

Copyright is owned by the Author of the thesis. Permission is given for a copy to be downloaded by an individual for the purpose of research and private study only. The thesis may not be reproduced elsewhere without the permission of the Author.

**ACCURATE THERMAL SENSING WITH MODERN
CMOS INTEGRATED CIRCUITS**

A THESIS PRESENTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE

OF
DOCTOR OF PHILOSOPHY
IN
ENGINEERING

AT
MASSEY UNIVERSITY
AUCKLAND
NEW ZEALAND

ROBERT PATRICK FISK
B. ENG.

2010

ABSTRACT

Digital control systems can be found performing a wide range of duties throughout modern society. These systems demand accurate, low cost interfaces to physical parameters of interest, one of the most common being temperature. A ‘smart’ sensor takes advantage of modern integrated circuit technology to create a sensor and analog-to-digital converter on the same silicon chip. Smart temperature sensors are widely available offering simple digital interfaces, high reliability, low power consumption and low cost. The primary weakness of these devices is the low inherent accuracy of on-chip thermal sensors.

This thesis presents a smart thermal sensor design that improves upon current technology by employing a modern 0.13 μm CMOS process and circuit-level techniques to reduce sensor size and power consumption while increasing digital converter resolution. Data is presented that shows uncalibrated sensor accuracy can be increased by using correlated device characteristics to compensate for random inter-device variation. The research findings guide the construction of future smart thermal sensors with uncalibrated accuracy levels exceeding that of any currently available design.

ACKNOWLEDGEMENTS

I would firstly like to thank my chief supervisor Dr. Rezaul Hasan for his endless encouragement, and supervisor Dr. Tom Moir for timely and friendly advice. Thanks also goes to the staff and students at the School of Engineering for their friendship and words of encouragement over the past years. The support of the MOSIS academic research program is gratefully acknowledged for covering the cost of prototype fabrication. This research also would not have been possible without the support of Massey University in the form of a three-year Doctoral Scholarship.

I would like to pay respect to the memory of my parents, who made every effort to give me the academic foundation necessary to reach this point. And ultimately, I would like to thank the founder of Falun Dafa, Mr Li Hongzhi, for teaching me the Chinese proverb “After passing the shady willow trees, there will be bright flowers and another village ahead!” ☸

DECLARATION

The author declares that this is his own work except where due acknowledgement has been given. It is being submitted for the PhD in Engineering to Massey University, New Zealand.

This thesis describes the research carried out by the author at the School of Engineering, Massey University, New Zealand from February 2005 to January 2010, supervised by Dr. Rezaul Hasan and Dr. Tom Moir.

Candidate's signature:

Robert Patrick Fisk

10th March 2010

TABLE OF CONTENTS

Abstract	ii
Acknowledgements	iii
Declaration	iv
Table of Contents	v
Table of Figures	ix
Table of Tables	xv
Definitions and Abbreviations	xvi
1 Introduction	1
1.1 Research Objectives	3
1.2 Contributions to Knowledge.....	4
1.3 Structure of the Thesis	4
2 Literature Review	5
2.1 On-Chip Sensing Elements.....	5
2.2 History of Bipolar IC Temperature Sensors	11
2.3 Curvature Compensation for Bipolar Sensors	13
2.4 History of CMOS IC Temperature Sensors.....	17
2.5 Analog Techniques for Accurate CMOS Circuitry	22
Chopping	22
Dynamic Element Matching	24
2.6 History of Smart Temperature Sensors	25
2.7 Optimal Filtering for Incremental Delta-Sigma ADCs	32
2.8 Discrete Techniques for Accurate $\Delta\Sigma$ Modulators	34
Switched-Capacitor Integrators.....	35
Correlated Double-Sampling	38
2.9 A Case Study	40
2.10 Process Compensation	43
Anatomy of Bipolar Process Sensitivity	44
Correlated Device Characteristics.....	45
2.11 Summary	46
3 Design Methodology	48
3.1 Design Evaluation Criteria	48

