

Direct Digital Synthesis by Analogue Interpolation


Alistair McEwan

Supervised by Dr Steve Collins

University College

Trinity Term, 2004

This thesis is submitted to the Department of Engineering Science, University of Oxford, in
fulfilment of the requirements for the degree of Doctor of Philosophy



Abstract

An improvement in efficiency of direct digital frequency synthesis (DDFS) systems is demanded for low power frequency synthesis in wireless communications. Concurrently a reduction in cost is important for disposable, low resolution frequency synthesis in biomedical instrumentation systems. To meet both these needs a new ROM-less architecture is presented here that uses less than half the circuit area of previous state of the art systems and improves the efficiency by operating at up to a tenth of the power consumption. The main contribution presented in this thesis is a novel, efficient method of interpolation for DDFS that uses the nonlinear response of the CMOS differential switch already present in the high speed current steering DAC. The nonlinear response provides a smooth transition between the conventional, quantised DAC output. This interpolation may be performed with the conventionally discarded phase bits leading to highly compact and efficient DDFS architectures for application in instrumentation and communications systems.

DDFS systems typically consist of a large overflowing accumulator to generate the phase, a ROM lookup table to convert the phase to amplitude and a DAC to perform the digital to analogue conversion. Approximations are often used to reduce the size of the ROM, however the most efficient DDFS systems remove the ROM completely and calculate the phase to amplitude conversion directly or store the conversion in a non-linear DAC. State of the art, high speed CMOS DACs consisting of thermometer decoded arrays of current steering cells are often used to reduce non-ideal effects that cause unwanted transients leading to a degradation in spectral purity (SFDR).

A novel ROM-less technique is introduced here that uses the non-linear response of a current cell consisting of an ideal current source and differential current switch to interpolate between the output levels of a non-linear DAC. Using this technique two architectures are developed. A compact architecture using only four or six current cells suitable for instrumentation applications and a thermometer decoded architecture using 64 current cells for communications applications that require better spectral purity. The compact architecture is 100% efficient as all the bias current is used to form the output. The only additional component is a small linear phase DAC. One compact system with a nonlinear DAC of four current cells achieved an SFDR of -40dBc up to output frequencies of 1MHz for dielectrophoresis consumed only $5\mu\text{W}/\text{MHz}$ and a second compact system with a six cell nonlinear DAC for electrical impedance spectroscopy, achieved an SFDR of -48dBc for output frequencies up to 1MHz and consumed only $8\mu\text{W}/\text{MHz}$. As an extension to improve the SFDR a segmented system with 64 current cells was developed. The larger number of current cells required the use of a modified thermometer decoder that had the added benefit of improving the spectral purity by linearising the response of each cell. The total active area was 0.6mm^2 , less than half of state of the art ROM-less DDFS systems that include a DAC. Although measurement results of the 64 cell system were disappointing, simulations suggest that these problems may be solved in a future chip that should be able to achieve -70dBc SFDR at 100MHz . Despite the loss in performance from simulation to measurement, the measured 64 cell system still meets the spectral purity requirements of UMTS and Bluetooth, -60dBc SFDR.

Acknowledgements

Firstly, I would like to thank Dr Steve Collins, an attentive supervisor and mentor who greatly enhanced my knowledge of analogue circuits. I would also like to acknowledge the support of University College and Mansfield College and funding from Analogue Devices, the EPSRC and the Overseas Research Scholarship. Finally, I would like to thank my wife, Alexandra Lobb and my family for their constant encouragement.

Glossary

AWG	Arbitrary Waveform Generator.
BSIM	Berkeley IFGET Model.
BT	Bluetooth communications standard.
DAC	Digital to Analogue Converter.
dB	Decibel scale.
dBc	Proportional power to fundamental or carrier in dB.
dBm	Decibel scale with reference to 1 milliwatt = 0dBm.
DDFS	Direct Digital Frequency Synthesis.
DEP	Dielectrophoresis.
DFF	D-type Flip Flops.
DNL	Differential Non-Linearity.
DSP	Digital Signal Processor.
EIS	Electrical Impedance Spectroscopy.
EKV	The compact MOS simulation model proposed in [1].
FG	Function Generator.
IFGET	Insulated Gate Field Effect Transistor.
INL	Integral Non-Linearity.
LPF	Low Pass Filter.
NCO	Numerically Controlled Oscillator.
OPAMP	OPerational Amplifier.
PCB	Printed Circuit Board.
PPA	Parallel Pipelined Accumulator.
PTL	Pass Transistor Logic.
PW	Phase Word (input to the DDFS).
PLL	Phase Lock Loop.
ROM	Read Only Memory.
SFDR	Spurious Free Dynamic Range.
THD	Total Harmonic Distortion.
UMTS	Universal Mobile Telephony.
VCO	Voltage Controlled Oscillator.

Contents

1	Introduction and Background	1
1.1	Thesis Statement	1
1.2	Thesis Overview	2
1.3	Synthesisers	2
1.4	Direct digital frequency synthesis (DDFS)	3
1.5	Applications of DDFS	7
1.6	Spurious signals in DDFS	13
1.7	Summary	26
2	Nonlinear interpolation technique	28
2.1	Introduction	28
2.2	Nonlinear interpolation DACs	29
2.3	Segmented systems	37
2.4	Summary	42
3	Effects of parameter variations	44
3.1	Introduction	44
3.2	Mismatch	45
3.3	Current source mismatch	47
3.4	Current switch mismatch	51

3.5	Summary	56
4	Circuit simulations and design of compact systems	57
4.1	Introduction	57
4.2	Architecture of compact systems	59
4.3	Low frequency circuit simulations of compact systems	60
4.4	Phase DAC	64
4.5	High frequency and mixed signal circuit simulation	67
4.6	Power dissipation	70
4.7	Summary	72
5	Circuit simulations and design of segmented systems	74
5.1	Introduction	74
5.2	Architecture of segmented system	75
5.3	Low frequency circuit simulations	84
5.4	High frequency and mixed signal circuit simulation	86
5.5	Power dissipation	89
5.6	Summary	89
6	Layout and design of the prototype	91
6.1	Introduction	91
6.2	Layout techniques	92
6.3	Layout of four and six cell systems	97
6.4	Layout of segmented current cells	97
6.5	Design and test plan	100
6.6	Layout verification	100

6.7 PCB and test bench design 102

6.8 Summary 102

7 Measurement results 104

7.1 Introduction 104

7.2 Low frequency measurements of the compact systems 105

7.3 Transient measurements of the compact systems 109

7.4 Low frequency measurements of the segmented system 114

7.5 Re-fabrication of segmented system 115

7.6 Transient measurements of the re-fabricated segmented system 119

7.7 Re-simulation of the 64 cell system using new process parameters 121

7.8 Summary 122

8 Conclusion and further work 124

8.1 Conclusion 124

8.2 Further work 128

A A low power parallel pipelined accumulator 131

A.1 Introduction 131

A.2 The architecture 131

A.3 Low power optimisation 134

A.4 Simulation results 136

A.5 Layout of the accumulator 137

A.6 Clock generation sub circuit 140

A.7 Testing of the accumulator 140

A.8 Summary 140

B	Derivations of circuit models	142
B.1	Derivation of the response of a single current cell	142
B.2	Derivations of the effect of variations	143
C	MATLAB Code	145
C.1	Single, four or six nonlinear current cells	145
C.2	Segmented system(linear)	147
C.3	Generation of phase DAC output to program function generator	149
C.4	Generation of four different phase outputs from one accumulator	150
D	Verilog code for mixed signal simulation	152
D.1	13-bit accumulator	152
D.2	Conventional thermometer decoder	154
D.3	Three state thermometer decoder	154

List of Figures

1.1	The architecture of a conventional direct digital frequency synthesizer.	4
1.2	The output of a DDFS for a normalized frequency of 3/16.	5
1.3	A differential current switch.	17
1.4	The position of dominant spurs in the output of a DDFS.	18
1.5	Quarter wave symmetry.	20
1.6	Segmented DDFS of Sunderland and co-workers.	22
1.7	An example of ROM compression by using a Taylor expansion.	22
1.8	The ROM-less DDFS proposed by Mohieldin and coworkers.	25
2.1	A differential current switch.	30
2.2	The current flowing out of the right branch (I_{out2}) of a single current cell in response to a triangular input voltage.	31
2.3	Dominant spurs of a single current cell.	32
2.4	A system based on an array of four current cells.	34
2.5	The combined output current for an array of four cells.	35
2.6	The dominant spurs of a system of four current cells.	36
2.7	The dominant spurs of six current cells.	37
2.8	SFDR of systems with various numbers of cells against bias current.	38
2.9	Linearisation of the current cell response.	40
2.10	The performance of an array of 64 cells versus efficiency simulated in MATLAB.	42

3.1	Mismatch of NMOS transistors of different sizes in a submicron process.	46
3.2	The effect of parameter truncation on the output spectrum of the 64 cell system.	49
3.3	The estimated yield against percentage error in the I_{bias} parameter for various minimum allowable SFDR for compact systems.	50
3.4	The estimated yield against percentage error in the I_{bias} parameter for various minimum allowable SFDR for 64 cell segmented systems.	50
3.5	The effects of variation between the β of the two transistors in the current switch on yield for compact systems.	53
3.6	The effects of variation between the β of the two transistors in the current switch on yield for the 64 cell segmented systems.	53
3.7	A comparison of the effect of β variation on the current cell transfer function.	54
3.8	The estimated yield against percentage error in the phase parameter for various minimum allowable SFDR for the compact systems.	55
3.9	The estimated yield against percentage error in the phase parameter for various minimum allowable SFDR in the 64 cell segmented system.	55
4.1	The architecture of a compact system.	58
4.2	Circuit simulation of an array of four cells.	61
4.3	Circuit simulation of an array of six cells.	62
4.4	Modeling differences between MATLAB and Cadence.	63
4.5	Settling time of the phase DAC signal when driving the four cell system.	69
4.6	Mixed signal simulation results of the four cell system.	71
4.7	Mixed signal simulation results of the six cell system.	71
5.1	The architecture of the segmented system.	75
5.2	A 64-cell thermometer decoded array.	78
5.3	A single current cell with conventional local decoding logic.	78
5.4	The decoding logic necessary for the interpolating thermometer decoder.	79

5.5	The two switch cell design.	80
5.6	The three-branch current switch design.	81
5.7	A conventional current cell with a three state analogue switch.	82
5.8	MATLAB simulations of the system based on an array of 64 current cells after truncation of the parameters.	85
5.9	Circuit simulations of the system based on an array of 64 current cells in Cadence.	85
5.10	The 64 cell system detailing the use of all 13 phase bits required to avoid phase truncation.	86
5.11	Settling time of the phase DAC in the 64 cell segmented system.	88
6.1	Cross section of the layers available in the process showing the formation of a PMOS and NMOS device.	92
6.2	Guard ring schematic.	94
6.3	Layout of the four cell system.	96
6.4	Layout of the six cell system.	96
6.5	Layout of the generic current cell of the 64 cell system.	98
6.6	Description of the layout of the 64 cell system.	99
7.1	Cadence simulation of the four cell system design from section 4.3.	106
7.2	Chip results from a single array of four cells on a typical chip.	106
7.3	Cadence simulation of the six cell system design from section 4.3.	108
7.4	Chip results from a single array of six cells on a typical chip.	108
7.5	Measured variation for the four cell system.	109
7.6	Measured variation for the six cells system.	110
7.7	Measured output waveform of $4kHz$ from the six cell circuit on a typical chip.	111
7.8	Frequency spectrum of the measured output waveform of $4kHz$ from the six cell circuit on a typical chip.	112
7.9	The settling times of the phase DAC signal emulated by the FG.	113

7.10	Variation in the SFDR against the bias voltage of the current sources for 64 cells for the second chip run.	117
7.11	The measured transient performance of the segmented system on a typical chip.	120
8.1	Differential current switch with 15V outputs.	129
8.2	A system capable of generating four phase shifted outputs at 0°, 90°, 180° and 270°. . .	130
8.3	Cadence simulation results of the 15v outputs phase shifted by 0°, 90°, 180° and 270°. All signals had a SFDR of $-42dBc$	130
A.1	A 24bit, 3 stage pipelined accumulator	132
A.2	(a) A 24-bit, 3 stage parallel pipelined accumulator (b) Parallel adder block	133
A.3	Delay element (D_1), positive edge triggered	135
A.4	Hybrid full adder (1-bit). (a) Sum stage. (b) Carry stage.	136
A.5	Delay element D_2	137
A.6	Final low power parallel pipelined accumulator.	138
A.7	Layout of the delay element D1.	139
A.8	Layout of the delay element D1bar.	139
A.9	The clock generation sub circuit.	141
B.1	The differential current switch.	142

List of Tables

1.1	Alternative frequency synthesisers [10].	3
1.2	Minimum SFDR and number of phase, amplitude and accumulator bits required by each application.	14
1.3	State of the art low power ROM-less DDFS systems.	24
2.1	The MATLAB predicted, optimum SFDR for systems containing different numbers of current cells.	38
2.2	MATLAB results showing the expected best SFDR for systems containing different numbers of current cells.	40
3.1	Truncated $10 \times I_{bias}$ for the four cell array.	47
3.2	Truncated $30 \times I_{bias}$ for the six cell array.	48
3.3	Truncated $1000 \times I_{bias}$ for the segmented system of 64 cells.	48
3.4	Allowable percentage error required to ensure a yield of better than 90% of systems. .	56
4.1	Summary of circuit simulation parameters for the compact systems.	60
4.2	The transient performance of six cells.	66
4.3	The transient performance of four cells.	67
4.4	Power dissipation breakdown of the compact systems.	70
5.1	The actual dimensions of the current switches in the 64 cell segmented system.	83
5.2	Performance of the 64 cell system as f_{clk} is increased.	88
5.3	Power dissipation breakdown of the 64 cell system.	89

6.1 Prototype test plan. 101

8.1 A comparison between the work presented here and recently published, state of the
art low power ROM-less DDFS systems. 126

A.1 Comparison of the development stages of the low power accumulator. 132

Chapter 1

Introduction and Background

The ability to create and finely control the frequency, phase, and amplitude of a signal is important in many engineering systems including radio communications and instrumentation [2, 3]. For example, instrumentation systems with sensors that measure impedance by comparison of signals with fine phase and frequency resolution benefit from well controlled AC signal sources or frequency synthesizers [4, 5]. In some instrumentation applications low power, inexpensive synthesizers are desirable to create integrated arrays of disposable sensors [6, 7]. Another example of a synthesizer application is the recent interest in spread spectrum techniques that alleviate congestion in the radio frequency band by using fast frequency hopping and good frequency resolution to share a limited bandwidth among many channels [8]. In mobile and wireless devices low power operation is of importance to conserve battery power.

1.1 Thesis Statement

A new ROM-less direct digital frequency synthesis (DDFS) architecture is proposed to show that it is possible to create compact and efficient DDFS systems through carefully controlling the nonlinear response of the CMOS differential current switch by use of the excess input bits of a conventional DDFS to meet the need for more compact and efficient integrated DDFS systems.

The main contribution presented in this thesis is a novel, efficient method of interpolation for DDFS that uses the nonlinear response of the CMOS differential switch already present in the high speed current steering DAC. The nonlinear response provides a smooth transition between the conven-

tional, quantised DAC output. This interpolation may be performed with the conventionally discarded phase bits leading to highly compact and efficient DDFS architectures for application in instrumentation and communications systems.

1.2 Thesis Overview

The structure of this thesis is as follows. In the remainder of this chapter the theoretical basis of DDFS and some application examples are described. State of the art DDFS systems are reviewed focusing on power and area efficient designs. Chapter 2 presents the novel nonlinear interpolation technique and MATLAB simulations are used to identify two architectures for the targeted applications. Further MATLAB simulations investigate the effect of parameter variations in Chapter 3. Chapter 4 introduces the first architecture that is based on a compact array of current cells and confirms the MATLAB predictions through circuit simulations. Chapter 5 outlines the second architecture that allows for arrays of larger current cells and again circuit simulations are used to confirm performance. Chapter 6 discusses layout techniques and details the layout and test design of the prototype. The measurement results of this prototype are given in Chapter 7. Chapter 8 concludes this thesis.

1.3 Synthesisers

A frequency synthesizer is an active electronic device which takes a reference frequency and, based on a control mechanism, produces one or more frequencies which share the accuracy and stability of the reference. These are preferred to single frequency sources due to the ability to tune frequencies for either modulation or channel selection.

A simple, tuneable frequency source is based on a resonator circuit where the characteristics are varied to set the frequency, for example in a voltage controlled oscillator. These circuits are used in frequency synthesisers to improve the performance and broaden the operating range. Two popular frequency synthesisers, DDFS and phase locked loops (PLL) are compared in Table 1.1. Currently in CMOS, PLLs are faster and consume less power than DDFS and this is the reason why they are dominant in wireless transceivers. However DDFS has many advantages that fit well with spread spectrum communications, namely high frequency resolution and very fast and continuous switching over a larger tuning range. This advantage is mostly due to the direct nature of DDFS compared to

the indirect, feedback based approach in PLL. There is great interest in improving the efficiency and bandwidth of CMOS based DDFS to bring these advantages to wireless and portable communications and instrumentation systems.

Attribute	DDFS	PLL
Frequency Resolution	mHz	kHz
Switching Time	ns	ms
Switching Effect	phase continuous	phase errors
Tuning Range	DC to Nyquist	limited by feedback loop
Phase Noise	determined by f_{clk}	worse then f_{clk}
Bandwidth	up to 1GHz	up to 100GHz
Power Dissipation	Historically High	micro watts

Table 1.1: Alternative frequency synthesisers [10].

Advantages of DDFS over other frequency synthesisers are the ability to rapidly change frequency (hop) over a wide frequency range and a very fine frequency resolution. In comparison to phase lock loops, DDFS has no stability problems as it does not use feedback and can therefore support fast and accurate frequency hopping. The stability and phase noise of the output is determined by the reference frequency which is typically provided by a crystal resonator. This agile output, combined with a natural digital interface, suggests DDFS is ideal for modern communications. DDFS is also a good candidate for a instrumentation applications as the spectral resolution is directly determined by the number of bits used in the system. By using a large number of bits it is possible to achieve a resolution of mHz in a bandwidth of MHz .

1.4 Direct digital frequency synthesis (DDFS)

All conventional DDFS systems are based on the Modulo-N architecture of Tierney, Rader and Gold [9]. The architecture consists of an overflowing accumulator, a read only memory (ROM) and a digital to analogue converter (DAC), Figure 1.1. The function of the accumulator is to generate an overflowing digital ramp that represents the instantaneous phase of the output. This digital phase word is then converted to the corresponding digital amplitude word by a lookup table stored in the ROM. Finally, the digital amplitude word is converted to the equivalent analogue signal by the DAC.

Consider the example depicted in Figure 1.2. If the frequency control word or phase word, k is chosen to be 3 and the accumulator size is $L = 4$ bits, the consecutive values of the accumulator output for the first 7 cycles of the reference clock will be 0, 3, 6, 9, 12, 15, 18. The size of the phase accumulator is $2^4 = 16 < 18$ so an overflow will occur with remainder of 2. Rather than 18 the accumulator




Figure 1.1: The architecture of a conventional direct digital frequency synthesizer. First proposed by Tierney, Rader and Gold [9].

contents will correspond to 2. $k = 3$ is continually added to this remainder and the accumulator output will be the sequence: 2, 5, 8, 11, 14, 17. Again an overflow occurs, this time with a remainder of 1 and the next sequence is: 1, 4, 7, 10, 13, 16. There is now no remainder. The process repeats with the sequence: 0, 3, 6....etc. (see Figure 1.2). If the reference clock is $1Hz$ then 3 overflows have been completed in 16 cycles of the reference clock and the frequency $3/16Hz$ has been generated.

For an L -bit accumulator working at a clock frequency of f_{ref} , with an input word k the DDFS creates an output at the frequency:

$$f_{out} = f_{ref} * k/2^L \quad (1.1)$$

Equation 1.1 shows that the stability of the output is based on the reference clock and that the output can be controlled accurately and quickly, by changing k . More importantly the frequency resolution is $f_{ref}/2^L$, indicating that high frequency resolution is possible by using a large accumulator. Furthermore, changes between these frequencies only take one clock cycle.

Consider again the conventional DDFS architecture shown in Figure 1.1. The accumulator output, which represents a phase in the range of 0 to 2π , is used to address a ROM lookup table that stores the nonlinear phase to amplitude conversion. Each memory address in the ROM contains the amplitude of the sine wave at the corresponding phase. Therefore the output of the ROM is a digital representation of the sine wave amplitude. This digital signal is then fed through a digital to ana-




Figure 1.2: The output of a DDS for a normalized frequency of $3/16$. Adapted from [3].

logue converter (DAC) to create the required analogue sine wave at the same accumulator overflow frequency.

The bandwidth of the output signal is limited to 1/2 the reference frequency due to the Nyquist sampling rate. However in real systems unwanted spurs are introduced into the Nyquist band due to constraints when building a real ROM and DAC. These degrade the spectral purity of the output signal and are the focus of most DDFS research. The origin of the unwanted spurs will be discussed in section 1.6.

The key measures of performance for DDFS systems are commonly specified in the frequency domain and relate to the quality of the signal generated. The frequency accuracy and resolution are directly determined by the number of bits in the accumulator, Equation 1.1. Conventionally the noise floor and total harmonic distortion (THD) are the figures of interest for a synthesized frequency. The noise determines a lower bound on signal power while the distortion or linearity determines an upper bound. The noise floor of the output of a DDFS is dominated by two effects, phase and amplitude quantisation, and the phase noise of the reference clock. The phase noise of the reference clock is reduced at the output as the output frequency is always less than the clock. This leads to an improvement of $20 \times \log(f_{out}/f_{clk})$ dB [10, 11]. Hence quantisation will typically dominate the noise floor as the reference clock in a DDFS is at least four times faster than the output and hence the phase noise is 12dB less than the phase noise of the reference. The effect of phase and amplitude quantisation are introduced in Section 1.6 and are typically chosen so that both have an equal effect on the noise floor, i.e. two more phase bits than amplitude bits are typically used.

A quantitative measure of spectral purity, the spurious free dynamic range (SFDR) has become popular in giving an overall performance figure of merit for DDFS, DACs and amplifiers by describing the tradeoff between noise floor and distortion, or linearity in one figure.

The SFDR is defined to be the difference in dB between the power at the desired output frequency and the largest spurious frequency in the spectrum (labelled dBc). The spurious tone may or may not be harmonic and therefore the SFDR is a more faithful figure of the available dynamic range than THD [12]. Many sources quote narrow-band SFDR [13, 14, 15] which assumes that a high order low pass filter will be provided to remove spurs at higher frequencies. For inexpensive, low power systems the whole bandwidth should be considered to avoid the need for additional filters which will increase power consumption and cost.

To reduce the power consumption and cost both the ROM and the DAC are minimised. This intro-

duces two sources of spurious signals within DDFS systems which lead to a reduction in the SFDR. Firstly the ROM size is minimised by using the minimum number of phase and amplitude bits. However this leads to approximation errors in the output. Secondly, an important source of error is the imperfections in the operation of switches within the DAC and DAC nonlinearity which causes harmonic distortion. A well-designed, compact, low power system needs to balance both types of error to ensure neither the ROM nor the DAC consume an unnecessary amount of power whilst ensuring the system achieves the spectral purity required by the application of interest. Many techniques have been proposed to deal with quantisation errors and these will be discussed in section 1.6.3 after considering some DDFS applications.

1.5 Applications of DDFS

Frequency synthesis has many applications. Those considered here are its use in communications technologies and instrumentation systems. Communication technologies use frequency synthesis to describe a common signal that can be accurately generated at the transmitter and receiver. One of the advantages of DDFS for communications is the ability to quickly and accurately change the carrier frequency (hop) to efficiently allocate many channels in a limited bandwidth by using spread spectrum techniques. Instrumentation systems commonly use frequency synthesis as a standard, reproducible source that can be used to guarantee the accuracy of a measurement. One of the advantages of DDFS in instrumentation is its good frequency resolution that allows fine trimming or calibration to remove variations in the synthesizer itself or in the rest of the measurement system.

1.5.1 Communications

The frequency synthesizer is one key component of communications transceivers. DDFS provides a convenient interface between the analogue communications channel and the digital baseband processing that has enabled Rofougaran and co-workers to integrate a complete transceiver on a single CMOS chip [16]. CMOS is an attractive technology for this application due to its low cost and wide use in highly integrated digital systems.

Equation 1.1 shows that the output frequency is very stable, can be controlled very accurately and changed within one clock period simply by changing the frequency control word or phase word (PW), k . DDFS offers a wide bandwidth by using the output of a large accumulator, leading Mitola to predict

fundamental changes in the architecture of radio transceivers [8]. More baseband processing may be performed digitally and the front end analogue electronics minimised leading to a software defined radio architecture where many channel characteristics, such as the carrier frequency, may be changed dynamically [17, 18, 19]. The potential benefits of DDFS characteristics to wireless communications systems has attracted significant interest [20, 21, 22, 13, 23, 24].

Bluetooth

A recent wireless communication standard in the ISM (Instrumentation, Scientific and Medical) radio frequency band is Bluetooth (BT). The standard is aimed at replacing cables for communications over 0 – 100m. Commercially successful applications of BT have appeared in peripheral devices around personal computers, wireless hands free interfaces with mobile phones and sensors in automobiles. In many of these applications cost is a critical factor. One method of keeping costs low is a single chip BT transceiver, preferably in low cost CMOS, with a minimum of external components.

One of the most important and challenging components of the BT transceiver to design is the frequency synthesizer. The BT standard [25] allows for most of the BT transceiver to be powered down when not in use to reduce power consumption. However the receiver is required to be turned on for long periods of time to find a synchronisation burst sent from other BT devices, therefore low power frequency synthesis is highly desirable.

The channel frequencies span 80MHz and continuously hop between channels at 1600hops/s. The channel spacing is 1MHz with frequency accuracy of better than 1ppm. A DDFS frequency synthesiser for BT would require an accumulator of 20 bits and would be need to span the entire 80MHz as each connection hops through the entire 80 channels. This may be achieved using reference frequency of at least 160MHz. The time division duplex (TDD) ¹ communication used in Bluetooth enables the decoupling of the transmit architecture from the receive architecture as the transmitter and receiver are never active at the same time. However to reduce both power consumption and circuit size, the transmit and receive circuits should share the same synthesizer.

The transmitter specifications are defined in three classes. Class 1 has the highest output power of 20dBm while Class 2 has an output power of 4dBm and Class 3 is the default class with the lowest output power of 0dBm. The transmit spectrum is defined by a mask that requires second harmonic

¹Time-division duplex (TDD) refers to duplex communications links where the up-link is separated from the down-link by the allocation of different time slots in the same frequency band.

spurs to be less than $-20dBm$ and other spurs less than $-40dBm$. This sets the SFDR requirement of the transmit side for Class 1, $-60dBc$ for Class 2, $-44dBc$ and for Class 3, $-40dBc$.

Bluetooth devices are likely to be used in environments where there are a large number of devices, for example the large number of peripheral devices that can connect to a PC. For this reason the specification emphasises inter-modulation from interference between devices rather than range or noise limitations. The SFDR of the receiver is set by the inter-modulation distortion of two devices transmitting on adjacent channels. Low power Bluetooth devices will probably use the default, low power output class 3 with output power of $0dBm$ ($1mW$). The specifications allow the signal to attenuate $69dBm$ over $10m$ leading to a receiver sensitivity requirement of $-69dBm$. Demodulation at the receiver requires a signal to noise ratio of $18dBm$ for the carrier to ensure a bit error rate of 0.1%, also part of the specifications. These set the noise floor of the receiver (N_{rx}) to $-87dBm$. As shown by Lee, SFDR is the interference level when the third order inter-modulation products of the interfering devices reach the same level as the receiver noise floor, in decibels [12]. (The noise floor is taken to be at the same level as the highest harmonic or non-harmonic spur.)

$$BT_{SFDR} = 2/3(IIP_3 - N_{rx}) \quad (1.2)$$

The $2/3$ factor is introduced by the third order of the inter-modulation products [12] (IIP_3) which is given as $-16dB$. This leads to SFDR requirement of $-50dBc$ using Equation 1.2, [26].

UMTS: The universal mobile telephone standard

The universal mobile telephone standard (UMTS) [27] is the basis of the 3rd generation mobile phone system to allow enriched data services over mobile devices such as video and multimedia. UMTS uses spread spectrum to efficiently allocate the spectrum between a large number of users. These services will be provided in conjunction with existing second generation standards. The UMTS standard recommends a frequency resolution of $200kHz$, a bandwidth of $60MHz$ and a switching time of $250\mu s$. The $60MHz$ bandwidth has been allocated to different service providers and the maximum bandwidth required for a single operator is $15MHz$. The SFDR set by the transmit mask levels for the standards is $-58dBc$ [28]. The frequency accuracy is set by a requirement of less than 0.5 degree phase error in the transmitter. For the $200kHz$ frequency resolution this translates to a frequency error of $278Hz$. To achieve this without placing a significant burden on the clock frequency precision an 18-bit accumulator is required (Equation 1.1) in an integrated DDFS system.

1.5.2 Instrumentation

A review of the instrumentation literature shows the use of discrete, commercially available DDFS in a wide range of instrumentation systems: AC metrology [29], determination of dielectric loss in lossy capacitors [30], capacitance microscopy [31], blood volume meters [32], force sensors [33], audiometers [34] and acoustic wave sensors [35]. A common theme is the use of DDFS to provide an easily variable frequency source with good spectral purity and direct interface to digital electronics. Low cost, integrated direct frequency synthesis is particularly useful in instrumentation systems that require the generation of many phase related signals and low power is important when the systems are battery powered. Two examples given here are dielectrophoresis and electrical impedance spectroscopy.

Dielectrophoresis

A potential application for frequency synthesis systems is to generate the AC electric fields required for AC electrokinetic phenomena. There are three different AC electrokinetic phenomena, dielectrophoresis, (DEP), traveling wave dielectrophoresis, (TWD) and electrorotation, (ROT), that can be used to move, separate and characterise small particles such as cells within microsystems that are variously known as laboratory-on-chip or biofactory-on-chip systems.

Dielectrophoresis (DEP) is the motion of suspended, electrically neutral particles in response to a non-uniform electric field. This motion is caused by the interaction of the non-uniform field with the poles of polarized, but otherwise neutral particles. Any electric field gradient causes a force

$$F_{dep} = 2\pi r^3 \epsilon_m \text{Re}(K(\omega)) \nabla E^2 \quad (1.3)$$

which depends upon the real part of the Clausius-Mossotti factor

$$K(\omega) = \frac{\epsilon_P^* - \epsilon_M^*}{\epsilon_P^* + 2 * \epsilon_M^*} \quad (1.4)$$

where $K(\omega)$ represents the difference between the complex permittivity of the particle (ϵ_P^*) and the suspension medium (ϵ_M^*). Since $\epsilon^* = \epsilon - j\sigma/\omega$ this force depends upon the permittivity and the conductivity of the particle and the medium. Over the past decade several groups have demonstrated that it is possible to separate mixtures of different cells by careful selection of the frequency of the applied electric field. For example in 1994 Becker and co-workers showed that it is possible to collect

both blood and leukaemia cells in the high field regions of a micro-electrode array by applying a frequency of 200 kHz to a sample. However, when the frequency is reduced to 80 kHz, the blood cells became detached from the high field region. They could then be swept away from the microelectrode array by a gentle flow of the suspension medium and thus separated from any leukaemia cells [6]. More recently Li and co-workers have shown that viable canola cells experience different forces to nonviable cells over the frequency range of $10Hz - 5kHz$ [7] and Green and co-workers have shown that rod-shaped viral cells experience DEP forces at frequencies of up to $1MHz$ [36]. Hence the useful bandwidth of the DEP phenomena is from close to DC up to a $1MHz$.

The dielectrophoretic phenomena can be exploited in systems which employ DEP to concentrate, separate and characterise cells [37]. The result is a system that can be designed to rapidly and automatically perform a wide range of complex diagnostic tasks on small volume samples in a single, miniaturised, low-cost package. This type of system has a wide range of applications in areas such as medical and single-cell diagnostics, chemical detection and water quality control.

This type of system could be created using commercially available, discrete DDFS integrated circuits for applications that require an AC electrokinetic signal with a SFDR of better than $-50dBc$. However, the AC electrokinetic forces are proportional to the square of the generated electric field, Equation 1.3. This means that a voltage signal with a SFDR of better than $-30dBc$ will have a force at the fundamental frequency that is a thousand times larger than that created by the largest spur. Furthermore, the characteristics of cells exploited in AC electrokinetic phenomena vary relatively slowly with frequency. Hence large frequency changes, such as the change from 200 kHz to 80 kHz used by Becker and co-workers, are more important than absolute frequency accuracy. This relaxes the frequency resolution requirements, reducing the number of bits required in the accumulator. Finally, many applications will require a disposable system and a primary feature for any disposable system must be cost and for an electronic component this equates to a compact, high yield integrated circuit. The ideal DDFS system for dielectrophoretic based diagnosis is therefore one which combines a SFDR of better than $-30dBc$ with compactness and a high yield.

Electrical Impedance Spectroscopy

Electrical Impedance Spectroscopy (EIS) is a method of determining the impedance of the contents of an object by measuring the impedance changes over frequency on the surface of the object. Typically currents or voltages are induced into the object at known electrode locations on the boundary and

the resulting voltage or current is measured. If a sufficient number of electrodes are used, and a large number of independent measurements are taken, the impedance of the contents of the object can be reconstructed from the measured electrical response. This impedance information can be used to form an image of inside the object in electrical impedance tomography (EIT). The pixels in the image represent conductivity changes, $\Delta\sigma$ over time or frequency. These conductivity changes are related to the measured voltages, ΔV , in Poisson's equation, expressed in matrix form:

$$\Delta V = A \Delta\sigma \quad (1.5)$$

The sensitivity matrix, A , is calculated either using an analytical model of the shape of the object or using a finite element mesh. The matrix is inverted using a technique such as singular value decomposition, and used to calculate, or reconstruct changes in conductivity from the measured voltage changes:

$$\Delta\sigma = A^{-1} \Delta V \quad (1.6)$$

Impedance changes may be obtained from the reciprocal of the change in conductivity. As Holder and co-workers point out it is important to collect a large number of independent measurements because the reconstruction process, Equation 1.6 is ill-posed [38].

The advantages of imaging EIS over other imaging methods such as X-ray are low cost and safety. EIS and electrical capacitance tomography have been successfully applied to wide range of industrial processes [39]. There is active research on the use of EIS in the clinical environment [38] however it has not yet been shown to be better than existing methods. Two examples where EIS would have a large potential are in imaging of the breast where there are safety fears of using ionising radiation such as x-ray or imaging of the brain as a low cost and portable alternative to MRI for diagnosis of tumors or stroke.

Existing EIS systems use 12-bit data converters and digital signal processors (DSPs) for waveform generation. The use of EIS in a clinical environment is limited by two sources of parasitic impedance. Firstly the need for long electrode leads introduces parasitic capacitance and secondly the bioelectrical interface at the skin introduces a complex impedance that has a frequency dependant capacitance [40]. By integrating the electronics within the electrode in a wireless 'active electrode', instrumentation errors can be reduced and patient comfort may be increased by the removal of wires. One of the major limitations is power consumption and hence this technique would benefit from a low power frequency synthesizer. The conventional method of reducing skin impedance is to remove the dead

skin by using abrasion. There are safety concerns over the transmission of blood born viruses such as hepatitis using this technique and thus a low cost disposable active electrode would be particularly advantageous.

