.

Another feature of the ADC is the ability to support both differential and single-ended modes of operation. In situations where energy efficiency is of primary concern and reduced robustness to power supply and common mode disturbances can be

tolerated, single-ended mode can be enabled to half the DAC power, further improving the energy efficiency of the ADC.

With the aforementioned techniques applied, the simulated ADC achieves a minimum FOM of 39-fJ/conversion step in 10-bit, differential mode from a supply of 1-V. Referring to Table 1.2, this is competitive with most current state-of-the-art high energy efficiency ADCs, with the exception of [13] and [16]. In addition to being very energy efficient in 10-bit mode, the proposed ADC is designed to maintain its efficiency over all resolution settings. When configured in 5-bit, single-ended mode at 0.4-V (without power gating), the FOM increases by only 3× to 118-fJ/conversion step. However, this FOM can be further improved by applying power gating.

Lastly, it is worthwhile to estimate the effect of the resolution scalable DAC as well as voltage scaling on the FOM as resolution is scaled. Although it is not possible to fully account for differences in dynamic linearity at different voltages and sample rates, reasonable assumptions can be made and order of magnitude estimates in the FOM can be considered. Figure 6-1 shows the FOM vs. resolution for four cases:

1. Voltage scaling and resolution scalable DAC enabled: Implemented in the proposed ADC.
2. Voltage scaling only: This case assumes that bit cycling begins with the MSB capacitor of the DAC and stops at the required resolution. Also, the sample rates are reduced with voltage as in case 1.
3. Resolution scalable DAC only: This case accounts for the savings from the use of the resolution scalable DAC, but the voltage and sample rate are kept constant at 1-V and 1-MS/s respectively over resolution.
4. No voltage scaling or resolution scalable DAC: This case assumes operation at 1-V and 1-MS/s, where the digital output is truncated to achieve resolution scaling.

It can be seen that the use of a resolution scalable DAC with voltage scaling is especially effective at lower resolutions. At 5-bits, the FOM for cases 2 and 3 are 7× worse than in case 1, and the FOM for case 4 is a full 10× worse than case 1.

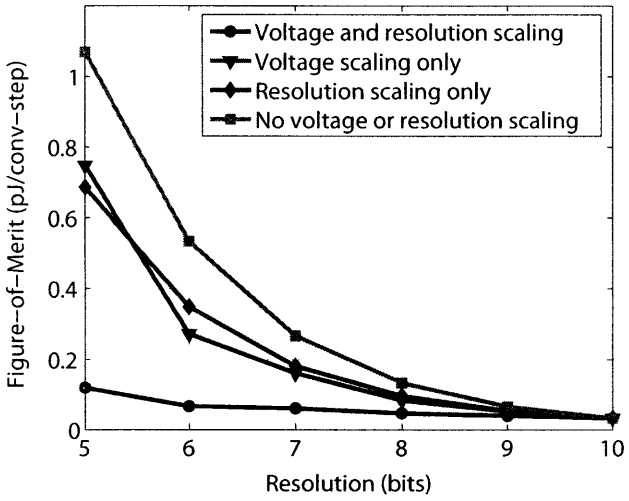

Figure 6-1: The effect of a resolution scalable DAC and voltage scaling on the FOM of the ADC.

The prototype has been submitted for fabrication and chip testing is scheduled for Fall 2009. A printed circuit board (PCB) has been design and details regarding the test setup can be found in Appendix B.

6.2 Future Work

This section will propose future work that can be done to improve the design presented in this thesis with respect to system integration and usability, power efficiency and scalability.

6.2.1 On-Chip Dynamic Voltage Scaling

Since this ADC relies heavily on voltage scaling to achieve high energy efficiency, it follows that an efficient on-chip DC-DC converter capable of delivering power at voltages from 1-V down to 400-mV is needed in order to make the ADC a more integrated subsystem. An example of such a DC-DC converter can be found in [44], where a switched capacitor DC-DC converter capable of delivering power in the μW to mW range at load voltages from 1.1-V down to 300-mV is reported.

Assuming that such an ADC is used in a medical monitoring system with a digital processor, digital control of the supply voltage is desirable because it enables feedback from either the user or the processor, making the ADC a dynamic voltage scalable system.

Lastly, if the voltage can be dynamically reduced during *SLEEP* mode, the idle mode leakage power can be reduced even further.