3.2	Desirable Design Characteristics.....	50
3.3	High-Level System Design.....	52
3.4	Bipolar Core Features	56
3.5	PTAT Bias Generator Features	58
3.6	Process Compensation Features	62
3.7	Delta-Sigma Modulator Features	64
	‘Implicit Reference’ Input.....	64
	$\Delta\Sigma$ Modulator Loop Structure.....	66
	Accurate Integration.....	68
3.8	Digital Filter Features	70
3.9	Summary.....	71
4	Design Implementation	73
4.1	Operational Amplifiers	74
	PMOS Well Connection.....	75
	Single-Ended Folded-Cascode Op-amp.....	76
	Differential Folded-Cascode Op-amp	78
	Common-Mode Feedback.....	78
	Gain-Boosted Differential Folded-Cascode Op-amp.....	80
	Amplifier Bias Voltage Generator	82
4.2	PTAT Bias Current Generator	83
	Bias Generator Control Logic	84
4.3	Bipolar Core	85
	Bipolar Core Control Logic	89
4.4	Process Compensation.....	90
	Implementable Compensation Circuit	94
4.5	$\Delta\Sigma$ Modulator	97
	Loop Structure.....	97
	Noise Transfer Function.....	98
	Capacitor Values	99
	Comparator.....	103
	ΔV_{BE} Gain Factor	104
	$\Delta\Sigma$ Modulator Control Logic	106
	Clock Frequency Analysis	108
4.6	Digital Decimation Filter.....	108

4.7	Layout.....	112
	Chip I/O.....	113
	Power and ESD Considerations	114
4.8	Software Tools.....	116
5	Experimental Setup & Results	117
5.1	Testing Methodology.....	117
5.2	Electrical Testbed	118
	Data Capture.....	118
	Sensor Operational Control.....	119
	Precision Current Source.....	120
	External Process-Compensation.....	121
	Testbed Power Supplies	123
5.3	Thermal Testbed.....	124
5.4	Functional Verification.....	128
5.5	Numerical Results	132
6	Analysis & Discussion	135
6.1	Sensor Output Linearity	135
6.2	Sources of Inaccuracy.....	137
	High-Temperature Junction Leakage.....	138
	Process Parameter Variations.....	138
	Residual Curvature.....	139
	Measurement Error.....	140
	Summary of Error Contributions	143
6.3	Pinch-Base Correlation with V_{BE}	143
	V_{BE} Characterisation	144
	R_P Characterisation	145
	Correlation Analysis.....	146
6.4	Pinch-Base Correlation with Sensor Output.....	148
6.5	Discussions	150
	Evaluation of Typical Smart Sensor Features	151
	Evaluation of High-Accuracy Analog Techniques	152
	Evaluation of Process Compensation.....	153
7	Conclusions	154
7.1	Summary of Research.....	154

Objective 1	154
Objective 2	155
Objective 3	156
7.2 Recommendations for High-Performance Smart Sensor Designs.....	156
7.3 Opportunities for Further Research	158
References	159
Appendices	I
Conference Proceedings I	II
Conference Proceedings II	VII
Conference Proceedings III.....	XIII

TABLE OF FIGURES

Fig 2.1	Temperature dependence of polysilicon resistors with different doping levels [Rasmussen, 1994].	6
Fig. 2.2a	Fig. 2.2b Simplified cross-section of PN diode junction (2.2a) and schematic symbol (2.2b)	6
Fig. 2.3a	Fig. 2.3b Simplified cross-section of an NPN bipolar transistor (2.3a), and equivalent schematic (2.3b)	7
Fig. 2.4	A PTAT voltage produced from the difference of two V_{BE} , and the derived bandgap reference voltage V_{REF} .	9
Fig. 2.5	Techniques for generating a bandgap reference and an intrinsically referenced temperature sensor [Rasmussen, 1994].	10
Fig. 2.6	Early bandgap-reference generator circuit [Widlar, 1971].	11
Fig. 2.7	An improved bipolar bandgap reference generator circuit [Brokaw, 1974].	12
Fig. 2.8	Residual curvature in V_{BE} with various bias current temperature dependencies [Rasmussen, 1994].	14
Fig. 2.9	Circuit technique to produce a linearised V_{BE} and bandgap reference [Meijer et al., 1982].	14
Fig. 2.10	Difference in V_{BE} curvature when biased with constant (CTAT) and linearly temperature-dependent (PTAT) currents	15
Fig. 2.11	Comparison of curvature compensated and non-compensated bandgap references [Meijer et al., 1982].	16
Fig. 2.12a	Fig. 2.12b Simplified cross-section of an NMOS transistor (2.12a), and equivalent schematic (2.12b).	18
Fig. 2.13	Simplified cross-section of a lateral PNP transistor on a CMOS process.	19
Fig. 2.14	Simplified cross-section of a substrate PNP transistor on a CMOS process.	20
Fig. 2.15a	Fig. 2.15b Current mirror (2.15a) and opamp-based (2.15b) CMOS bandgap references [Razavi, 2001]	21
Fig. 2.16	Implementation of a chopped differential op-amp and the associated frequency spectrum at various circuit nodes.	22
Fig. 2.17	Implementation of the chopping block of Fig. 2.16 using MOS switches.	23