In existing systems the signal to noise ratio is $40dB$ and hence a DDFS with spectral purity better than $-40dBc$ would suffice [40]. Similar to the DEP application a wide frequency range is more important than absolute frequency accuracy as the impedance differences in biological tissue occur slowly over a wide frequency range from near DC to $1MHz$.

Hence EIS and DEP are two instrumentation application areas for DDFS that have been identified both with relaxed spectral purity and frequency resolution requirements. These provide new needs for very low cost, disposable DDFS systems.

1.6 Spurious signals in DDFS

A large phase accumulator capacity is required for high frequency resolution (for example over 18 bits to meet the UMTS frequency stability requirement of $200Hz$ for a bandwidth of $200MHz$, see section 1.5.1). The ROM would be very large if all 18 phase bits were used. For example, if an 8-bit word is used to represent the amplitude of the sine wave, then a memory of $64MB$ would be needed.

Various ROM strategies may be used to implement a high speed, low power $64MB$ ROM in CMOS. The conventional ROM is based on a two dimensional array of NOR gates with pull down transistors in parallel or NAND gates based on pull down transistors in series. NOR arrays are the most popular used as they are faster than NAND arrays. To increase speed most ROMs precharge during one phase of the clock and evaluate in the other. The main source of power consumption is in precharging the drain and bit line capacitance. Using the figures from a survey of low power ROM techniques by de Angel and coworkers [41], a low power NOR ROM would consume in the order of $100\mu A$ at $100MHz$ using selective pre-charge, voltage supply scaling and buffer sizing. By using charge recycling and charge sharing techniques this may be reduced by an order of 10 to less than $8mW$ at $100MHz$ in an area of $0.51mm^2$ [42]. This is a high power and size consumption considering that an analogue output DDFS would also require an accumulator and DAC.

The conventional solution is to limit the ROM size by reducing the number of phase and amplitude bits. However this can only be achieved by approximations that introduce distortion into the output signal. This distortion appears as undesirable, spurious frequency components in the output that

in turn degrade the spectral purity (SFDR). Additionally non-ideal DAC characteristics degrade the spectral purity.

1.6.1 Approximation errors

In order to limit the size of the ROM, the number of output bits from the phase accumulator, L is truncated to the first, W bits and the remaining, $L - W$ bits are discarded. The consequence is a phase modulation of the output. An estimation of the level of the largest spurious signals due to the truncation of the phase bits is [43]:

$$Phase\ Truncation = 20\log\left(\frac{a_{sp}}{a_{car}}\right) \approx -6W\ dBc \quad (1.7)$$

where a_{sp} is the power of the spur and a_{car} is the power of the carrier, [3].

Equation 1.7 sets the number of separately addressable amplitude values that are stored as words in the ROM. The number of amplitude bits are also quantised to a word-length of S and the consequence is a spurious amplitude modulation of the output signal. The effects of amplitude quantisation can be estimated as [43]:

$$Amplitude\ Quantisation = -6S - 13.5\ dBc \quad (1.8)$$

Equations 1.7 and 1.8 can be used to estimate the ROM size required to achieve a target SFDR. As a general rule, there are 2 more phase bits than amplitude bits to make the highest phase truncation spur component smaller than the amplitude quantisation noise floor [43]. The minimum required phase bits and amplitude bits for the applications discussed in section 1.5 are listed in Table 1.2.

Application	SFDR(dBc)	min phase bits	min amplitude bits	min accumulator bits
BT	-50	9	7	20
UMTS	-60	10	8	18
DEP	-30	5	3	5 +
EIS	-40	7	5	7 +

Table 1.2: Minimum SFDR and number of phase, amplitude and accumulator bits required by the applications introduced in section 1.5. + The number of accumulator bits are set by the phase bit requirements.

1.6.2 Imperfections in high speed DACs

Spurious components in the DDFS outputs can also be introduced by imperfect operation of the DAC. The output of the DAC is degraded by both static and dynamic errors. Static errors within high speed DACs and the analogue circuit design techniques to overcome them have been well documented in the literature [44], [45], [46]. All dynamic errors depend on bandwidth and are currently one of the hardest problems in high speed CMOS DAC design and are an active area of research [47], [48]. Dynamic errors important for DDFS applications include delays, settling times, and clock feedthrough.

Static errors can be defined in terms of differential non-linearity (DNL) and integral non-linearity (INL). DNL is the difference between two successive digital codes, measured in fractions of an LSB while INL is the error between the analogue DAC output and the ideal output. INL can also be thought of as the cumulative DNL error.

The most widely used DAC architectures are based on the binary code and the thermometer code. A binary coded DAC uses a different number of devices for each bit. The N^{th} device is given the weight 2^N where N is the number of bits. One of the drawbacks is that for a large number of bits the difference between the MSB and LSB is large and hence the DAC becomes sensitive to parameter variations. The thermometer decoder helps to reduce these effects by having a large number of unused input codes. In a thermometer code the number of 1's represents the decimal value or the magnitude of the input in a similar way to the level of a thermometer represents a temperature. The lowest number is represented by all bits equal to 0. If the input increases, more significant bits will turn on and if it decreases less significant bits will turn off. In a thermometer decoded DAC the number of 1s corresponds to the number of devices that are on and the 0s the devices that are turned off in a similar way to a linear string of LEDs that represent the magnitude or level of an input. The weights of all devices are equal implying that the number of devices requires is $2^N - 1$ and only one change between equal devices is made between two adjacent codes.

Monotonicity can be guaranteed if the deviation from the best-fit straight line (INL) is less than half an LSB which implies that the DNL error is less than one LSB. Non-monotonic behaviour is usually seen at the MSB transitions in binary-weighted DACs. Some DAC architectures such as thermometer coded DACs are monotonic by design.

Dynamic errors

Dynamic errors are mainly the result of delays between switching signals caused by parasitic impedance [49], clock feedthrough and digital crosstalk. These errors or 'glitches' will cause the output to be incorrect for an instant and the sum of these errors will reduce the spectral purity of the output. As the clock frequency is increased dynamic errors become a larger portion of the period and the spurs associated with these errors increase in power leading to a degraded SFDR.

In a binary encoded DAC the glitch is usually worse when more devices switch simultaneously. For example in a binary encoded 3-bit DAC, a change in the input code from 3 (011) to 4 (100) will result in all the current sources switching. The devices representing the first and second bits, 2^0 and 2^1 , will be turned off and the devices representing 2^2 will be turned on simultaneously. In this example only three devices are changing whereas in a large DAC many more will switch. In the example of a 3-bit DAC, if the devices controlled by the first two LSBs turn off before the MSB device turns on the output will temporarily fall to zero. Alternatively, if the MSB device turns on before the LSB devices turn off, the output will temporarily rise to its full value. The most popular way to reduce these glitches is to minimise the number of devices which need to be switched simultaneously by using a thermometer decoder.

Current mode DACs

For high bandwidth applications current mode converters are commonly used as they can drive a resistive load directly at high speed without the need for output buffers. A simple current source is a single transistor where the gate voltage is directly connected to the digital input so that the transistor is switched completely on or off. The output is directly connected to the drain and the output current will depend directly on the drain voltage because the output resistance is small. An improvement is a cascode current source where a switch transistor is placed in series with the current source to increase the output resistance of the current source. Problems with this design are the large clock swings required to completely turn off the output current path. Large voltage swings lead to large glitches at the output due to capacitive coupling.

A preferred current cell uses a differential pair which can switch currents with minimum drain voltage excursions, Figure 1.3. In this circuit the current I_{bias} , is either steered to the output or steered to the power rails depending on the difference between the gate voltage of each transistor, V_{in} and V_{ref} . A

particular advantage of the design is that the swing of the control voltage, $V_{low} < V_{in} < V_{high}$, may be reduced minimising the voltage change required to switch current. Reduced voltage swings lead to a reduced size of the resulting glitch. For this reason most high speed, CMOS DACs use this type of current cell.




Figure 1.3: A differential current switch. The current is steered between the two transistors by changing the input voltage.

Clock feedthrough and digital crosstalk

When MOS switches are used to switch analogue signals with a digital clock there is capacitive coupling between the digital switching signal and the sensitive analogue signal. This is referred to as clock feedthrough and occurs at the rising and falling edges of the switching signal. As clock feedthrough occurs every clock cycle, in the frequency domain there will be a spur at the clock frequency, f_{clk} . Clock feedthrough has more effect for some phase words when the output frequency is increased as the aliased frequencies interfere constructively in the Nyquist band as depicted in Figure 1.4 [3].

In the time domain the effect will be similar to a glitch in current mode DACs where it causes an error in the final value. The clock feedthrough can be reduced by decreasing the size of the coupling capacitances or by decreasing the clock swing.

In current mode DACs the coupling capacitances are the parasitic capacitances of the MOS devices, predominantly the gate to source capacitance, C_{gs} . Therefore using minimum size MOS devices for the switch should reduce clock feedthrough. However the on resistance of the switches increases with smaller transistors which may decrease performance due to increased settling time. Where possible a good compromise is the use of transmission gates where the on resistance will be low over most

voltage levels. The NMOS device has a high resistance at low voltage levels but this is reduced by the parallel combination of the low resistance of the PMOS device. At high voltage levels the situation is reversed. The charge feedthrough will also tend to cancel as the PMOS transistor absorbs the charge that the NMOS repels as it goes into its cut-off region [50].




Figure 1.4: The position of dominant spurs in the output of a DDFS. (a), (b) and (c) show the effect of clock feedthrough alias as the phase word is increased. Most DDFS systems report frequencies up to $F_{clk}/4$ (b). When the phase word is increased above an output frequency of $F_{clk}/4$ the spurs are increased due to the mixing of clock aliases (a). Adapted from [51]. NB SFDR is Narrowband SFDR and WB SFDR is Wideband SFDR, 'a' denotes the alias of the fundamental, i.e. '9th a' is the alias of the 9th harmonic.

Techniques in the physical layout of the DAC, such as separate analogue and digital ground and guard rings between the digital and analogue circuits, prevent charge leakage through the substrate and can reduce clock feedthrough and digital crosstalk leakage to the analogue components. An effective de-glitching circuit that also removes clock feedthrough is a sample and hold at the DAC output [46].

When a new output level is formed the DAC output current is initially diverted into ground while the glitch occurs. After the output transition has occurred the sampler switches state and directs the current to the output load. The disadvantage of this technique is the requirement of a sampling clock that operates at a frequency of at least twice the frequency of the maximum data rate.

Even with good design and layout the glitch or clock feedthrough energy can never be completely reduced. High speed CMOS DACs which employ the techniques described here are limited in SFDR to -70 dBc due to remaining transient errors. Low-power versions are limited further to less than -60 dBc due to the use of lower bias current. For example Van den Bosch and co-workers [48] achieve a SFDR of between -60 dBc to -70 dBc for a CMOS DAC with bias current of 1 mA to 10 mA , while for bias currents less than 1 mA the SFDR remains worse than -60 dBc . Therefore it is expected that CMOS DDFS systems will be limited to an SFDR of worse than -70 dBc .

1.6.3 ROM reduction techniques

Non-idealities in the DAC degrade the SFDR, however the digital signal at the output of the ROM already contains unwanted spurious frequencies caused by minimisation of the ROM. With truncation of both phase and amplitude the ROM is still the largest and most power hungry element within the DDFS system. Hence there has been interest in techniques to reduce the ROM size further [52, 23, 21, 53, 54, 55, 56] and a good summary is given by Vankka [2]. For example the symmetry of the sinusoid² wave can be used to compress the ROM without any loss in phase or amplitude information, however most methods are based on approximations that degrade the SFDR. All methods require additional circuitry that in some cases significantly limits the throughput and frequency switching/hopping speed. A common theme is the tradeoff of additional circuitry overhead for improved ROM compression. For example most techniques store the difference between the amplitude and phase in the ROM and form the amplitude using an extra set of adders. One additional adder per output bit is not a significant overhead compared to the logic required by the accumulator and therefore the so called amplitude-phase difference method is used by all ROM compression techniques except CORDIC. The CORDIC technique is an algorithm that can perform the entire phase to amplitude conversion without a ROM³, however the overhead circuitry includes adders, shifters, multiplexers and registers. It is not suitable for high speed, low power applications due to the power consumption and iterative nature of the algorithm and will not be considered further here.

²A sinusoid is a waveform shaped like the sine function. The waveforms of cosine and sine are sinusoids[57].

³In practice a small ROM is commonly used to store the known arc tangent to improve speed [2]

Sinusoid symmetry




Figure 1.5: The quarter wave symmetry of the sine wave may be used to reduce the size of the ROM by four times. Adapted from [2].

Most ROM based DDFS systems store only a quarter ($\sin(0)$ to $\sin(\pi/2)$) of the full period ($\sin(0)$ to $\sin(2\pi)$) by using the symmetry of the sine wave, Figure 1.5. This is achieved by using the first two MSBs to reform the full period. The first MSB is used to determine the sign of the amplitude, ($\pm\sin(\phi)$) at the output of the ROM before it is passed to the DAC and is referred to as half wave symmetry. The second MSB is used to determine if the phase, ϕ , is increasing or decreasing and is used before the ROM lookup table. Using both the first and second MSBs is referred to as quarter wave symmetry. The quarter wave symmetry of the sinewave allows a compression of four times without degradation of the SFDR.

Sine-phase difference

Any further compression requires an approximation. The sine-phase difference method was proposed by Nicholas and co-workers [53] as a linear approximation to the sine wave. A ROM lookup table that contains $\sin(\phi\pi/2) - \phi$ is used with an additional adder to calculate the amplitude, $\sin(\phi\pi/2)$. The maximum amplitude term is reduced by over four times:

$$\max(\sin(\phi\pi/2) - \phi) \approx 0.21\max(\sin(\phi\pi/2)) \quad (1.9)$$

Leading to a saving of two amplitude bits. Vankka finds that the additional adder does not provide a significant burden to the DDFS system compared to the computational logic required by the accumulator [2].

Interpolation and trigonometric approximation

An interpolation based system by Hutchinson [52] first proposed segmentation of the ROM. The ROM size is halved by segmenting the amplitude information into a coarse approximation ROM then a fine ROM is used to interpolate between the coarse values. The fine ROM uses smaller word lengths than the coarse ROM because its dynamic range is reduced. The most popular approximation technique that forms the basis for many segmented designs is that originally proposed by Sunderland and co-workers [23]. In the original design 14 phase bits are used from the accumulator. The first two are used for quarter wave symmetry and the next twelve are segmented into three four-bit blocks; $\alpha < \pi/2$, $\beta < (\pi/2) * (2^{-4})$ and $\gamma < (\pi/2) * (2^{-8})$ such that:

$$\sin(\alpha + \beta + \gamma) = \sin(\alpha + \beta)\cos(\gamma) + \cos(\alpha)\cos(\beta)\sin(\gamma) - \sin(\alpha)\sin(\beta)\sin(\gamma) \quad (1.10)$$

Due to the relative magnitudes of α, β and γ Equation 1.10 can be simplified to ⁴

$$\sin(\alpha + \beta + \gamma) \approx \sin(\alpha + \beta) + \cos(\alpha)\sin(\gamma) \quad (1.11)$$

A saving of up to 11 times in ROM size can be made in implementing this function due to the small magnitudes of β and γ relative to α . The first term $\sin(A + B)$ is implemented in a coarse ROM table and the sine-phase difference method is used to further reduce its size. The second term $\cos(\alpha)\sin(\gamma)$ is implemented in a fine ROM table that effectively interpolates between the values in the coarse ROM, Figure 1.6.

Taylor approximation

An early chip by Qualcomm [58] used a Taylor series to compress the amplitude information. The phase address ϕ is divided into an upper block α and lower block $\phi - \alpha$. The first three terms of the Taylor series of the upper block is stored in three ROMs. Other inaccuracies in the DDFS system

⁴ $\cos(\gamma) \approx 1 \approx \cos(\beta)$ and $\sin(\gamma) < \sin(\beta) \ll \text{other terms}$ hence $\sin(\gamma)\sin(\beta) \approx 0$




Figure 1.6: Segmentation is a popular method to reduce the ROM size. This example shows the use of a trigonometric approximation by Sunderland and co-workers, [23]. Note that the first term uses the sine-phase difference method. A,B,C are the number of bits used to represent α,β,γ .

override the finer accuracy provided by including any more terms [2].

$$\sin(\phi\pi/2) = \sin(\alpha\pi/2) + k_1(\phi - \alpha)\cos(\alpha\pi/2) - \frac{k_2(\phi - \alpha)^2 \sin(\alpha\pi/2)}{2} \tag{1.12}$$




Figure 1.7: An example of ROM compression by using a Taylor expansion. Adapted from [2]. Note that the first term effectively uses the sine-phase difference method. a is the number of bits used to represent α .

Figure 1.7 shows how the first three terms of the Taylor expansion are combined to give the upper

and lower blocks and hence the amplitude. The size of the ROM has been reduced by 64 times. However the method relies on two additional adders and a multiplier that negate the power savings from compression of the ROM. This drawback was alleviated in a recent improvement of the Taylor series method by Bellaouar and co-workers, [21]. Here linear interpolation is used to remove the need for the storing the third coefficient of the Taylor expansion. The compression improved to 300:1 and low power logic methods are used to reduce the power consumption of the additional digital components significantly.

Guess/correction methods

A guess/correction method is the so called sine-phase difference method introduced previously that saves two amplitude bits. It is a straight line approximation $y_i(\phi) = 2\phi/\pi$ for $0 < \phi < \pi/2$ where the output of the phase accumulator is used as the initial guess.

Two improvements of the sine-phase difference method are firstly the use of the 3rd MSB by Yamagishi and co-workers [54] to improve the initial guess and save three amplitude bits; and secondly the parabolic approximation by Sodagar and co-workers [55] to save four amplitude bits. These techniques require overheads of extra addition, 2's complement and digital multiplication circuits to implement the improved compression.

These sine-phase difference techniques are generalised in the piece-wise linear approximation by Langlois and co-workers [56]. Here n straight line approximations defined by two co-efficients are used to form the sine wave. As with the Sodagar and co-workers parabolic technique, 4 bits of memory word length are saved, but here the circuitry is reduced and simplified. Two adders and two multiplexers are required rather than a full multiplier and 2's complement.

1.6.4 ROM-less systems

A recent improvement for low power DDFS is the removal of the ROM entirely with the phase to amplitude conversion performed in digital circuitry or within the DAC. In the place of a ROM, Langlois and co-workers propose a hardware optimised ROM-less phase to amplitude converter with equal piecewise linear segments where

$$s_i(x) = m_i \times (x - 1/8) + y_i, \quad (1.13)$$

The best slope (m_i) and offset (y_i) of each segment is evaluated by using a genetic algorithm. A multiplier is avoided by representing the parameters of each segment as a power of two so that adders and bit shifts may be used in the place of the multipliers. By deriving the Fourier series of the output of a DDFSs whose sine amplitude converters use linear interpolation, Langlois and co-workers show that the upper bound on the SFDR is

$$SFDR \approx 24 + 40 \log s \tag{1.14}$$

where s is the number of segments. Hence the SFDR improves by approximately $12dB$ for each doubling of the number of segments. The Fourier series also reveals that the dominant spur is the $(4s - 1)^{th}$ harmonic of the fundamental.

Author	Segments	Power	F_{clk}	SFDR	On chip
		(mW/MHz)	(MHz)	(dBc)	DAC
Lee resistor string [59]	96	0.4	230	-55	Yes
Lee current mode [20]	192	0.8	300	-62	Yes
Mohieldin [60]	8	0.08	130	-59	Yes
Langlois [56]	8	0.1	125	-60	No
Bellaouar [21]	16point ROM	0.3	30	-60	No

Table 1.3: State of the art low power ROM-less DDFS systems. The design of Bellaouar is not completely ROM-less as it contains a small ROM to correct for errors. It is included for comparison as it is the most efficient DDFS system with a ROM with over 300 times compression from the conventional system.

Using this method a system with 8 linear segments can achieve $-60dBc$. For this system, power consumption was estimated to be less than $10mW$ at $f_{clk} = 100MHz$ ($0.1mW/MHz$) in simulation, making this design one of the lowest power DDFS systems ever designed. Moheildin and coworkers, [60] propose an alternative design that includes a DAC and uses a nonlinearly weighted network of resistors to provide the slopes of the piecewise linear segments, Figure 1.8. The technique has the same SFDR performance as Langlois and co-workers, the SFDR improves by $12dB$ for each doubling of the number of segments. However the power consumption is reduced to $8mW$ at $100MHz$, a significant saving as this system also includes the DAC.

The design uses the quarter wave symmetry of the sine wave so that the resistor networks only need store a piecewise linear approximation of the first quadrant of the sine wave, $\sin(0 : \pi/2)$. The remaining phase bits are segmented into two parts (Figure 1.8). α is the MSB part, a bits long which corresponds to the segment number. β is the LSB part, b bits long which is sent to a linear DAC. There is a weighted resistor network, demultiplexer switches and buffers in each of the switched weighted sum (SWS) blocks whose output is the voltage sum of the combination of resistors. There are two




Figure 1.8: The ROM-less DDFS proposed by Mohiaddin and coworkers [60]. The two differential voltage outputs combine to form a sine wave using a differential amplifier.

SWS blocks one for each of the positive/negative half cycle of the sine wave. The output of the DAC $V_{MIN} < V_{DAC} < V_{MAX}$ is sent to the each of the SWS blocks that store the slope of the piecewise linear segments, k such that the output of the SWS block when enabled ($E=1$) is given by

$$V_{SWS} = V_{min} \quad E = 0 \quad (1.15)$$

$$V_{SWS} = \sum_{n=0}^{\alpha-1} k_n V_{min} + k_{\alpha} V_{DAC} + \sum_{n=\alpha+1}^{2^{\alpha}-1} k_n V_{max} \quad E = 1 \quad (1.16)$$

A system with 8 segments achieves an SFDR of $-59dBc$ and requires only a 5-bit linear DAC. However one critical part of the design is the matching of the resistors as any mismatch will introduce distortion in the output signal. A normalized unit resistance $\Delta R/R_{max}$ of 0.4% is determined in MATLAB simulations to ensure the SFDR is not degraded by resistor mismatch.

Another critical design parameter is the settling time of the linear DAC which is dependant on the load impedance of the switches. An analogue demultiplexer is used to switch the linear DAC output to buffers that drive the required number of resistors in a weighted resistor network. A fundamental problem that limits the final bandwidth of the system is the effect of the load resistance of the buffers between the linear DAC and the weighted resistors. The output of the buffers is finite but the resistive loading is different for each buffer. Therefore each buffer is designed to have the same ratio of output resistance to load resistance by using a unit device approach where a sub-buffer is the unit

cell and the number of parallel unit cells in each buffer is inversely proportional to the desired output resistance. Consequently the total current drain in all buffers is inversely proportional to the total resistance, R_T in the weighted sum network. However R_T also determines the response time of the circuit. Increasing R_T will increase the load on the linear DAC. Thus to keep the power consumption low the design of R_T was optimised for $f_{clk} = 100MHz$. The power consumption was $8mW$ at $100MHz$ or $0.08mW/MHz$.

An alternative, faster ROM-less method that does not use any additional circuitry is proposed by Lee and co-workers [59]. A nonlinear DAC provides both the phase to amplitude mapping and digital to analogue conversion. Lee proposed two nonlinear implementations, a resistor string DAC which had low power consumption but was limited to $f_{clk} = 25MHz$ and consumed a large area to ensure matching between the resistors. The second implementation used a current mode DAC to increase the bandwidth to $f_{clk} = 300MHz$ at the expense of increased power consumption. A segmented architecture based on the trigonometric approximation used in the Sunderland technique, Equation 1.11 was used to reduce the power consumption [20] to that comparable to state of the art systems, Table 1.3.

Comparison of the power consumption between DDFS systems is difficult due to the different implementations of parts of the architecture. Most authors reduce the power consumption of the accumulator by using efficient architectures but do not quote separate power dissipation figures for the accumulator, digital logic, ROM and possible DAC. Details from published state of the art systems are shown in Table 1.3, where it is clear that the interpolation using linear segment approach of Mohieldin and Langlois is the most efficient.

1.7 Summary

DDFS is an attractive method of producing spectrally pure signals with high frequency resolution and fast frequency hopping for application in communications and instrumentation. Two instrumentation systems that would benefit from using low power and inexpensive integrated DDFS to perform accurate measurements from DC to $1MHz$, have been identified and described. A low power dielectrophoresis based sensor for cell characterisation would require a spectral purity of better than $-30dBc$ and a low power, wireless active electrode for electrical impedance spectroscopy would require a spectral purity of better than $-40dBc$ to ensure accurate measurement.

Another application area is mobile communications where low power operation is important to preserve battery life in handheld devices. DDFS is well suited to spread spectrum techniques which require high frequency hopping speed, good frequency resolution and spectral purity with low phase noise. Specifications for an integrated CMOS DDFS synthesiser in the third generation standards (UMTS) require a spectral purity of $-60dBc$, frequency accuracy of $5ppm$ and minimum bandwidth of $15MHz$. The low power output classes of the recent short range wireless protocol, Bluetooth have more relaxed spectral purity of $-50dBc$ with frequency accuracy of $1ppm$ and a bandwidth of $80MHz$.

The spectral quality of the DDFS output is theoretically unlimited. However to reduce the cost and power consumption the size of the ROM must be minimised, causing approximation errors and hence spurious components in the output spectrum. In addition spurious errors are created by phase and amplitude quantisation and imperfect operation of the DAC. A thermometer decoder can reduce both static and dynamic DAC errors, although high speed, state of the art, CMOS DACs are still limited in spectral purity to $-70 dBc$ due to parasitic capacitances within the current switch. Existing direct digital frequency synthesis (DDFS) systems use trigonometric identities, quarter wave symmetry and interpolation to reduce the size of the ROM, however state of the art systems still have considerable power consumption. The most efficient method is to completely remove the ROM and use a nonlinear DAC.

Chapter 2

Nonlinear interpolation technique

2.1 Introduction

Direct digital frequency synthesis is well suited to a wide area of application due to its fast frequency hopping, good frequency resolution and digital interface. The need to provide a large ROM imposes a relatively large power and area consumption and limits the bandwidth as large ROMs have slow access times. One area of improvement is to reduce the size of the ROM. However as shown in Chapter 1 most ROM compression algorithms are based on approximations that lead to spurious errors due to quantisation. Approximations can reduce the ROM by up to 300 times, and the most efficient way to miniaturise a DDFS is to take the approach of Lee and co-workers who remove the ROM completely by combining phase to amplitude conversion and digital to analogue conversion in a nonlinear DAC.

ROM-less direct digital frequency synthesizers that are based on conventional nonlinear DACs require 2^N current cells to produce an output with spectral purity limited to $-6NdBc$ without any approximation. State of the art, high speed, CMOS DACs are limited to a $-70dBc$ SFDR by transient effects. To meet this performance a DDFS based on a conventional nonlinear DAC would require 2^{12} current cells. Lee has shown that the number of current cells can be reduced by using approximations, however these techniques limit the spectral purity to $-60dBc$ [12].

An alternative system is introduced in this chapter that aims to reduce the number of current cells while improving spectral purity. This new method takes advantage of the nonlinear transfer function of a differential current switch to interpolate between the contributions of each current cell. By

carefully choosing the parameters of the differential switch transfer function a small array of current cells can be used as a small nonlinear DAC with spectral purity that is comparable to state of the art systems. This technique may be extended to improve the spectral purity further by using more current cells and a modification of the conventional thermometer decoding technique. This chapter introduces the novel nonlinear interpolation technique and provides predicted spectral purity using MATLAB simulations.

2.2 Nonlinear interpolation DACs

The starting point for the proposed system is Lee's system [59]. In a nonlinear current switched DAC for cosine¹ wave output DDFS the phase to amplitude conversion information is encoded within the DAC bias currents. Exploiting the half-wave symmetry of the output means that the DAC only needs to represent the output signal over the range between 0 and π . To represent a cosinusoidal output, the bias current of the i^{th} circuit should be

$$I_{biasi} = \cos((i - 1)\pi/N) - \cos(i\pi/N) \quad (2.1)$$

This is a real number and the accuracy to which the DAC cell can represent it will have an effect on the amplitude resolution and spectral purity of the system. Lee scales the amplitude (I_{biasi}) to 2^N , for different values of N , where N is the number of current cells, and rounds the real number to its nearest integer [59]. The rounding allows for good matching between DAC cells and ease of layout. However it will decrease the SFDR. Lee presented a quadrature DDFS system incorporating segmented nonlinear DACs that achieved a SFDR of -55 dBc with power dissipation of 92 mW for synthesis of low frequencies (less than 100 kHz). More importantly with the first 2 MSBs used for quarter wave symmetry, Lee's design required $N = 8$ or 256 DAC cells for 10-bit phase resolution. The method presented here aims to reduce the number of DAC cells by using each current cell efficiently.

2.2.1 Nonlinear interpolation

In the current cell of a conventional DAC, V_{in} is either high or low and the cell output is one of two values. However, by addressing the current switch with intermediate input voltages between V_{high}

¹A cosine function is described from here on as it simplifies the explanation of the novel nonlinear interpolation DAC.

and V_{low} the current steering ability of the differential switch may be used to set the current output at intermediate values between 0 and I_{bias} , Figure 2.1.

The intermediate current values are obtained by holding the gate of one of the transistors (M2) in the differential pair at a constant reference voltage, (V_{ref}), while the input voltage is applied to the gate of the other transistor (M1). When the input voltage is low ($< V_{low}$), the bias current flows through M2, and when the input voltage is high ($> V_{high}$), the bias current flows through M1. The input voltage swing ($V_{high} - V_{low}$) should be small to reduce the glitches at the output (see section 1.6.2). To reduce the swing V_{low} is chosen so that M1 is just off, and V_{high} is chosen to just turn M2 off. The two transistors are of equal size and the input voltage swing is chosen symmetrically around the reference voltage. Where possible the current I_{out2} is chosen in preference to I_{out1} as V_{ref} is constant and hence I_{out2} is less effected by glitches and clock feedthrough.




Figure 2.1: A differential current switch. The current is steered between the two transistors by changing the input voltage. When V_{in} is less than V_{ref} most of the current flows through the right (V_{ref}) branch (M2) and I_{out1} is less than I_{out2} . As V_{in} is increased the current in M2 decreases and the current in the left (V_{in}) branch (M1) increases until $V_{in} = V_{ref}$ where equal current ($I_{bias}/2$) flows in each branch. As V_{in} is increased further most of the current flows in M1 and I_{out1} is greater than I_{out2} . Eventually all the bias current flows through M1 and $I_{out1} = I_{bias}$.

In order to derive a relationship between the input voltage and output current assume that the two transistors of the current switch are matched and in saturation ($V_{ds} > (V_{gs} - V_t)$ where V_{ds} is the drain to source voltage, V_{gs} is the gate to source voltage and V_t is the threshold voltage.). If the current source is ideal, the current flowing through M2 is:

$$I_{out2} = \frac{\beta}{2}(V_{gs} - V_t)^2 \quad (2.2)$$

The bias current is shared between the two switch devices, $I_{bias} = I_{out1} + I_{out2}$. Following the derivation in Appendix B:

$$I_{out2} = \frac{\beta}{2} \left[\frac{(V_{ref} - V_{in})}{2} - \sqrt{\frac{I_{bias}}{\beta} - \frac{(V_{in} - V_{ref})^2}{4}} \right]^2 \quad (2.3)$$

where V_{in} is the input voltage, V_{ref} a constant reference voltage, I_{bias} is the current provided by the ideal, constant current source and β is the transconductance parameter of M1 and M2, (Figure 2.1).

The current switch is off ($I_{out2} = 0$) or fully on ($I_{out2} = I_{bias}$) when $(V_{in} - V_{ref})^2 = 4I_{bias}/\beta$ (See Appendix B). Therefore the condition $|V_{in} - V_{ref}| > \sqrt{4I_{bias}/\beta}$ determines the minimum differential input voltage swing required for the cell to operate as a digital switch. However as expected smaller differential input voltages will steer part of the current to the output. By controlling the fraction of the current which is steered to the output, the cell can be used to generate a nonlinear interpolation function between an output of 0 and I_{bias} , (Figure 2.1).




Figure 2.2: The current flowing out of the right branch (I_{out2}) of a single current cell in response to a triangular input voltage. (a) When the bias current is small the differential switch acts as a digital switch where the current is either flowing through the left or the right branch. (b) If the amount of available bias current is increased slightly the output resembles a sinusoid. (c) By increasing the amount of available bias current to the switch the output becomes more linear.

When the differential input voltage, $|V_{in} - V_{ref}|$, is a triangular wave the output can resemble a cosine wave as shown in Figure 2.2b. The triangular wave voltage input represents the phase, V_{phase} .




Figure 2.3: How the harmonic spurs change when the bias current is scaled for a single cell. At low bias currents the harmonics are at similar magnitude to those of a square wave. This is when the total bias current (in arbitrary units) is 0.01 and corresponds to the waveform shown in Figure 2.2a. The optimum SFDR of -45dBc is at a total bias current of 0.4 and this point corresponds with the waveform shown in Figure 2.2b. At high bias currents the input voltage does not turn the current switch fully on or off and the output is linear, similar to a triangular wave. This occurs for total bias current of greater than 2 which corresponds to the waveform shown in Figure 2.2c.

The fraction of bias current steered to the output can be controlled by changing the differential input voltage swing (V_{phase}) or by changing the amount of bias current available. At low current values, $V_{phase} > \sqrt{4I_{bias}/\beta}$ the output current saturates to either 0 or I_{bias} for most of each period. The result shown in Figure 2.2a resembles a square wave. Increasing the bias current, while keeping both V_{phase} and β constant, reduces the time for which the output is saturated. Eventually a bias current is reached at which the output current saturates at the minimum and the maximum input voltage, Figure 2.2b. The output signal resembles one period of a cosine wave.

If the current is increased beyond this point, so that $\sqrt{4I_{bias}/\beta} > V_{phase}$, the input voltage range is too small to saturate the output and the input voltage range corresponds to the linear part of the response of the circuit. At high current levels the output signal resembles a triangular wave, Figure 2.2c.

The effect of the change in the shape of the output signal on its spurious free dynamic range can be seen in more detail in Figure 2.3 where the dominant spurious frequencies are plotted over a range of bias currents. These spurs are determined from the FFT of the response of Equation 2.3 calculated in MATLAB for different total bias currents. As the results presented in this chapter are all based

on MATLAB simulations the current and voltages are in arbitrary units. For these output signals the dominant frequencies are the lower odd harmonics. As expected for small bias currents, when the output resembles a square wave, the dominant harmonics are the 3rd harmonic at $-10dBc$ (1/3) and the 5th harmonic at $-14dBc$ (1/5).