6.2.2 Comparator Offset Compensation

In this design, comparator offset was digitally cancelled offline through the use of dual 4-bit capacitor banks. No special effort was made to solve the problem of comparator offset through either automatic on-chip mixed-signal calibration [45], or analog offset cancellation techniques [12].

In future versions, an offset calibrating latch such as the one found in [12], or an on-chip digital-to-offset feedback scheme [45] should be used.

6.2.3 Improved Sampling Network

Section 4.2.1 described how boosted switches were used to improve sampling linearity. At the 10-bit level, simulations showed that a constant boosted voltage was sufficient. However, if future versions are to include scalability to higher resolutions, constant-V_{GS} boosting can be used to further improve linearity.

6.2.4 Optimization of Digital Power Gating

Section 4.4.1 described how a portion of the digital state machine was power gated during *SLEEP* mode. In this version of the ADC, a significant portion of the digital state machine was not power gated due to the interface between the shift register and the DAC switches. The percentage of digital circuits being power gated can be increased through optimization of the interface, perhaps by employing registers which can hold its state even when being power gated. Increasing the use of power gating will result in even more leakage reduction at low sample rates.

6.2.5 Complete ADC Subsystem

In applications such as medical monitoring where power consumption and form factor are primary concerns, it is important to maximize energy efficiency and minimize the size of the system. An effective way to minmize the size of the system is to integrate as many components on-chip as possible. In terms of the ADC subsystem, this means that it is desirable to include digitally programmable reference generation and clock generation on-chip, in addition to the scalable DC-DC converter mentioned earlier in Section 6.2.1.

6.2.6 Ultra Scalable Hybrid ADCs

The reconfigurable resolution range of this design is from 5 to 10-bits. If the range is to be extended into the noise limited regime at 12-bits or more, challenges such as the trade-off between capacitor sizing at high resolutions and energy efficiency at low resolutions will have to be solved. Quite likely, the approach used in [27] can be leveraged, where two different ADC architectures can be used for low and high resolution settings. Designs with an extremely high degree of scalability can benefit in terms of chip area through a high degree of reuse of on-chip components such as capacitors.

6.2.7 Reconfigurable Analog Front End

Lastly, referring back to Figure 1-1, a reconfigurable ADC can be integrated with a fully reconfigurable analog front-end with programmable gain, bandwidth and noise settings. The amplifier gain can be made programmable in many ways. In the case of a closed loop amplifier, the gain can be digitally programmed by setting a feedback ratio of capacitance. A more analog approach to gain adjustment may include tuning the transconductance of an operational transconductance amplifier (OTA). The front end bandwidth can be digitally set with a simple capacitor bank. Finally, the instrumentation amplifier can be adaptively biased to meet varying noise requirements.

Appendix A

Voltage Scaling Analysis

This appendix will present analysis on how the supply voltage should be scaled such that the ADC achieves a constant FOM as resolution is scaled. For the ADC described in this thesis, the power scaling relationships for each block of the ADC with resolution, frequency and supply voltage is listed in Table A.1.

In the following analysis, we first state three assumptions:

1. Assume that all blocks of the ADC scale quadratically with V_{DD}, which is a good assumption for a fully dynamic ADC with no static bias currents.

2. Assume that any components of power that does not scale quadratically with V_{DD} is small. That is, assume that the leakage power is much less than the active power consumption.

3. Assume that the fractions d_E and d_L stay relatively fixed over resolution. This

Table A.1: ADC Power Scaling Relationships

Block Power	**Resolution**	**Frequency**	V_{DD}
Digital Logic	Linear	Linear	Quadratic
Comparator	Linear	Linear	Quadratic
DAC	$\sim$ Exponential	Linear	Quadratic
Clocks	Linear	Linear	Quadratic
Charge Pumps	Linear	Linear	Quadratic
Leakage	$\sim$ Constant	Constant	$\sim$ Linear

is the weakest assumption, however, this can be achieved by turning off circuits that aren't needed as resolution is reduced.

Also, let any parameter with the subscript H be associated with an arbitrary high resolution mode, and any parameter with the subscript L be associated with an arbitrary low resolution mode. We define the following:

- Let b_H and b_L be the high and low resolution modes respectively.
- Let $V_{DD,H}$ and $V_{DD.L}$ be the supply voltages when operating at a resolution of b_H and b_L respectively.
- Let $P_{TOT,H}$ and $P_{TOT,L}$ be the total power consumption in the high and low resolution modes respectively.
- Let d_L and d_E be the fractions of the total power with linear and exponential power relationships with resolution respectively (see Table A.1). Note, we are assuming that $d_L + d_E \approx 1$.