Fig. 2.18 Chopping implemented on a simple CMOS op-amp with single-ended output.	23
Fig. 2.19a Fig. 2.19b A ratioed current source (2.19a), and accurate circuit-level implementation (2.19b).....	24
Fig. 2.20 Dynamic element matching swaps the single unit source between all possible sources.....	24
Fig. 2.21 Block diagram of early smart temperature sensor [D. van Maaren et al., 1982].	26
Fig. 2.22 An improved smart sensor with frequency output [Meijer et al., 1988].....	26
Fig. 2.23 A simple one-bit first-order delta-sigma ADC.	27
Fig. 2.24 A discrete z-domain model of the first-order $\Delta\Sigma$ modulator in Fig. 2.23.....	28
Fig. 2.25 Spectrum of the input signal and shaped quantisation noise of a typical $\Delta\Sigma$ modulator [Schreier and Temes, 2005, Chapter 1].	29
Fig. 2.26 First-order sigma-delta modulator implemented on bipolar technology [Riedijk and Huijsing, 1991].....	30
Fig. 2.27a Fig. 2.27b Simplified ΔV_{BE} (2.27a) and V_{BE} (2.27b) generator circuits (bias current generator not shown) [Pertjjs, Niederkorn et al., 2005].....	31
Fig. 2.28 The second-order $\Delta\Sigma$ modulator used by [Pertjjs, Niederkorn et al., 2005], showing the connection to the analog bias currents of Fig. 2.27.	31
Fig. 2.29 Block diagram of a second-order symmetric sinc decimation filter.....	32
Fig. 2.30 Impulse response of first- and second-order digital filters with filter length L $= 100$	33
Fig. 2.31a Fig. 2.31b Simple switched-capacitor integrator (2.31a) and switch clock phases (2.31b)	35
Fig. 2.32 Parasitic capacitances present at node A affect the accuracy of the simple SC integrator.	36
Fig. 2.33 Parasitic-insensitive SC integrator.....	36
Fig. 2.34 Charge injection from a MOS switch caused by channel charge and parasitic gate capacitances.....	37
Fig. 2.35a Fig. 2.35b Parasitic-insensitive SC integrator modified to reduce switch charge injection (2.35a), and clock waveforms (2.35b).....	38
Fig. 2.36 Differential charge injection-insensitive parasitic-insensitive SC integrator.	38
Fig. 2.37 A parasitic-insensitive SC integrator with correlated double-sampling.	39
Fig. 2.38 Block diagram of the smart sensor of [Pertjjs, Makinwa et al., 2005].	40

Fig. 2.39 Bipolar sensor core for an accurate temperature sensor [Pertijs, Makinwa et al., 2005].	41
Fig. 2.40 A second-order differential switched-capacitor delta-sigma modulator [Pertijs, Makinwa et al., 2005]	43
Figure 2.41a Figure 2.41b Cross-section of a bipolar transistor (2.41a) and base pinch resistor (2.41b) in a bipolar manufacturing process [Hastings, 2006]	45
Fig. 2.42 Conceptual circuit diagram of V_{BE} process-compensation circuit [Amador et al., 1998].....	46
Fig. 3.1 Block diagram of the proposed smart temperature sensor.....	53
Fig. 3.2a Fig. 3.2b Conceptual circuits to generate V_{BE} (3.2a) and ΔV_{BE} (3.2b)	57
Fig. 3.3 Circuit for generating either V_{BE} or ΔV_{BE}	57
Fig. 3.4 Simplified diagram of PTAT bias current generator circuit.	59
Fig. 3.5 Bias generator with chopping to reduce offset errors.	61
Fig. 3.6 Complete PTAT bias generator with chopping and DEM.....	62
Fig. 3.7 Conceptual operation of process compensation circuit.	63
Fig. 3.8 Complete circuit diagram of process compensation scheme.	64
Fig. 3.9 Summing junction at input of a generic delta-sigma modulator	65
Fig. 3.10 Multiplexer to replace the summing junction of Fig 3.9.	65
Fig. 3.11 Cascade of Integrators, FeedBack (CIFB) delta-sigma modulator structure..	66
Fig. 3.12 Cascade of Integrators, FeedForward (CIFF) delta-sigma modulator structure.	67
Fig. 3.13 CDS switched-capacitor integrator [Pertijs, Makinwa et al., 2005].	68
Fig. 3.14 Parasitic capacitances of a MOS switch.	69
Fig. 3.15 Improved CDS integrator with negligible charge leakage.....	69
Fig. 3.16 Complete switched-capacitor modulator employing CIFF structure.....	70
Fig. 3.17 Block diagram of a third-order sinc digital filter.....	71
Fig. 4.1 Block diagram of the proposed smart temperature sensor.....	73
Fig. 4.2a Fig. 4.2b Conceptual output stage of a transconductance op-amp (4.2a) and transistor-level implementation (4.2b)	75
Fig. 4.3 Overdrive and bias voltages of the standard current-source stack.....	75
Fig. 4.4 Folded-cascode amplifier used in the bias generator circuit.....	77
Fig. 4.5 Simulated AC magnitude (top) and phase (bottom) of the amplifier in Fig. 4.4.	77
Fig. 4.6 Differential folded-cascode amplifier used in the $\Delta\Sigma$ modulator.	78