For a very small range in bias current (when $V_{phase} = \sqrt{4I_{bias}/\beta}$ and $I_{bias} \approx 0.4$) the magnitude of the harmonics drop dramatically to less than $-40dB$ and the SFDR of the output is optimum. This is the region where the bias voltage resembles a sine wave, Figure 2.2b. At higher bias currents the spectrum of the output signal approaches that of a triangular wave with the 3rd harmonic at $-20dB$ (1/9), the 5th at $-28dB$ (1/25) and the 7th harmonic at $-34dB$ (1/49). Overall, the results in Figure 2.3 show that the circuit in Figure 2.1 can create a surprisingly good approximation to a sinusoidal output at the optimum bias condition. However the optimum operating range will be sensitive to changes in the bias current. For example small changes in the bias current reduces the spectral purity of the single current cell to $-25dBc$. To achieve a high yield any analogue circuit design must be able to allow for the inevitable variations of each device from the typical behaviour. Hence a more robust system is pursued.

2.2.2 Compact systems

DDFS systems use more current cells to improve the spectral purity of the output waveform generated by the DAC. The circuit architecture of an array of four current cells, each of which contributes a segment of the cosine wave is shown in Figure 2.4. The reference voltages were chosen to be uniformly spaced (ΔV_{ref}) and all four current cells share the input voltage. The input voltage is chosen to span from half the reference voltage difference ($\Delta V_{ref}/2$) below V_{ref1} (to ensure all current is steered to I_{out2}), to $\Delta V_{ref}/2$ above V_{ref4} (to ensure all current is steered to I_{out1}). Hence $V_{low} = V_{ref1} - \Delta V_{ref}/2$ and $V_{high} = V_{ref4} + \Delta V_{ref}/2$.

When the input voltage is V_{low} all the reference transistors are on, I_{out2} is maximum and I_{out1} is minimum. As the input voltage increases from V_{low} the first current source begins to steer current from the reference transistor to the input transistor. As the input reaches $V_{ref1} + \Delta V_{ref}/2$ the reference transistor in the first current cell turns off and the input transistor turns on. I_{out2} has dropped and I_{out1} has increased. As the input continues to increase each of the other current sources steer current from I_{out2} to I_{out1} in turn. To ensure that four current sources represent half a period the bias current of each current cell must represent the changes in output amplitude over an eighth of a period, using




Figure 2.4: A system based on an array of four current cells. The current outputs are summed together using current summation at the output nodes. The input voltage is common to all cells and the reference voltages are equally spaced between the current cells. The bias currents are chosen so that the outputs form a sinusoid.

Equation 2.1:

$$\begin{aligned}
 I_{bias1} &= \cos(0) - \cos(\pi/4) \\
 I_{bias2} &= \cos(\pi/4) - \cos(\pi/2) \\
 I_{bias3} &= \cos(\pi/2) - \cos(3\pi/4) \\
 I_{bias4} &= \cos(3\pi/4) - \cos(\pi)
 \end{aligned}$$

The quarter wave symmetry of the cosine wave (about $\pi/2$) means that the parameters are duplicated for half of the current cells ($I_{bias1} = I_{bias4}$ and $I_{bias2} = I_{bias3}$). Furthermore, the β parameter for the switch devices in each current source is chosen to ensure that (I_{bias}/β) is the same in each current cell. This is a requirement of this method as V_{phase} is constant and optimum operation is known to occur when $V_{phase} = \sqrt{4I_{bias}/\beta}$ from section 2.2.1.

The response of the array of four cells to a triangular wave input voltage, V_{phase} , is shown in Figure 2.5. At low current levels the output current from each circuit saturates for small differential input voltages. As a result, the output from each current source saturates before the output from its neighbour starts to change. This introduces the eight flat regions per period clearly visible in Figure 2.5a.

As the current level is increased these steps begin to disappear and the output becomes more sinusoidal. When the input swing is the minimum for the switch to operate as a digital switch, $V_{phase} = \sqrt{4I_{bias}/\beta}$, the response of the cells is near optimum with spectral purity of $-40dBc$ (Figure 2.5b). As the total current increases further the response of each current source becomes more linear

and the output begins to resemble a triangular wave Figure 2.5c. It is clear that the current sources are not saturating here as the full scale of the output has decreased.




Figure 2.5: The combined output current for an array of four cells. (a) At low current levels the contribution of each cell is clearly visible as steps in the output when each cell saturates. (b) At the optimum current level the spectral purity is better than $-40dBc$. (c) At higher than optimum current levels the output begins to resemble a triangular wave and no cells are saturating.

As with the single cell, the changes in the shape of the output of four cells (Figure 2.5) will be reflected in changes to its power spectrum and hence its SFDR. The behaviour of the dominant spurious harmonics at different currents is shown in Figure 2.6. At low bias currents the eight steps per period create large spurious outputs at the 7th and 9th harmonics. The width of the steps decreases with increasing current and the power at these spurious frequencies reduces.

There are eight steps per period, corresponding to the four cells used to create each half waveform. The waveform is the same as if it had been sampled and held to create the flat regions with a sampling frequency of eight times per period. It is well known that sampling creates spurs at image frequencies at $f_s - f_{out}$ and $f_s + f_{out}$, here $f_s = 8$ and $f_{out} = 1$.

At high bias currents the output approaches a triangular wave and the relative power levels of the lower odd harmonics increase in power. The optimum SFDR is $-40dBc$ and is better than $-35dBc$ over a wide range of bias current levels.

The important difference between the results from a single cell, (shown in Figure 2.3), and those for a system of four cells, (shown in Figure 2.6) is that the larger system is expected to be more robust over




Figure 2.6: Relative power levels of the dominant harmonics in the output spectrum of an array of four cells. At low total bias current of 0.1 (again arbitrary units) the dominant harmonics are that of a waveform that has been sampled and held at a rate of eight times per period as in the trace shown in Figure 2.5a. The optimum SFDR of a system of four cells is better than -40dBc at a total bias current of 2 and this corresponds to the trace in Figure 2.5b. The performance is better than -35dBc over a wide range of bias current levels. When the bias current is increased the harmonics approach that of a triangular wave. The trace shown in Figure 2.5c is for a total bias current of 3.

a range of bias current levels. The optimum performance of the four cell system of -40dBc just meets the requirements of the EIS instrumentation system without any margin for error. To investigate if the spectral purity could be improved along with the robustness by using more current cells, a system based on an array of six cells was examined. In a similar fashion to the four cell system a common input voltage V_{in} is shared amongst six cells by using six uniformly spaced reference voltages V_{refi} . Each of the current cells is assigned a bias current according to Equation 2.1.

The effect of the change in shape of the output signal on the spectral purity of the six cell system is shown in Figure 2.7. At low bias current the individual responses of each current cell saturates and the image frequencies at the 11^{th} and 13^{th} harmonics limit the SFDR of the six cell system, Figure 2.7. As with the system of four cells, the position of these harmonics is related to the number of current cells used to form the output. The six cells are addressed twice per output period and therefore the dominant harmonics appeared adjacent to a sampling frequency of $f_s = 12$. The SFDR is an improvement over the systems with one or four cells by at least 5dB over a wide range of bias current levels. Comparing Figure 2.7 with Figure 2.3 and Figure 2.6 shows that the six cell system is also more robust than the smaller systems.




Figure 2.7: Relative power levels of the dominant harmonics in the output spectrum of an array of six current cells. The optimum SFDR of a system of six cells is better than -48dBc over a wide range of bias currents.

2.2.3 Performance scaling

The spectral purity of a single current cell, -45dBc , is surprisingly good. However analogue circuits require systems that are robust to variations in parameters and small changes in bias current degrades the performance of a single cell dramatically to approximately -25dBc . Arrays of four and six current cells are more robust than a single cell to variations in bias current. Four cells are expected to achieve better than -35dBc and six cells better than -40dBc reliably over a wide range of bias current.

The MATLAB simulation results in Figure 2.8 show that the SFDR improves for systems with more current cells. The results in Table 2.1 indicate that systems with as few as 8 devices are comparable to state of the art systems described in section 1.6.4. These systems are expected to be limited at high frequencies to -70dBc SFDR due to transient effects, as this is the limit of state of the art, low power, high speed, CMOS DACs. Therefore an array of 64 current cells, with a predicted SFDR of -76dBc is of interest to test the optimum performance of this type of system.

2.3 Segmented systems

A change in architecture to a segmented system is required for a system of 64 current cells. While the compact systems used a small number of current cells and a small number of voltage references,




Figure 2.8: SFDR of systems with various numbers of cells against bias current.

Number of current cells	SFDR
4	- 40 dBc
6	- 48 dBc
8	- 54 dBc
16	- 59 dBc
32	- 69 dBc
64	- 76 dBc
128	- 82 dBc

Table 2.1: The MATLAB predicted, optimum SFDR for systems containing different numbers of current cells.

for a system with 64 cells, too many voltage references within the limited voltage range are required. Hence the popular thermometer decoding method is modified to create an architecture that is able to interpolate between 64 current cells using only one reference voltage. This new architecture is referred to as a segmented system as the contribution to the output of each cell represents a particular segment of the output waveform, whereas the contributions of the cells in the compact systems were allowed to overlap.

To combine the interpolation technique with thermometer decoding (section 1.6.2), a segmented address system based on conventional thermometer decoding can be used to select the appropriate cell to interpolate for each phase input. An advantage of this approach is that all current cells use the same reference voltage, rather than equally spaced reference voltages, removing the headroom requirement to keep all the reference voltages within the range offered by the voltage supply. For example the creation of 64 or more precise and independent reference voltages within a 3.3V voltage

range is a significant design problem. Therefore a change in architecture is important to create a working system.

The architecture is described in detail in Chapter 5 and its function is summarised here. From Equation 1.7, to achieve an SFDR of $-76dBc$ at least 13 phase bits are required to avoid phase quantisation. The most significant bit is used to identify the half wave symmetry of the cosine wave. The next 6 most significant phase bits are used to thermometer decode the array of 64 devices and the remaining 6 phase bits are used to create the analogue interpolation voltage (V_{phase}). In a conventional DDFS with 64 current cells these last 6 bits would be discarded. All current cells have the same V_{ref} . The 64 current cells represent half a period and therefore each is assigned a bias current using Equation 2.1. The quarter wave symmetry of the sinusoid means that the parameters are duplicated for half the current cells, hence there are 32 independent current, (I_{bias}), parameters. As in compact systems, the slope, I_{bias}/β is kept constant and hence there are also 32 independent gain, (β), parameters.

For any given phase input the thermometer decoder determines how many current cells will be switched fully on. Conventionally the rest of the current cells will be switched off. However to use the cells more effectively, the analogue interpolation voltage, V_{phase} is fed to a particular current cell that is at the intersection of the cells that are on and the cells that are off. Conventionally this particular current cell is off. In the new architecture V_{phase} is used to partially turn it on. V_{phase} causes this current cell to interpolate between the amplitude value of the current cell preceding it and its own full amplitude value, represented by its bias current. The circuit architecture and associated logic to achieve this operation is described in Chapter 5.

2.3.1 Linearisation

One important feature of the segmented architecture is that the interpolation cell is only connected to the analogue input of a single cell at any one time as each current cell receives the input voltage over a limited phase range. By limiting the input voltage range of each current cell the proportion of the nonlinear shape that contributes to the output can be modified. The response of the current cell may be linearised by either decreasing the voltage difference $|V_{in} - V_{ref}|$, or by increasing the I_{bias}/β ratio, Figure 2.9. For example when increasing I_{bias} the I_{bias}/β ratio increases, and the slope of the interpolation decreases. Over a smaller input voltage range the response of the cell is more linear. Reducing input voltage swing will make the system more sensitive to variations in the reference

voltage and input voltage, making the system less reliable. Therefore linearisation is achieved here by increasing the I_{bias}/β ratio by increasing the amount of current in each circuit.




Figure 2.9: Linearisation of the current cell response. By increasing the I_{bias}/β ratio (small I_{bias}/β to large I_{bias}/β) or by reducing the voltage range from ΔV to δV for the large I_{bias}/β curve, the contribution of current δI is more linear.

Number of current cells	Best SFDR Nonlinear interpolation	Best SFDR Linear interpolation	SFDR from [22] Optimised linear
1	- 45 dBc	-	-
2	- 29 dBc	- 19 dBc	- 24dBc
4	- 40 dBc	- 34 dBc	- 36dBc
6	- 48 dBc	- 42 dBc	- 43dBc
8	- 54 dBc	- 47 dBc	- 48dBc
16	- 59 dBc	- 60 dBc	- 60dBc
32	- 69 dBc	- 72 dBc	- 72dBc
64	- 76 dBc	- 84 dBc	- 84dBc
128	- 82 dBc	- 96 dBc	- 96dBc

Table 2.2: MATLAB results showing the expected best SFDR for systems containing different numbers of current cells. The nonlinear interpolation results are at a power efficiency of 100%, while the linear results are at a power efficiency of 10%.

The performance of the systems with a linearised current cell response is shown in the third column of Table 2.2. The linear interpolation results are significantly better than nonlinear interpolation for systems with 32 cells or more. The improvement may be explained by the observation that a sinusoid is more linear during the areas of greater change. The systems with less than 32 current cells were better with nonlinear interpolation as the approximation in the maximum and minimum regions of the sinusoid are poor with few linear segments. This is not surprising as the linearised response of the one and two cell systems is a triangular wave.

An interesting observation emerges when the results in Table 2.2, column 3, are compared to the results reported by Langlois and Al-Khalili [22], (see section 1.6.4) that are listed in Table 2.2, column

4, taking into account that they used quarter wave symmetry and hence the number of segments is half those used here. Langlois and Al-Khalili used a detailed parameter optimisation to determine the parameters, however a comparison with the results determined here shows little difference. The results for more than two cells differ by at most $2dB$. The similarity is surprising considering the detailed optimisation method of Langlois and Al-Khalili when compared to the parameter selection procedure presented here. This suggests the optimisation is not effective in improving the performance. More importantly, systems with eight current cells or less have a better SFDR with nonlinear interpolation.

2.3.2 Power efficiency

The results in Table 2.2 show that the spectral purity may be improved by linearisation for systems with 32 current sources or more. When using nonlinear interpolation all the available bias current is used to form the output waveform and hence the nonlinear DAC in the compact system is 100% power efficient. By taking the linear part of the nonlinear response a fraction of the bias current is not used to form the output and hence this system is less efficient than the compact systems. The gain in performance is achieved at the expense of a reduction in power efficiency. Power efficiency can be measured as the output current amplitude divided by the total bias current used to create the output.

$$efficiency = I_{out} / \Sigma I_{bias} \quad (2.4)$$

The performance of an array of 64 cells versus efficiency is shown in Figure 2.10. There is a clear tradeoff between the amount of current used to generate the output and the spectral purity performance. The dominant harmonic is always the 127^{th} harmonic in agreement with the Fourier analysis of Langlois and co-workers [22]. In particular these results show that in the SFDR is $-72dBc$ when the 64 cell system is operated at an efficiency of 50%. This is likely to be limited by the performance of state of the art low power, high speed CMOS DACs of $-70dBc$ SFDR. This system would also be able to meet the spectral purity requirements of the UMTS and Bluetooth specification at an efficiency of 80% for the high output power class ($-60dBc$) and at 100% for the lower output power Bluetooth classes ($-50dBc$). However these results are based on a first order model of transistor response and do not take into account non-ideal effects of real circuits. Therefore the circuits are simulated in more detail in Chapters 4 and 5 using the Cadence circuit simulator.




Figure 2.10: The performance of an array of 64 cells versus efficiency simulated in MATLAB. Efficiency is the ratio of output current used to form the output to the total bias current used by the cells. The performance decreases with increases in efficiency providing a tradeoff between the amount of current used and the spectral purity.

2.4 Summary

MATLAB simulations have shown how the nonlinear response of a conventional differential current switch can be used to produce spectrally pure outputs from compact ROM-less DDFS systems consisting of four or six current cells. An analogue input is used to partially turn on the current cells that subsequently provide intermediate values to interpolate between the nonlinearly weighted current sources of the conventional nonlinear DAC. All the available bias current of the nonlinear DAC is used to form the output and hence these systems are 100% efficient. The results of these compact systems look promising for the instrumentation systems introduced in Chapter 1 that require SFDR of better than $-30dBc$.

Further MATLAB simulations show that by using more current cells the spectral purity improves. A different architecture that is based on the conventional thermometer decoder has been suggested as a method to build a system to meet the spectral purity requirement of the UMTS and Bluetooth communications standards. This architecture is referred to as segmented because the thermometer decoder divides the nonlinear DAC into N segments where N is the number of current cells.

The conventional thermometer decoder reduces transient glitches in high speed current mode DACs by ensuring the minimum number of cells switch. This circuit is modified to a segmented ther-

momenter decoder to ensure that only one current cell is interpolating for any given phase input. Segmenting the current cell array allows linear interpolation which improves the spectral purity of the output by up to $8dB$ when 64 current cells are used. The drawback of using a linear interpolation response is a loss in power efficiency as a fraction of the total bias current is used to form the output.

There is a clear tradeoff between the spectral purity improvement and power efficiency from $-52dBc$ at 100% to $-84dBc$ at 10%. The transfer functions of the current switches are linear when only 10% of the bias current is used to form the output and hence there is little improvement by decreasing the efficiency further than this point. It is expected that an array of 64 current cells in the segmented architecture will achieve an SFDR of $-72dBc$ at an efficiency of 50%. The efficiency can be improved to 80% and the SFDR degrades accordingly to $-60dBc$ to meet the Bluetooth and UMTS SFDR requirements.

Chapter 3

Effects of parameter variations

3.1 Introduction

The promising results so far lead to an interest in a DDFS system with analogue interpolation. All analogue circuits should be designed to be robust to variations inherent in the manufacturing process. These create unwanted differences or mismatch between transistors. The output signal of the ROM-less DDFS presented in the previous chapter will be distorted by the effect of mismatch and the distortion will appear as spurs in the output spectrum, degrading the SFDR. Mismatch and hence the unwanted spurs can be reduced by using large devices as the random variations are averaged out over a larger area [61]. However larger devices will also increase the cost and gate capacitance thereby limiting the speed of operation.

In this chapter the methods of modelling mismatch between transistors are reviewed and the effect of parameter mismatch on the performance of the nonlinear interpolation system is simulated in MATLAB. The results of these simulations will be used in a prototype design to determine the accuracy of the transistors in the differential switch and current source. The viability of a resulting system will be based on the yield, that is, the percentage of circuits that pass the specifications. The specifications of interest are determined from the discussion in Chapter 1. Yield predictions will be given for the systems that target the instrumentation applications requiring a spectral purity of better than -30dBc and -40dBc . Predictions will also be given for a DDFS system that targets the performance of state of the art, high speed, CMOS DACs, -70dBc SFDR.

In this chapter the conventional approaches to minimising mismatch through circuit design are in-

roduced. A common approach to ensuring that the ratio between two current sources is robust to mismatch is to use combinations of unit devices with exactly the same geometry and orientation. Although, this creates a design which is robust to mismatch it is only easy to implement if the ratio of the two bias currents within the circuit can be approximated by the ratio of two small integers. MATLAB is used to find appropriate integer values for this method.

3.2 Mismatch

Mismatch between devices in the nonlinear DAC will introduce errors in the approximation to a cosine wave and cause distortion in the output waveform. The most widely used model of the area dependence of mismatch is the Pelgrom model [61]. The Pelgrom model defines a mismatch parameter for each of the two parameters of the first order MOS model, threshold voltage (V_t) mismatch, A_{V_t} and transconductance (β) mismatch, A_β . Both are combined in a model of current (I_{bias}) mismatch, σ_I/I . For transistors in the linear region [61], and saturation [62] the relationship between mismatch and size of two identically drawn transistors with area (W.L) is:

For threshold voltage mismatch:

$$\sigma_{V_t}^2 = \frac{(A_{V_t})^2}{W.L} \quad (3.1)$$

where A_{V_t} is an absolute matching parameter of typically $5mV\mu$ for an NMOS transistor.

For β mismatch:

$$\frac{\sigma_\beta^2}{\beta} = \frac{(A_\beta)^2}{W.L} \quad (3.2)$$

where A_β is a relative matching parameter of typically $1\%\mu$ for an NMOS transistor.

Mismatch will cause errors in the bias currents and the offset voltages. Both β and V_t mismatch contribute to variation in the output current of a transistor [62].

$$\sigma_I^2 = \left(\frac{\delta I}{\delta \beta} \sigma_\beta \right)^2 + \left(\frac{\delta I}{\delta V_t} \sigma_{V_t} \right)^2 \quad (3.3)$$

and

$$I = \frac{\beta}{2}(V_{gs} - V_t)^2 \quad (3.4)$$

leads to:

$$\frac{\sigma_I^2}{I^2} = \frac{\sigma_\beta^2}{\beta^2} + \frac{4\sigma_{V_t}^2}{(V_{gs} - V_t)^2} = \frac{A_\beta^2}{W.L} + \frac{4A_{V_t}^2}{W.L.(V_{gs} - V_t)^2} \quad (3.5)$$

Hence the threshold mismatch has more influence at lower gate voltages where $(V_{gs} - V_t)$ is small. As the gate voltage is increased, the β mismatch begins to dominate. This effect can be seen in the relationship between σ_I^2/I^2 against V_{gs} , for various transistor sizes, Figure 3.1.

A convenient choice is to choose the gate to source bias voltage when the β and V_t variation equal [48]. For $A_{V_t} = 5mV\mu$, $A_\beta = 1\%\mu$ and $V_t \approx 0.5V$ the gate to source voltage must be $V_{gs} \approx 1.5V$ for the contribution of β and V_t variation to be equal from Equation 3.5. However for the model of the current cell in Equation 2.3 to be valid the current source and the current switch must be saturated ($V_{ds} > V_{gs} - V_t$). If the current source gate voltage is $1.5V$ this leaves insufficient voltage for the current switch V_{ds} (which must include six independent voltage references in the six cell system) and a voltage drop across the output load. Hence it is desirable to lower the V_{gs} requirement for the current source to approximately $1V$ to allow more voltage headroom for the current switch and output load. From Figure 3.1 this corresponds to moderately sized transistors of $8\mu^2$ for current matching of better than 1%.




Figure 3.1: Mismatch of NMOS transistors of different sizes in a submicron process.

3.3 Current source mismatch

Mismatch between the current sources will cause errors between the bias currents of each current cell. This is a problem in all current mode digital to analogue converters. Most use unit current sources to efficiently implement large area transistors to reduce mismatch. Other techniques are common centroid layout and calibration, however these techniques are difficult to arrange for a high speed current mode nonlinear DAC as it is desirable to have the current switch close to the current source to reduce delays. Common centroid layout of a nonlinear DAC would result in higher parasitic impedances and a more complex routing scheme. Calibration requires a large ROM to store the correct values and the overhead of the associated ROM access time. Whereas the goal of this design is to avoid using a ROM in a DDFS system. Hence calibration and common centroid are not used in this design, instead unit devices with sufficient area are used to reduce mismatch. To use unit current sources the number of units required to implement each parameter of the nonlinear parameter set and the minimum size of each unit device are determined using MATLAB.

3.3.1 I_{bias} determination

The current (I_{bias}) parameters, determined in MATLAB in Chapter 2, are scaled and truncated to form integer numbers of unit devices to reduce mismatch. The β parameters are chosen to keep I_{bias}/β constant. Truncation was performed in MATLAB by rounding up or down $N \times I_{bias}$, depending on the cumulative sum of the parameter, to ensure that truncation errors are not accumulated. The scaling factor of N was arbitrarily chosen to provide suitable small numbers of unit devices. Starting with rounding the first current cell parameter to the nearest integer, the cumulative truncated current, ΣI_{trunc} is compared with the cumulative un-truncated current, ΣI . If larger, ($\Sigma I_{trunc} > \Sigma I$) the next current cell is rounded down to the nearest integer, otherwise, ($\Sigma I_{trunc} \leq \Sigma I$) and the next current cell is rounded up.

Cell	I_{bias}
1 and 4	3
2 and 3	7

Table 3.1: Truncated $10 \times I_{bias}$ for the four cell array.

The truncated values for systems of four, six and 64 cells are shown in Tables 3.1, 3.2 and 3.3. The values represent the number of unit devices needed within each current cell. Hence a system of four cells requires 20 unit current sources, a system of six cells, 60 unit sources and a system of 64

Cell	I_{bias}
1 and 6	4
2 and 5	11
3 and 4	15

Table 3.2: Truncated $30 \times I_{bias}$ for the six cell array.

Cell	I_{bias}	Cell	I_{bias}	Cell	I_{bias}	Cell	I_{bias}
1 and 64	2	9 and 56	20	17 and 48	35	25 and 40	46
2 and 63	3	10 and 55	23	18 and 47	38	26 and 39	47
3 and 62	6	11 and 54	24	19 and 46	38	27 and 38	47
4 and 61	8	12 and 53	26	20 and 45	41	28 and 37	48
5 and 60	11	13 and 52	28	21 and 44	41	29 and 36	48
6 and 59	13	14 and 51	30	22 and 43	42	30 and 35	49
7 and 58	16	15 and 50	32	23 and 42	44	31 and 34	48
8 and 57	17	16 and 49	34	24 and 41	45	32 and 33	50

Table 3.3: Truncated $1000 \times I_{bias}$ for the segmented system of 64 cells.

cells, 2000 unit devices.¹ These large numbers of unit cells emphasize the importance of finding the smallest area of devices that are robust to mismatch and meet the specifications.

By truncating the parameters further, the number of unit devices could be reduced, however this would have an adverse effect on the spectral purity. The level of truncation of the compact systems was chosen at a level where the spectral purity did not degrade below their optimum, $-40dBc$ and $-48dBc$ for the four and six cell systems respectively.

However the target performance of the 64 cell system is to meet an SFDR better than $-70dBc$, the performance limit of state of the art, low power, high speed, CMOS DACs. The segmented architecture allows the linearisation of the response that drops the efficiency to 10% and the SFDR is $-84dBc$. A truncation level of the current parameter was chosen where the SFDR degrades to $-72dBc$. However there were substantial changes in the output spectrum, Figure 3.2. While the spectrum of the response of the 64 cell system before truncation showed spurs above the noise floor only at the 127th and 129th harmonics (and the harmonics of these spurs, namely at 254, 258 and 381, 387). Truncation of the current parameter introduced quantisation spurs throughout the spectrum, the dominant spurs being the lower odd harmonics which degraded the linear SFDR to $-72dBc$.

¹A direct comparison between the number of unit current sources required here and other designs can not be made as the published work of other designs only report the number of nonlinearly related devices. How these devices are implemented to avoid mismatch is unclear.




Figure 3.2: The effect of parameter truncation on the output spectrum of the 64 cell system. The circles show the spectrum of the parameters before truncation (the noise floor is -160dBc). The lines show the spurs that are introduced when the current parameters are truncated.

3.3.2 Modelling I_{bias} variations in MATLAB

To quantify the effect of variations in the current source (I_{bias}) parameter, a series of MATLAB simulations were performed. In each simulation a random number generator with a normal distribution was used to determine the percentage error in the I_{bias} parameter. For each percentage error the SFDR of 100 different systems was evaluated to determine the statistical behaviour of a population of DDFS systems, and hence the expected yield for different performance criteria. A yield of better than 90% was targeted.

Figure 3.3 shows the effect of variations in the I_{bias} parameter for the systems of four and six cells. The yield decreases with an increase in mismatch error as expected. To achieve a SFDR of $-40dBc$ for the instrumentation systems described in section 1.5 with the four cell system the variation in the I_{bias} parameter needs to be controlled to within 1%. This can be relaxed to 3% if the six cell system is used. The variation in the I_{bias} parameter needs to be controlled to within 1% to achieve a SFDR of $-45dBc$ in the six cell system. From Figure 3.1 transistors with an area of $8\mu^2$ are sufficient to ensure 1% matching at a bias (V_{gs}) voltage of 1V.

Figure 3.4 shows the effect of variation in the I_{bias} parameter for the segmented system of 64 current cells. The target specification for Bluetooth and UMTS is better than $-60dBc$ which can be achieved




Figure 3.3: The estimated yield against percentage error in the I_{bias} parameter for various minimum allowable SFDR for compact systems.




Figure 3.4: The estimated yield against percentage error in the I_{bias} parameter for various minimum allowable SFDR for 64 cell segmented systems.

with 1.5% variation in I_{bias} , however the variation must be reduced to less than 0.4% to achieve a SFDR of better than $-70dBc$, the performance of state of the art, low power, high speed, CMOS DACs. Hence considerably larger current cells are required in the segmented system, although the area is not excessive, approximately $35\mu^2$ at $V_{gs} = 1V$, from Figure 3.1.

3.4 Current switch mismatch

The I_{bias}/β ratio is encoded in each current cell by sizing the W/L ratio of the switch transistors. Mismatch between the current switches of one current cell with the other current cells will introduce mismatch in the I_{bias}/β ratio between current cells, causing distortion in the output. The technique also relies on good matching between the two devices of the differential current switch. Variations in the transconductance (β), and threshold voltage (V_t) between the two switch transistors within each current cell will also cause variations in the output of each current cell.

3.4.1 Modelling β variations in MATLAB

Switch transistor β mismatch

To study the effect of within cell β mismatch on the combined output of all current cells, the transfer function of the current cell can be expressed with a separate variable representing the variation in β , $\Delta\beta$.

$$I_{out} = \frac{\bar{\beta} + \Delta\beta}{2} \left[\frac{(V_{ref} - V_{in})}{2} - \frac{\Delta\beta (V_{in} - V_{ref})}{\bar{\beta}} - \sqrt{\frac{I_{bias}}{\bar{\beta}} - \frac{(V_{in} - V_{ref})^2}{4} \left(1 - \left(\frac{\Delta\beta}{\bar{\beta}} \right)^2 \right)} \right]^2 \quad (3.6)$$

Where V_{in} is the input voltage, V_{ref} is a constant reference voltage, I_{bias} the current provided by a constant current source. $\bar{\beta}$ is the average transconductance of the two switch transistors and $\Delta\beta$ the variation or mismatch between the two switch transistor geometries. The equation can be simplified

by observing that $\Delta\beta/\beta \ll 1$ to become:

$$I_{out} = \frac{\bar{\beta} + \Delta\beta}{2} \left[\frac{(V_{ref} - V_{in})}{2} - \frac{\Delta\beta (V_{in} - V_{ref})}{\bar{\beta}} - \sqrt{\frac{I_{bias}}{\bar{\beta}} - \frac{(V_{in} - V_{ref})^2}{4}} \right]^2 \quad (3.7)$$

As expected this equation simplifies to Equation 2.3 when $\Delta\beta = 0$.

Figure 3.5 shows the effect of variation in the β parameter for the systems of four and six cells. The yield decreases with an increase in mismatch error as expected. To achieve a SFDR of $-40dBc$ for the instrumentation systems described in section 1.5 with the four cell system the variation in the β parameter needs to be controlled to within 0.3%. The variation requirement can be relaxed to 3% if the six cell system is used. The variation in the β parameter needs to be controlled to within about 0.6% to achieve a SFDR $-45dBc$ in the six cell system.

Figure 3.6 shows the effect of variation in the β parameter for the segmented system of 64 current cells. The segmented system is much more robust to β mismatch than the other systems, due to the use of linear rather than nonlinear interpolation, Figure 3.7. To meet the Bluetooth and UMTS requirement of a SFDR of $-60dBc$ the variation in β needs to be controlled to within 5%. To meet the performance of state of the art low power, high speed, CMOS DACs of $-70dBc$ SFDR the variation in β needs to be controlled to within 4%.

Switch transistor V_t mismatch

To study the effect of V_t mismatch of the switch transistors the transfer function of the current switch can be expressed with a separate variable representing the variation in V_t , ΔV_t .

$$I_{out} = \frac{\beta}{2} \left[\frac{(V_{ref} - V_{in})}{2} + \frac{\Delta V_t}{2} - \sqrt{\frac{I_{bias}}{\beta} - \frac{(V_{in} - V_{ref} + \Delta V_t)^2}{4}} \right]^2 \quad (3.8)$$

This equation simplifies to Equation 2.3 when $\Delta V_t = 0$.

Again to quantify the effect of V_t mismatch in the current switches a series of MATLAB simulations were performed. In each simulation a random number generator with a normal distribution was used to determine the percentage error in V_t . From Equation 3.8, ΔV_t also represents the percentage




Figure 3.5: The effects of variation between the β of the two transistors in the current switch on yield for compact systems.




Figure 3.6: The effects of variation between the β of the two transistors in the current switch on yield for the 64 cell segmented systems.




Figure 3.7: A comparison of the effect of β variation on the current cell transfer function, Equation 2.3. For a 20% variation in β the nonlinear response (from 0V to 2V) varies more than the linear response (from 0.6V to 1.4V). The reference voltage was 1V.

error in the reference voltage, V_{ref} and the percentage error in the phase input voltage, V_{phase} .

Figure 3.8 shows the effect of variation in the V_t parameter for the systems of four and six cells. To achieve the SFDR of $-40dBc$ for the instrumentation systems described in section 1.5 with the four cell system the variation in the V_t parameter needs to be controlled to within 0.3% of V_{phase} . This can be relaxed to 2% if the six cell system is used. The variation in the V_t parameter needs to be controlled to within about 0.5% of V_{phase} to achieve a SFDR of $-45dBc$ in the six cell system.

Figure 3.9 shows the effect of variation in the V_t parameter for the segmented system of 64 current cells. Again the segmented system is more robust than the compact systems due to linear interpolation. As expected, all systems are less robust to V_t mismatch than β mismatch as the V_t term is included in the square expression of the first order MOS equation of drain current (Equation 2.2). For the segmented system to meet the Bluetooth and UMTS specification of better than $-60dBc$ SFDR the V_t mismatch need to be controlled to within 4% of V_{phase} and the control requirement tightens significantly to 1% of V_{phase} in order to meet the $-70dBc$ SFDR of state of the art low power, high speed, CMOS DACs.




Figure 3.8: The estimated yield against percentage error in the phase parameter for various minimum allowable SFDR for the compact systems.




Figure 3.9: The estimated yield against percentage error in the phase parameter for various minimum allowable SFDR in the 64 cell segmented system.

Parameter	I_{bias}	β	V_t
Compact			
Four cells	1%	0.3%	0.3%
Six cells			
$> -40dBc$	3%	3%	2%
$> -45dBc$	1%	0.6%	0.5%
Segmented			
$> -60dBc$	1.5%	5%	4%
$> -70dBc$	0.4%	4%	1%

Table 3.4: Allowable percentage error required to ensure a yield of better than 90% of systems.

3.5 Summary

The matching requirements of the transistors in the current sources and current switches have been estimated using MATLAB to reduce the effect of mismatch and guarantee a good yield. To easily implement the parameters they have been truncated to ratios of small integers. Truncation caused no detrimental effect on the SFDR of the compact systems. For the linear 64 cell system the SFDR degraded to $-72dBc$ due to truncation in the switch devices. However, this is still within the desired requirement of meeting the spectral purity of $-70dBc$ of high speed CMOS DACs established in Chapter 1 and offers a margin of error over the UMTS and Bluetooth requirements of $10dB$.

As expected, local variations caused by mismatch between transistors in the same circuit degrade the performance of all systems. The effect of mismatch on the compact systems suggest a required accuracy of 1% for the current sources. This is achievable using transistors of a small area. The segmented system requires more accurate current sources with matching of better than 0.4%. This will require larger devices, however this accuracy requirement is not too demanding. The segmented system was more robust to variations in the gain, β and threshold voltage, V_t of the switch transistors. This is due to the use of linear interpolation in the segmented system. The simulated variation in V_t also captures the percentage variation in V_{ref} and V_{phase} . All systems are less robust to variations in V_t than β because the V_t term appears in the square of the first order model of the MOS drain current, Equation 2.2. The accuracy requirements of the switch transistors are summarised in Table 3.4 and will be used to determine the transistor sizes in the following two chapters.