We wish to determine what the supply voltage $V_{DD,L}$ should be when operating at b_L bits of resolution to maintain a constant FOM when compared to operation at b_H bits. In the ideal case where the power of all blocks scale exponentially with resolution (i.e. $d_E = 1$) such that no voltage scaling is required, we expect that the power at the low resolution mode, $P_{TOT,L,ideal}$, be given by Equation A.1.

$$P_{TOT,L,ideal} = \frac{P_{TOT,H}}{2^{(b_H - b_L)}} \tag{A.1}$$

However, in cases where the power of some blocks scale only linearly with resolution at a constant V_{DD}, the power in the low resolution mode, $P_{TOT,L,constVDD}$, is actually given by Equation A.2.

$$P_{TOT,L,constVDD} = \frac{d_E \cdot P_{TOT,H}}{2^{(b_H - b_L)}} + \frac{d_L \cdot P_{TOT,H}}{b_H / b_L} \tag{A.2}$$

Therefore, to achieve a constant FOM, an additional factor of power savings is required. This is given by the ratio shown in Equation A.3.

$$\frac{P_{TOT,L,constVDD}}{P_{TOT,L,ideal}} = \frac{\frac{d_E}{2^{(b_H-b_L)}} + \frac{d_L}{b_H/b_L}}{\frac{1}{2^{(b_H-b_L)}}} \tag{A.3}$$

We can achieve this additional factor of power savings through supply voltage scaling. Since we have assumed that the power of all blocks scale quadratically with V_{DD}, then we know that for two arbitrary voltages $V_{DD,1}$ and $V_{DD,2}$, the relationship to the power P_1 and P_2 respectively is given by Equation A.4.

$$\frac{P_1}{P_2} = \left(\frac{V_{DD,1}}{V_{DD,2}}\right)^2 \tag{A.4}$$

Rearranging Equation A.4 for $V_{DD,2}$, we have $V_{DD,2} = \frac{V_{DD,1}}{\sqrt{P_1/P_2}}$. Here, P_1/P_2 is exactly the additional factor in power savings that is required. Substituting Equation A.3 for P_1/P_2, and letting $V_{DD,1} = V_{DD,H}$ and $V_{DD,2} = V_{DD,L}$, we arrive at an expression for $V_{DD,L}$ which is given in Equation A.5.

$$V_{DD,L} = \frac{V_{DD,H}}{\sqrt{\frac{\frac{d_E}{2^{(b_H-b_L)}} + \frac{d_L}{b_H/b_L}}{\frac{1}{2^{(b_H-b_L)}}}}} = \frac{V_{DD,H}}{\sqrt{d_E + d_L \cdot 2^{(b_H-b_L)} \cdot \frac{b_L}{b_H}}} \tag{A.5}$$

Note that this derivation assumes that leakage is a small percentage of the total power consumption. At low supply voltages where the bandwidth is limited, this assumption breaks down and leakage starts to limit the energy efficiency of the ADC. Further discussion surrounding Equation A.5 is given in Section 3.2.4.

Appendix B

PCB Design and Test Setup

A 4-layer, 8.4" × 6.9" PCB was designed and fabricated for ADC testing. The layout is shown in Figure B-1. The signals were routed on the top and bottom layers, with a supply and ground plane sandwiched in between. The test board supports a differential analog input from either a XLR connector or two SMA connectors. A high quality input transformer can be used to separate the DC level between the signal source and the input to the ADC. On board *RC* filters were used to limit any high frequency noise to the ADC input.

Digital inputs will be generated with a Tektronix TLA7PG2 pattern generator. Resistive dividers were used to level shift down to the ADC supply voltage. Digital outputs are level shifted up using off-chip LMV72139 comparators from National Semiconductor. The digital outputs will be captured by a Tektronix TLA7NA3 logic analyzer.

All supplies are heavily de-coupled on the test PCB as well as on-chip. The chip will be packaged in a 48-pin, 0.5-mm pitch TQFP package, to be placed into a socket on the test PCB. The chip bonding diagram is shown in Figure B-2.