Fig. 4.7 Conceptual operation of SC-based CM feedback for differential op-amps.....	79
Fig. 4.8 Alternative analog switch implementations using MOS devices.	79
Fig. 4.9 Final SC-based CM control block for a differential op-amp.....	80
Fig. 4.10 Gain-boosted differential folded-cascode op-amp.....	81
Fig. 4.11a Fig. 4.11b Auxiliary amplifiers A_1 (4.11a) and A_2 (4.11b) used in the gain-boosted op-amp of Fig. 4.10	82
Fig. 4.12 A simple constant- g_m bias generator circuit.	82
Fig. 4.13 Complete constant- g_m bias generator circuit.	83
Fig. 4.14 Complete PTAT bias generator with chopping and DEM.....	84
Fig. 4.15 Control logic block for the PTAT bias generator circuit.....	85
Fig. 4.16 Bipolar core generating either V_{BE} or ΔV_{BE}	86
Fig. 4.17 Bandgap reference V_{REF} vs. temperature, produced by the bipolar core with α $= 16$ and ΔV_{BE} bias ratio = 4:1.....	88
Fig. 4.18 Deviation from linear fit in D_{OUT} when using V_{REF} from Fig. 4.17.	88
Fig. 4.19 Maximum deviation from linear fit of D_{OUT} from Fig. 4.18.....	89
Fig. 4.20 Control logic block for the bipolar sensor core.	90
Fig. 4.21 Conceptual operation of process compensation circuit.	91
Fig. 4.22 The desired relationship between I_S and B , and the actual relationship implemented by the circuit of Fig. 4.20.	93
Fig. 4.23 Complete circuit diagram of process compensation scheme.	93
Fig. 4.24 Fully-adjustable process compensation circuit realised with off-chip amplifiers.....	95
Fig. 4.25 Complete schematic of the on-chip analog front-end – bias generator, process compensation interface, and bipolar core.....	96
Fig. 4.26a Fig. 4.26b Second-order CIFB (4.26a) and CIFF (4.26b) $\Delta\Sigma$ loop architectures with ‘implicit reference’ front-end.	98
Fig. 4.27 Complete switched-capacitor modulator employing CIFF structure.....	100
Fig. 4.28 A z-domain model of the CIFF integrator in Fig. 4.27.....	100
Fig. 4.29 Operating principle of a latching comparator.	103
Fig. 4.30 Clocked comparator implemented in the $\Delta\Sigma$ modulator.	104
Fig. 4.31 Complete switched-capacitor modulator with multi-clocking first integrator.	105
Fig. 4.32 Details of the modulator clock phases used in Fig. 4.31.	106
Fig. 4.33 Non-overlapping clock generator with delayed falling edges.	107