Chapter 4

Circuit simulations and design of compact systems

4.1 Introduction

MATLAB simulations of compact nonlinear DACs consisting of arrays of four or six current cells have shown promising spectral purity results. Mismatch will not affect the yield if the variations between the devices is controlled to within better than 1%. These accuracy figures will be used to determine the minimum transistor sizes and the required bias voltages.

The MATLAB model used in the preceding chapters, is based on the first order model of transistor behaviour which fails to take into account second order effects that may introduce distortion into the output at all frequencies. The Cadence circuit simulation software is used in this chapter to simulate the system in more detail. The BSIM 3 CMOS model [63] within Cadence uses over 50 equations to model various second order effects accurately, using parameters provided by the circuit manufacturer. In addition the measurement data collected by the circuit manufacturer describes the statistical behaviour of mismatch using the Pelgrom model, and this allows the yield predictions of the previous chapter to be confirmed by Monte Carlo Cadence circuit simulations.

Of the other MOS simulation models (EKV, HSPICE level 28, MOS 9) that might be used to examine the second order effects, EKV is the most popular and widely used. EKV [1] was designed to better model low power MOS circuits by using models of device behaviour based on the underlying physics. This leads to a model with only extracted 18 parameters rather than nearly 100 required by BSIM.

EKV is also better at modelling the transition between above threshold and subthreshold which spans nearly four decades of drain current [64]. Terry and coworkers find that both models are deficient at modelling output resistance over a range of transistor lengths and BSIM is better at modelling output conductance at short gate lengths. Hence there is no clear superior model. Unfortunately the process parameters for the EKV simulation were not available at the time and hence all simulations presented here use the BSIM model.

This chapter also presents the architecture of prototype circuits to test the nonlinear interpolation DDFS concept for compact arrays of four and six current cells. The architecture consists of an accumulator, phase DAC and nonlinear DAC. The only superfluous addition to the system is the phase DAC that provides the analogue interpolation voltage. The phase DAC is required to drive the gate capacitance of the switch transistors in the nonlinear DAC. Hence one design choice is the current output required by the phase DAC to ensure that there are no unwanted transient effects such as glitches that degrade the output. This is simulated using the BSIM 3 model which includes parasitic capacitances associated with the devices to enable more realistic modelling of the transient effects.

In this chapter the Cadence simulations of the compact systems at low frequencies are compared to the MATLAB simulation results. Dimensions of the transistors used in the nonlinear DACs and the choice of bias levels are given. The design of the phase DAC is then addressed by using Cadence simulations at higher frequencies. Finally the entire DDFS system, including the accumulator is simulated in mixed signal Cadence simulations to examine the spectral purity at higher output frequencies.




Figure 4.1: The architecture of a compact system. Nine phase bits are used to ensure that the system is not limited by phase resolution

4.2 Architecture of compact systems

The architecture of the compact systems with either four or six cells (Figure 4.1) is based on the ROM-less DDFS architecture that was discussed in Section 1.6.4. The four cell or six cell compact array of current cells is represented by the nonlinear DAC. A linear phase DAC is used to generate the analogue input voltage that represents the phase, V_{phase} . To ensure the SFDR of the six cell system is not limited by phase quantisation 9 phase bits are used, (Equation 1.7). The half wave symmetry of these bits is determined by the XOR block that uses the first MSB from the accumulator to take the 1's complement of the next 8 MSB bits using the process explained in section 1.6.3.

4.2.1 Circuit simulation conditions

To determine the transistor sizes and reference voltages for the nonlinear DAC the following requirements need to be considered. The voltage rails in the manufacturing process are limited to 3.3V. Within this voltage range a saturated current source, current switch and output voltage swing must be allocated. In Chapter 2 a V_{gs} of approximately 1V was allocated for the current source. Hence the input voltage may be allowed to move over a range of approximately 2V. However for the systems with more than one current cell the input voltage swing is required to span the voltage range of all the current cells. For this reason the four and six cell systems were simulated with a 200mV voltage range for each current cell. These voltages were chosen with a compromise of limited phase accuracy for voltage range.

4.2.2 Device sizes

The dimensions of the devices in the nonlinear DAC were chosen to ensure a high yield as a prototype for the instrumentation systems described in section 1.5. With a V_{gs} of 1V the current source devices within the six cell system need to be at least $8\mu^2$ to reduce $\delta I/I$ to 1% (Figure 3.1, Table 3.4). An area of $10\mu^2$ was chosen and the dimensions of the unit device are $W = 1\mu m$ and $L = 10\mu m$. The four cell system has a reduced input voltage range, therefore the gate voltage can be higher and the current source devices are allocated a smaller area. An area of $3\mu^2$ was chosen and the dimensions of the unit device are $W = 1\mu$ and $L = 3\mu$. Devices with small W/L ratios were chosen as they have better matching than devices with large W/L ratios [65].

In the four and six cell systems the V_t mismatch dominates the effect of mismatch in the switch transistors and hence determines the area. The four cell system has the smaller input voltage range of $800mV$ with a mismatch requirement of 0.3% (Table 3.4). From Equation 3.1 and a typical A_{V_t} of $5mV\mu$ the area of the switch devices for both the four and the six cell system need to be more than $5\mu^2$. The sizes are given in Table 4.1.

current cell	number of unit	switch transistors	V_{ref_i}
index (i)	current sources	$W/L (\mu m)$	reference voltage (V)
four cells			
1	3	4.3/2.2	1.9
2	7	10/2.2	2.1
3	7	10/2.2	2.3
4	3	4.3/2.2	2.5
six cells			
1	4	4.05/3.5	1.7
2	11	11.5/3.5	1.9
3	15	16/3.5	2.1
4	15	16/3.5	2.3
5	11	11.5/3.5	2.5
6	4	4.05/3.5	2.7

Table 4.1: Current source, switch transistor sizes and reference voltages used in circuit simulations of the compact systems. For ease of layout a constant length was chosen for the switch devices in each system.

The total active area of the four cell system was $116\mu m^2$ and the six cell system $656\mu m^2$. For reference the area of a single pad is $10000\mu m^2 (0.01mm^2)$, demonstrating that the four and six cell circuits are extremely compact and suggesting that DDFS systems based on these DACs will be very small.

4.3 Low frequency circuit simulations of compact systems

Initial circuit simulations at low frequency were used to quantify the effects of transistor behaviour such as limited output resistance (channel length modulation), body effect, leakage current, short channel, and subthreshold operation on the performance of the circuits. An ideal piecewise linear voltage source was used in these simulations to generate the analogue input voltage. This ensured that any distortion was solely caused by the four or six cell circuits. As in the MATLAB simulations of Chapter 2, the output current from the reference branch of the array current cells is chosen in preference to output current from the input branches as the reference branch is less effected by glitches and clock feedthrough.




Figure 4.2: Circuit simulation of an array of four cells. The results show similar trends to the MATLAB simulations. The 5th harmonic showed a similar trend to the 3rd harmonic at a lower power level and is not shown for clarity.

In agreement with the results from the MATLAB model, at high bias current values, when the output is a triangular wave, the relative importance of the third harmonic is dominant in the Cadence simulation (Figure 4.2). However, the results at lower bias currents are quite different from MATLAB. In particular, unlike the MATLAB results, the circuit simulation results suggest that the 3rd harmonic reduces to less than -40dBc at low bias currents whereas the MATLAB simulations predicted that it reduced to the noise floor of -100dB (see 2.7). The MATLAB simulations also predicted that the 2nd harmonic is below the noise floor, however it is above -60dBc in the circuit simulations.

The circuit simulation in Cadence of the six cell system showed more deviation from the MATLAB simulations, (Figure 4.3). The best case SFDR was -52dBc , 4dB better than the best results from MATLAB and the SFDR is better than -40dBc for most of the bias range. At low bias values, the 11th and 13th harmonics become dominant and at high bias values the 3rd and 5th become dominant in agreement with the MATLAB. However, in the best bias range, between 0.7V and 1.1V , the 5th and 7th harmonics are dominant in Cadence whereas they were never dominant in MATLAB.

One possible cause of the difference between Cadence and MATLAB is that they employ different fast Fourier transform (FFT) algorithms. Cadence samples the output data uniformly by interpolation and then applies the direct Fourier transform [66] while MATLAB uses a radix-2 FFT algorithm for




Figure 4.3: Circuit simulation of an array of six cells. The results are significantly different from MATLAB simulations.

vectors with a length of a power of two. To investigate this inconsistency, the output waveform from Cadence was saved and the FFT calculated in MATLAB. Identical frequency spectrums were computed by Cadence and MATLAB indicating that the anomaly must be caused by another effect.

Another effect not represented in MATLAB is the V_{ds} dependence of the current sources that is caused by channel length modulation in MOS devices (finite output resistance). This difference between Cadence and MATLAB was further investigated by removing the V_{ds} dependence from the Cadence simulations by using ideal current sources instead of MOSFET transistors. This reduced the power of the 2nd harmonic in both four and six cell simulations, indicating that limited output resistance causes some second harmonic distortion. This is expected as the current cells that are addressed by lower reference voltages will have lower V_{ds} across their current sources. At lower bias voltages these current sources will leave the saturation region breaking the quarter wave symmetry of the output leading to second harmonic distortion in the output. The second harmonic is not dominating the SFDR and therefore the output resistance is sufficient. The other dominant harmonics maintained their power levels and hence must be caused by another effect that is modelled in Cadence but not in MATLAB.

4.3.1 Subthreshold operation

The performance improvement seen in the Cadence simulations of the six cell system at bias voltages between $0.7V$ and $1.1V$ may be due to a smoothing effect of subthreshold operation (see Figure 4.4). The equations used in MATLAB fail to take into account the subthreshold region of operation of a MOSFET at low source-drain currents. It is therefore not surprising that the simple MATLAB model fails to predict the exact shape of the nonlinear response of each circuit.




Figure 4.4: Modeling differences between MATLAB and Cadence. The difference between the two nonlinear interpolation curves is shown exaggerated here but is more than 1% in simulation. This is significant in a signal with a spectral purity of $-40dBc$. The MATLAB trace has a hard discontinuity whilst the Cadence trace starts and saturates more smoothly.

The ability of the circuit simulator to model subthreshold current causes a small difference in the interpolation function of the differential current switch. The difference in the response of one of these cells is demonstrated in Figure 4.4. A larger deviation occurs when either the reference or input transistor is turning off. The function used in MATLAB has a hard discontinuity whilst the Cadence data starts and saturates more smoothly. This hard discontinuity arises from an approximation used in the MATLAB model. In the MATLAB model the current is set to 0 or I_{bias} when the expression within the square root of Equation 2.3 is negative, (see the MATLAB code in Appendix C for details). The model used in BSIM 3 for subthreshold operation provides a more detailed continuity at this point by using the second order model. The transistors enter the subthreshold region of operation for a small amount of the output period shown exaggerated in Figure 4.4. A 1% variation in current

is important in signals of spectral purity better than -40dBc (see section 3.3.2). When the total amount of bias current is decreased, the transistors are in subthreshold for a larger fraction of the input voltage swing. Hence there is more discrepancy between the MATLAB and Cadence simulations when the total bias current is decreased.

The most important conclusion of comparing the MATLAB simulation results to the Cadence simulation results is that they agree that an ideal four cell system will have a SFDR of -40dBc and an ideal six cell system will have a SFDR of better -48dBc . The Cadence simulations use a more detailed model of transistor behaviour which explains the differences in the results between the two types of circuit simulation. In this case the MATLAB simulations appear to be an underestimate of the SFDR.

4.3.2 Confirmation of the effect of variations

The Cadence and MATLAB simulations suggest that the spectral purity performance of the systems is as good or better than the MATLAB predictions. However the choice of transistor sizes will determine how robust the performance is to variations inherent in the manufacturing process. The control of parameter variations by sizing the transistors appropriately using the mismatch analysis of Chapter 3, was confirmed in Cadence. These simulations used the BSIM 3 model and variation parameters of the Pelgrom model (Section 3.2), which have been measured from manufacturing runs by the circuit manufacturer. The results of 100 Monte Carlo simulations at the optimum bias voltage for the six cell system indicated a yield of 79% for spectral purity better than -50dBc . For spectral purity better than -45dBc the yield is 100%. The four cell system had a yield of 70% for systems with a SFDR of better than -40dBc and a yield of 100% for systems with a SFDR of better than -35dBc . In agreement with the earlier Cadence simulations at this operating point the dominant harmonic in the four cell system is the 3^{rd} harmonic, while the 5^{th} harmonic is the dominant spur in the six cell system. These results confirm high yields for the compact four and six cell systems with SFDRs expected from the MATLAB simulations.

4.4 Phase DAC

In the compact architecture (Figure 4.1) the nonlinear DAC is small and the current cells are 100% efficient as all of the bias current can be used to form the output signal. A linear phase DAC is required to form the analogue input signal. This is the only additional part of the system and therefore requires

a careful design.

4.4.1 Phase DAC device sizes

A binary coded current switched architecture was chosen for the phase DAC for small size, high speed operation and simplicity [44]. The phase DAC generates the analogue input voltage that drives the gate of switch transistors in the current cells within the nonlinear DAC. Therefore the size of the phase DAC current sources were determined by considering the effect of variations in threshold voltage, V_t , input voltage, V_{phase} and reference voltages, V_{ref} that were investigated in Chapter 3. The variations analysis suggested that the analogue phase voltage at the output of the phase DAC must be accurate to within 0.5% for the compact systems. Hence the phase DAC current sources need to be matched to within 0.5%. This corresponds to a unit current source of area $8\mu^2$ biased with a gate voltage of 1.5V (see Figure 3.1). Higher gate bias voltages would enable smaller unit current sources, to reduce cost, however the voltage headroom is limited by the input voltage range required by the interpolation system of six cells (1.2V). The switch transistors are of minimum size to reduce the parasitic capacitances that cause clock feedthrough and increase glitches. One drawback of using small switches is the increased on resistance that will increase the settling time of phase DAC output. NMOS transistors are used as switches as they have higher charge mobility and hence lower on resistance and clock feedthrough.

There are N current cells within a binary encoded DAC. The first current cell is assigned a weight of 1 and the N^{th} current cell is assigned a weight of 2^N . Hence the number of unit current cells needed in a binary DAC is $2^{(N-1)}$ and an 8-bit binary encoded DAC requires 255 current sources which require an active area of $2040\mu^2$. The nonlinear array of six cells uses an active area of $656\mu^2$ hence the phase DAC is approximately four times larger than the six cell array. A current switched phase DAC requires a resistor to change its output from a current to a voltage. In the process used, a resistor is required to be at least $5\mu^2$ to achieve matching of 0.5%. For comparison, a resistor string DAC would require 255 unit resistors of $5\mu^2$ or a total area of $1275\mu^2$. However the resistor string DAC would not be as fast as the current switched DAC [44]. The phase DAC is required to form the analogue input signal and will have a large impact on the size, power efficiency and bandwidth of the entire system.

4.4.2 Phase DAC power dissipation

The power dissipation and output current of the phase DAC are determined by the speed requirement to drive the input capacitance at the gates of the differential switch transistors of the current cells in the nonlinear DAC. In the compact architecture all cells are driven by the phase DAC at the same time so the total capacitance will determine the bandwidth. The signal is required to drive at least six switch devices, each of which occupy an active area of $7\mu^2$. These will have a combined capacitance of less than $0.2pF$ (C_{ox} is less than $5fF/\mu^2$). The current required to drive this load is determined by $i = C \times \delta V / \delta t$ where δV is $1.2V$, the maximum voltage swing, and δt is the rise time (2% of the period¹) of the fastest moving signal. This means that at least $10\mu A$ is required to drive the capacitance at $10MHz$ and $100\mu A$ to drive at $100MHz$.

<i>LSB</i>	SFDR	SFDR	SFDR
	$10\mu A$	$100\mu A$	$200\mu A$
0.1MHz	- 52.0 dBc	- 52.0 dBc	- 52.0 dBc
1MHz	- 51.4 dBc	- 52.0 dBc	- 52.0 dBc
10MHz	- 42.7 dBc	- 51.4 dBc	- 52.0 dBc
100MHz	- 30.1 dBc	- 41.7 dBc	- 51.3 dBc
200MHz	- 21.0 dBc	- 38.2 dBc	- 41.2 dBc
500MHz	- 21.0 dBc	- 34.0 dBc	- 35.3 dBc

Table 4.2: The transient performance of six cells.

Full analogue Cadence circuit simulations of the six cell system were used to determine the amount of current required in the phase DAC of each system and hence give an indication of the required rise time. In these initial simulations the linear phase DAC was driven by ideal voltage inputs that represented an overflowing accumulator with a phase word of 1 loaded into the phase register. In these simulations the fastest changing signal is the LSB. The transient performance is given in Table 4.2, columns two, three and four. As expected the results show that there is not enough current to charge the capacitance when using a phase DAC with only $10\mu A$ of output current when the LSB frequency is $10MHz$ or more. This is due to slowed rise times causing distortion in the analogue phase voltage. This distortion introduces a strong second harmonic effect which is dominant over any other transient effects. For a particular SFDR the current and LSB frequency are proportional. Hence by increasing the current to $200\mu A$, the bandwidth of the system is increased to LSB frequencies up to $100MHz$ for SFDR better than $-50dBc$.

The four cell system requires at least 7 bits of phase resolution to ensure an SFDR of at least $-42dBc$

¹2% of the period is an arbitrary specification and the required risetime is determined by increasing the frequency in simulation.

<i>LSB</i>	SFDR	SFDR	SFDR
	$10\mu A$	$50\mu A$	$100\mu A$
0.1MHz	-40.0dBc	-40.0dBc	-40.0dBc
1MHz	-40.0dBc	-40.0dBc	-40.0dBc
10MHz	-38.4dBc	-40.0dBc	-40.0dBc
100MHz	-32.0dBc	-38.0dBc	-39.1dBc
200MHz	-28.3dBc	-36.8dBc	-38.0dBc
500MHz	-22.9dBc	-33.7dBc	-36.7dBc

Table 4.3: The transient performance of four cells.

the MSB for half wave symmetry, and a 6-bit phase DAC. The combined gate capacitance of the four cell system is less than the six cell system as there are less transistors used, the analogue voltage range is smaller and the spectral purity is relaxed by approximately $10dB$. Table 4.3 shows that to maintain an SFDR of $-38dBc$, the phase DAC that drives the four cell system is required to draw at least $10\mu A$ of current for LSB frequencies up to $10MHz$, $50\mu A$ for LSB frequencies of $100MHz$ and $100\mu A$ for LSB frequencies of $200MHz$.

4.5 High frequency and mixed signal circuit simulation

The simulations in previous sections were performed using a fixed phase word. The transient performance results in Table 4.2 and Table 4.3 are from full analogue simulations with a set of eight or nine ideal voltage sources representing the output of an accumulator with a phase word of one. It is difficult to specify ideal voltage sources to represent an accumulator overflowing with phase words greater than one as all 2^N combinations need to be coded by hand. Cadence provides the capability of simulating some parts of a system with a digital simulator such as Verilog and other parts with the analogue simulator (Spectre). These mixed signal simulations provide more flexibility in specifying the digital components of the system. The effect of generating higher output frequencies was therefore studied using the mixed signal simulator to generate the output of an accumulator with various phase words. One limitation of the mixed signal simulator was a restriction within Cadence to clock frequencies of $5MHz$, $50MHz$ and $500MHz$.

4.5.1 Mixed signal simulation

In the six cell system at least a 9-bit accumulator is required to provide 9 phase bits to ensure the SFDR is not limited by phase truncation. By removing the need for a ROM and reducing the number

of current cells significantly, the accumulator consumes most of the power and area in the DDFS system [67] and [68]. For applications where better frequency resolution is required the low power accumulator presented in section A.1 could be used. At the interface with the accumulator the first MSB from the accumulator is used as an input in a bank of XOR gates which take the 1's complement of the next 8 most significant bits. This allows the system to use the half wave symmetry of the cosine wave to halve the number of required current sources. The eight bits from the XOR block are used to generate the input waveform using the small, binary encoded 8-bit linear phase DAC. By using 9 phase bits from the accumulator to generate the input to the nonlinear DAC the system is not limited by phase truncation in the following simulations.

A 9-bit accumulator can be specified in Verilog as follows:

```
module accum9 (A, B);
output [0:8] A;
input [0:8] B;
reg [0:8] A;
initial A = 0;
always
#2 A = A + B;
endmodule
```

The phase word can then be specified as the input B . The output A is used to address the phase DAC.

4.5.2 Settling time of the phase DAC

The output of the phase DAC was simulated using the mixed signal simulator when driving the four cell system and the results are shown in Figure 4.5 for clock frequencies of $5MHz$, $50MHz$ and $500MHz$. The effect of quarter wave symmetry can be seen in the height of the largest glitch at the second MSB transition. As expected it is positive in the first half period and negative in the second half period. The settling time is $5ns$ when $f_{clk} = 50MHz$ and the degradation of the SFDR with increasing phase words is acceptable. However when the clock period is reduced to $2ns$, i.e. $500MHz$ the settling time is longer than the clock period and the DAC signal never settles to its desired value. This limits the clock frequency to less than $f_{clk} = 500MHz$. Despite the poor quality of the phase DAC output at $500MHz$ the four cell system attained a SFDR of $-36dBc$ and the six cell




Figure 4.5: Settling time of the phase DAC signal when driving the four cell system. (a) $f_{clk} = 5MHz$, (b) $f_{clk} = 50MHz$, (c) $f_{clk} = 500MHz$.

system $-41dBc$, the reason for the loss of $10dB$ in SFDR is evident from the poor phase DAC output signal in Figure 4.5(c).

4.5.3 Simulations at $f_{clk} = 50MHz$

The results of mixed signal simulations with $f_{clk} = 50MHz$ suggest that the four cell system is capable of producing better than $-35dBc$ for f_{out} up to $1MHz$. The simulations were performed over a range of output frequencies from $f_{clk}/2^9$ to $(f_{clk}/4) - 1$. The performance falls with an increase in output frequency as expected, Figure 4.6. In a similar fashion to all DDFS systems the clock alias spur ($f_{clk} - f_{out}$), and its harmonics degrade the inband spur due to clock feedthrough, see section 1.6.2. This effect degraded the performance of the system at output frequencies of greater than $3MHz$. In a real circuit a low pass filter is expected to be formed at the output by the output pin capacitance and load resistor. The capacitance of the circuit pad was $1.3pF$ while the load of the scope probe was $1.5pF$ using a times 10 probe, and the input capacitance of a buffer is $1.2pF$. With the $20k\Omega$ output resistor the two measurement arrangements of scope probe loading or a gain buffer, form a low pass filter with cutoff of $3-4MHz$ and hence the clock and other spurs above this frequency will be attenuated in the measurement. A capacitor of $1pF$ and a resistor of $20k\Omega$ were included in the

simulations to represent these parasitics and the filtering effect improved the SFDR to $-40dBc$ at $1MHz$. Hence the load acts as a good enough filter to reject the aliased output.

A similar effect was observed in mixed signal simulations of the six cell systems with an identical phase DAC used to generate the interpolation voltage, Figure 4.7. The unfiltered spur degrades from the optimum SFDR at $100kHz$ due to the larger gate capacitance of the six cell array. The filtering effect of the parasitic output improved the performance to $-50dBc$ at $1MHz$ output frequency. As the phase word is increased the dominant harmonics of the clock alias spur begin to constructively interfere with the dominant harmonics of f_{out} . For example there was degradation in SFDR to $-48dBc$ when the phase word was increased to 51 with $f_{clk} = 5MHz$ ($f_{out} = 0.498MHz$) while at $f_{clk} = 50MHz$ and $PW=5$ ($f_{out} = 0.488MHz$) there was no degradation in SFDR. Therefore increasing f_{out} by increasing the phase word degrades the SFDR more than increasing f_{out} by increasing the clock frequency. This is expected as increases in the phase word cause larger voltage changes at the gates of the current switches as well as faster changes, while increases in clock rate only increase the frequency of changes, not the magnitude. This suggests that the performance is limited by an effect of loading on the phase DAC, however the phase DAC uses a large amount of current and hence was not modified further.

4.6 Power dissipation

Component	Current (μA at $100MHz$)	Voltage Supply (V)	Power Dissipation ($\mu W/MHz$)
Four cell system			
Nonlinear DAC	20	3.3	0.66
Linear Phase DAC	100	2.6	2.6
8-bit XOR Complementor	4	3.3	0.1
8-bit Accumulator	40	3.3	1.3
Total			4.7
Six cell system			
Nonlinear DAC	30	3.3	0.99
Linear Phase DAC	200	2.8	5.6
9-bit XOR Complementor	4.5	3.3	0.1
9-bit Accumulator	45	3.3	1.4
Total			8.1

Table 4.4: Power dissipation breakdown of the compact systems. The total power dissipation is over 10 times less than state of the art systems (see Table 1.3).

The bias current for best spectral purity in the four cell system was $20\mu A$ and in the six cell system




Figure 4.6: Mixed signal simulation results of the four cell system simulated in Cadence with a 9-bit accumulator, 8 XOR gates, 8-bit DAC and $f_{clk}=50MHz$. The in band spur, ($< f_{clk}/2$) has a constant magnitude up to output frequencies of $1MHz$. The clock alias spur, ($f_{clk} - f_{out}$) dominates above $500kHz$ and is reduced by the low pass filtering. The filtered output, ($(< f_{clk}/2)_{filt}, (f_{clk} - f_{out})_{filt}$) is better than $-40dBc$ for f_{out} up to $1MHz$ and the narrowband SFDR is better than $-35dBc$ for f_{out} up to $10MHz$




Figure 4.7: Mixed signal simulation results of the six cell system simulated in Cadence with a 9-bit accumulator, 8 XOR gates, 8-bit DAC and $f_{clk} = 50MHz$. The in band spur is below $-50dBc$ up to output frequencies of $1MHz$. The clock alias spur dominates and may be reduced by using a low pass filter. The filtered output is better than $-50dBc$ for output frequencies up to $1MHz$, and the filtered, narrow band SFDR better than $-40dBc$ for output frequencies up to $10MHz$.

30 μA . A phase current of 100 μA was required to operate the four cell system at 100MHz and 200 μA to operate the six cell system at the same frequency due to the increased capacitive loading and improved spectral purity. A larger current for the phase DAC is acceptable as it operates from a lower supply voltage, see Table 4.4. Low power accumulators that use paralleling and pipelining (see Appendix A) are able to operate with less than 0.05 $\mu W/MHz$ per bit. Hence the expected power dissipation is 0.005mW/MHz for the four cell system and 0.008mW/MHz for the six cell system (Table 4.4). Despite using most of the power in the phase DAC, these figures are over 10 times less than state of the art systems, Table 1.3.

4.7 Summary

Detailed circuit simulations of the compact systems are encouraging and predict an improved spectral purity over the MATLAB simulations. These simulations suggest that the SFDR of the six cell system can be improved by 2dB to better than $-50dBc$ SFDR due to the smoothing effect of transistors operating in subthreshold. The transistors of the current cell are in subthreshold when the current cell is turning on or off. In the first order MATLAB model the transistors turned on or off with a sharp discontinuity.

The only additional component needed in the compact systems is a small binary encoded phase DAC. This DAC is used to create the analogue input phase voltage required for interpolation. The size of the phase DAC is equivalent to the size of the six current cells combined. This large size is required to ensure the linearity of phase DAC output is not affected by mismatch variations. Unwanted transient effects caused by the phase DAC can be controlled by increasing the current output of this DAC so that it is able to drive the gate capacitance of the current cells of the nonlinear DAC. To drive the system at clock frequencies of 100MHz the current dissipation in the phase DAC is greater than the output current, however the total power dissipation in the four cell system is less than 5 $\mu W/MHz$ and in the six cell system 8 $\mu W/MHz$.

As expected transient effects at the output of the phase DAC degrade the spectral purity at high output frequencies where the harmonics of the clock alias spur combine constructively with the dominant spurs. However appropriate use of low pass filters enables the spurs to be suppressed. In the four cell system a low pass filter at 8MHz, formed by the output load and the parasitic capacitance of the pad pin, reduced these spurs to less than $-35dBc$ when the output frequency is 1MHz and hence these systems are suitable for the instrumentation systems discussed in Chapter 1 with a margin of error

of $5dB$ over their SFDR requirements. For the six cell system better than $-40dBc$ SFDR is expected with a yield of 100% at clock frequencies up to $f_{clk} = 50MHz$, suitable for integration in active sensors for electrical impedance spectroscopy. The four cell system is expected to operate up to the same frequency range with SFDR better than $-35dBc$ and 100% yield, suitable for dielectrophoresis application for cell characterisation.

The yield of six cell systems with spectral purity better $-50dBc$ SFDR is expected to be 79%. This spectral purity is within the requirement of the Bluetooth specification for the low output power classes 2 and 3. However at output frequencies of greater than $1MHz$, transient effects degrade the performance. Hence this system is not appropriate to generate the $80MHz$ bandwidth required by the Bluetooth specification. The 64 cell segmented architecture will be investigated for this application in the next chapter.

Chapter 5

Circuit simulations and design of segmented systems

5.1 Introduction

The ability to digitally synthesise sinusoid waveforms with good spectral purity at low power consumption is important in wireless handheld devices. It was predicted in Chapter 2 that the nonlinear interpolation system with 64 or more current cells should be able to generate signals with a SFDR of better than -70dBc . This is the target of the segmented system presented in this chapter. Nonlinear interpolation systems with more than six current cells require a change in architecture from the compact system to a segmented architecture where each cell represents a particular segment of the output waveform.

This change in architecture is required due to four limiting effects. Firstly the limited voltage range available in low power CMOS restricts the number of current cells that can be arranged within the power rails whilst keeping the transistors within the current cells saturated. The bias voltage of the current source requires some headroom to keep the variation within the 0.4% requirement determined in Chapter 3. Secondly the limited accuracy in setting reference voltages restricts the number of cells that may be individually selected within a single input voltage range. Thirdly the effect of mismatch on the threshold voltage of switch devices of each current cell places a lower limit on the voltage swing per current cell. In Chapter 3 the allowable variation in threshold voltage of the switch devices was determined to be 1% of the phase voltage swing. Finally transient effects

introduce glitches which cause high frequency spurs in the output spectrum. Design techniques used in high speed current switched CMOS DACs such as thermometer decoding, reduced voltage swings and latches will be used to combat these unwanted transient effects.

A possible architecture, introduced in Chapter 2, uses a modified thermometer decoding technique to address an array of current cells. The phase bits that are normally discarded in a conventional system are used with the phase DAC to create the analogue phase input. A 64 cell system was chosen for further study to demonstrate an improvement over state of the art DDFS systems with a SFDR of better than $-70dBc$. This chapter examines the impact of second order and transient effects on the spectral purity of the 64 cell segmented system by using the Cadence circuit simulation tool. The circuits required to implement the 64 segmented cell architecture are a phase DAC, a modified thermometer decoder and a modified current cell for the nonlinear DAC. These circuits are introduced in this chapter and the entire system is simulated using the Cadence mixed signal circuit simulator. Statistical information from the circuit manufacturer is used to confirm the correct sizing of the transistors to achieve the accuracy suggested by work in Chapter 3 for a high yield of working circuits.

5.2 Architecture of segmented system




Figure 5.1: The architecture of the segmented system. A conventional thermometer decoded ROM-less DDFS is shown by the unbroken lines. The dashed lines show the additional components required for the segmented interpolation architecture.

The defining difference between the compact and segmented systems is the architecture. While the compact systems used a small number of current cells and a small number of voltage references, for a system with 64 cells, too many voltage references within the limited voltage range are required. Hence the popular thermometer decoding method is modified to create an architecture that is able to interpolate between 64 current cells using only one reference voltage.

Thermometer decoding reduces glitches by ensuring that the switching is monotonic. The thermometer decoder can be viewed as a linear string of N current cells which are indexed by the digital phase input. For the output to represent a cosine function, the i^{th} current cell should represent $\cos(i\pi/N) - \cos((i-1)\pi/N)$ for phase in the range 0 to π . The cosine function in the range π to 2π can be generated by exploiting the symmetry of the cosine function. This can be achieved by taking the ones complement of the phase bits before the input to the DAC by using the first MSB of the phase bits and a bank of XOR gates.

A nonlinear DAC for DDFS that meets the performance specification of having a SFDR better than -72dBc requires 13 phase bits to ensure the system is not limited by phase truncation. In a conventional architecture, using the first MSB for half wave symmetry, the nonlinear DAC would require $2^{(13-1)} = 4096$ current cells. The segmented method presented in section 5 suggests that each current cell can be used more efficiently by using interpolation to reduce the number of current cells to as few as 64.

To understand how interpolation can be used to reduce the size of a nonlinear DAC consider the first P bits of the phase accumulator (Figure 5.1). The first MSB is used to determine the half-period and the remaining bits are used to address a conventional nonlinear DAC. In a conventional system the number of phase bits are truncated to $(W-1)$ bits to reduce the number of current cell units. Alone this will limit the SFDR to $-6W$ dBc due to phase truncation. To improve the SFDR the remaining phase accumulator bits $(P-W)$ may be used to create an analogue voltage (V_{phase}) that partially turns on a single current cell inside a copy of the conventional DAC labelled the interpolation DAC¹ in Figure 5.1. The combined output of the two DACs is an improvement in SFDR of up to $-6(P-W)$ dBc. The analogue phase interval voltage may be produced by using a small linear phase DAC that is addressed with the $P-W$ accumulator bits.

To determine which current cell in the interpolation DAC should receive the analogue phase voltage consider the operation of the conventional DAC. The most significant bit is used for half wave symmetry and the next 6 most significant bits for the thermometer decoder. At any one instant only

¹This DAC is interfaced with both a digital address (the thermometer decoder) and the analogue input function, V_{phase} .

current cells 1 to J may be contributing to the output. In a system with 64 cells the combined sum of the outputs should represent the amplitude for a phase of $J\pi/64$. Using the conventional thermometer decoder the maximum SFDR for a conventional system is ($-6W = -42dBc$ for $W = 7$). The conventional method of improving the SFDR is to increase W , and increase the number of current cells. However if we drive the current cell that is just about to turn on, i.e. the $(J + 1)^{th}$ current cell, of the interpolation DAC with an analogue phase input, (formed using the remaining $(P-W)$, phase bits), then this current cell will contribute a fraction of its current to the output. In effect this cell is performing interpolation between the absolute outputs available from a conventional DAC. Therefore only the $(J + 1)^{th}$ current cell in the nonlinear DAC should receive the analogue phase voltage. In this way the phase truncation limit on SFDR for 64 current cells is increased up to $-6P = -78dBc$ for $P = 13$. The cells addressed in the conventional and nonlinear arrays are mutually exclusive and hence may be combined if an appropriate switch is used, see section 5.2.2.