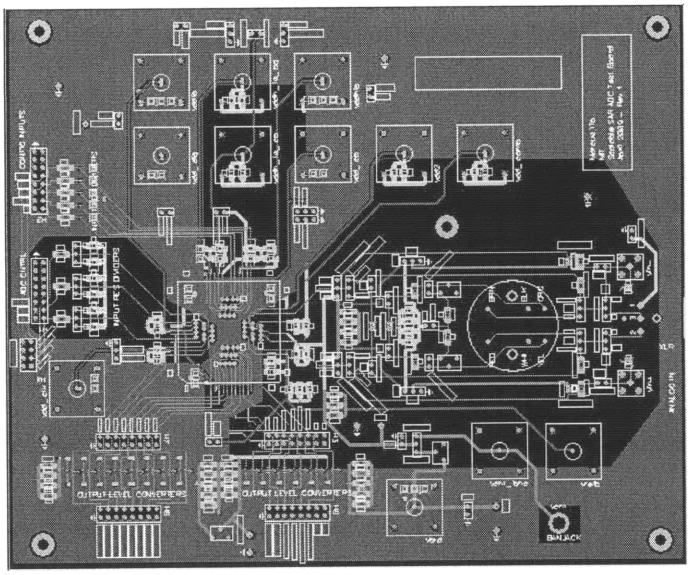

Figure B-1: Layout of the fabricated test PCB.

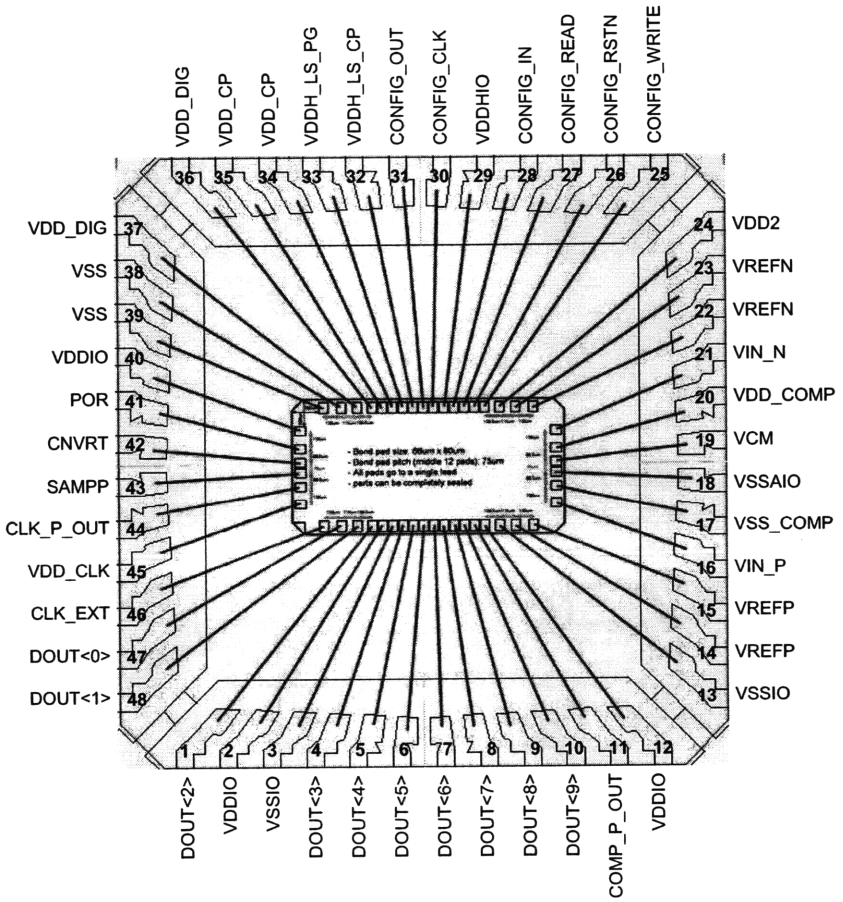

Figure B-2: Bonding diagram of the ADC test chip.

Bibliography

[1] B. H. Calhoun *et al.*, "Design considerations for ultra-low energy wireless microsensor nodes," *IEEE Trans. on Computers*, vol. 54, no. 6, pp. 727–749, Jun. 2005.

[2] M. Yin and M. Ghovanloo, "A low-noise preamplifier with adjustable gain and bandwidth for biopotential recording applications," in *Proc. IEEE Int. Symp. Circuits and Systems*, May 2007, pp. 321–324.