Fig. 4.34 Complete $\Delta\Sigma$ modulator control block.	107
Fig. 4.35 Plot of conversion resolution achieved by 2 nd and 3 rd -order symmetric sinc filters, versus number of modulator cycles to produce the result.	109
Fig. 4.36 Block diagram of a third-order sinc digital filter.	110
Fig. 4.37 Complete circuit diagram for third-order sinc decimation filter.	111
Fig. 4.38 Layout of N-well pinched-base resistor test devices.	112
Fig. 4.39 External connections to sensor on a 28-pin DIP package.	114
Fig. 4.40a Complete layout of active circuitry for the smart temperature sensor.	115
Fig. 4.40b Chip photo of active die area.	115
Fig. 5.1 DUT clock and reset control circuit for single-conversion operation.	120
Fig. 5.2 Precision current source for measurement of on-chip pinched-base resistors.	121
Fig. 5.3 External process-compensation circuitry.	122
Fig. 5.4 Testbed power supply circuitry.	123
Fig. 5.5 Implemented circuitry for the electrical testbed.	124
Fig. 5.6 Example of temperature data captured from datalogger.	125
Fig. 5.7 Method used to attach DUT and reference thermometer to isothermal aluminium block.	126
Fig. 5.8 Schematic of complete thermal and electrical testbed.	127
Fig. 5.9 Test equipment used in the thermal testbed.	128
Fig. 5.10 Typical filtered output bitstream captured on DSO.	129
Fig. 5.11 Raw $\Delta\Sigma$ modulator bistream multiplied by 3 rd -order sinc filter impulse.	130
Fig. 5.12 Filtered ADC output bitstream with binary weighting.	131
Fig. 5.13 Output of ADC versus temperature.	133
Fig. 5.14 Substrate bipolar voltage V_{BE} versus temperature.	133
Fig. 5.15 N-well resistor R_1 versus temperature.	134
Fig. 5.16 Base-pinch resistor R_2 versus temperature.	134
Fig. 6.1 Deviation from average linear fit for the 6 tested smart sensors.	136
Fig. 6.2 Deviation from average linear fit of 5 smart sensors.	137
Fig. 6.3 Correlation between gradient and intercept of linear approximations of individual device output curves.	137
Fig. 6.4 Deviation from individual linear fit for the 6 tested smart sensors from 0°C to 100°C.	139

Fig. 6.5	Deviation from individual 2 nd -order polynomial fit for the 6 tested smart sensors from 0°C to 100°C.....	140
Fig. 6.6a	Sequential ADC readings obtained from device #3 at room temperature. ...	141
Fig. 6.6b	Sequential ADC readings obtained from device #3 at 100°C.	141
Fig. 6.7	Deviation from average linear fit of V_{BE} generated by the 6 test devices.	144
Fig. 6.8	Correlation between gradient and intercept of linear approximations of individual V_{BE} curves.....	145
Fig. 6.9	Diagram showing possible correlations between on-chip resistors and V_{BE} characteristics.....	147
Fig. 6.10	Correlation between value of R_2 @ 80°C and V_{BE} 's gradient.....	147
Fig. 6.11	Diagram showing possible correlations between on-chip resistor R_2 , V_{BE} , and ADC output characteristic.....	149
Fig. 6.12	Correlation between V_{BE} single-dimensional gradient representation and ADC output gradient.	149
Fig. 6.13	Diagram showing significant correlations between on-chip resistor R_2 , V_{BE} , and ADC output characteristic.....	150

TABLE OF TABLES

Table 3.1	Characteristic quantities for a generic temperature sensor, and the direction of movement that will produce a better sensor.....	49
Table 3.2	The strengths and weaknesses of a typical smart integrated temperature sensor.	50
Table 4.1	Summary of design variables affecting curvature in D_{OUT}	87
Table 4.2	Preliminary $\Delta\Sigma$ loop gain values.....	101
Table 4.3	Final $\Delta\Sigma$ loop gain values.	101
Table 4.4	Final values for $\Delta\Sigma$ loop capacitances in Fig. 4.26.....	103
Table 4.5	Complete list of input and output ports on the implemented design.	113
Table 6.1	Summary of error sources in the measured smart sensor response.....	143

DEFINITIONS AND ABBREVIATIONS

$\Delta\Sigma$	Delta-Sigma
A-D	Analog-to-Digital
AC	Alternating Current
ADC	A-D Converter
CDS	Correlated Double Sampling
CIFB	Cascade of Integrators, FeedBack
CIFF	Cascade of Integrators, FeedForward
CM	Common-Mode
CMC	CM Control
CMOS	Complementary MOS
CSV	Comma-Separated Values
CTAT	Constant To Absolute Temperature
D-A	Digital-to-Analog
DAC	D-A Converter
DC	Direct Current
DEM	Dynamic Element Matching
DMM	Digital Multimeter
DRC	Design Rule Check
DSO	Digital Storage Oscilloscope
DUT	Device Under Test
ENOB	Effective Number Of Bits
ESD	Electrostatic Discharge
FIR	Finite Impulse Response
IC	Integrated Circuit
IIR	Infinite Impulse Response
I/O	Input / Output
MIM	Metal-Insulator-Metal
MOS	Metal Oxide Semiconductor
NC	No Connection