5.2.1 Thermometer decoding

Thermometer decoding is popular in the design of DACs to reduce static and dynamic errors. One example of decoding logic for an array of 64 thermometer decoded current cells splits 6 address bits into 3 MSBs and 3 LSBs which address a two dimensional array of 64 current sources, shown in Figure 5.2. Generally the conventional thermometer decoder has two decoding stages, the first n MSBs are decoded to 2^n row flags, the remaining m MSB's are decoded to 2^m column flags. The flags address an array of 2^n by 2^m current cells (see Figure 5.2 for $n = m = 3$). Decoding logic in each cell switches the current source with the following rules: If the two adjacent rows are both high (row_i and row_{i+1}), the current source is turned on, if the adjacent rows are both low, the current cell is turned off. If the two rows are different, the column signal (col_j) is used to determine if the cell should be turned on, for example, see Figure 5.3 for the 8x8 cell system. The logic at the cell is (active low):

$$CELL\ ON = Hi = (row_{i+1}\ OR\ col_j)\ NAND\ row_i$$

$$CELL\ OFF = Lo = NOT((row_{i+1}\ OR\ col_j)\ NAND\ row_i)$$

In most cases more than one cell switches for a given change in input phase so a latch is used within each cell to ensure that the logic is synchronized with the clock otherwise a glitch will be produced whenever two or more cells change asynchronously. The power rails on the latch are often used to decrease the voltage swing [48]. This reduces the voltage change at the source of the switch

transistors and leads to a smaller glitch at the output. The nonlinear interpolation system can be implemented using a thermometer decoding logic similar to the conventional thermometer decoded DAC. The logic within the cell is modified to allow for the analogue phase voltage, V_{phase}




Figure 5.2: A 64-cell thermometer decoded array. A binary input of 30 will turn on the first 30 cells




Figure 5.3: A single current cell with conventional local decoding logic. The latch is used to ensure cells switch in synchrony.

In the previous section it was shown that in a set of thermometer decoded current sources, where sources 1 to J are contributing to the output, only the $(J + 1)^{th}$ current source should be fed the analogue phase input to partially turn on. Therefore the logic within the current cell needs to be modified to receive the analogue phase input when it is the $(J + 1)^{th}$ current source.

To achieve this, the row and column decoders of a conventional decoder were used (Figure 5.2), and




Figure 5.4: The decoding logic necessary for the interpolating thermometer decoder.

the logic at the cell was modified to include a state called AN . The cell receiving the analogue phase voltage is the cell where both the row and column flags are different. This requires an XOR function between both adjacent row flags and column flags.

$$AN = (row_i \text{ XOR } row_{i+1}) \text{ AND } (col_j \text{ XOR } col_{j+1})$$

In addition, the Hi (cell fully on) and Lo (cell fully off) states need to be redefined to ensure that they are mutually exclusive to the AN state:

$$Hi = (row_i \text{ AND } row_{i+1}) \text{ OR } (row_i \text{ XOR } row_{i+1})col_j$$

$$Lo = NOT((row_i \text{ AND } row_{i+1}) \text{ OR } (row_i \text{ XOR } row_{i+1})col_{j+1})$$

With this logic the switches are turned on row by row so at any one time there is a row that is partially on, $(row_i \text{ XOR } row_{i+1})$. All other rows are either fully on or off determined by $(row_i \text{ AND } row_{i+1})$. For the row that is partially on, cells in columns that are flagged as on (col_j) will be fully on. The cell having a different adjacent column flags will receive the analogue phase input (V_{phase}), and will be turned partially on. While all other cells in this row will be fully off. Figure 5.4 shows the necessary decoding logic.

5.2.2 The interpolation current cell

In a segmented system a current cell design is required which can accept three inputs consisting of a high and low digital signal or the analogue phase signal (V_{phase}) which interpolates between the

two digital values. Three designs of a current cell were considered. In the first design two current switches are used for each segment. One switch takes the digital signal and the other the analogue phase signal. The second design is a modification where the reference transistor is shared to create a three branch current switch. This enables the current source to also be shared with the aim to reduce power consumption. In the third design a three state analogue voltage switch is used to create the input to an otherwise conventional current cell design.

Two switch cell design




Figure 5.5: The two switch cell design. All four current switch transistors are the same size and the two current source transistors have the same number of unit transistors to ensure the current is equal. The phase voltage is drawn to represent a phase word of 1 for clarity here. For most phase words the phase voltage will be more discontinuous.

In the first cell design two current sources are used, Figure 5.5, one to perform the conventional digital step transition and the other to interpolate between the two levels of the digital step. The advantage of this system is that a nonlinear DAC with a conventional thermometer decoder can be used to generate the digital step transition. The analogue interpolation can be introduced by a separate current source and switch, addressed by the modified thermometer decoder. In effect this method combines the two DACs of Figure 5.1 into a single nonlinear array. The current switches share the same current and voltage reference to minimise mismatch. The output is combined to form I_{out} . The major drawback of this type of current source is that it consumes twice the necessary power and area due to its dual current source design. Therefore the maximum efficiency achievable is 50% and more efficient current cell designs were considered.




Figure 5.6: The three-branch current switch design. The response to a phase word of one is shown for clarity. Most phase words would result in a more discontinuous response.

Three branch cell design

To reduce the size and power dissipation of the current cell a modified design with a three-branch current switch was considered, Figure 5.6. This design shares the reference transistor and the current source between the conventional and interpolation circuits. One branch takes the analogue phase input, another the digital input and the third is held at a reference. A switch is inserted in the current path of the analogue branch to prevent the digital and analogue controls competing (Figure 5.6). The switch requires a large voltage swing to completely turn off the analogue branch. However this voltage is coupled onto the common source by the parasitic capacitance, V_{gs} of the switch devices, causing glitches at the output. The three branch cell suffers from combining devices in parallel. Voltage swings associated with current switches need to be small to reduce glitches, however unless currents are switched completely off, the two current paths will compete and cause errors in the output.

A three state analogue switch

A design was sought to merge the analogue phase and digital on/off voltages to create an input which was as close as possible to the ideal inputs used in the MATLAB simulations and the design requirement. The digital and analogue voltages can be combined on a single differential current switch by using an analogue switch based on transmission gates (TGs) (Figure 5.7). The transmission

gate is favoured over single NMOS or PMOS gates for its low on resistance to signals near half the supply voltage [50]. TGs are also less susceptible to clock feedthrough as the PMOS device will absorb the charge that the NMOS device repels as it goes into its cut-off region, (see section 1.6.2). Three TGs are used to select between the digital voltages (V_{high} and V_{low} and the analogue phase signal (V_{phase})). The maximum of V_{phase} is equal to V_{high} and the minimum of V_{phase} is equal to V_{low} . This ensures that there are minimum glitches due to voltage differences between transitions of the gate input. The analogue switch was implemented using minimum size transmission gates to reduce clock feedthrough [44]. The digital control signals Hi, An, and Lo are provided directly from the modified thermometer decoder (Figure 5.4) and are synchronized by using latches.




Figure 5.7: A conventional current cell with a three state analogue switch. The sequence shown in figure is: Initially Lo is enabled, An and Hi disabled. An is enabled to allow analogue to pass through, then Hi is enabled to hold the output high. Only one input is switched to the current cell at any one time. The response of the cell when the phase word is one is shown for clarity, most phase words result in a more discontinuous output.

A drawback of this circuit is that when switching between states at low frequency, all three transmission gates may be temporarily off and hence the voltage on the current switch will start to fall as charge leaks across the gate capacitance. This charge leakage effect is reduced by using static latches on the thermometer decoded digital signals to ensure that they change synchronously.

It is important that the analogue switch is placed close to the current cell so that the voltage line between the switch and the gate of the current switch is short and the load on the analogue voltage switch is minimised. This reduces the settling time by reducing the load. The transmission gates used in the analogue voltage switch are of minimum size and are not connected to the power rails and therefore have a negligible effect on static power consumption. By introducing three TGs the current cell has not appreciably increased in size and power dissipation. As the minimum number of

analogue transistors are used, mismatch has not been increased and the number of current cells may be compared to other nonlinear DAC methods. This current cell is the preferred design that is used in the following circuit simulations.

5.2.3 Device sizes

The dimensions of the transistors are determined by the allowable variations from work in Chapter 3. For a variation in I_{bias} of less than 0.4%, a device area of $35\mu^2$ is required at a gate voltage of 1V. Devices that are long and thin exhibit less mismatch hence unit devices of $W/L = 1\mu/35\mu$ were used in each current cell. The number of unit devices in each current cell are the truncated I_{bias} parameters given in Table 3.3.

In the switch devices the V_t mismatch dominates the effect of β mismatch and hence the requirement for V_t uniformity determines the area of the switch devices. Using Equation 3.1 an area of $6.25\mu^2$ is required to reduce V_t mismatch to 1% of the minimum input voltage swing. The voltage swing is chosen to be 200mV, as discussed in section 4.2.1. The sizes of all the devices are given in Table 5.1.

Cell	$W/L \mu m$	Cell	$W/L \mu m$	Cell	$W/L \mu m$	Cell	$W/L \mu m$
1 and 64	1/4.5	9 and 56	2.85/1.5	17 and 48	5.1/1.5	25 and 40	4.3/0.95
2 and 63	1.2/3.5	10 and 55	3/1.5	18 and 47	5.1/1.5	26 and 39	4.4/0.95
3 and 62	1.8/2.5	11 and 54	3.45/1.5	19 and 46	5.55/1.5	27 and 38	4.4/0.95
4 and 61	2.4/2.5	12 and 53	3.6/1.5	20 and 45	5.55/1.5	28 and 37	4.4/0.95
5 and 60	2/1.95	13 and 52	4.05/1.5	21 and 44	5.7/1.5	29 and 36	4.6/0.95
6 and 59	2.4/1.95	14 and 51	4.2/1.5	22 and 43	4/0.95	30 and 35	4.5/0.95
7 and 58	3/1.95	15 and 50	4.35/1.5	23 and 42	4.1/0.95	31 and 34	4.6/0.95
8 and 57	3.2/1.95	16 and 49	4.8/1.5	24 and 41	4.2/0.95	32 and 33	4.5/0.95

Table 5.1: The actual dimensions of the current switches in the 64 cell segmented system. The values change between current cells to ensure the ratio I_{bias}/β ratio is constant as a constant V_{phase} is used and the technique is based on Equation 2.3. Again it proved difficult to keep all the devices the same area and for the layout it was more practical to base all devices on six different lengths.

The active area of the 64 cell segmented system was $0.4mm^2$ over 500 times larger than the compact systems due to the overhead of the thermometer decoding logic and latches. To put this in perspective, this system is of interest as a significant improvement over state of the art frequency synthesizers that typically require in the order of $1mm^2$ active area [20, 60] where a 16 bit accumulator usually consumes 20% of the active area [21]. Thus the active area of the a DDFS system based on 64 cells would be approximately $0.6mm^2$.

5.3 Low frequency circuit simulations

In Chapter 3 MATLAB was used to investigate truncation of the parameters of the nonlinear interpolation model based on Equation 2.3 to determine the number of unit current cells in each current cell. The truncation degraded the SFDR to $-72dBc$. An array of 64 current cells was simulated in Cadence at low frequency with the truncated parameters of Table 3.3 and Table 5.1. To ensure any distortion was solely caused by non-ideal transistor behaviour in the current cells, 64 ideal, piecewise linear, input voltage sources were used to create the input voltage in these initial circuit simulations. Each input voltage swing was chosen to be $200mV$ as a tradeoff for phase accuracy and limited voltage supply. All current cells shared the same reference voltage of $2V$. This simulation produced a result of SFDR, $-72dBc$, with the 5^{th} harmonic the dominant spur at a bias voltage of $1V$.

5.3.1 Comparison with MATLAB simulations

The cells of the 64 cell system can be linearised due to the thermometer decoding architecture. The tradeoff of performance and efficiency shown in the MATLAB simulations in section 2.3.2 was tested by increasing the bias current with the input voltage range fixed. Unlike the MATLAB simulations of Chapter 2, there is no longer a smooth tradeoff of performance and efficiency. This is caused by the truncation of the parameters in Chapter 3 that determined the number of unit current sources and equalised the area of the current switches. The effect on the trends of the dominant harmonics with bias current can be seen in MATLAB simulation after truncation, Figure 5.8.

The decrease of the 127^{th} harmonic with increased bias current corresponds to the efficiency trade-off seen in the MATLAB simulations before truncation, Figure 2.10. Increased bias current causes linearisation and hence less efficiency. When the bias voltage is low ($0.5V$) the efficiency is high (95%), when the bias current is increased the efficiency decreases, for example at $0.8V$ the efficiency has dropped to 50%. When the bias currents are increased further the approximation error from truncation becomes dominant and the 3^{rd} and 5^{th} harmonics increase in power. The optimum SFDR is $-75dBc$ in MATLAB simulations, this corresponds to an efficiency of 35%. In the Cadence circuit simulations the dominant harmonics follow a similar trend and the best case SFDR is $-72.4dBc$ at $V_{bias} = 1$, Figure 5.9. The power efficiency of the system was 95% at a bias voltage of 0.5 , SFDR $-53dBc$, 75% at a bias voltage of $0.7V$, SFDR $-60dBc$, and 48% at a bias voltage of $1V$ where the SFDR was $-72.4dBc$. Hence there is broad agreement between the Cadence and MATLAB simulations. However the effect of truncation of the parameters restricts the range of a trade-off




Figure 5.8: MATLAB simulations of the system based on an array of 64 current cells after truncation of the parameters. Unfortunately the approximation error introduced due to truncation introduces 3rd and 5th distortion at bias voltages greater than 0.8V, there is no longer a smooth tradeoff of performance and total bias current (or efficiency). The tradeoff still exists when the bias voltage is low (0.5V) where the efficiency is high (95%). The dominant harmonic is the same (127th) seen in the MATLAB simulations before truncation. When the bias current is increased the efficiency decreases, for example at 0.8V the efficiency has dropped to 50%. The second harmonic is not shown as it was below the noise floor of $-120dBc$.




Figure 5.9: Circuit simulations of the system based on an array of 64 current cells in Cadence. The power efficiency of the system was 95% at a bias voltage of 0.5V, SFDR $-53dBc$, 75% at a bias voltage of 0.7V, SFDR $-60dBc$, and 48% at a bias voltage of 1V where the SFDR was $-72.4dBc$.

between SFDR performance and power efficiency for a particular design. Therefore the prototype circuit was designed for operation at the optimum point after truncation, at a bias voltage of 1V, SFDR $-72.4dBc$ and efficiency 48%.

5.3.2 Confirmation of the effect of variations

The control of parameter variations by appropriate sizing of transistors was confirmed in Cadence with variation parameters based on the Pelgrom model (see Section 3.2), which have been measured by the manufacturer from processing runs. 100 Monte Carlo simulations were run with a current source bias voltage of 1V. The results indicate a yield of 95% for systems with SFDR $-70dBc$ and 100% for $-65dBc$ SFDR. The 127th harmonic is the most powerful as variations in the threshold voltage will cause mismatch in the interpolated output current from each cell, introducing the 128 steps per period that is related to the number of current cells. This is similar to the sampling effect seen in the compact systems in Chapter 2. In agreement with MATLAB, the 5th harmonic is the next most powerful. This was caused by accumulated errors in the truncated parameter set.

5.4 High frequency and mixed signal circuit simulation




Figure 5.10: The 64 cell system detailing the use of all 13 phase bits required to avoid phase truncation.

The low frequency simulation results using an ideal phase DAC are promising and match the performance achieved in MATLAB. To examine the influence of transient effects and the effect of loading on the phase DAC the full system was simulated using the mixed signal simulator. The complete architecture to implement the 64 cell system, including all peripheral circuits, is illustrated in Figure 5.10 where the three state analogue switch allows the digital phase address and analogue phase voltage to be used with a single nonlinear DAC. The 7 MSBs are used to select the nonlinear DAC cells that contribute to the output. The next 6 phase accumulator bits are used in conjunction with a linear phase DAC to construct an analogue voltage that drives a particular current cell of the nonlinear DAC, (see section 5.2).

The 6-bit linear phase DAC is similar to those considered in section 4.4.2 and is required to bring the total number of phase bits up to 13 to ensure that the system is not limited by phase truncation. The phase DAC is required to drive only one current cell at a time hence the gate capacitance is less than in the compact systems, however due to the larger circuit size more track capacitance is expected. The areas of the transistors in the switch devices were designed to be approximately equal to ensure that they represent an approximately equal load to the phase DAC. In the process used C_{ox} is $0.5 fF/\mu^2$ hence the capacitance of one switch and the parasitic track impedance was estimated at less than $0.1 pF$. A binary encoded 6-bit phase DAC was used in the following simulations with a bias current of $200 \mu A$ as this current level is found to be sufficient to drive a load of $0.2 pF$ in section 4.4.2. The settling time of the analogue interpolation voltage signal was examined and found to limit the bandwidth to less than $500 MHz$ in this design. The settling time is shown in Figure 5.11 and is approximately $2.5 ns$ which is 25% of the clock frequency at $100 MHz$.

Mixed signal simulation results of the complete systems with an estimated load of $0.1 pF$ on the phase DAC are in good agreement with the low frequency and MATLAB predictions. For simplicity the accumulator was initially programmed with a phase word of 1, so that the system generated a synthesised frequency of $f_{clk}/8192$. With a clock frequency of $1 MHz$, a SFDR of $-72.2 dBc$ was obtained.

The frequency spectrum of the output showed a dominant spur at the 5th harmonic in agreement with MATLAB and the earlier Cadence simulations. As expected the SFDR of the system decreased at higher clock frequencies when glitches become the dominant source of errors. The results in Table 5.2 suggest that the design is robust for f_{clk} up to $100 MHz$. A prototype system was therefore designed to allow the system to be characterised over a wider range of input words. Although there is a slight degradation in SFDR, the SFDR is better than existing systems (Table 1.3) by approximately $10 dB$




Figure 5.11: Settling time of the phase DAC in the 64 cell segmented system. (a) $f_{clk} = 5MHz$, (b) $f_{clk} = 50MHz$, and (c) $f_{clk} = 500MHz$.

at these higher frequencies. The true influence of glitches on performance can only be known after a prototype chip has been fabricated and measured.

f_{clk}	circuit SFDR
1MHz	- 72.2 dBc
50MHz	- 70.3dBc
100MHz	- 70.1dBc
500MHz	- 69.2dBc
1GHz	- 68.7dBc

Table 5.2: Performance of the 64 cell system as f_{clk} is increased. The accumulator is loaded with a phase word of 1 for these simulations. The output frequency is $f_{clk}/8192$.

In a similar way to the compact systems the clock alias spur, $f_{clk} - f_{out}$ adds to the dominant harmonics, degrading the SFDR when the phase word is increased. This mechanism is the clock feedthrough introduced in section 1.6.2. For example generating an output of $f_{out} = 100kHz$ with $f_{clk} = 5MHz$ and a phase word of 333 was limited by a 7th harmonic spur at $-69dBc$ below the fundamental. While an SFDR of $-72.5dBc$ was maintained for $f_{out} = 100kHz$ by increasing the clock frequency to $f_{clk} = 50MHz$ and decreasing the phase word to 33. These simulations took 5 days to complete on a Ultra 60 workstation with 2GB of memory, hence a prototype was built to examine the effects more easily.

5.5 Power dissipation

Component	Current	Voltage Supply	Power Dissipation
	(μA at $100MHz$)	(V)	($\mu W/MHz$)
Nonlinear DAC	200	3.3	6.6
Linear Phase DAC	200	2.6	5.2
13-bit XOR Complementor	6.5	3.3	0.2
13-bit Accumulator	65	3.3	2.1
Thermometer Decoder	100	3.3	3.3
Total			17.4

Table 5.3: Power dissipation breakdown of the 64 cell system. The total power dissipation is more than four times less than state of the art systems (see Table 1.3).

The total bias current for best spectral purity in the 64 cell system was $200\mu A$ to operate the system at $100MHz$. The large current for the phase DAC is acceptable as it can operate from a lower supply voltage. Low power accumulators that use paralleling and pipelining (see Appendix A) are able to operate with less than $0.05\mu W/MHz$ per bit. Hence the expected power dissipation is $0.018mW/MHz$ (Table 5.3), less than a quarter of state of the art systems, Table 1.3.

5.6 Summary

Circuit simulations of the segmented systems are encouraging. A small binary encoded phase DAC can be used to create the analogue input phase voltage required for interpolation. In the case of the segmented system this voltage can be created by using phase bits that are discarded in conventional DDFS systems to improve the spectral purity of the output. Simulations suggest that transient effects may be controlled by appropriate use of thermometer decoding, latches and reduced voltage swings. A modified thermometer decoder has been presented to ensure that the correct cell is interpolating at any given time. Latches have been used in each current cell to ensure synchronised switching. While reduced voltage swings are used to reduce the magnitude of glitches.

Three possible current cell designs were reviewed. The most straightforward design consists of an array of conventional current cells and a separate array of interpolation current cells. This arrangement is inefficient as only one current cell is interpolating at any given time. Hence a design was sought to combine the two functions into one current cell. The first attempt shared a single current source between three current switch transistors. This design was unsuccessful as the source voltage of the switch transistors was required to be gated depending on the mode of the current cell.

The second attempt was more efficient and was based on the conventional current cell. A three state analogue switch based on transmission gates is used on the gate of the input transistor to pass the phase voltage when required and otherwise to pass the conventional 'on' and 'off' values. This technique ensures that small voltage changes are made at the gate of one of the switches rather than larger voltage changes in the bias current path. This method is shown to reduce the inevitable output glitch.

Analogue and mixed signal simulations suggest that the system is able to meet the limit of state of the art, low power CMOS DACs, $-70dBc$ SFDR at clock frequencies up to $100MHz$. The power dissipation is estimated to be $18\mu W/MHz$, less than a quarter of state of the art, low power, ROM-less DDFS systems. The yield of the system with spectral purity of $-70dBc$ is 100% and the power efficiency is 50%. This gives a $10dB$ margin of error in SFDR between the requirements of the UMTS and Bluetooth communication standards. Due to the truncation of parameters required to implement unit devices the tradeoff of power efficiency for SFDR for bias voltages above $1V$ is lost, however at low bias voltages the tradeoff still exists. To meet the Bluetooth low output power class SFDR requirement of $-50dBc$ the efficiency of the system can be improved to 95% at a bias voltage of 0.5 . Simulations are too time consuming to simulate the effect of all possible input conditions and transitions. It is expected that SFDR performance will degrade as the output frequency is increased as the effect of the clock alias spur and its harmonics interfere constructively in the band of interest. The most efficient way to test this effect is to build a prototype system that can be tested with greater ease and in real time.

Chapter 6

Layout and design of the prototype

6.1 Introduction

The simulation results of the compact four and six cell, and segmented 64 cell DDFS systems are promising and demonstrate the performance improvements possible over existing DDFS systems. The simulations predict performances that meet the requirements of the applications introduced in Chapter 1. A prototype system is required to demonstrate the circuit functionality and to measure the performance over a wider range of operating conditions than was possible in simulation.

A two dimensional physical description of the prototype called a circuit layout must be designed for prototype manufacture. A careful layout is required for the analogue part of the circuit to ensure that undesirable effects such as mismatch and crosstalk are minimised. The proposed design contains a number of novel circuits and techniques, and therefore the layout was not straightforward. The exact details of the layout will be described in this chapter.

The prototype integrated circuit design has certain constraints. Due to the number of circuits on the prototype this integrated circuit design is limited by the number of pins available. Pins that do not carry a critical current may be shared between circuits and bi-directional pads may be used to share digital pins. Of the circuits that are included only the nodes that are brought out to a pin may be measured after manufacture. Hence it will be difficult to discover any flaws in the circuits that are not directly connected to the output pins. For this reason test structures will also be included on the chip. Due to the ability to share pins these test circuits will not require a substantial number of pins.

This chapter introduces layout techniques of mixed signal circuits such as reduction of bias voltage mismatch, device orientation and placement, and guarding between digital and analogue circuits. Details of the layout of the circuits that were included on the prototype chip are given. These are based on the device sizes determined in Chapters 4 and 5 and Tables 4.1 and 5.1. The circuits to be included in the prototype are determined by a test plan that ensures that a sufficient combination of circuits are included.

6.2 Layout techniques

The layout was designed in a $0.35\mu m$, N-well, 3-metal, CMOS process for manufacture by Austria Microsystems. In an N-well process NMOS transistors are formed in the P-substrate. The PMOS transistors are created inside an N-well and the additional process step can introduce more variations between PMOS transistors [49] . Hence NMOS transistors were used for the sensitive devices such as the current source and switches. The process has two polysilicon layers, one for forming the gate of the active devices, (POLY1) and another for forming large impedance resistors in a small area or capacitors between the polysilicon layers, (POLY2). There are three metal layers to enable tight signal routing, the first two metal layers (M1, M2) have the same sheet resistance, while the third metal layer, (M3) has half the sheet resistance of the other two layers.




Figure 6.1: Cross section of the layers available in the process showing the formation of a PMOS and NMOS device.

6.2.1 Process tolerances

Absolute device dimensions can not be determined accurately due to the finite resolution of the process. The resolution of the width and length of each transistor is limited by effects in the manufacture such as over-etching of the polysilicon used to form the gate, lateral diffusion when forming

the channel and increased substrate doping at the edges of the channel. These errors cause differences between the drawn width and length (W and L) and the effective or actual width and length (W_{eff} and L_{eff}), [44]:

$$W_{eff} = W + \Delta W \quad (6.1)$$

$$L_{eff} = L + \Delta L \quad (6.2)$$

The smaller the devices chosen, the more impact ΔW and ΔL will have. In the process chosen $\Delta W = 0.027\mu m$ and $\Delta L = 0.008\mu m$.

6.2.2 Bias voltage drop

Bias nodes of matched devices are often assumed to be at the same voltage. In reality the wire that connects these nodes has a non-zero impedance. The current flowing through the bias line will create a voltage drop and hence the bias voltage at remote devices will be different to that intended. In an array of current sources the power supply nodes, ground, V_{dd} , and the output nodes, are the most affected by voltage drops as they carry the majority of the current. The bias voltages at the gate experience little voltage drop as MOS devices consume little current through the gate. To reduce the parasitic impedance the width of the sensitive bias lines is increased for the higher current lines and their length is minimised. The least resistive metal layer, (M3) is reserved for the longer wires, including the ground path, and the width of these paths is chosen to ensure that bias voltage drop is less than the tolerable variation in V_{ref} or V_{phase} .

6.2.3 Device orientation and placement

One cause of local variations is tilt-induced variations in carrier mobility. To reduce the impact of this effect all critical transistors are arranged in the same orientation so that any direction dependent variation effects all transistors. Mismatch can also be caused by temperature and oxide thickness gradients across the wafer and variations in the surrounding silicon. To reduce gradient effects all critical devices are positioned as close as possible to each other. Devices can also be arranged in a common centroid pattern to share or average out variations. However this will complicate the routing of switching signals and is not used for this design.

6.2.4 Dummy devices

Devices at the edges of large arrays are more affected by variations [49], therefore dummy devices are placed at the edges of the array to ensure that transistors at the edges of the array are effected by directional variations in the same way as transistors in the middle of the array. Dummy devices also reduce mismatch caused by etch rate variations by artificially increasing the size of a device array, thereby pushing any edge effects to the dummy perimeter.

6.2.5 Guard rings




Figure 6.2: The analogue circuit is protected from the digital substrate by a guard ring that encourages current to flow to grounds rather than through the substrate. The three different ground nodes, DGND, GRD-GROUND and AGND are connected off the chip to prevent the creation of a ground loop.

In mixed signal designs, digital noise coupled through the substrate induces noise in the analogue circuits. The analogue and digital substrate regions are isolated from each other using guard rings. The separate substrate areas are connected to different pads and supplied by independent power supplies to improve isolation. A large proportion of the substrate current flows in the surface of the substrate as this is the path of the least impedance. In an N-well process a region of N-well will provide a discontinuity in the P-substrate surface. This can be used to increase the impedance between the analogue and digital substrate. On either side of the N-well guard, P+ regions are used to reduce any currents flowing in the surface by reducing the impedance to ground, thereby encouraging the unwanted currents to flow to ground rather than through the substrate. The P+ and N-well are connected to a separate ground pin to complete the guard ring and ensure the digital and analogue grounds are isolated within the chip [49], Figure 6.2.

6.2.6 Mixed signal design

A high speed, mixed signal design must take into account the effect of all impedances, including inductance, capacitance and resistance of every connection in the system. This includes the metal track connecting devices, the bonding pads, bonding wire, the package pins and PCB tracks and connections. The inductive component of track impedance is less than the resistive component at frequencies of less than 1GHz and therefore is usually ignored and the resistance taken as the worst case impedance. Also, measurement data of inductance was not available from the manufacturer for the process or package used here therefore inductance was not considered in this design.

Capacitance will limit the frequency performance considerably as there is parasitic capacitance between all tracks and between each track and the substrate. Combined with the resistance of the tracks, the track capacitance forms a low pass filter that slows down the edges of digital signals, in particular the clock signal (clock skew). Due to the reliance on synchronous switching signals in the current cells, the clock edges and rise and fall times of the digital control signals are important. Any slew or decrease in rise/fall time or delay will introduce mismatch in the timing of the control signals, causing glitches in the response of the current cells, degrading the output of the DDFS system.

The coupling capacitance of the metal tracks to substrate is proportional to their surface area, hence minimum width tracks were used. However the resistance of the tracks is proportional to their aspect ratio, that is long thin tracks have a higher resistance than long and wide tracks or short and thin tracks. Therefore tracks that are as short as possible are desired.

The large numbers of unit current cells required to reduce variations between the current sources set the size of each current cell and hence the 64 cell array. This results in very long thin metal tracks with considerable resistance. Hence a compromise was made to introduce additional latches and buffers into the design to reduce the impact of track impedance. For each row of the eight current cells, the digital bus contains two row select lines, eight column select lines and the clock. The entire digital bus was buffered at each row using inverters. Additionally, the digital control logic at each current cell was latched to ensure synchronous operation of the three state analogue switch.




Figure 6.3: Layout of the four cell system. Signal lines are across the top, switch transistors are in the middle (two different sizes) and the 22 unit current sources are along the bottom. Dummy unit cells are at the edge of the array to prevent edge effects. To minimise bias voltage drop the whole circuit was covered in a sheet of metal 3.




Figure 6.4: Layout of the six cell system. The top half has a similar placement to the four cell system, see Figure 6.3. The entire top half of the circuit is mirrored to ensure that the unit current sources are placed close together to improve matching. As in the four cell system the whole circuit was covered in a sheet of metal 3 to minimise bias voltage drop.

6.3 Layout of four and six cell systems

The physical layout of the circuits was drawn using the Cadence layout tool, Virtuoso. The four cell system requires 20 unit devices and the six cell system, 60 unit devices (Table 4.1). To reduce complexity in the four cell system 22 unit devices (including two dummy devices) were placed in a row with the gates connected together and the sources connected to ground (Figure 6.3). In the six cell system, unit devices were implemented in two rows of 32 unit devices (including two dummy devices) to reduce the length of wires (Figure 6.4). The parameters in Table 4.1 were used to size the switch devices and the transistors were placed as close as allowed by the design rules to minimise mismatch. The drain connection of a unit device at either end of the row was connected to ground to form a dummy device and the source node of the switch transistors was connected to the required number of unit devices to form each current source.

Minimum width metal wires were used for the voltage bias lines as little current is drawn through the transistor gates and hence there is minimum voltage drop across the wires. However the power supply (V_{dd} and G_{nd}) and output current I_{out} buses were drawn in $2\mu m$ wide metal to reduce the resistance and hence the voltage drop. The resistance of these lines is approximately 3Ω corresponding to a supply voltage drop of less than $0.6mV$ when $I_{out} = 200\mu A$. Similar wires connect V_{dd} , G_{nd} and I_{out} to their respective pads to ensure minimal voltage drop. At least one substrate connection was placed next to each device to ensure a low resistance to the substrate to reduce the body effect. All signal routing was done in the first two layers of metal, reserving the third metal layer for covering the entire circuit to reduce the resistance between the ground connections seen by each device. A single ground wire was connected from the circuit to the ground pad. This ensures that the maximum difference in ground resistance experienced by all devices is less than 0.4Ω . The entire area, including metal wires, of the four cell circuit was 80μ by 30μ and the six cell circuit, 120μ by 70μ .

6.4 Layout of segmented current cells

To minimise the complexity of the segmented 64 cell system layout, a generic cell was created containing common digital logic, guard rings, the maximum number of unit current sources, the largest switch transistors and a digital and analogue bus for the common input and output connections. A single generic cell is shown in Figure 6.5. The cell includes the three-state analogue switch introduced in section 5.2.2, Figure 5.7 which allows the system to be implemented with the minimum




Figure 6.5: Layout of the generic current cell of the 64 cell system.

number of current sources.

At the top of the generic cell are the digital bus lines and the within-cell thermometer decoding logic (see Figure 5.4). The digital part of the cell is isolated using a guard ring (section 6.2.5). The guard ring of each cell interconnects to form rows of interchanging digital and analogue substrate. Typical CMOS DACs use a simpler design where the digital switch matrix is a separate array to the array of current sources. This approach was not in this design as it would increase the wire length and hence capacitive load of the phase DAC.

The switch devices are in the centre with a large area of unused silicon as the number of unit current sources sets the cell width. The switch device in each cell is different and the parameters are taken from Table 5.1. The lower half of the cell is occupied by the unit current sources and an analogue signal bus runs along the bottom of the cell. Another guard ring is beneath the analogue signal bus to separate this cell from the digital substrate of other nearby cells.

There are 52 unit current source devices placed in each copy of the generic cell. A different number are connected in each current cell to form the current source using the parameters in Table 3.3. The remaining unit current sources are connected to ground to form dummy devices. This generic method was used to simplify the complexity of the array of 64 generic current cells, see Figure 6.6. A more compact design is possible by using only the required number of current sources in each cell,

however the layout of such a system would be more complex.




Figure 6.6: Description of the layout of the 64 cell system. 64 copies of the generic current cell (Figure 6.5) were layed out in an 8 by 8 array. The number of connected unit current source devices and the geometry of the current switch devices of each cell was then modified corresponding to the parameters in Tables 3.3 and 5.1. The order of the current cells was chosen to simplify the thermometer decoding logic.

The 64 cell array contains two sets of identical 32 current cells. The first set increases in order from cell 1 through to cell 32 across the top four rows and the second set decreases in order from cell 32 through to cell 1 in the bottom four rows. Cell 1 has the lowest current source weighting and smallest switch transistor sizes (β), while cell 32 has the highest weighting and largest β .

The final chip layout was verified against the simulated schematic using the Layout Versus Schematic (LVS) software tool. The size of the cell including wires was 150μ by 100μ . The 64 current cells in an 8 by 8 array was 1200μ by 800μ . A substrate contact was placed next to all devices and a sheet of metal 3 was placed over the analogue elements and connected to ground to ensure a low resistance ground connection. The largest mismatch in ground resistance seen by any current cell was 3Ω which corresponds to a voltage drop of $0.5mV$ or a variation of 0.05% in V_{bias} . This is within the 0.4% requirement established in Chapter 3.

6.5 Design and test plan

The test plan of the prototype chip was specified considering the available chip area, pads and equipment. The circuits of interest were the four and six cell compact systems and the 64 cell segmented system. The test plan was composed to provide methods of determining faults within the system. Therefore a test cell and several different architectures were fabricated (Table 6.1).

The decoder test cell contains the local thermometer decoding logic. (Figure 5.4) and a three state analogue switch, (Figure 5.7). This circuit is within each cell of the 64 cell segmented system. If it were to fail the entire system would not work and the fault would be difficult to discover. Hence a separate decoder test cell was included to isolate faults in the local thermometer decoder or analogue switch that would be difficult to discover by only examining the full 64 cell system.

Prototype circuits of the compact four and six cell systems were included on the chip without a phase DAC or accumulator due to limited silicon area. The analogue phase voltage and accumulator function will be provided by external equipment.