[3] T. Denison *et al.*, "A 2 μW 100 nV/rtHz chopper-stabilized instrumentation amplifier for chronic measurement of neural field potentials," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 12, pp. 2934–2945, Dec. 2007.

[4] M. Bach *et al.*, "Standard for pattern electroretinography," *Documenta Ophthalmologica*, vol. 101, pp. 11–18, 2000.

[5] S. Borromeo *et al.*, "A reconfigurable, wearable, wireless ECG system," in *Proc. IEEE EMBS*, Lyon, France, Aug. 2007, pp. 1659–1662.

[6] G.-Y. Jeong and K.-H. Yu, "Design of ambulatory ECG monitoring system to detect ST pattern change," in *Proc. SICE-ICASE Int. Joint Conf.*, Busan, Korea, Oct. 2006, pp. 5873–5877.

[7] D. A. Johns and K. Martin, *Analog Integrated Circuit Design*, 1st ed. Canada: John Wiley & Sons, 1997.

[8] R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE Journal on Sel. Areas Commun.*, vol. 17, no. 4, pp. 539–550, Apr. 1999.

[9] B. Murmann, "A/D converter trends: Power dissipation, scaling and digitally assisted architectures," in *Proc. IEEE Custom Integrated Circuits Conference*, San Jose, CA, Sep. 2008, pp. 105–112.

[10] H.-C. Hong and G.-M. Lee, "A 65-fJ/conversion-step 0.9-V 200-kS/s rail-to-rail 8-bit successive approximation ADC," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 10, pp. 2161–2168, Oct. 2007.

[11] J. Craninckx and G. Van der Plas, "A 65fJ/conversion-step 0-to-50MS/s 0-to-0.7mW 9b charge-sharing SAR ADC in 90nm digital CMOS," in *IEEE Int.*

Solid-State Circuits Conf. Dig. Tech. Papers, San Francisco, CA, Feb. 2007, pp. 246–247.

[12] N. Verma and A. P. Chandrakasan, "An ultra low energy 12-bit rate-resolution scalable SAR ADC for wireless sensor nodes," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 6, pp. 1196–1205, Jun. 2007.

[13] M. van Elzakker *et al.*, "A 1.9μW 4.4fJ/conversion-step 10b 1MS/s charge-redistribution ADC," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2008, pp. 244–245.

[14] V. Giannini *et al.*, "An 820μw 9b 40MS/s noise-tolerant dynamic SAR ADC in 90nm digital CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2008, pp. 238–239.

[15] B. Verbruggen *et al.*, "A 2.2mW 5b 1.75GS/s folding flash ADC in 90nm digital CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2008, pp. 252–253.

[16] G. Van der Plas and B. Verbruggen, "A 150MS/s 133μW 7b ADC in 90nm digital CMOS using a comparator-based asynchronous binary search sub-ADC," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2008, pp. 242–243.

[17] L. Brooks and H.-S. Lee, "A 12b 50MS/s fully differential zero-crossing-based ADC without CMFB," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2009, pp. 166–167.

[18] M. Boulemnakher *et al.*, "A 1.2V 4.5mW 10b 100MS/s pipeline ADC in a 65nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2008, pp. 250–251.

[19] B. P. Ginsburg and A. P. Chandrakasan, "Highly interleaved 5-bit 250-MSample/s, 1.2-mW ADC with redundant channels in 65-nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 12, pp. 2641–2650, Dec. 2008.

[20] R. H. M. van Veldhoven, R. Rutten, and L. J. Breems, "An inverter-based hybrid $\Sigma\Delta$ modulator," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2008, pp. 492–493.

[21] Y.-S. Shu, B.-S. Song, and K. Bacrania, "A 65nm CMOS CT $\Sigma\Delta$ modulator with 81dB DR and 8MHz BW auto-tuned by pulse injection," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2008, pp. 500–501.

[22] C.-Y. Chen, M. Le, and K. Y. Kim, "A low power 6-bit flash ADC with reference voltage and common-mode calibration," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, HI, Jun. 2008, pp. 12–13.

[23] S.-K. Shin *et al.*, "A fully-differential zero-crossing-based 1.2V 10b 26MS/s pipelined ADC in 65nm CMOS," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, HI, Jun. 2008, pp. 218–219.