NTF	Noise Transfer Function
OSR	Oversampling Ratio
PDK	Process Development Kit
PID	Proportional, Integral, Differential
PTAT	Proportional To Absolute Temperature
Recursion	See 'Recursion'
RTD	Resistance Temperature Detector
SC	Switched Capacitor
STF	Signal Transfer Function
ZIF	Zero Insertion Force

CHAPTER 1

INTRODUCTION

At the microscopic level, all matter in this universe is in motion. This motion is random in nature, and the average magnitude of this motion causes the macroscopic phenomenon known as temperature. Temperature can be measured on several scales; most common in scientific and technical disciplines are degrees Celsius ($^{\circ}\text{C}$) and Kelvin (K). The theoretical temperature at which all thermal motion ceases is known as *absolute zero*, and is 0 on the Kelvin scale.

The Celsius scale is defined by the temperatures at which ordinary water changes phase from solid to liquid and from liquid to gas – 0°C and 100°C respectively. The gradients of the Celsius and Kelvin scales are identical, i.e. a 1°C change corresponds to a change of 1K. Measured on the Celsius scale, absolute zero occurs at -273.15°C . Degrees Celsius can thus be converted to Kelvin by adding 273.15 (Eq. 1.1)

$$K = ^{\circ}\text{C} + 273.15 \quad (1.1)$$

The earliest temperature-measuring devices (thermometers) were built in the 1600's [Fraden, 2004, Chapter 3]. They exploited the property common to many materials that volume changes with temperature. The modern descendent of these is the ubiquitous mercury-glass thermometer. Thermal sensors with an output in the electrical domain are available in a wide variety of designs and target applications. They typically employ one of two sensing methods: thermoelectric and thermoresistive.

A junction between two dissimilar metals generates a small voltage proportional to the junction temperature. Thermal sensors employing this principle are known as thermocouples, and employ the thermoelectric effect to directly convert thermal energy into electricity. Different metals can be used to form the sensor junction, and the choice of constituent metals determines the thermocouple's sensitivity, linearity, and operating temperature range. Thermocouples are popular due to their small thermal mass, simple and rugged construction, and wide operating range. A 'type K' thermocouple

constructed with a chromel-alumel junction can operate from -200°C to 1250°C [Carstens, 1993].

Another class of thermal sensors exhibit a change in electrical resistance with temperature, and can be constructed with a wide range of materials. *Resistance temperature detectors* (RTDs) are a family of thermal sensors constructed with a wound length of pure metal wire. This produces a sensor with a relatively low resistance and a positive temperature coefficient. Chemically stable metals exhibit greater linearity and less long-term drift. Platinum wire-wound RTDs are by far the most common type where high accuracy and repeatability is desired. Another type of thermally sensitive resistor known as a thermistor is usually constructed with a semiconducting metal-oxide material. Devices with both positive and negative temperature coefficients are available. Thermistors are less linear and operate over a narrower temperature range than RTDs, but exhibit greater temperature sensitivity. High sensitivity and low cost means that thermistors are widely used in non-demanding sensor applications.

Continuing advancements in the capability of microelectronics has led to increased levels of signal processing and intelligent management of electronic systems throughout modern society. The bandwidth of commodity digital systems continues its skyward march, while mass production and intense competition drives prices down. Advanced control systems place increased demands on their interface with the real world. Accurate and cheap measurement of temperature is demanded in the fields of industrial process control and automation, environmental monitoring, agriculture, healthcare, robotics and the automotive industry, to name only a few.

All of the discrete thermal sensors discussed above have an output in the electrical domain, yet using the generated signal for further electronic processing requires suitable buffering and amplification circuitry. Although the sensor elements themselves may be cheap and simple to manufacture, the completed sensor system may not. Enter the smart sensor.

A smart sensor takes advantage of the same advancements in integrated circuit technology that creates the increasing demand for sensing capability. A thermal sensor and an analog-to-digital converter (ADC) are manufactured on the same integrated

circuit to allow the sensor to directly and easily interface with digital circuitry. The sensor might implement a serial bus interface such as I²C (Inter-IC) to allow a microprocessor to address the desired sensor, initiate a temperature conversion, and read out the resulting digital value. By eliminating the need for discrete signal-processing components, reliability is increased and integration into the target application is eased. Mass production on standard IC manufacturing processes drives down unit cost.

A typical example of a modern low-cost smart temperature sensor is the TCN75 manufactured by Microchip Technology Inc. [Microchip, 2010]. It has a typical accuracy of $\pm 1^{\circ}\text{C}$ over the range of -40°C to 125°C . The integrated $\Delta\Sigma$ ADC has a selectable resolution of up to 12 bits (0.06°C), and the device interfaces over the I²C bus. Active current consumption is $200\mu\text{A}$. It is available in standard 8-pin surface-mount plastic packages, and costs around US\$0.90 in bulk quantities.