Two copies of the 64 cell system are included on the prototype. The first copy is a conventional nonlinear ROM-less DDFS, *conv* circuit where interpolation is not used and hence the SFDR is phase limited to 7 bits. This circuit was included to allow confirmation of the correct operation of the thermometer decoding logic and current cell sizing. The second copy of the 64 cell system, *int*, provides the interpolation required to improve the phase limitation to 13bits by using a phase DAC. When driven by the phase DAC and the digital low signal, the *int* circuit provides the interpolation voltages between the quantised output of the *conv* circuit. In this way the two nonlinear DACs are used to form the interpolation system (Figure 5.1). The *int* system may also be operated without the *conv* system by applying both the digital high and digital low signals, to form a full interpolation DDFS using the minimum number of current cells. This method implements the system described in Figure 5.10.

6.6 Layout verification

The Cadence verification tool, DIVA LVS was used to compare the layout with the schematics used in simulation. The LVS tool generates two netlists containing all the device sizes and electrical connections from both the layout and schematic. The schematics were entered as a hierarchical design,

Experiment	Circuit	Equipment	Pads
Check	Test decoder	Digital source	4x digital input
decoding		Clock source	1xclock
logic and		Oscilloscope	3x Analogue output
analogue switch		Bias supply	4x Analogue bias/reference
Compact	Four and	Function generator	1x Analogue input
systems	six cells	Oscilloscope/Spectrum analyser	2x Analogue output
		Bias supply	6x Analogue bias/reference
Segmented	64 cells <i>conv</i>	NCO	13x digital input
system		Clock source	1x clock
conventional DAC		Oscilloscope/Spectrum analyser	1x Analogue output
		Bias supply	4x Analogue bias/reference
Segmented	64 cells <i>int</i>	NCO	13x digital input
system	and phase DAC	Clock source	1x clock
conventional +		Oscilloscope/Spectrum analyser	1x Analogue output
interpolation DAC		Bias supply	5x Analogue bias/reference

Table 6.1: Prototype test plan. This test plan was composed to make efficient use of chip area, pads and equipment while also ensuring that faults could be determined without using a probing station. All circuits were manufactured on a single chip with 64 pads by sharing most of the digital signals.

where subcells represent the components of the DDFS systems such as the nonlinear DAC, phase DAC and accumulator. Each of these subcells use the power supply pins which were referenced as global or virtual power supply nodes (vdd! and gnd!). This means that the schematic subcells do not display a pin for these subcells as the simulator and verification tool interpret global pins to be common for all subcells. This method simplifies schematic entry and simulation, however it is up to the designer to ensure that the power supply connections are made in the layout. The design was verified at each stage of the hierarchy from device level (discrete transistors, resistors and capacitors) to chip level for each of the circuits listed in Table 6.1 and the pads. However at each of these stages the connections between the subcells were tested and the virtual power pins were in fact not verified as Cadence assumes that virtual power pins are always connected. There are two methods to ensure the design is correct. One is to flatten both the schematic and layout of the entire chip to device level. Another is to verify or to simulate the netlist extracted from the layout. The DIVA tool has reported problems with the first method and in this design thousands of errors were reported, most of which were in fact not errors. The second method was used to simulate each hierarchy up the chip level. One advantage of simulating the netlist extracted from the layout is that the parasitic impedances of the wire interconnections can be included in the simulations. The design kit automatically extracted the capacitance of the interconnect but not the resistance or inductance. Simulations of the extracted layout of the compact systems agreed with schematic simulations. However the large number of devices in the segmented system prevented the simulation of an entire output period as the simulation

was too lengthy. The extracted layout of each component of the segmented system was simulated instead and the interconnections between the components checked manually.

6.7 PCB and test bench design

A six layer mixed signal PCB was designed to test the prototype chip. The six layers provided isolation between the analogue and digital supplies and dual supply layers for the digital and analogue test chip sets. The digital chip set included line drivers to convert the 5V signals from the data acquisition kit to 3.3V digital signals, an AEL1200 crystal based oscillator for clock generation, and a HSP45106 numerically controlled oscillator chip to emulate the accumulator. The analogue test chip set included a LM317 voltage regulator to guarantee 3.3V supply and OPA655 current amplifiers for buffering the outputs. Decoupling capacitors ($10\mu F$ and $100\mu F$) were placed close to each of the chips to provide power supply filtering of all bias signals. For all analogue signals 50Ω transmission lines were used to reduce the attenuation of the analogue RF signals. Surface mount devices were used to reduce the length and hence the impedance of the chip pins. For the prototype chip a surface mount socket was used to minimise capacitance while also allowing the testing of the 10 samples manufactured.

6.8 Summary

A physical representation of the compact and segmented architectures has been drawn to form a layout for prototype fabrication using the Cadence Virtuoso tool. The novel analogue circuits were designed carefully to ensure that the sensitive parameters were not adversely effected by processing variations, mismatch and bias voltage drop.

In the segmented system a generic current cell was used to simplify the layout and ensure that the signal paths of the common signals were equal. Guard rings and latches were used to reduce the effect of noise and delay. These might have caused unwanted transients at high frequency. The chip area and number of pads were used efficiently by sharing circuits between test configurations. The areas of the four and six cell compact systems are very small, approximately a quarter of the size of a single pad. The use of the generic current cell in the 64 cell segmented system caused this circuit to be considerably larger, the entire 8 x 8 array was 1200μ by 800μ .

Test cells were included on the prototype chip to ensure that faults could be identified. The design and test plan is given in Table 6.1 in detail with the equipment, peripheral circuits and the number of pads required by each circuit. This integrated circuit is pad limited and hence it was important to identify which pads could be shared between designs. The final chip contained five different circuits in a chip with 64 pads. To enable testing of the chip a six layer printed circuit board (PCB) was designed to host the chip during testing.

Chapter 7

Measurement results

7.1 Introduction

The simulation results of the nonlinear interpolation DDFS systems are encouraging. The compact circuits which are based on nonlinear arrays of four and six current cells are able to meet the requirements of various instrumentation systems. The segmented system uses a nonlinear array of 64 current cells and a novel, three state analogue switch to exceed the performance of state of the art DDFS systems where the high frequency performance is limited by the performance of low power CMOS DACs.

All simulations are limited in predicting the performance of a circuit design, and hence it is important to measure a real circuit to confirm a working system. A prototype chip with various implementations of the novel DDFS systems has been designed with particular attention to minimising variations between the devices, delays and crosstalk between the analogue and digital parts of the system. The prototype chip was fabricated in the Austria Microsystems $0.35\mu m$ CSX process and packaged in a JLCC (J shaped lead chip carrier) 68 pin ceramic package.

This chapter describes the measurement method required to generate bias and reference voltages and reports the measurement results. Each circuit is first tested at low frequency to confirm that the bias and reference voltages have been set correctly, variations have been controlled sufficiently, and the layout is functionally correct. Then each of the ten manufactured chips are measured to give an indication of the yield. The yield is compared to that predicted in MATLAB and Cadence simulations to give an indication of the ability to reliably manufacture working circuits. Finally the

transient performance of each circuit is measured by increasing the clock frequency and the results are compared with the predictions of the circuit simulations.

7.2 Low frequency measurements of the compact systems

7.2.1 Measurement method

In all tests a data acquisition system (PowerDAQ PDIO-210 from United Electronic Industries) was connected to a PC running HPVee and used to generate all bias voltages and timing signals for the test circuits. At low frequencies a HP 4155B semiconductor parameter analyser was used to generate the analogue voltage signals and measure the output current. The acquired data was transferred to a PC over a GPIB bus. This provided a large increase in data storage over that provided by the limited internal memory of the HP 4155B and allowed the measurements to be automated. MATLAB was used to calculate the FFT of the acquired waveform ¹.

A source measure unit on the HP4155B parameter analyser was used to measure the output current of the four and six cell compact systems on all 10 manufactured prototype chips. The measurement time of the parameter analyser is limited by the need to measure current and source voltage alternatively for each sample. The fastest measurement time of the source measurement unit on the HP4155B is $640\mu s$ and 800 samples were taken per output cycle to measure an output frequency of $2Hz$.

As the compact systems on the prototype chip do not include a phase DAC, a voltage output unit on the HP4155B was used to emulate a phase DAC and accumulator with a phase word of 1. With the phase word set to one the analogue input is a triangular wave. The transfer function of each chip was recorded at various bias voltages and the result stored on a PC. MATLAB was used to calculate the spectrum of each response to compare with the Cadence simulations. The bias range of interest was from 0.5V to 1.1V as this is the region where the performance is best.

7.2.2 Low frequency measurement results

Figures 7.1 and 7.2 show similar trends in the four dominant spurs between the Cadence simulation and the chip measurements for a typical four cell system. Both simulation and measurement show

¹MATLAB uses the FFTW C library to calculate the FFT. This library allows arbitrary size transforms, where the size is not restricted to 2^N [69].




Figure 7.1: Cadence simulation of the four cell system design from section 4.3.




Figure 7.2: Chip results from a single array of four cells on a typical chip. The 2nd harmonic on this chip is an improvement from the typical mean suggested by the Cadence simulations. The 3rd, 7th and 9th harmonics show similar trends to the Cadence simulations with the optimum SFDR at $V_{bias} = 0.65V$. The 5th harmonic is not dominant and is not shown for clarity.

that at lower bias voltages the 7th harmonic dominates and at higher bias voltages the 3rd harmonic dominates. The small shift in the bias voltage is possibly due to process variation in the threshold voltage. This possibility is consistent with a shift in the same direction for the six cell system on the same, typical chip, Figures 7.3 and 7.4. The six cell system also shows similar trends in the harmonics to simulations. Other chips behaved in a similar manner to the typical chip as they have all been manufactured in the same process run.

The largest difference between the simulated trends and measurement is the second harmonic in both systems. The Cadence simulations suggested that the second harmonic is introduced when the quarter wave symmetry of the output is disrupted either by mismatch or the limited output resistance of the current sources. The low frequency measurement results are made without an output load as the HP4155B is able to measure current directly. Hence the second harmonic measured for the four and six cell systems on the typical chip is the distortion introduced by mismatch and is an improvement from the typical mean condition that was simulated in Cadence. In some of the other chips the second and the third, fifth harmonics increased in power. This may be caused by an adverse effect from variations in the parameters that increase these harmonics. The effect of the variability of the harmonics on the SFDR was investigated further by examining the yield of the ten manufactured chips.

7.2.3 Yield measurements of the compact systems

The trends of the dominant harmonics of all ten fabricated chips are shown in Figure 7.5. The four cell system had a measured yield of 80% for systems with SFDR of better than $-40dBc$ at a bias voltage of $0.65V$ for the ten manufactured chips. This is in good agreement with the variation analysis performed in MATLAB (Chapter 3) and the yield simulations in Cadence (section 4.3.2). For all 10 chips the trends in the dominant harmonic are similar and vary slightly at the optimum causing a variation in the optimum SFDR obtained for each system. The optimum bias voltage is between $0.6V$ and $0.7V$ and here the dominant harmonic is different for different chips due to parameter variations caused by mismatch, (Figure 7.5). For chip 1 and 8 the 2nd harmonic was dominant due to mismatch, in the other chips the 3rd and 7th harmonics are dominant indicating that mismatch is not limiting the SFDR in these chips. At a bias of greater than $0.8V$ the 3rd harmonic dominates and the variation and performance are reduced.

The measurement results of the six cell system in Figure 7.6 also showed similar trends with Cadence




Figure 7.3: Cadence simulation of the six cell system design from section 4.3.




Figure 7.4: Chip results from a single array of six cells on a typical chip. The 3rd, 5th, 7th and 9th harmonics show similar trends to the Cadence simulations with the optimum SFDR at $V_{bias}=0.9V$.




Figure 7.5: Measured variation for the four cell system.

for seven of the ten manufactured prototypes. In the bias voltage range where the optimum performance occurs, between $0.8V$ and $1.1V$, the SFDR of chips 1, 3 and 10 is bias independent while the SFDR of other chips reach an optimal SFDR. In this range chips 1, 3 and 10 are limited by the 2^{nd} harmonic spur while the other chips are limited by the 5^{th} and 7^{th} harmonic spurs. This is expected from simulation as the 2^{nd} harmonic spur is one of the most susceptible spurs to parameter variations. At a bias voltage of $1V$, all chips have a SFDR of better than $-45dBc$ and only two of the chips were within $1dB$ of the $-45dBc$ requirement (chip10 was $-45.3dBc$ and chip3 was $-45.6dBc$). At a bias voltage of $0.9V$ the yield of chips with SFDR better than $-50dBc$ is 70%. The system is more sensitive to variations at lower bias voltages as the output currents are smaller and measurement error is greater. Figure 7.6 shows that the SFDR of the six cell system varies by over $5dB$ when the bias voltage is $0.7V$ and the SFDR varies by less than $1dB$ when the bias voltage is $1.2V$.

7.3 Transient measurements of the compact systems

7.3.1 Measurement method

The HP 4155B is limited to low frequency measurements. At high frequencies an Agilent 54642D oscilloscope was used to capture complete periods of the output of each system. The oscilloscope MATH




Figure 7.6: Measured variation for the six cells system.

function was used to compute the FFT and the spectrum recorded from the oscilloscope screen. The effect of transients at higher frequencies was measured using an Agilent 33250A function generator (FG) to emulate the output of the accumulator and phase DAC. The measured SFDR is restricted to -48dBc due to the limited size of the FG RAM. The output current was converted into a voltage using a $20\text{k}\Omega$ resistor and the resulting voltage measured using the oscilloscope. The combination of the output pin capacitance (about 1pF) and load resistor form a low pass filter that reduces the power of signals and spurs above 8MHz .

The output voltage from the six cell system at the optimum bias voltage of 0.85V was 1.2V peak to peak, while the output voltage of the four cell system at the optimum bias voltage of 0.65V was only 200mV peak to peak. The resolution of the oscilloscope is 2mV , corresponding to a noise floor of $2\text{mV}/1200\text{mV} = -55\text{dBc}$ for the six cell system and $2\text{mV}/200\text{mV} = -40\text{dBc}$ for the four cell system. Hence the resolution of the oscilloscope was inadequate to measure the output of the four cell system directly. Therefore an OPA655 amplifier was used to amplify the output voltage so that the spectrum could be measured on a HP8590L spectrum analyser.

The suitability of the OPA6355 was judged by the nearest operating conditions reported in the OPA655 datasheet, a second harmonic distortion of -55dBc with a load resistance of 50Ω , $V_{out} = 2\text{V}_{pp}$ at a frequency of 5MHz and a gain of 5. The measurement setup used a gain of 2 where the

distortion should be less. This was confirmed by measuring the distortion of the OPA655 directly using a filtered $200mV$ sine wave from the FG. The results were a SFDR of $-57dBc$ for frequencies up to $5MHz$, limited by the second harmonic. At $10MHz$ the third harmonic distortion dominated at $-48dBc$. Hence the OPA655 is suitable for measuring the response of the compact systems as the SFDR is limited to $-48dBc$ by the FG RAM.

The FG was programmed with a waveform generated in MATLAB by simulating the accumulator overflow, phase truncation and half wave symmetry (see appendix C for the MATLAB code). The FG has a 256×65536 RAM to store the programmed waveform which limits the emulated phase DAC to 8bits in amplitude resolution. This limits the measurable SFDR to $-48dBc$ for either the four or six cell systems due to phase truncation. There are 65536 time points in the FG RAM and hence the 8-bit phase DAC can be scanned 256 times for each cycle of the FG clock. While the clock frequency of the FG is limited to $1MHz$, the effective clock rate of the phase DAC represented by the FG can be increased to $256MHz$. This effective clock frequency is actually limited to less than $100MHz$ by the sampling rate of the FG which is $200MSa/s$.




Figure 7.7: Measured output waveform of $4kHz$ from the six cell circuit on a typical chip.

An example of the output waveform from the six cell system, driven by the FG with a phase word of 1 and recorded by the oscilloscope is shown in Figure 7.7. For this signal the MATH function of the oscilloscope measured a SFDR of $-48dBc$. This result was confirmed by transferring the waveform to the PC and calculating the FFT in MATLAB, (Figure 7.8). This result is consistent with the phase

resolution of the 8-bit phase DAC that is emulated by the waveform stored in the RAM of the FG.




Figure 7.8: Frequency spectrum of the measured output waveform of $4kHz$ from the six cell circuit on a typical chip. The SFDR is $-48dBc$, limited by the size of the FG RAM.

7.3.2 Transient results

The compact systems were measured with the FG programmed with a ramp waveform ($PW=1$) and f_{clk} increased to examine the transient performance. This allowed the system to be measured with small voltage changes at the input. One difference between measurement and simulation is the phase DAC. A binary coded current mode phase DAC was used in simulation and was assumed to be on chip where the dominant load was the input resistance of the four and six cell systems. However due to chip area restrictions a phase DAC was not provided on chip for these systems. In the measurement setup the FG represents a voltage output phase DAC and is required to drive a 1m long coax cable, a PCB track and the input pin of the chip.

When the effective clock frequency was increased with a phase word of one the 2^{nd} harmonic increased in power, dominating the output at effective clock frequencies greater than 1MHz. At $f_{clk} = 1MHz$ the spectral purity of the four and six cell systems were equivalent to their DC performance, $-40dBc$ and $-48dBc$. At $f_{clk} = 10MHz$ the second harmonic had degraded the performance of the four cell system SFDR to $-38dBc$ and the six cell system to $-40dBc$. By $f_{clk} = 50MHz$ the SFDR of the four cell system was reduced to $-33dBc$ and the six cell system to $-36dBc$. These results

are worse than simulations.

The voltage output of the FG is the input to the device under test and hence was also measured at various effective clock frequencies. The results are shown in Figure 7.9. Distortion and hence a degraded SFDR was detected in simulation when the settling time of the FG signal becomes more than 25% of the clock period. The poor settling time will introduce nonlinearities into the input signal that will subsequently degrade the spectral purity of the output. This occurred at a clock frequency of greater than 50MHz in simulation, however parasitic impedances in the measurement system limit the bandwidth further. The parasitic impedances are due to capacitances and resistances in the chip socket, PCB tracks and cables. By comparing the measured and simulated settling times, the parasitic impedance in the measurement system is approximately a five fold increase over that used in the simulations. The settling time of the phase DAC is the same proportion of the clock period for $f_{clk} = 10\text{MHz}$ in measurement and $f_{clk} = 50\text{MHz}$ in simulation. With current equipment the system is frequency limited to $f_{clk} = 10\text{MHz}$ with an SFDR of -40dBc . However, with an integrated phase DAC -48dBc SFDR at $f_{clk} = 50\text{MHz}$ is expected.




Figure 7.9: The settling times of the phase DAC signal emulated by the FG. These were measured at the input to the four cell system using the oscilloscope. The settling time is approximately 25ns. Hence at (a) $f_{clk} = 1\text{MHz}$ it is 2.5% of the clock period, at (b) $f_{clk} = 10\text{MHz}$ it is 25% of the clock period (the equivalent of $f_{clk} = 50\text{MHz}$ in simulation) and at (c) $f_{clk} = 50\text{MHz}$ it is longer than the clock period so the input signal never reaches its final value, causing severe second harmonic distortion in the output signal.

7.4 Low frequency measurements of the segmented system

7.4.1 Measurement method

As with the measurements of the compact systems the PowerDAQ data acquisition system was used to set the bias and reference voltages in all measurements of the segmented system. This data acquisition system also provided the 13-bit digital accumulator output for the low frequency measurements. Agilent Vee software was programmed to generate the accumulator function. The HP4155B was used to source the output voltage and measure the output current which was transferred to the PC and the spectrum calculated in MATLAB. Again this measurement system was frequency limited to 2Hz due to the limited sampling rate of the source measurement unit of the HP4155B.

7.4.2 Thermometer decoder and three state analogue switch

A test cell containing a single instance of the thermometer decoding logic and the three state analogue switch (Figure 5.7) was included on the chip to test the operation of this novel circuit independently of the rest of the system. Correct operation of the test cell was confirmed by measurement at frequencies up to 1MHz by driving the analogue input with the FG and the digital lines with a clock generator and observing the response on the oscilloscope. The three modes 'on', 'off' and 'interpolate' were confirmed by changing the row and column address using switches on the printed circuit board.

7.4.3 Conventional thermometer decoded nonlinear DAC

The segmented system was designed in such a way that it could be configured as a conventional nonlinear DAC (*conv*, see Table 6.1) by holding the analogue input with the V_{high} voltage so that the system does not interpolate. The segmented system in this configuration was shown to produce a sine wave output with SFDR of -42dBc at 2Hz using the HP4155B parameter analyser. This result is consistent with a phase limitation of 7-bits and confirms that the row and column (global) thermometer decoding cells were operating correctly. This result indicates that the current cell sizes were matched and correctly sized to a percentage error of better than 10% with the spectral purity of -42dBc (see the yield estimates calculated in MATLAB, Figure 3.4). The dominant harmonic was the 127^{th} harmonic in agreement with simulation.

7.4.4 Novel interpolation nonlinear DAC, segmented system

A second copy of the segmented system (*int*, see Table 6.1) was included on the chip to implement the system described in Figure 5.1 where the output of the conventional DAC (*conv*) is improved by summing with the output of the interpolation DAC (*int*). However, when the *int* DAC circuit was connected no improvement was found in the SFDR suggesting that the phase DAC was working incorrectly. The phase DAC signal was not available at an output pin so it could not be checked directly. Changing each of the 6 least significant bits made no change in output. This suggested a serious error in the phase DAC layout.

When the layout design was manually checked by inspection it was found that the phase DAC used in the interpolation circuit was not internally connected to the power supply. This error explained the observed limitation of the segmented system. The error was caused by one missing via (connection between two metal layers) out of more than one thousand vias in the layout. The error was not found before fabrication by the layout versus schematic tool because the power rail of the phase DAC was connected to a virtual power pin, see section 6.6. To avoid these errors in the future virtual pins should not be used in sub cells to ensure that every connection in a sub cell connects to a circuit higher up in the hierarchy.

7.5 Re-fabrication of segmented system

The layout error was corrected by fabricating the segmented system a second time on a multi-project chip shared with 3 other projects. The number of pads available was limited to 20 so only the interpolation circuit could be included in the second chip. The second chip was fabricated 6 months after the first. During this time the manufacturer changed the process from CSX to C35. The information provided by the fabrication service required changing some of the digital parts of the design to comply with design rules in the new process such as minimum track widths and minimum spacing rules. Adhering to the new design rules was the only change requested by the manufacturer.

7.5.1 Measurement setup

To reduce the pin count on the second chip only one interpolation nonlinear DAC (segmented system) was included with two internal phase DACs, one to generate the interpolation voltage and a copy

with its inputs tied to $2^6/2$ to generate the reference voltage, V_{ref} . To ensure that V_{ref} was always in the middle of the interpolation voltage range, ($V_{phase_{min}}$ to $V_{phase_{max}}$) both phase DACs shared the same current source bias voltage, supply voltage and used equal value resistors. The measurement method was the same as described in section 7.4.1.

Two critical bias voltages for this design are the external V_{high} and V_{low} voltages which have to be set to match the lowest and highest output from the phase DAC, $V_{phase_{min}}$ and $V_{phase_{max}}$. Any difference will create a discontinuity in the output signal at each cell transition. The discontinuity appears as a spur in the output spectrum at the frequency $2N - f_{out}$ where N is the number of current cells², 64 in this design. Hence the $2N - f_{out}$ spur can be used to identify mismatch between ($V_{phase_{min}}$ and V_{low}) and ($V_{phase_{max}}$ and V_{high}). The phase DAC is a current mode DAC with the current output converted to a voltage using an on chip resistor between the output node and a pin, $V_{dd_{DAC}}$. Hence the full scale range of the phase DAC output is set by the bias voltage that is connected to the gate of its current sources and the voltage levels are set by $V_{dd_{DAC}}$. To ensure the discontinuity was minimised the highest output of the phase DAC, $V_{dd_{DAC}}$ was set equal to $V_{phase_{max}}$. This allowed the voltage $V_{phase_{min}}$ to be trimmed independently from $V_{phase_{max}}$ by changing the bias voltage of the phase DAC. The trimming was performed automatically using the PowerDAQ data acquisition system and Agilent Vee software. This bias voltage was trimmed by observing the level of the 127th harmonic as this is the dominant spur predicted in simulations.

7.5.2 Measurement results

Suppression of the discontinuities by trimming the phase DAC bias voltage pushed the 127th spur to $-72dBc$. For a typical chip the optimum SFDR was $-64.5dBc$ at a bias voltage of $0.75V$ with the 2nd harmonic dominant. The power efficiency was 65%. To investigate the effect of linearity and the efficiency and performance tradeoff the bias voltage was reduced to $0.7V$. The SFDR degraded to $-60dBc$ and the power efficiency improved to 72%.

To investigate how well the effect of variations has been managed the yield of the systems of the second chip run was measured. The variation between the measured responses of the 10 chips from the second fabrication run is shown in Figure 7.10. At a bias voltage of $0.75V$ chip5 is has a performance of better than $-65dBc$ SFDR, however the other chips are limited by the second harmonic to worse than $-65dBc$ SFDR. The yield is 90% for chips with an SFDR better than $-60dBc$. The best region

²This is predicted by the Fourier analysis of Langlois [22]

of bias voltage, 0.75V is a quarter of a volt different to the designed value used in simulations of 1V. This suggests that there is a large discrepancy between simulations and measurement.

At the optimum bias level the 2nd harmonic is dominant in the measurement results, while in both MATLAB and Cadence simulations the 5th harmonic was dominant. However the trend of the dominant harmonics is similar to that simulated in Cadence (compare with Figure 5.9). In agreement with Cadence the 3rd harmonic dominates at higher bias voltages, and the 127th harmonic dominates at lower bias. The power efficiency for SFDR tradeoff is also evident. At a bias voltage of 0.5V the efficiency is 88%, at 0.7V the efficiency is 72% and at 1V the efficiency is 46%. The loss of the region where the 5th harmonic dominates and the SFDR is optimal is a further indication that there is a difference between the fabricated chip and simulations that is significant enough to degrade the SFDR by 5dB. Nevertheless the system can achieve better than -60dB, the requirements of the low output power Bluetooth classes and UMTS and the efficiency can be traded for SFDR by changing the bias voltage.




Figure 7.10: Variation in the SFDR against the bias voltage of the current sources for 64 cells for the second chip run. The trends of the dominant harmonics are similar to those simulated in Cadence and the power efficiency for SFDR tradeoff is also evident. At a bias voltage of 0.5V the efficiency is 88%, at 0.7V the efficiency is 72% and at 1V the efficiency is 46%.

7.5.3 Mismatch distance effect

One observation is that the second harmonic distortion is greater than $10dB$ more powerful in the measurements than in simulation. One cause of second harmonic distortion discovered in simulations is a loss of the symmetry of the output wave caused by mismatch. Hence the increase in the second harmonic seen in the measurements may be caused by the distance (correlation) mismatch effect that is not included in the mismatch simulation within Cadence. The 64 cell system is significantly larger than the compact four or six cell systems and therefore the transistors are placed further apart. The 64 cell system contains two sets of 32 current cells that are placed in an 8×8 array in increasing order in terms of size, see section 6.4. This was done to minimise signal routing and simplify the thermometer decoding logic. However this pattern places the smaller current cells at a large distance apart on the chip. To reduce second harmonic distortion due to mismatch it is desirable to have the smaller current cells closer together as these cells are more susceptible to gradient errors due to their small size.

The circuit simulations assume that the transistors are placed close together, at a distance that is less than their dimension (measured from centre to centre) which is impossible in this nonlinear DAC design. The mismatch effect is also difficult to model at greater distances, changing with transistor size and orientation [70]. Pelgrom highlighted this distance effect problem and included a distance parameter in his model. However in this design no indication of how distance effects mismatch was given by the circuit manufacturer and therefore could not be included in simulations.

To test the hypothesis that the distance mismatch effect is causing the difference between simulation and measurement, a full waveform period was synthesised in MATLAB by folding the first quadrant of the measured waveform four times. This process reduced the second harmonic power to $70dB$ below the fundamental, and increased the third harmonic by $3dB$ to $-63dBc$ and the fifth harmonic by $10dB$ to $-65dBc$. This pattern of dominant harmonics is closer to that predicted by simulation, supporting the hypothesis that mismatch between the two sets of 32 current cells is limiting the SFDR. Unfortunately the only way to discover the influence of the distance effect is to measure a fabricated chip and in this case it appears that the distance effect is limiting the performance of the manufactured prototype. This indicates that it is important to design a system that places the more sensitive current cells closer together or uses a method that reduces gradient errors such as a common centroid layout. Nevertheless, given the influence of the distance effect the measurement results are within $10dB$ of the designed performance of $-70dBc$ SFDR and 90% of the chips meet the Bluetooth and UMTS spectral purity requirements of $-60dBc$.

7.6 Transient measurements of the re-fabricated segmented system

7.6.1 Measurement method

For measurements at higher frequencies, the re-fabricated segmented system was driven with the numerically controlled oscillator (NCO) at clock frequencies of up to 33MHz (the frequency limitation of the NCO chip). The NCO chip was programmed as an accumulator using a set of switches on the PCB and the phase word was set via the PowerDAQ data acquisition system connected to the PC. Both the NCO and segmented system (device under test) shared the same clock signal which was generated by the FG (function generator). The 16-bit digital output of the NCO was monitored using the digital channels on the oscilloscope. As with the previous measurements the data acquisition system was used to set the bias and reference voltages. The output was acquired using the oscilloscope triggered to the LSB of the NCO output to remove as much timing jitter as possible between the NCO output and the clock. This was important as jitter is a major source of error in the transient results.

7.6.2 Transient results

In agreement with the measurements at low frequency the SFDR of the transient measurement setup was limited by the second harmonic to -63dBc up to $f_{clk} = 1\text{MHz}$. As the clock frequency was increased above 1MHz the fundamental and spurs decreased in signal power with the second harmonic remaining dominant. The SFDR was constant at -63dBc up to a clock frequency of $f_{clk} = 10\text{MHz}$ where the noise floor was less than -70dBc . The SFDR was better than -60dBc at $f_{clk} = 33\text{MHz}$ with the phase word set to 1. At this clock frequency substantial jitter between each of the digital output lines of the NCO increased the noise floor to -60dBc masking the effect of the dominant harmonics and degrading the SFDR by 3dB .

The degradation in SFDR as the phase word was increased at $f_{clk} = 10\text{MHz}$ is shown in Figure 7.11. The SFDR is better than -60dBc at output frequencies up to 200kHz , $f_{clk}/55$ and the SFDR degrades at higher phase words. The drop in performance due to increasing the phase word is similar to that found in simulation, section 4.5.2 due to long settling time of the phase DAC. In simulation the settling time was found to be 2.5ns and this caused significant nonlinearities in the input signal at frequencies greater than $f_{clk} = 50\text{MHz}$ where the settling time is 20% of the clock period. The measurement results suggest a settling time of 20ns in the manufactured system as this is 20% of the clock period at $f_{clk} = 10\text{MHz}$. This error is likely to be due to an underestimated value for the

impedance that the phase DAC is required to drive. The estimation of the impedance considered the parasitic capacitance between the track and substrate and the gate capacitance of one current cell as the phase DAC drives only one current cell at a time. However the parasitic coupling capacitance between adjacent tracks and the parasitic capacitance within the transmission gates of the 64 tri-state switches will add significantly to the phase DAC load. Adding in the resistivity of the long track and it is possible that the load on the phase DAC could be eight times that estimated³. Re-simulation with the load on the phase DAC increased eight times confirmed the measured results.




Figure 7.11: The measured transient performance of the segmented system on a typical chip.

To determine if the degradation with increasing phase words is seen across all clock frequencies, the experiment of increasing the phase word was repeated with a $f_{clk} = 1\text{MHz}$, Figure 7.11. The measured SFDR was better than -60dBc for phase words up to $f_{clk}/4 = 2047$, suggesting that either the clock jitter or the phase DAC loading is limiting the performance at $f_{clk} = 10\text{MHz}$. The SFDR of the segmented system may degrade at phase words greater than 55 due to insufficient current of the phase DAC to drive the larger voltage changes at the input of the nonlinear DAC.

The limitation in transient performance is similar to that simulated and measured with the compact systems (see sections 4.5.3 and 7.3.1). The important difference between these two systems is that the phase DAC is external to the compact systems and its ability to drive the input gates of the compact systems could be measured directly. In the segmented system the phase DAC is located on

³The design kit used for process in this design did not extract any resistance values for the tracks

chip and hence its output voltage cannot be measured directly. Improving the transient performance of this system could be achieved in two ways in subsequent designs. Firstly the output current of the phase DAC could be increased to allow it to drive the track capacitance and switch transistors of the nonlinear DAC. Secondly an accumulator could be provided on chip to allow testing at clock frequencies greater than $10MHz$ to overcome the problem of timing jitter.

7.7 Re-simulation of the 64 cell system using new process parameters

The measurement results of the 64 cell system meet the spectral purity requirements of Bluetooth and the UMTS specifications however they are disappointing, worse than the simulation results. The discrepancy between the dominant harmonic, 2^{nd} or 5^{th} , can be explained by the distance mismatch effect (section 7.5.3), however the resulting SFDR is still $5dB$ worse than in simulation and there is a large voltage shift in the optimum bias voltage.

During the design of the second chip the manufacturer advised that the difference between the two processes (CSX for the first chip and C35 for the second chip) only affected the design rules. However, after the second run it was discovered that this was not the case. The process parameters had also been changed and had not been released before the second chip fabrication. The digital cells were unaffected by this change, however the key parameters in the analogue cells changed substantially. For instance there was a 7% decrease in V_t and 4% increase in β . These together with other changes caused a 2.5% change in I_{ds} above threshold and a 10% change in I_{ds} in subthreshold. More importantly the difference between drawn and electrical geometry was worse. In the new process, C35, $\Delta L = 0.09\mu m$ and $\Delta W = -0.05\mu m$ whereas in the original process, CSX, $\Delta L = 0.008\mu m$ and $\Delta W = 0.027\mu m$. This change means that the lengths chosen for the switch devices will be less accurate leading to distortion of the output signal.

The low frequency simulations of the segmented system were repeated with the C35 process parameters. The best SFDR was $-66.6dBc$ at a bias voltage of $0.8V$, limited by the 5^{th} harmonic. This was a degradation in SFDR of approximately $7dB$ due to the changes in the process parameters. This shows that the seemingly unimportant changes have a substantial detrimental effect on the performance of the 64-cell nonlinear interpolation DDFS system. These results are closer to those measured and explain the remaining discrepancy between measurement and simulation. The distance effect (section 7.5.3) explains why the 2^{nd} harmonic was dominant in the measurements and the change in process parameters explain why the measured result was approximately $5dB$ worse.

To show that the design can be re-adjusted to the change in process parameters the system was re-simulated in Cadence using the following adjustments. The current source bias voltage was reduced by 7% to account for the decrease in V_t . The increase in β was accounted for by reducing the output voltage swing of the phase DAC. The degradation in drawn resolution was accounted for by doubling the width of the switch devices to ensure that the difference between drawn and effective geometry caused less error in the β parameters. This improved the SFDR to $-73dBc$, limited by the 5th harmonic, and efficiency of 45% in better agreement with MATLAB and previous Cadence simulation results. The improved SFDR could not be measured using the fabricated chip as the device sizes need to be changed and simulation results are the best method to show the performance that could be achieved with this system. A third chip could not be manufactured due to time constraints on this project. Simulations agree with the measured systems and show that this design can be improved. High frequency simulations were repeated with the C35 process parameters using the conditions described in section 5.4. The SFDR was $-73dBc$ at clock frequencies up to $50MHz$, $-70dBc$ at $f_{clk} = 100MHz$ and $-68dBc$ at $f_{clk} = 500MHz$. These results demonstrate that this system is suitable for the intended communications applications, Bluetooth and UMTS.