[24] B. P. Ginsburg and A. P. Chandrakasan, "A 500MS/s 5b ADC in 65nm CMOS," in *Proc. IEEE Symp. VLSI Circuits*, Honolulu, HI, Jun. 2006, pp. 140–141.

[25] K.-W. Hsueh *et al.*, "A 1V 11b 200MS/s pipelined ADC with digital background calibration in 65nm CMOS," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2008, pp. 546–547.

[26] I. Ahmed and D. Johns, "A 50MS/s (35mW) to 1kS/s (15μW) power scaleable 10b pipelined ADC with minimal bias current variation," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2005, pp. 280–281.

[27] K. Gulati and H.-S. Lee, "A low-power reconfigurable analog-to-digital converter," *IEEE Journal of Solid-State Circuits*, vol. 36, no. 12, pp. 1900–1911, Dec. 2001.

[28] R. Sarpeshkar, *Low-Power Analog VLSI - 6.376 Course Notes*, Massachusetts Institute of Technology, Cambridge, MA, 2007.

[29] J. Y. S. Kwong, "A sub-threshold cell library and methodology," M.S. Thesis, Massachusetts Institute of Technology, Cambridge, MA, Jun. 2006.

[30] M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," *IEEE Journal of Solid-State Circuits*, vol. 24, no. 5, pp. 1433–1440, Oct. 1989.

[31] M. J. M. Pelgrom, H. P. Tuinhout, and M. Vertregt, "Transistor matching in analog CMOS applications," in *Int. Electron Devices Meeting Tech. Digest*, Dec. 1998, pp. 915–918.

[32] B. Razavi, *Design of Analog CMOS Integrated Circuits*, 1st ed. New York, NY: McGraw-Hill, 2001.

[33] A. Wang, A. P. Chandrakasan, and S. Kosonocky, "Optimal supply and threshold scaling for sub-threshold CMOS circuits," in *IEEE Computer Society Annual Symposium on VLSI*, Apr. 2002, pp. 7–11.

[34] J. McCreary and P. Gray, "All-MOS charge-redistribution analog-to-digital conversion techniques I," *IEEE Journal of Solid-State Circuits*, vol. 10, no. 6, pp. 371–379, Dec. 1975.

[35] S. Mutoh *et al.*, "A 1-V power supply high-speed digital circuit technology with multi-threshold voltage CMOS," *IEEE Journal of Solid-State Circuits*, vol. 30, no. 8, pp. 847–854, Aug. 1995.

[36] J. Kao, S. Narendra, and A. P. Chandrakasan, "MTCMOS hierarchical sizing based on mutual exclusive discharge patterns," in *Proc. ACM/IEEE Design Automation Conference*, Aug. 1998, pp. 495–500.

[37] S. Narendra *et al.*, "Scaling of stack effect and its application for leakage reduction," in *Proc. Int. Symp. Low Power Electronics and Design*, Huntington Beach, CA, Aug. 2001, pp. 195–200.

[38] N. Ickes, "A micropower DSP for sensor applications," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, May 2008.

[39] B. P. Ginsburg and A. P. Chandrakasan, "An energy-efficient charge recycling approach for a SAR converter with capacitive DAC," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 1, May 2005, pp. 184–187.

[40] B. Fotouhi and D. A. Hodges, "High-resolution A/D conversion in MOS/LSI," *IEEE Journal of Solid-State Circuits*, vol. SC-14, no. 6, pp. 920–926, Dec. 1979.

[41] N. Verma, "An ultra low power ADC for wireless micro-sensor applications," M.S. Thesis, Massachusetts Institute of Technology, Cambridge, MA, Jun. 2005.

[42] D. C. Daly and A. P. Chandrakasan, "A 6b 0.2-to-0.9V highly digital flash ADC with comparator redundancy," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, Feb. 2008, pp. 554–555.

[43] *IEEE standard for terminology and test methods for analog-to-digital converters.* IEEE Std 1241-2000.

[44] Y. K. Ramadass and A. P. Chandrakasan, "Voltage scalable switched capacitor DC-DC converter for ultra-low-power on-chip applications," in *Power Electronics Specialists Conference*, Jun. 2007, pp. 2353–2359.

[45] S. Gambini and J. Rabaey, "Low-power successive approximation converter with 0.5v supply in 90 nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 42, no. 11, pp. 2348–2356, Nov. 2007.