The TCN75 highlights the primary advantages and disadvantages of smart IC-based thermal sensors. They offer low-cost and low-power sensing, and interface easily with microprocessor-based systems. The temperature sensing range defined by the limits of the IC manufacturing process is more restrictive than discrete thermal sensors, and the absolute accuracy of the on-chip thermal sensor is low.

1.1 Research Objectives

Entitled “**Accurate Thermal Sensing with Modern CMOS Integrated Circuits**”, this thesis is concerned with advancing the state of the art of temperature measurement with smart IC-based sensors. The objectives of this thesis are thus:

1. To thoroughly review relevant publications and literature, establish design techniques used to implement state-of-the-art smart sensors, and develop a design methodology that incorporates identified shortcomings and opportunities for improvement within the reviewed literature.
2. To design, simulate, lay out and fabricate a novel smart sensor IC using the developed methodology that implements design improvements identified in (1).

3. To evaluate the performance of this novel smart sensor against its own predicted performance as well as against previously published designs, thereby providing recommendations and insights into the future direction of IC smart sensors.

1.2 Contributions to Knowledge

The research work described in subsequent chapters of this thesis makes original contributions to knowledge in the field of smart temperature sensors by:

- Presenting a design for a fully-integrated smart temperature sensor on a modern 0.13 μm manufacturing process.
- Demonstrating circuit-level techniques to reduce smart sensor layout area by optimising $\Delta\Sigma$ modulator capacitor values.
- Improving the accuracy-conversion time tradeoff over previous designs by using a 3rd-order decimation filter with the on-chip incremental $\Delta\Sigma$ analog-to-digital converter.
- Presenting experimental data that indicates uncalibrated sensor accuracy can be improved by compensating for process-induced variations in on-chip thermal sensors with ‘pinch-base’ resistors on modern CMOS processes.
- Presenting designs for pinch-base process-compensation circuitry that interfaces with the implemented high-accuracy analog sensor front-end.

1.3 Structure of the Thesis

The structure of this thesis can be determined by parsing the *table of contents* (TOCTM) provided for the reader’s convenience at the beginning of this document. Both random and sequential data access is supported through the use of an advanced page numbering system that correlates with indexing numerals embedded within the TOCTM.

CHAPTER 2

LITERATURE REVIEW

The design of an analog integrated circuit is inextricably tied to the manufacturing process it will be implemented on. Critical device characteristics change between processes, minimum dimensions continue to shrink, and supply voltages have reduced from ten volts to one volt. Circuit techniques that work well on one process generation are often superseded on the next. Smart temperature sensors, with their need for accurate sensors and on-chip analog-to-digital converters (ADCs), are certainly no exception to this trend. This chapter will introduce the methods used to create on-chip temperature sensors, and describe the evolution of these techniques as manufacturing technology has progressed. The various ADC designs suitable for on-chip temperature conversion are also described, and a state-of-the-art sensor design published in 2005 is analysed. The chapter concludes with research that suggests it is possible to eliminate the need for individual sensor calibration.

2.1 On-Chip Sensing Elements

For the sensing element in a temperature sensor to be considered accurate, it should exhibit high linearity, high repeatability, and low manufacturing-induced variation. Furthermore, a low-cost sensor manufactured on a silicon chip is limited to the devices (active and passive) supplied as standard on that process. And finally, with the integration of an A-D converter on the same integrated circuit, a constant reference is required with the same accuracy as the sensor itself. This assortment of requirements greatly restricts the choice of thermal sensors available to the designer.

Many measurable device characteristics on an integrated circuit are influenced by temperature, yet few offer the linearity or process-insensitivity required for an accurate sensor. One notable possibility is the oxide-isolated polysilicon resistor. Lightly-doped polysilicon exhibits a strong negative temperature coefficient of resistance, with significant non-linearity (Fig. 2.1). High doping levels produce resistors with a smaller




Fig 2.1 Temperature dependence of polysilicon resistors with different doping levels [Rasmussen, 1994].

positive temperature coefficient, and greatly reduced non-linearity [Rasmussen, 1994]. The main drawback of these resistors as temperature sensors is their strong dependence on manufacturing process conditions. Fortunately, there is a better choice.