7.8 Summary

The measured results of the compact systems based on arrays of four and six current cells are in agreement with both the MATLAB and Cadence simulations from DC to frequencies up to $10MHz$. The SFDR is better than $-40dBc$ for the four cell system and better than $-48dBc$ for the six cell system at $f_{out} = 1MHz$. The measurements of these circuits were limited $f_{clk} = 10MHz$ by the settling time of the arbitrary function generator. The simulations indicate that these systems are useful at clock frequencies up to $100MHz$ for the instrumentation applications. No power dissipation measurements could be made as all the circuits share power supply pins to reduce the pin count.

The 64 cell system used a novel three state analogue switch. This was confirmed to operate correctly at frequencies up to $1MHz$. A conventional nonlinear DAC that does not use interpolation was also confirmed to operate correctly and was limited to $-42dBc$ SFDR due to phase truncation as expected. However a layout error meant that the configurations of the 64 cell segmented system that used nonlinear interpolation were unusable. A second chip was manufactured with the 64 cell system included. The second design contained a working system, however the performance was not as good as expected. The best SFDR achieved was $-65dBc$ for one of the ten chips fabricated at

a bias voltage of $0.75V$ with a power efficiency of 65%. At the same operating point nine of the chips were better than $-60dBc$. After checking the accuracy of the experiment, the circuits were re-simulated using the parameters for the actual manufacturing process. These new results matched the experiments. It was therefore concluded that despite the reassurances of the manufacturer a late change in the manufacturing process (from CSX to C35) has significantly changed the performance of the system.

Transient measurements of the 64 cell segmented system were limited to $10MHz$ due to the jitter limitations of the off chip accumulator (NCO). The SFDR was constant at $-63dBc$ for clock frequencies up to $10MHz$. By increasing the phase word it was discovered that loading of the phase DAC is the reason for the limited performance. From simulation it was expected that the bandwidth would be limited by the settling time of the phase DAC to clock frequencies of up to $500MHz$, however this result could not be confirmed due to limitations in the measurement setup. Measurement at higher frequencies could be achieved in future designs by including an accumulator on chip and increasing the output drive of the phase DAC.

The important conclusion from the measurements that could be made of the segmented system is that it is robust to errors caused by changes in the manufacturing process. Amid the unexpected changes in process parameters the system was still able to achieve better than $-60dBc$ with a power efficiency of 65%, the SFDR requirement for the Bluetooth and UMTS communications standards. In agreement with simulations the efficiency could be traded off for SFDR, at a bias voltage of $0.5V$ the efficiency is 88% and the SFDR was $-50dBc$, the requirements of the low output power Bluetooth classes. Further simulation results using the C35 parameters have shown that this system can perform in at least two operating regimes. Spectral purity of $-73dBc$ with efficiency of 45% in close agreement with the circuit simulations predictions of the CSX process and the efficiency can be improved to 75% with a drop in performance to SFDR, $-60dBc$, both at $500MHz$, meeting both the required spectral purity and bandwidth of the Bluetooth and UMTS communication standards.

Chapter 8

Conclusion and further work

8.1 Conclusion

The main contribution presented in this thesis is a novel, efficient method of interpolation for DDFS that uses the nonlinear response of the CMOS differential switch already present in the high speed current steering DAC. The nonlinear response provides a smooth transition between the conventional, quantised DAC output. This interpolation may be performed with the conventionally discarded phase bits leading to highly compact and efficient DDFS architectures for application in instrumentation and communications systems.

Low power direct digital frequency simulation (DDFS) has many applications. Two particular applications that have been described are instrumentation systems and wireless communications. Two examples of instrumentation systems that require wideband frequency synthesis from near DC to $1MHz$ are dielectrophoresis (DEP) and electrical impedance spectroscopy (EIS). DEP can be employed to concentrate, separate and characterise cells based on their response to a non-uniform electric field and would benefit from low power integrated DDFS systems with spectral purity better than $-30dBc$. EIS measures the spectral characteristic of objects and would benefit from low power active sensors that use DDFS systems with spectral purity better than $-40dBc$. Both DEP and EIS measure effects that vary slowly over a large frequency range and hence the frequency resolution requirements are low as are the bandwidth and spectral purity. These two instrumentation applications identify a new application area for compact, low power, low resolution and low cost, disposable frequency synthesisers. The system presented here is one example of how this need may be met.

Two wireless digital communications standards were of interest, UMTS that forms the air interface standard for the third generation of mobile phone standards and Bluetooth the new low power protocol for short range peripheral devices. These standards use frequency hopping to spread the spectrum of a limited bandwidth to many channels. Low power operation of synthesizers in these mobile devices is of importance to conserve battery power. DDFS is attractive for this application as the resolution and output frequency are digitally defined. UMTS requires a bandwidth of at least 15MHz and spectral purity of better than -60dBc . Bluetooth requires a bandwidth of 80MHz and a spectral purity better than -50dBc SFDR. Spurious frequency dynamic range (SFDR) is the commonly used measure of spectral purity and is defined as the difference in dB between the wanted frequency (carrier) and the most powerful, unwanted spurious frequency or spur.

Conventional DDFS systems consist of an overflowing accumulator to generate the phase, a ROM lookup table to convert the phase to amplitude and a DAC to perform the digital to analogue conversion. These systems suffer from the approximations used to reduce the size of the ROM and non-ideal effects of the DAC. High speed, low power CMOS DACs achieve -70dBc SFDR using a current steering architecture with differential switching and thermometer decoding to minimise the glitches resulting from unsynchronised transitions and unwanted clock feedthrough.

Minimisation of the size of the ROM is the most active area of DDFS research where an increase in additional digital circuitry is traded off against a reduced ROM size. Two popular techniques lead to a saving of two ROM bits by utilising quarter wave symmetry and a saving of a further two bits by storing the difference between the amplitude and phase. Further approximations using Taylor series, interpolation and trigonometric identities reduce the ROM by up to 300 times, however the most efficient method is to remove the ROM entirely.

Langlois and co-workers approximate the sinusoid using piecewise linear segments that are generated digitally. Moheildin and co-workers use the same approximation but generate the linear segments using a small linear DAC and array of nonlinearly weighted resistors. Lee and co-workers store the phase to amplitude conversion information in a nonlinear current steering DAC. The conclusions following this review indicated that the most efficient form of DDFS minimisation is to approximate the sinusoid output using a small number of segments and to implement the design using current steering DACs to allow the highest bandwidth operation achievable in CMOS.

The novel DDFS design used in this design utilises the first order nonlinear response of a current cell consisting of an ideal current source and differential current switch. A single such cell was able to achieve an SFDR of -45dBc when driven by a triangular wave. This surprisingly good result led

to the investigation of the performance of nonlinear DACs based on small numbers of current cells. Two architectures were identified, a compact architecture for the instrumentation applications and a segmented architecture for the communications applications. These systems used the current cells in the nonlinear DAC more efficiently than previous designs by carefully controlling the nonlinear response of each current cell by using phase bits that are discarded in a conventional DDFS. The compact cells are 100% efficient as all the bias current is used to form the output. The only extra component was a linear DAC and hence the power consumption was reduced by up to 10 times over state of the art designs, Table 8.1. The compact systems consisted of a nonlinear DAC, a small linear phase DAC and an overflowing phase accumulator. MATLAB simulations of the compact architecture suggested two systems based on the nonlinear interpolation technique for the instrumentation applications. One system with a nonlinear DAC of four current cells achieved an SFDR of $-40dBc$ up to output frequencies of $1MHz$ for DEP and consumed only $5\mu W/MHz$. A system with a six cell nonlinear DAC achieved $-48dBc$ for output frequencies up to $1MHz$ for EIS and consumed only $8\mu W/MHz$. The linear phase DAC was the dominant source of power and area consumption of these systems contributing half of the power consumption and over four times the area. Previous ROM-less DDFS are summarised in Table 8.1 along with the systems presented in this work.

Author	Segments	Power (mW/MHz)	F_{clk} (MHz)	SFDR (dBc)	On chip DAC	Active area (mm^2)
This work (6 cells)	12	0.008	100	-50	Yes	<0.01
This work (64 cells)	128	0.018	100	-72	Yes	0.6
Mohieldin [60]	8	0.08	100	-59	Yes	1.4
Langlois [56]	8	0.1	125	-60	No	0.011
Bellaouar [21]	16point ROM	0.3	30	-60	No	0.9
Lee resistor string [59]	96	0.4	230	-55	Yes	1.6
Lee current mode [20]	192	0.8	300	-62	Yes	1.4

Table 8.1: A comparison between the work presented here and recently published, state of the art low power ROM-less DDFS systems. The active area comparisons include a phase accumulator with a phase input word of at least 16 bits which accounts for 20% of the area in some state of the art systems [60, 21]. Except for the 6 cell system that does not include a phase accumulator as it is not intended for applications that require fine frequency resolution. Based on state of the art systems a 16-bit phase accumulator would add an area of $0.2mm^2$, [21].

The nonlinear interpolation technique was extended to a segmented architecture with 64 current cells. The result was a system with a SFDR of better than $-70dBc$. The larger number of current cells required the use of a modified thermometer decoder. The additional benefit of the thermometer decoder was the ability to make the response of each cell more linear which improves the spectral purity when more than 16 current cells are used. This reduced the power efficiency as all the bias current was no longer used to form the output. A target efficiency point of 50% with SFDR of $-72dBc$

was chosen as state of the art low power CMOS DACs are limited to $-70dBc$ SFDR. This system was over 200 times larger than the compact systems however the area consumption is not excessive when compared to state of the art DDFS systems (Table 1.3), particularly noting that [56, 21] do not include a DAC.

These systems were sensitive to the unwanted mismatch inherent in CMOS processes as they relied on matching both the nonlinear weighted current sources and the nonlinear weighted gains of the current switches in the nonlinear DAC. Mismatch was controlled using the Pelgrom model to specify the minimum allowable transistor areas. The allowable mismatch was modelled in MATLAB and the device sizes confirmed in Cadence Monte Carlo simulations. Unit devices and gain parameters were selected using MATLAB and the result of this truncation of parameters reduced the linear performance of 64 cell system to $-72dBc$. Cadence simulations were used to analyse the effect of transients and glitches on SFDR and the performance was better than $-72dBc$ at $F_{clk} = 100MHz$ with $18\mu W/MHz$ power consumption.

A custom layout prototype was required to measure the real performance. One sample of each of the four and six cell circuits was included on the prototype. Each was about the size of a single pad and no phase DAC was included to simplify measurements. One copy of the 64 cell system was included with two matched phase DACs to reduce the number of reference voltages. A custom layout of the circuits was made with careful attention to reducing bias voltage drop and guard rings were used to minimise crosstalk.

The measurement results of the four and six cell systems confirmed the Cadence transient simulation results and showed that these systems are capable of implementing the instrumentation systems DEP and EIS with a high yield. The SFDR is better than $-40dBc$ for the four cell system and better than $-48dBc$ for the six cell system at $f_{out} = 1MHz$. The power efficiency of these systems is 100% as all the bias current is used to form the output current. The 64 cell, segmented system was measured at $f_{clk} = 10MHz$ with an optimum SFDR of $-60dBc$, with an efficiency of 65% for 9 out of the 10 chips manufactured. In agreement with simulations the efficiency could be traded off for SFDR when the efficiency is improved to 88% the SFDR degrades $-50dBc$. The loss in optimum performance was due to both the distance mismatch effect and a change in all the process parameters, particularly an increase in the difference between drawn and electrical geometry. Further simulations showed that it should be possible to build a system with SFDR of $-73dBc$ by increasing the size of the switch devices within the nonlinear DAC and using a layout structure such as a common centroid scheme that minimises the distance effect. Despite the loss in performance from simulation to measurement,

the measured system still met the UMTS and Bluetooth spectral purity requirements as intended.

Overall the results of the 64 cell system are disappointing because of the distance effect and a change in the process parameters. A more promising avenue is the new area of biomedical applications that call for compact, disposable and low power frequency synthesisers. The compact systems presented here fit the specifications of such systems, DEP and EIS at a tenth of the power consumption and size of state of the art systems that include a DAC.

8.2 Further work

The compact system based on an array of four current cells is suitable for application in dielectrophoresis (DEP). The requirements of the system are a spectral purity of better than $-30dBc$ and bandwidth of $1MHz$ output frequency (section 1.5.2). The compactness of the four cell system makes this circuit ideal for the manufacture of arrays of large numbers of DEP electrodes and the digital tunability of the output frequency means that the frequency can be changed accurately to characterise cells. Phase resolution requires that at least 5 phase bits are used from the accumulator and amplitude resolution requires at least 4-bit phase DACs for $-30dBc$. To allow a margin of error 7 phase bits should be used. The system is limited in size reduction by the requirement for an accumulator, phase DACs and the pads. If further development suggests that the system does not require the frequency resolution offered by the accumulator and phase DAC, these components could be replaced by a more compact dual slope generation circuit. The number of pads required by the system may be reduced by using a resistive ladder to generate the reference and bias voltages.

DEP requires the generation of electric fields from signals of greater than 10V peak to peak [6]. Low power CMOS processes are usually limited to voltage levels of less than 3.3V due to the small oxide thickness. However some processes (such as the AMS C35 process used for this design) have high voltage transistors with greater oxide thickness extending the voltage levels to 15V for the processes used in this design in particular. The high voltage transistors can be used to shift the output voltage to 15V, Figure 8.1. The 15V transistors are biased with a gate voltage of 5V and act as cascode transistors to ensure that the drains of the current switch transistors are at 3.3V so that the circuit operates as designed. The output currents generated by the current cell are I_{out1} and I_{out2} and are 180° out of phase. These currents flow through the buffer transistors and are generate a 10V voltage output across the $500k\Omega$ output resistors. Figure 8.1 shows a single current cell, however the same buffer circuit may be used to convert the currents flowing from an array of any number of current




Figure 8.1: Differential current switch with 15V outputs.

sources to a higher output voltage.

The traveling wave DEP (TWD) technique requires a phase gradient generated using rows of parallel electrodes. Each addressed with four sinusoidal voltages at the same frequency but phase shifted by 0° , 90° , 180° and 270° . The electrorotation (ROT) DEP method requires a rotating electric field that can be generated at the centre of a chamber surrounded by four electrodes if each electrode is connected to one of four sinusoidal voltages with relative phase shifts of 0° , 90° , 180° and 270° . The four cell array of current cells generates two outputs at the same frequency, one with a phase of 0° and the other 180° . One way of generating the other two output phases is by driving an identical copy of the phase DAC and four cell array with a digital phase signal shifted by 90° , Figure 8.2. In Figure 8.2 the top system uses the conventional method of forming a full wave from a half wave by taking the XOR of the MSB of the accumulator with the remaining bits. The bottom system performs a phase shift of 90° by replacing the MSB by the XOR of the most significant bit ($A < 0 >$) and the next most significant bit ($A < 1 >$) and inverting $A < 1 >$.

A mixed signal simulation of the method in Figure 8.2 confirmed that four sinusoidal voltages at the same frequency with relative phase shifts of 0° , 90° , 180° and 270° could be generated using one accumulator and two sets of a phase DAC and nonlinear array of four cells. The output voltages shown in Figure 8.3 are promising and further work could lead to an even more compact design.




Figure 8.2: A system capable of generating four phase shifted outputs at 0° , 90° , 180° and 270° .




Figure 8.3: Cadence simulation results of the 15v outputs phase shifted by 0° , 90° , 180° and 270° . All signals had a SFDR of $-42dBc$.

Publications related to this work

A. McEwan and S. Collins, "An Efficient CMOS DDFS Architecture for Bluetooth",
Bi – annual IEE Analogue Signal Processing Conference, 2002, Nokia best paper award.

A. McEwan, S. Shah and S. Collins, "A Direct Digital Frequency Synthesis System for Low Power Communications",
Proceedings European Solid State Circuits Conference, 2003

A. McEwan and S. Collins. "Analogue Interpolation Based Direct Digital Frequency Synthesis".
Proceedings of 2003 IEEE International Symposium on Circuits and Systems, 2003, pp621-624.

A. McEwan and M. Chappell, "A Low Power, High Speed Accumulator for DDFS Applications",
Proceedings of 2004 IEEE International Symposium on Circuits and Systems, 2004, pp797-800.

Appendix A

A low power parallel pipelined accumulator

A.1 Introduction

Power consumption in existing DDFS systems is dominated by the ROM required to perform phase to amplitude conversion. For example, even when ROM reduction techniques are employed to minimise power consumption in the ROM, it still consumes four times more power than the accumulator [21]. This has led to the development of systems in which the ROM is replaced by a nonlinear DAC [20].

The development of ROM-less DDFS systems [71] introduced in Chapter 1, means that the accumulator is once again a significant drain on the power source. A design is therefore required that reduces power consumption in the accumulator, whilst maintaining a high update rate.

In this appendix a parallel pipelined architecture is presented that was designed in collaboration with Michael Chappell as part of his undergraduate work. The aim was to use standard digital cells to create this auxiliary circuit. The power consumption is reduced by identifying and removing the redundant circuitry and introducing a dynamic delay element. The use of a hybrid adder to reduce power consumption is also introduced and detailed simulation results are discussed and used to justify the final accumulator design.

A.2 The architecture

A.2.1 Pipelining

Many existing DDFS systems include a pipelined accumulator architecture, Figure A.1, which improves throughput over a simple ripple-carry adder based accumulator [72]. The maximum update rate of a pipelined accumulator is limited by the time it takes for the carry to ripple through the largest adder in the system (T_{cr}). The size of the largest adders, and hence T_{cr} , can be reduced by subdividing the accumulator further. This is done at the cost of increasing the number of pipeline stages and the number of D-type flip flops (DFFs). DFFs refer to the standard cell D-type flip flops provided in the AMS digital library.




Figure A.1: A 24bit, 3 stage pipelined accumulator

In DDFS systems some power saving is made when the number of phase bits is truncated in line with the spurious free dynamic range (SFDR) requirement [3]. Phase truncation allows the removal of the post adder pipeline stages. For example, in a 24-bit accumulator with 13-bit output, 19 out of 96 DFFs may be removed. Despite this saving the DFFs still account for 95% of the total accumulator power consumption.

A.2.2 Paralleling and pipelining

To achieve a saving in power consumption the DFFs need to be used more efficiently. In a DDFS accumulator the input word usually changes less often than each clock cycle. Therefore several results can be calculated in parallel. The progression of states technique [73] can be used to develop

System	Flip-flop type	Adder type	Predicted power $\mu\text{W}/\text{MHz bit}$
1	DFF	CMOS	2.25
2	DFF	CMOS	1.14
3	DFF	CMOS	1.1
4	D_1	CMOS	0.55
5	$D_1, (+D_2)$	hybrid	0.50

Table A.1: Comparison of the development stages of the low power accumulator : 1 Pipelined (3 stage); 2 PPA; 3 PPA with unnecessary sum circuitry removed; 4 PPA employing D_1 elements; 5 PPA employing hybrid adders and D_1 elements, also using D_2 where appropriate.





Figure A.2: (a) A 24-bit, 3 stage parallel pipelined accumulator (b) Parallel adder block

a semi pipelined architecture where N parallel adder elements are used to calculate the next N results simultaneously. This allows the DFFs in the pipeline to be run N times slower providing more scope to reduce the overall power consumption. The parallel pipelined accumulator (PPA) proposed by Yang et al, [74] is depicted in Figure A.2 (a). Four adders calculate the next four results simultaneously. The four carry results are passed from the parallel LSB adders to the parallel MSB adders every $f_{clk}/4$ and the results are placed on an output bus that is wide enough to accommodate all four results simultaneously. The final part of the system is a multiplexer that is clocked at f_{clk} and presents each of the four results to the output in the correct sequence.

The internal architecture of the parallel adder element (Figure A.2), can be derived by considering the parallel accumulation of an 8-bit word within four accumulator clock cycles (f_{clk}). The first result, $O1$, is the previous result plus the PW. The second, $O2$, is the previous result with the PW doubled, i.e. shifted by 1-bit. The MSB, $I[0]$, is shifted out and the MSB from the lower next 8-bits is shifted into the input, $Low[0]$. The third, $O3$, could be computed by multiplication by 3, however it is more efficient to use the $O2$ result, creating a series cascade of two 8-bit adders. The fourth, $O4$, uses a left shift of two bits to multiply by four [74].

$$O1[0:7] = I[0:7] + O4'[0:7]$$

$$O2[0:7] = 2 \cdot I[0:7] + O4'[0:7] = [Low[0], I[1:7]] + O4'[0:7]$$

$$O3[0:7] = 3 \cdot I[0:7] + O4'[0:7] = O2[0:7] + O4'[0:7]$$

$$O4[0:7] = 4 \cdot I[0:7] + O4'[0:7] = [Low[0:1], I[2:7]] + O4'[0:7]$$

(where $O4'$ is $O4$ of the previous clock cycle)

By calculating four results simultaneously this architecture offers a better maximum output rate than a simple pipelined accumulator. The apparent increase in power required by the additional components is offset by running each component at a lower speed, in particular the DFFs. The power consumptions of different types of accumulator have been estimated using the power consumptions of the standard cells that would be required to create each system. To allow a comparison of the different architectures the estimated power consumptions for these systems are included in Table A.1, along with the estimated power consumptions of other systems that will be described later. In this table the pipelined architecture is system 1 and the parallel pipelined architecture, created from standard cells, is system 2. The results in Table A.1 show that the estimated power consumption of the parallel pipelined architecture is approximately half that of the pipelined architecture. These estimates show that although the parallel pipelined architecture requires more components, its power consumption is lower. This is because each component can be designed to have a slower response time and therefore consume less power.

A.3 Low power optimisation

Although the parallel pipelined architecture consumes less power than the pipelined architecture, further power savings are possible. The first power saving measure to be investigated was the trimming of the architecture to match the particular combination of input and output word lengths required by a DDFS system with 24-bits of amplitude resolution and 13-bits of phase resolution. Savings may be made by removing the sum circuitry from the majority of the outputs from the lower 11 (truncated) bits. This results in a small power saving of approximately 3% corresponding to the difference between the estimated power consumptions of systems 2 and 3 in Table A.1.

A.3.1 Low power flip-flops




Figure A.3: Delay element (D_1), positive edge triggered

In the PPA architecture 70% of the power is consumed by the DFFs. An alternative, low power, dynamic delay element (D_1) was therefore designed. This exploits the operation of an accumulator in a DDFS system whereby the flip-flops are refreshed every clock cycle. The DFF does not need to hold its value indefinitely, and a simple 'sample and hold' alternative is sufficient. The resulting clocked delay element, shown in Figure A.3, requires only 25% of the power used by the DFF in the standard cell library. The only drawback is that the delay element inherently produces an inverted output. Therefore, in paths with an odd number of delay elements an additional inverter is added to the path. The result is a small increase in the number of components which slightly offsets the reduction in power consumption achieved by using delay elements rather than the standard cell DFF. However, the estimates for the power consumption of the resulting system (system 4 in Table A.1), show an overall significant power saving.

A.3.2 Low power adders

With the delay element, (D_1) employed in the PPA, the adder elements will account for 40% of the overall power consumption. Further significant power savings may therefore be possible if a more efficient adder architecture is also employed.

Previous research has suggested the use of pass transistor logic (PTL) to create low power adders [21, 75]. PTL is suited to performing the XOR operation required to compute the sum, however it is not advantageous for the logic functions required to calculate the carry [76]. PTL structures are difficult to use as the output lines draw current directly from the input rather than from the power supplies as in CMOS. This can be alleviated by using buffers to restore signal levels, however this reduces the power saving sought from the use of PTL. For this system a hybrid full adder was designed, (Figure A.4), in which the sum is calculated using complementary pass transistor logic without level restoration [77]. The carry is calculated using standard CMOS. This is a good arrangement for an accumulator where the sum signals are immediately restored by delay elements while the carry




Figure A.4: Hybrid full adder (1-bit). (a) Sum stage. (b) Carry stage.

signals are required to quickly ripple to the next most significant adder.

The power consumption of the hybrid adder is only 50% of that required by the standard CMOS full adder. However the PTL sum section requires complementary inputs. The basic accumulator architecture was therefore modified. In particular an inverter is required at the output of some delay elements and the delay element that drives the feedback path is found to require level restoration at the input. This required the design of another delay element, D_2 , Figure A.5. The additional components in D_2 means that it will consume more power than the D_1 element. Unfortunately, the increased power consumption within the pipeline offsets the power savings in the adder. In fact a comparison of the estimated power consumption for systems 4 and 5 in Table A.1 suggests that the additional power consumption in the pipeline will almost totally negate the power saving within each adder element that has been achieved by using PTL.

A.4 Simulation results

The results in Table A.1 are estimates which fail to take into account the different loading of each component in the two systems which have very similar estimated power consumptions. In particular the estimated power consumption values are unable to fully account for the current drawn from the delay elements by the PTL. Consequently it was necessary to simulate both systems 4 and 5 of Table A.1 in detail.

For a typical input sequence simulation results from a Spectre circuit simulator suggest that the average accumulator power consumption when using a CMOS based adder is $0.46\mu W/MHz$ bit. This compares well with the result for the system that was modified to incorporate an adder based upon PTL which had an average power consumption for the same inputs of $0.52\mu W/MHz$ bit.




Figure A.5: Delay element D_2 providing level restoration and complementary outputs, positive edge triggered.

Simulation results therefore suggest that the best design for an accumulator, shown in Figure A.6 combines a pipeline based upon delay elements and adders based upon standard CMOS cells. Simulation results suggest that this system will operate at $800MHz$ with an average power consumption of $8.9mW$, when manufactured on a $0.35\mu m$ process and operated at $3.3V$. This corresponds to a power consumption of $0.46\mu W/MHz$ bit, which is less than half the power consumption of the recently proposed version of this architecture based on standard DFF elements [74].

A.5 Layout of the accumulator

An automatic layout tool called "Silicon Ensemble" that is part of the Cadence design tool was used to automatically generate the low power parallel-pipelined accumulator. The automatic tool places standard cells in rows and then determines a path for the wires between the cells to be routed, if possible. If the automatic tool cannot find a route for all connections, the designer must increase the area of the design to allow more space for routing. The automatic tool attempts to reduce the length of wires to ensure that the impedance of the wires is at a minimum. Timing constraints may be set on the critical paths such as the clock path.

The standard cells are designed by the manufacturer to generic rules to make automatic placement simpler and more area efficient. The standard cells are all of the same height with variable width, a $2\mu m$ voltage supply V_{dd} track runs across the top of the cell and a $2\mu m$ ground line runs across the bottom. The PMOS transistors are placed close to the top half and the NMOS transistors in the bottom half of the cell. The substrate contacts for these transistors are placed within the $4\mu m$ supply lines. This allows the top half of all cells to contain the N-well region, and all cells to be placed directly next to each other within rows of equal height and the rows to be abutted. This is possible




Figure A.6: Final low power parallel pipelined accumulator.

by vertically flipping every second row upside down so that the ground lines overlap the ground lines of the row above and the power supply V_{dd} lines overlap the power supply V_{dd} lines of the row below.

To minimise the area of the accumulator design, and aid the automatic layout tool, the custom cells were made the same height as the standard cells. This allows the automatic placement tool to place both standard cells and custom cells within the same height rows and allows the rows to abut each other.

The custom cells used in the accumulator were the $D1$ delay element and a $D1bar$ delay element which is identical to the $D1$ element with an inverter included to provide the complement output signal. The layouts of these cells are shown in Figures A.7 and A.8. On the left is the clock inverter, in the middle left the transmission gate switches and middle right the clocked output buffer (compare with Figure A.3). The large metal regions were drawn to provide a large area for the automatic router to connect interconnections between the cells.

To use timing constraints, the automatic layout tool uses critical timing information of each cell, such as the rise/fall times and delay of each standard cell. This information is provided by the manufacturing plant as a binary file that could not be edited to include the timing information of the custom cells, $D1$ and $D1bar$. Therefore timing constraint driven routing could not be used in this design. Rather the clock paths were routed as normal (minimum) width tracks, and widened after automatic routing to enable them to carry the required current for fast operation without electromigration (melting). This reduced the resistance of the clock tracks and increased the capacitance and was deemed acceptable in this design as the architecture of the parallel-pipelined accumulator allows the clock to be run at a quarter of the system clock speed. The capacitance of the tracks was estimated and an additional clock generation sub circuit designed to generate the clock signal and multiplexer signals to drive the estimated track capacitance.




Figure A.7: Layout of D1. Compare with Figure A.3. From left to right the layout consists of the clock inverter, transmission gates and the clocked inverter that generated the output. The substrate contacts are shown in the power rails across the top and the bottom of the cell. V_{dd} is at the top of the cell with the PMOS transistors and gnd runs along the bottom of the cell with the NMOS transistors.




Figure A.8: Layout of D1bar. From left to right the layout consists of the clock inverter, transmission gates, the clocked inverter that generates the output and an output buffer. The substrate contacts are shown in the power rails across the top and the bottom of the cell. V_{dd} is at the top of the cell with the PMOS transistors and gnd runs along the bottom of the cell with the NMOS transistors.

The accumulator requires an input of 24-bits. To reduce the number of pins required on the prototype chip, a shift register was used to take the 24-bit input serially from one pin and store the phase word for the accumulator. The shift register requires two pins one to load the serial data and an enable signal and therefore saves a total of 22 pins.

A.6 Clock generation sub circuit

Power dissipation is reduced by clocking the D latches at a frequency one quarter of the system clock, $clock_{4o} = f_{clk}/4$. The multiplexers used to form the accumulator output require one select signal at the system clock rate, and another at half the system clock rate. These signals are generated by a clock generation sub circuit. The clock generation sub circuit is shown in Figure A.9. The combination of the XOR gate and D-latch halves the clock frequency as the output of the D-latch only changes on the positive clock edge. The output of the XOR gate is always the complement of the output of the latch and therefore every positive clock pulse causes the output of the latch to toggle, creating a clock signal with twice the period of the latch clock. Two of these circuits are cascaded to generate a clock signal with a quarter of the system clock frequency (Figure A.9).

The first four buffers in each path are sized in exponential increments to ensure that the load each is required to drive is not large enough to slow the clock path [78]. The 8 strength buffers in parallel are required to drive the large impedance of the clock lines of the parallel-pipelined accumulator which was modelled as a 30Ω resistor and $0.4pF$ capacitor delay line. The series combinations of buffers are used to ensure the correct phase difference between the multiplexer control signals and the clock at very high frequencies. This phase difference is only required at frequencies of greater than $500MHz$. At lower frequencies, the response of the inverters are fast enough to provide rise and fall times that do not affect the output of the multiplexers. All cells within the clock generation circuit are standard cells and hence the automatic layout tool was used to create the output.

A.7 Testing of the accumulator

The accumulator was tested using a mixed signal oscilloscope at clock frequencies of up to $80MHz$. The first three MSBs operated independently from the rest of the accumulator suggesting a fault in the carry signal between the ten least significant bits and the three most significant bits. Apart from this anomaly the accumulator behaved as expected. No fault could be found in the layout between the between the third and fourth most significant bit, however one drawback of an automatically generated layout is an over complicated mix of signal lines that are hard to follow to find errors. The layout of the accumulator had passed the layout versus schematic check that ensures the connectivity in the layout is exactly identical to that in the schematic. The layout was also successfully simulated with the same result as the schematic. A possible fault that could not be checked is incorrect use of the standard cells.

A.8 Summary

An accumulator optimized for use in direct digital frequency synthesis systems with 24-bits of frequency resolution and 13-bits of phase resolution was presented. Update rates in excess of $800MHz$




Figure A.9: The clock generation sub circuit. The number inside the buffer symbol represents the strength of the buffer the xN figure underneath some of the 8 strength buffers represents a serial chain of 8 strength buffers which are used to generate the required delay. These numbers were discovered experimentally in simulation.

are achieved by careful choice of architecture. Power dissipation is reduced by removing redundant circuitry and by using a delay element suitable for accumulators within a DDFS system. The static, D-type flip flops used in the pipeline architecture are the dominant source of power dissipation. Hence the refresh-every-cycle operation of accumulators in DDFS is exploited by using dynamic delay elements to reduce power consumption.

Attempts to reduce power consumption further by using pass transistor logic within the parallel adders in the accumulator were frustrated by the need to modify the pipeline structure. Simulation results suggest that employing an efficient pass transistor adder design actually increases the overall power consumption. The most efficient design is therefore a parallel pipelined accumulator which employs a delay element in the pipeline and an adder based upon standard cells.

Appendix B

Derivations of circuit models

B.1 Derivation of the response of a single current cell




Figure B.1: The differential current switch.

In order to derive a relationship between the input voltage and output current assume that the two transistors of the current switch are matched and in saturation ($V_{ds} > V_{gs} - V_t$). The current flowing through each current cell using the first order model of MOSFET behaviour:

$$I_{out1} = \frac{\beta_1}{2} (V_{in} - V_s - V_{t1})^2 \quad (B.1)$$

and

$$I_{out2} = \frac{\beta_2}{2} (V_{ref} - V_s - V_{t2})^2 \quad (B.2)$$

where V_{in} is an input voltage, V_{ref} a constant reference voltage, V_s is the voltage at the source of the switch transistors, V_{t1} is the threshold voltage of M1, V_{t2} is the threshold voltage of M2, β_1 is the transconductance parameter of M1 and β_2 is the transconductance parameter of M2. I_{bias} is the current provided by the ideal constant current source :

$$I_{bias} = I_{out1} + I_{out2} \quad (B.3)$$

Assuming that the two switch transistors, M1 and M2, are well matched and have equal β and V_t then:

$$(V_{in} - V_{st})^2 + (V_{ref} - V_{st})^2 = \frac{2I_{bias}}{\beta} \quad (B.4)$$

where $V_{st} = V_s - V_t$ Expanding:

$$(V_{in}^2 - 2V_{in}V_{st} + V_{st}^2) + (V_{ref}^2 - 2V_{ref}V_{st} + V_{st}^2) = \frac{2I_{bias}}{\beta} \quad (B.5)$$

or:

$$2V_{st}^2 - 2V_{st}(V_{in} + V_{ref}) + V_{in}^2 + V_{ref}^2 - \frac{2I_{bias}}{\beta} = 0 \quad (B.6)$$

solve for V_{st} :

$$V_{st} = \frac{(V_{in} + V_{ref})}{2} \pm \sqrt{\frac{I_{bias}}{\beta} - \frac{(V_{in} - V_{ref})^2}{4}} \quad (B.7)$$

When $V_{in} = V_{ref}$ then $V_{st} = V_{in} \pm \sqrt{\frac{I_{bias}}{\beta}}$. V_{st} should be less than V_{in} and hence the negative part of the result is correct. To solve for I_{out2} substitute B.2 into B.7:

$$I_{out2} = \frac{\beta}{2} \left[\frac{(V_{ref} - V_{in})}{2} - \sqrt{\frac{I_{bias}}{\beta} - \frac{(V_{in} - V_{ref})^2}{4}} \right]^2 \quad (B.8)$$

Which reduces to $I_{out2} = I_{bias}/2$ when $V_{in} = V_{ref}$ as expected. The solution is complex when $\frac{I_{bias}}{\beta} - \frac{(V_{in} - V_{ref})^2}{4} < 0$. I.e. when $I_{bias} > \beta \frac{(V_{in} - V_{ref})^2}{4}$. This is when the current cell is turning fully on, ($I_{out2} \approx I_{bias}$) or fully off, ($I_{out2} \approx 0$) and is valid for both M1 and M2. The condition is just met when $(V_{in} - V_{ref})^2 = \frac{4I_{bias}}{\beta}$.

B.2 Derivations of the effect of variations

The variation in transconductance, β is represented as $\Delta\beta$ and the average is represented by $\bar{\beta}$ ($\beta_1 = \bar{\beta} + \Delta\beta$ and $\beta_2 = \bar{\beta} - \Delta\beta$). From Equation B.4:

$$(1 + \Delta\beta/\bar{\beta})(V_{in} - V_{st})^2 + (1 - \Delta\beta/\bar{\beta})(V_{ref} - V_{st})^2 = \frac{2I_{bias}}{\bar{\beta}} \quad (B.9)$$

where $V_{st} = V_s - V_t$ Expanding:

$$(1 + \Delta\beta/\bar{\beta})(V_{in}^2 - 2V_{in}V_{st} + V_{st}^2) + (1 - \Delta\beta/\bar{\beta})(V_{ref}^2 - 2V_{ref}V_{st} + V_{st}^2) = \frac{2I_{bias}}{\bar{\beta}} \quad (B.10)$$

or:

$$2V_{st}^2 - 2V_{st}((V_{in} + V_{ref}) + \Delta\beta/\bar{\beta}(V_{in} - V_{ref})) + (V_{in}^2 + V_{ref}^2) + \Delta\beta/\bar{\beta}(V_{in}^2 - V_{ref}^2) - \frac{2I_{bias}}{\bar{\beta}} = 0 \quad (B.11)$$

solve for V_{st} :

$$V_{st} = \frac{(V_{in} + V_{ref}) + \Delta\beta/\bar{\beta}(V_{in} - V_{ref})}{2} \pm \sqrt{\frac{I_{bias}}{\bar{\beta}} - \frac{(V_{in} - V_{ref})^2}{4}(1 - (\Delta\beta/\bar{\beta})^2)} \quad (B.12)$$

Following Equation B.8:

$$I_{out2} = \frac{\bar{\beta} + \Delta\beta}{2} \left[\frac{(V_{ref} - V_{in}) - \Delta\beta/\bar{\beta}(V_{in} - V_{ref})}{2} - \sqrt{\frac{I_{bias}}{\bar{\beta}} - \frac{(V_{in} - V_{ref})^2}{4}(1 - (\Delta\beta/\bar{\beta})^2)} \right]^2 \quad (B.13)$$

Which reduces Equation B.8 when $\Delta\beta = 0$.

From Equation B.1 and Equation B.2 variation in the threshold (V_t), reference (V_{ref}) or input voltage (V_{in}) can be expressed as ΔV_t as the output currents can be written:

$$I_{out1} = \frac{\beta}{2}(V_{in} - V_s - (\bar{V}_t + \Delta V_t))^2 \quad (B.14)$$

and

$$I_{out2} = \frac{\beta}{2}(V_{ref} - V_s - (\bar{V}_t - \Delta V_t))^2 \quad (B.15)$$

and hence Equation B.8 becomes:

$$I_{out} = \frac{\beta}{2} \left[\frac{(V_{ref} - V_{in})}{2} - \frac{\Delta V_t}{2} - \sqrt{\frac{I_{bias}}{\bar{\beta}} - \frac{(V_{in} - V_{ref} + \Delta V_t)^2}{4}} \right]^2 \quad (B.16)$$

This equation simplifies to Equation 2.3 when $\Delta V_t = 0$.

Appendix C

MATLAB Code

C.1 Single, four or six nonlinear current cells

```
%Programme to generate parameters for nonlinear interpolation
%change N_segments to 4 or 6 for four and six cells.
N_segments=1
L=8; %Accumulator phase bits
S=8; %ROM amplitude bits (not used)
W=8; %ROM phase bits
B=L-W; % phase truncated
Cycles_min=1;
Delta_cycles=1;
Nfreqs=1;
gt80=1:1:Nfreqs;
gt70=1:1:Nfreqs;
gt60=1:1:Nfreqs;
gt50=1:1:Nfreqs;
spur_pos=1:1:Nfreqs;
peak_pos=1:1:Nfreqs;
max_spur=1:1:Nfreqs;
Npoints=2^L;
N_lut=2^(W-1);
index=1:1:N_lut;
LUT=1:1:N_lut;
diff_LUT=1:1:N_lut;
Final_LUT=1:1:2*N_lut;
Whole_sine_wave=1:1:2*N_lut;
Full_target=1:1:2*N_lut;
cos_wave=cos(pi*index/(N_lut));
target_wave=1-cos_wave;
Delta_phi=pi/N_segments;
bias=1:1:N_segments;
gain=1:1:N_segments;
phi_ref=1:1:N_segments;
beta=1:1:N_segments;
```

```

current=1:1:N_segments;
for i=1:1:N_segments
    phi_ref(i)=(i-1)*Delta_phi
    bias(i)=1-cos(phi_ref(i));
    gain(i)=(cos(phi_ref(i))-cos(phi_ref(i)+Delta_phi))/Delta_phi;
    current(i)=(cos(phi_ref(i))-cos(phi_ref(i)+Delta_phi));
    beta(i)=0.45*4*current(i)/(Delta_phi*Delta_phi);
end
current=current*curr_scale/1;
dphi=pi/N_lut
for i=1:1:N_lut-1;
    phi=i*dphi;
    segment_index=fix(phi/Delta_phi)+1;
    LUT(i)=bias(segment_index)+gain(segment_index)*(phi-phi_ref(segment_index));
end %Linear
for i=1:1:N_lut-1;
    diff_LUT(i)=0;
    phi=i*dphi;
    for j=1:1:N_segments;
        phi_r=phi_ref(j)+Delta_phi/2;
        diff=phi-phi_r;
        diff2=diff*diff;
        arg=current(j)/beta(j)-diff2/4;
        test=current(j)/beta(j)-diff2/2;
        contribution=0;
        if test >= 0.00
            bracket=diff/2-sqrt(arg);
            contribution=current(j)-beta(j)*bracket*bracket/2;
        else
            if diff > 0
                contribution=current(j);
            else
                contribution=0;
            end
        end
        diff_LUT(i)=diff_LUT(i)+contribution;
    end %for j
end % for i
diff_LUT(N_lut)=diff_LUT(N_lut-1);
LUT=diff_LUT;
for i=1:1:2*N_lut-1;
    if i<= N_lut
        Final_LUT(i)=LUT(i);
    else
        Final_LUT(i)=LUT(2*N_lut-i);
    end
end
end

```

C.2 Segmented system(linear)

```
%Program for linear interpolation
N_segments = 64;
L=16;
S=10;
W=13;
B=L-W;
Cycles_min=1;
Delta_cycles=1;
Nfreqs=1
gt80=1:1:Nfreqs;
gt70=1:1:Nfreqs;
gt60=1:1:Nfreqs;
gt50=1:1:Nfreqs;
spur_pos=1:1:Nfreqs;
peak_pos=1:1:Nfreqs;
max_spur=1:1:Nfreqs;
Npoints=2^L;
N_lut=2^(W-1);
index=1:1:N_lut;
LUT=1:1:N_lut;
diff_LUT=1:1:N_lut;
Final_LUT=1:1:2*N_lut;
Whole_sine_wave=1:1:2*N_lut;
Full_target=1:1:2*N_lut;
cos_wave=cos(pi*index/(N_lut));
target_wave=1-cos_wave;
Delta_phi=pi/N_segments
bias=1:1:N_segments;
gain=1:1:N_segments;
phi_ref=1:1:N_segments;
beta=1:1:N_segments;
current=1:1:N_segments;
for i=1:1:N_segments
    phi_ref(i)=(i-1)*Delta_phi;
    bias(i)=1-cos(phi_ref(i));
    gain(i)=(cos(phi_ref(i))-cos(phi_ref(i)+Delta_phi))/Delta_phi;
    current(i)=(cos(phi_ref(i))-cos(phi_ref(i)+Delta_phi));
    beta(i)= 4*current(i)/(Delta_phi*Delta_phi);
end
dphi=pi/N_lut
for i=1:1:N_lut-1;
    phi=i*dphi;
    segment_index=fix(phi/Delta_phi)+1;
    LUT(i)=bias(segment_index)+gain(segment_index)*(phi-phi_ref(segment_index));
end %Linear
for j=1:1:N_segments;
    phi_r=phi_ref(j)+Delta_phi/2;
    phi=phi_ref(j);
```

```

diff=phi-phi_r;
diff2=diff*diff;
arg=current(j)/beta(j)-diff2/4;
bracket=diff/2-sqrt(arg);
cont_min(j)=current(j)-beta(j)*bracket*bracket/2;
phi=phi_ref(j)+Delta_phi;
diff=phi-phi_r;
diff2=diff*diff;
arg=current(j)/beta(j)-diff2/4;
bracket=diff/2-sqrt(arg);
cont_max(j)=current(j)-beta(j)*bracket*bracket/2;
end
for i=1:1:N_lut-1;
    diff_LUT(i)=0;
    phi=i*dphi;
    for j=1:1:N_segments;
        phi_r=phi_ref(j)+Delta_phi/2;
        diff=phi-phi_r;
        diff2=diff*diff;
        arg=current(j)/beta(j)-diff2/4;
        test=current(j)/beta(j)-diff2/2;
        contribution=0;
        if test >= 0.00
            bracket=diff/2-sqrt(arg);
            contribution=current(j)-beta(j)*bracket*bracket/2;
        else
            if diff > 0
                contribution=current(j);
            else
                contribution=0;
            end
        end
        if contribution > cont_max(j)
            contribution = cont_max(j);
        end
        if contribution < cont_min(j)
            contribution=cont_min(j);
        end
        diff_LUT(i)=diff_LUT(i)+contribution;
    end %for j
end %for i
diff_LUT(N_lut)=diff_LUT(N_lut-1);
LUT=diff_LUT;
for i=1:1:2*N_lut-1;
    if i<= N_lut
        Final_LUT(i)=LUT(i);
    else
        Final_LUT(i)=LUT(2*N_lut-i);
    end
end
Final_LUT(2*N_lut)=Final_LUT(2*N_lut-1);

```



```

t=1:1:Npoints;
for k=1:1:Nfreqs;
    N_cycles=Cycles_min+(k-1)*Delta_cycles;
    accum=N_cycles*t;
    accum=rem(accum,2^L);
    rom_address=fix(accum/2^B);
    output=rom_address;
    for i=1:Npoints;
        index=rom_address(i)+1;
        output(i)=Final_LUT(index)-mean(Final_LUT);
    end
    out_spec=fft(output);
    out_power=20*log10(abs(out_spec)/max(abs(out_spec)));

```

C.3 Generation of phase DAC output to program function generator

```

%Creates phase DAC output for Function Generator
clear all;
L=13; % L-bit accumulator
P=8; % P-bit DAC
PW=33 % Specify desired phase word
steps=2^L;
out=2^L;
out_xor=2^P; % before half wave symm.
out_xor2=2^P;
v1=steps %maximum voltage
v2=0 %minimum voltage (scale later)
stepsize=(v1-v2)/steps*PW
    for i=1:steps
        out(i)=mod(stepsize*i,v1-v2)+v2;
    end
out_xor=out; %allocate
k=0; %take half wave symmetry
for j=1:2^L
    bin=dec2bin(out_xor(j),L); % convert to binary
        for i=1:P %only take P amplitudebits
            if xor(bin2dec(bin(1)),bin2dec(bin(i)))
                bin_xor(i)='1';
            else bin_xor(i)='0';
            end
        end
    end
%bin_xor(1)=(bin(1));
out_xor2(j)=bin2dec(bin_xor);
end %convert back to decimal
v1=1; v2=-1;
scale=(v1-v2)/max(out_xor2)
output=scale*out_xor2-1;

```

C.4 Generation of four different phase outputs from one accumulator

```
%Shows that four outputs may be generated by a single accumulator and
%four nonlinear DAC (four cells)
clear all;
L=6; %L-bit accumulator
P=5; %P-bit phase DAC
PW=1 %phase word
steps=2^L;
out=2^L;
out_xor=2^P; % before half wave symm.
out_xor2=2^P;
v1=steps
v2=0
stepsize=(v1-v2)/steps*PW
    for i=1:steps
        out(i)=mod(stepsize*i,v1-v2)+v2;
    end
k=0; %take HWS
for j=1:2^L
    %convert to binary
    bin1=dec2bin(out(j),L); bin2=dec2bin(out(j),L);
    bin3=dec2bin(out(j),L); bin4=dec2bin(out(j),L);
    % shift by 90
    if bin1(2)=='0'
        bin2(2)=='1'; %flip second MSB
    else bin2(2)=='0';
        if bin1(1)=='1'
            bin2(1)=='0';
        else bin2(1)=='1';
        end %flip first MSB
    end; %flip second MSB
    %shift by 180
    if bin1(1)=='0'
        bin3(1)=='1';
    else bin3(1)=='0';
    end; %flip first MSB
    %take the 1st and shift by 270
    if bin1(1)=='0'
        bin4(1)=='1';
    else bin4(1)=='0';
    end; %flip first MSB
    if bin1(2)=='0'
        bin4(2)=='1'; %flip 2nd
    else bin4(2)=='0';
        if bin1(1)=='0'
            bin4(1)=='0';
        else bin4(1)=='1';
        end; %flip 1st MSB
    end
end
```

```

for i=1:P %only take P amplitude bits
    if xor(bin2dec(bin1(1)),bin2dec(bin1(i)))
        bin1_xor(i)='1';
    else bin1_xor(i)='0';
    end
    if xor(bin2dec(bin2(1)),bin2dec(bin2(i)))
        bin2_xor(i)='1';
    else bin2_xor(i)='0';
    end
    if xor(bin2dec(bin3(1)),bin2dec(bin3(i)))
        bin3_xor(i)='1';
    else bin3_xor(i)='0';
    end
    if xor(bin2dec(bin4(1)),bin2dec(bin4(i)))
        bin4_xor(i)='1';
    else bin4_xor(i)='0';
    end
end
out1(j)=bin2dec(bin1);out2(j)=bin2dec(bin2);
out3(j)=bin2dec(bin3);out4(j)=bin2dec(bin4);
out_xor1(j)=bin2dec(bin1_xor);out_xor2(j)=bin2dec(bin2_xor);
out_xor3(j)=bin2dec(bin3_xor);out_xor4(j)=bin2dec(bin4_xor);
end
figure;plot(out1);hold on;
plot(out2,'r'); plot(out3, 'k');plot(out4, 'g');hold off;
figure;hold on;
subplot(4,1,1);plot(out_xor1);
subplot(4,1,2);plot(out_xor2,'r');
subplot(4,1,3);plot(out_xor3, 'k');
subplot(4,1,4);plot(out_xor4, 'g');
hold off;

```

Appendix D

Verilog code for mixed signal simulation

D.1 13-bit accumulator

Simulates a 13-bit accumulator in Verilog

```
module stimulus;
wire clock, clock_bar ;
reg [0:12] B ; // make 13-bit
wire [0:11] A_xor ; // the accumulator output after bit flip
wire [0:12] A_out ; // accum output
wire [0:6] row; // row and column registers of them dec.
wire [0:6] col;
//assign clock=0;
initial
begin
    B = 13'b00000001000000;
    // Display a header
    $display("Ti 012345678 A_out  A col row clock clock_bar");
    // Prints the values anytime a value of A, B or C changes
    $monitor("%0d %b %b %b %b %b %b %b %b", $time, A_out[
0], A_out[1:6], ~A_xor[0:2], ~A_xor[3:5] ,col,row, clock, clock_bar);
    // after 10 simulation units increaseB
    // #10 B = 256;
    // #20 B = 20'b10000000000000;
    // #1 B = 20'b00001000000000;
    // Will stop the execution after another 20 simulation units.
    #60; $finish;
end
//main_process
// call accumulator
accum4 accum1_4(A_out, B);
//call xor
flip flip_1(A_xor, A_out);
//call thermometer decoder
therm therm_1(A_xor[0:5], col, row);
```

```

//call clock
clock clock_1(clock, clock_bar);
endmodule
module flip(out, A_out);
input [0:12] A_out;
output [0:11] out;
reg [0:11] out;
always #1 begin
if (A_out[0]) out[0:11] = A_out[1:12] ;
else out[0:11] = ~A_out[1:12];
// flip A if MSB is low as we want A_xor active low
end
endmodule
module accum4(A, B);
output [0:12] A;
input [0:12] B;
reg [0:12] A;
initial A = 0;
always #1 A = A + B;
endmodule
module therm(A, col, row);
output [0:6] col;
output [0:6] row;
reg [0:6] col;
reg [0:6] row;
input [0:5] A;
initial row=0;
initial col=0;
always
#1 begin
//thermometer decoder
//row[0] = 1;
row[0] = ~(~A[0] & ~A[1] & ~A[2]);
row[1] = ~(~A[0] & ~A[1]);
row[2] = ~(~A[0] & ~(A[1] & A[2]));
row[3] = A[0];
row[4] = A[0] & ~(~A[1] & ~A[2]);
row[5] = A[0] & A[1];
row[6] = A[0] & A[1] & A[2];
//row[8] = 0;
col[0] = ~(~A[3] & ~A[4] & ~A[5]);
col[1] = ~(~A[3] & ~A[4]);
col[2] = ~(~A[3] & ~(A[4] & A[5]));
col[3] = A[3];
col[4] = A[3] & ~(~A[4] & ~A[5]);
col[5] = A[3] & A[4];
col[6] = A[3] & A[4] & A[5];
//col[7] = 0;
end
endmodule
module clock(clock, clock_bar);

```

```

output clock, clock_bar;
reg clock, clock_bar;
initial clock=0;
always #1 clock = ~clock;
always #1 clock_bar = ~clock;
endmodule

```

D.2 Conventional thermometer decoder

```

module stimulus;
reg coljj, rowi, rowii;
wire En;
initial
begin
    rowi = 0; rowii=0; coljj=0;
    $display("Ti rowi rowl+1 colj+1 En");
    // Prints the values anytime a value of A, B or C changes
    $monitor("%0d %b %b %b %b", $time, rowi, rowii, coljj, En);
    // after 10 simulation units increaseB
    #1 coljj =1;
    #1 rowi=1;
    #1 coljj=0;
    #1 rowii=1;
    #1 rowii=1;
    #10 $finish;
end
//call clock
local local_1(rowi, rowii, coljj, En);
endmodule
module local(rowi, rowii, coljj, En);
output En;
reg En;
input rowi, rowii, coljj;
always #1 En = ~(~(rowi&coljj) & ~(rowi & rowii));
endmodule

```

D.3 Three state thermometer decoder

```

module stimulus;
reg colj, coljj, rowi, rowii;
wire En;
initial
begin
    rowi = 0; rowii=0; coljj=0; colj=0
    $display("Ti rowi rowl+1 colj colj+1 En");
    // Prints the values anytime a value of A, B or C changes

```



```

        $monitor("%0d  %b %b %b %b %b", $time, rowi, rowii, col
j, coljj, En);
        // after 10 simulation units increaseB
        #1 coljj =1;
        #1 colj =1;
        #1 rowi=1;
        #1 rowii=1;
        #1 coljj=0;
        #1 rowii=0;
        #1 rowi=0;
        #10 $finish;
end
//call clock
local local_1(rowi, rowii, colj, coljj, En);
endmodule
module local(rowi, rowii, colj, coljj, En);
output En;
reg  En;
input  rowi, rowii, colj, coljj;
always #1 En =  (rowi ^ rowii) & (colj ^ coljj);
endmodule

```

Bibliography

- [1] C. Enz, F. Krummenacher, E. Vittoz. An analytical mos transistor model valid in all regions of operation and dedicated to low-voltage and low-current applications. *Journal on Analog Integrated Circuits and Signal Processing*, pages 83–114, 1995.
- [2] J. Vankka, M. Waltari, M. Koseunen, and K. A. I. Halonen. A Direct Digital Synthesizer with an On-Chip D/A-Converter. *IEEE J. Solid-State Circuits*, 33(2):218–227, February 1998.
- [3] V. F. Kroupa. *Direct Digital Frequency Synthesizers*. Wiley-IEEE Press, 1998.
- [4] W. Helbach and H. Schollmeyer. Impedance measuring methods based on multiple digital generators. *IEEE Transactions on Instrumentation and Measurement*, IM-36:400–405, June 1987.
- [5] L. Callegaro and G. Galzerano and C. Svelto. A Multiphase Direct-Digital-Synthesis Sinewave Generator for High-Accuracy Impedance Comparison. *IEEE Transactions on Instrumentation and Measurement*, 50(4):926–929, August 2001.
- [6] F.F. Becker and X.B. Wang and Y. Huang and R. Pethig and J. Vykoukal and P.R.C. Gascoune. The removal of human leukaemia cells from blood using interdigitated microelectrodes. *Journal of Applied Physics D*, 27:2659–2662, 1994.
- [7] Y. Li, K. V. I. S. Kaler. Dielectrophoretic fluidic cell fractionation system. *Analytica Chimica Acta*, pages 151–161, 2004.
- [8] J. Mitola. The Software Radio Architecture. *IEEE Communications Magazine*, pages 26–38, May 1995.
- [9] J. Tierney, C. M. Rader and B. Gold. A Digital Frequency Synthesizer. *IEEE Trans. Audio Electroacoust*, AU-19:48–57, 1971.
- [10] Analog Devices. A technical tutorial on digital signal synthesis,. www.analog.com, 1999.
- [11] E. M. Mattison and L. M. L. Coyle. Phase Noise in Direct digital Synthesizers. *Proceedings of the IEEE International Frequency Control Symposium*, pages 352–356, 1988.
- [12] T. H. Lee. *The Design of CMOS Radio-Frequency Integrated Circuits*. Cambridge University Press, 1998.
- [13] A. Madisetti, A. Y. Kwentus, and A. N. Wilson. A 100MHz, 16-b Direct Digital Frequency Synthesizer with a 100dBc Spurious Free Dynamic Range. *IEEE Journal of Solid State Circuits*, 34(8):1034–1043, August 1999.
- [14] Analog Devices. AD7008 DataSheet RevB, 2001.
- [15] Analog Devices. AD9852 DataSheet, 2001.

- [16] A. Rofougaran, G. Chang, J. J. Rael, J. Y.C. Chang, M. Rofougaran, P. J. Chang, M. Djafari, M.-K. Ku, E. W. Roth, A. A. Abidi and H. Samueli. A Single-Chip 900-MHz Spread-Spectrum Wireless Transceiver in 1- μ m CMOS—Part I: Architecture and Transmitter Design. *IEEE J. Solid-State Circuits*, 33(4):515–534, April 1998.
- [17] A. Bateman, D. Haines and R. Wilkinson. Direct Conversion Linear Transceiver Design. In *5th International Conference on Mobile Radio and Personal Communications*, pages 53–56. IEE, 1989.
- [18] A. Bateman and D. M. Haines. Direct conversion transceiver Design for Compact Low-Cost Portable Mobile Radio Terminals. In *39th IEEE Vehicular Technology Conference*, volume 1, pages 57–62, 1989.
- [19] P. B. Kenington. Emerging Technologies for Software Radio. *Electronics and Communication Engineering Journal*, 11(2):69–83, April 1999.
- [20] J. Jiang and E. Lee. A ROM-less Direct Digital Frequency Synthesizer Using Segmented Nonlinear Digital-to-Analog Converter. In *IEEE Custom Integrated Circuits Conference*, pages 165–168, 2001.
- [21] A. Bellaouar, M. O’brecht, A. Fahim, and M. Elmasry. Low-Power Direct Digital Frequency Synthesis for Wireless Communications. *IEEE J. Solid-State Circuits*, 35(3):385–390, March 2000.
- [22] J.M.P. Langlois and D. Al-Khalili. Novel Approach to the Design of Direct Digital Frequency Synthesizers based on Linear Interpolation. *IEEE Transactions on Circuits and Systems II*, 50(9):567–578, September 2003.
- [23] D. Sunderland, R. Strauch, Wharfield S., H Peterson, and C. Cole. CMOS/SOS frequency synthesizer LSI circuit for spread spectrum communications,. *IEEE Journal of Solid State Circuits* , 19:497–505, August 1984.
- [24] H. T. Nicholas and H. Samueli. A 150-MHz Direct Digital Frequency Synthesizer in 1.25- μ m CMOS with -90-dBc Spurious Performance. *IEEE J. Solid-State Circuits*, 26(12):1959–1969, December 1991.
- [25] Bluetooth special interest group (from IEEE 802.11 working group). Bluetooth core specification version 1.2. In www.bluetooth.org. www.bluetooth.org, 2004.
- [26] J.C. Haartsen and S. Mattisson. Bluetooth-a new low-power radio interface providing short-range connectivity. *Proceedings of the IEEE*, 88(10):1651–1661, 2000.
- [27] ETSI 3rd Generation Partnership Project (3GPP). Umts is the european name for imt-2000, the international telecommunication union’s (itu) framework for third-generation standards. www.3gpp.org, 2004.
- [28] J. P. Castro. *The UMTS network and radio access technology: air interface techniques for future wireless systems*. Wiley, Chichester, NY, 2001.
- [29] L. Callegaro, and V. D’Elia. A synchronized two-phase sinewave generator for ac metrology system compensations. *IEEE Transactions on Instrumentation and Measurement*, 49(2):320–324, 2000.
- [30] G. Ramm and H. Moser. From the calculable ac resistor to capacitor dissipation factor determination on the basis of time constants. *IEEE Transactions on Instrumentation and Measurement*, 50(2):286–289, 2001.

- [31] T. Tran, D. R. Oliver, D.J. Thomson and G. E. Bridges. Zeptofraad resolution capacitance sensor for scanning capacitance microscopy. *Review of Scientific Instruments*, 72(6):2618–2623, 2001.
- [32] J. N. Lygouras, K. N. Tarchanidis and P. G. Tsalides. A digital optical fibre finger blood volume meter. *IEEE Instrumentation and Measurement Technology Conference*, pages 648–653, 2001.
- [33] K. N. Tarchanidis, J. N. Lygouras. Data glove with a force sensor. *IEEE Instrumentation and Measurement Technology Conference*, pages 380–385, 2001.
- [34] H. Wu, L. P. Yang, T. H. Chen, S. T. Young and T. S. Kuo. A low cost architecture of audiometer based on direct digital synthesizer. *Proceedings of the first joint BMES-EMBS Conference*, 1999.
- [35] J. D. Sternhagen, K. Mitzner, E. Berkenpas, M. Karlgaard, C. E. Wold and D. W. Galipeau. A direct digital synthesis system for acoustic wave sensors. *IEEE Sensors Journal*, 2(4):288–293, 2002.
- [36] N. G. Green and H. Morgan. Dielectrophoretic manipulation of rod-shaped viral particles. *J. Electrostatics*, 42:279–293, 1997.
- [37] R. Pethig, J. P. H. Burt, A. Parton, N. Rizvi, M. S. Talary and J. A. Tame. Development of biofactory-on-a-chip technology using excimer laser micromachining. *J. Micromech. and Microeng*, 8:57–63, 1998.
- [38] R. H. Bayford A. T. Tidswell, A. Gibson and D. S. Holder. Three-dimensional electrical impedance tomography of human brain activity. *Neuroimage*, 13:283–294, 2001.
- [39] T. York. Status of electrical impedance tomography in industrial applications. *Journal of Electronic Imaging*, 10(3):608–619, 2001.
- [40] B. H. Brown. Medical impedance tomography and process impedance tomography: a brief review. *Measurement Science and Technology*, (12):991–996, 2001.
- [41] E. de Angel and E. E. Swartzlander, Jr. Survey of low-power techniques for roms. *Proc. Int. Symp. Low Power Electronics and Design*, pages 7–11, 1997.
- [42] B.D. Yang and L.S. Kim. A low-power rom using charge recycling and charge sharing techniques. *IEEE Journal of Solid State Circuits*, pages 641–653, 2003.
- [43] Z. Papay. Numerical Distortion in Single Tone DDS. *Proceedings of IEEE Instrumentation and Measurement Technology Conference (IMTC 2001)*, pages 720–724, May 2001.
- [44] D. A. Johns and K. Martin. *Analog Integrated Circuit Design*. Wiley, 1997.
- [45] T. Miki and Y. Nakamura and M. Nakaya and S. Asai and Y. Akasaka and Y. Horiba. An 80-MHz 8-bit CMOS D/A Converter. *IEEE J. Solid-State Circuits*, 21(6):983–988, June 1986.
- [46] A. R. Bugeja and Bang-Sup Song. A self-trimming 14-b 100-MS/s CMOS DAC. *IEEE J. Solid-State Circuits*, 35(12):1841–1852, December 2000.
- [47] J. Jacob Wikner and Nianxiong Tan. Modeling of CMOS Digital- to- Analog Converters for Telecommunications. *Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions*, 46(5):489–499, May 1999.
- [48] A. Van den Bosch and M. A. F. Borreman. A 10-bit 1-GSample/s Nyquist Current Steering CMOS D/A converter. *IEEE J. Solid-State Circuits*, 36(3):315–324, March 2001.
- [49] A. Hastings. *The Art of Analog Layout*. Pretice Hall, 2001.

- [50] P. Horowitz and W. Hill. *The Art of Electronics*. Cambridge University Press, 1980.
- [51] L.E. Larson. *RF and Microwave Circuit Design for Wireless Communications*. Artech House., INC., 1997.
- [52] B.H. Hutchinson. *Frequency Synthesis and Applications*. IEEE Press, New York, 1975.
- [53] H. T. Nicholas and H. Samueli. An Analysis of the Output Spectrum of Direct Digital Frequency Synthesizers in the Presence of Phase-Accumulator Truncation. pages 495–502. 41st Annual Frequency Control Symposium, IEEE, 1987.
- [54] T. Tsukahara A. Yamagishi, M. Ishikawa and S. Date. A 2-V, 2-GHz Low-Power Direct Digital Frequency Synthesizer Chip-Set for Wireless Communications. *IEEE J. Solid-State Circuits*, 33(2):210–217, February 1998.
- [55] A. M. Sodagar and G. R. Lahiji. Mapping From Phase to Sine-Amplitude in Direct Digital Frequency Synthesizers Using Parabolic Approximation. *Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions*, 47(12):1452–1457, December 2000.
- [56] J.M.P. Langlois and D. Al-Khalili. Hardware optimized direct digital frequency synthesizer with 60dBc spectral purity. *IEEE International Symposium Circuits and Systems*, pages 361–364, May 2002.
- [57] B. Simmons. Mathwords: Terms and formulas from algebra to calculus. www.mathwords.com, June 2004.
- [58] Qualcomm 2368. Synthesizer data book, August 1997.
- [59] S. Morteza pour and E. Lee. Design of a Low-Power ROM-Less Direct Digital Frequency Synthesizer Using Nonlinear Digital-to-Analog Converter. *IEEE J. Solid-State Circuits*, 34(10):1350–1359, October 1999.
- [60] A. N. Mohieldin, A. A. Emira and S.E. Sanchez. A 100mhz 8mw rom-less quadrature direct digital frequency synthesizer. *IEEE Journal of Solid State Circuits*, 31(10):1235–1243, 2002.
- [61] M.J. Pelgrom and A.C.J. Duinmaijer and A.P.G. Welbers. Matching Properties of MOS transistors. *IEEE J. Solid-State Circuits*, 24(5):1433–1440, October 1989.
- [62] S. Shah. *Floating Gate Based Trimmable Current Sources*. Oxford University DPhil Thesis, 2002.
- [63] Foty D. P. *MOSFET MODELLING with SPICE. Principles and Practice*. Prentice Hall PTR, 1997.
- [64] S. C. Terry, J. M. Rochelle, D. M. Binkley, B. J. Blalock, D. P. Foty, and M. Bucher . Comparison of a bsim3v3 and ekv mosfet model for a 0.5 um cmos process and implications for analog circuit design. *IEEE TRANSACTIONS ON NUCLEAR SCIENCE*, 50(4), 2003.
- [65] P.G. Drennan and C.C. McAndrew. Understanding MOSFET Mismatch for Analog Design. *IEEE Journal of Solid State Circuits*, 38(3):450–456, March 2003.
- [66] A. Sangiovanni-Vincentelli K. Kundert, J. White. *Steady-State Methods for Simulating Analog and Microwave Circuits*. Kluwer Academic Publishers, Boston, 1990.
- [67] A. McEwan and S. Collins. A Direct Digital Frequency Synthesis System for Low Power Communications . *Proceedings European Solid State Circuits Conference, 2003*, 2003.
- [68] A. McEwan and M. Chappel. A Low Power, High Speed Accumulator for DDFS Applications . *Circuits and Systems, Proceedings of 2003 IEEE International Symposium on*, 1:797–800, 2004.

- [69] M. Frigo and S. G. Johnson. FFTW: An adaptive software architecture for the FFT. In *Proc. 1998 IEEE Intl. Conf. Acoustics Speech and Signal Processing*, volume 3, pages 1381–1384. IEEE, 1998.
- [70] U. Schaper, C. G. Linnenbank, and R. Thewes. Precise characterization of long-distance mismatch of cmos devices. *IEEE Transactions on Semiconductor Manufacturing*, 14(4):331–317, Nov. 2001.
- [71] A. McEwan and S. Collins. Analogue Interpolation Based Direct Digital Frequency Synthesis . *Circuits and Systems, Proceedings of 2003 IEEE International Symposium on*, 1:621–624, 2003.
- [72] F. Lu, H., Samueli, J., Yuan, C., Svensson. A 700-MHz 24-b pipelined accumulator in 1.2- μ m CMOS for application as a numerically controlled oscillator. *IEEE J. Solid-State Circuits*, 28(8):878–886, Aug. 1993.
- [73] M.Thompson. Low-Latency, High-Speed Numerically Controlled Oscillator Using Progression-of-States Technique. *IEEE J. Solid-State Circuits*, 27(1):113–117, Jan. 1992.
- [74] B.Yang, L.Kim and H.Yu. A High Speed Direct Digital Frequency Synthesizer Using a Low Power Pipelined Parallel Accumulator. *Circuits and Systems, Proceedings of 2002 IEEE International Symposium on*, pages 373–376, 2002.
- [75] A., H., Shams and M., A., Bayoumi. A Novel High-Performance CMOS 1-Bit Full-Adder Cell. *IEEE Transactions on Circuits and Systems - II*, 47(5):478–481, May. 2000.
- [76] R.Zimmermann and W.Fichtner. Low-Power Logic Styles: CMOS Versus Pass-Transistor Logic. *IEEE J. Solid-State Circuits*, 32(7):1079–1089, July 1997.
- [77] I.Abu-Khater, A.Bellaouar and M. Elmasry. Circuit Techniques for CMOS Low-Power High-Performance Multipliers. *IEEE J. Solid-State Circuits*, 31(10):1535–1546, Oct. 1996.
- [78] N.H.E. Weste and K. Eshraghian. *Principles of CMOS VLSI design*. Addison-Wesley Pub Co; 2nd edition, 1994.