Active devices on an integrated circuit consist of regions of doped N-type and P-type silicon (Fig. 2.2) The PN junction is thus extremely common, and the equations governing its operation are well-understood [Dimitrijevic, 2000, Chapter 3]. The current flowing through a PN diode junction is described by:

$$I_D = I_S \left(e^{\frac{V_D}{V_T}} - 1 \right), \tag{2.1}$$

as a function of applied voltage V_D .





Fig. 2.2a

Fig. 2.2b

Simplified cross-section of PN diode junction (2.2a) and schematic symbol (2.2b).

On an integrated circuit the PN diode typically appears as part of a more complex device. The base-emitter junction of a bipolar transistor is one example. Current gain is produced in a bipolar transistor by making the base region as thin and as lightly-doped as possible (Fig. 2.3). When a forward voltage is applied to the base-emitter junction,

the majority of electrons coming from the emitter pass straight through the base region and recombine in the collector. Thus a small base-emitter current produces a larger collector-emitter current.




Fig. 2.3a **Fig. 2.3b**
Simplified cross-section of an NPN bipolar transistor (2.3a),
and equivalent schematic (2.3b).

The Ebers-Moll model [Ebers and Moll, 1954] can be used to determine the terminal currents of a bipolar transistor. Considering the base-emitter diode junction in isolation, its behaviour differs from the ideal diode relationship of Eq. 2.1 due to the transistor's current gain β . In the forward active region, the base-emitter current can be simplified to [Dimitrijević, 2000, Chapter 6]:

$$I_B = \frac{1}{\beta} I_S \left(e^{\frac{V_{BE}}{V_T}} \right). \quad (2.2)$$

The current gain factor β in 2.2 introduces an additional unwanted temperature dependency due to recombination/generation currents within the base region [Rasmussen, 1994]. These unwanted currents only flow between the base and emitter terminals, so the transistor's collector current possesses the desired relationship of Eq. 2.1. In the forward active region of operation, $\exp(V_{BE} / V_T)$ is much greater than 1, therefore the Ebers-Moll equation for bipolar collector current can be simplified to:

$$I_C = I_S \left(e^{\frac{V_{BE}}{V_T}} \right). \quad (2.3)$$

Two temperature-dependent terms are present in Eq. 2.3. The thermal voltage V_T is linearly proportional to absolute temperature:

$$V_T = \frac{kT}{q}, \quad (2.4)$$

where k is Boltzmann's constant and q is the electron charge. The saturation current I_S can be expressed as:

$$I_S = BAT\bar{u}_n n_i^2, \quad (2.5)$$

where B is a process-dependent constant, and A is the emitter area. The terms \bar{u}_n and n_i of Eq. 2.5 exhibit further temperature dependence:

$$\bar{u}_n \propto T^{-m}, \quad (2.6)$$

$$n_i^2 \propto T^3 e^{\left(\frac{-V_{G0}}{V_T}\right)}. \quad (2.7)$$

It should be noted that 2.7 is simplified for the purposes of this discussion: the bandgap voltage V_{G0} is assumed to be linearly proportional to temperature, which is true only over a limited range of temperatures [Tsividis, 1980].

Introducing a simplifying constant C and exponent $\gamma = 4 - m$, rewriting Eq. 2.3 for V_{BE} results in:

$$V_{BE} = V_{G0} + V_T \ln\left(\frac{I_C T^{-\gamma}}{CA}\right). \quad (2.8)$$

The temperature dependence of V_{BE} is then fully defined when the temperature dependence of bias current I_C is specified. With a constant bias current, the typical base-emitter junction exhibits a temperature coefficient of around -2mV/K , with notable second-order curvature.

The value of V_{BE} at room temperature exhibits significant sensitivity to manufacturing conditions. Variations in doping profile and junction area cause variations in I_S (Eq. 2.5), which introduces variations in the gradient of V_{BE} . These gradient variations can be compensated by varying the magnitude of I_C . The ability to adjust the gradient of V_{BE} via bias current also provides a unique opportunity for accurate temperature measurement, as follows.

Consider two bipolar transistors on an integrated circuit. With careful layout the various process-dependent constants in 2.8 will be effectively identical, and thus the difference between their respective base-emitter voltages will be:

$$\Delta V_{BE} = V_{BE1} - V_{BE2} = V_T \ln\left(\frac{I_{C1} \cdot A_2}{I_{C2} \cdot A_1}\right). \quad (2.9)$$




Fig. 2.4 A PTAT voltage produced from the difference of two V_{BE} , and the derived bandgap reference voltage V_{REF} .

This neatly eliminates the various sources of non-linearity and process dependence present in 2.8, and leaves a voltage dependent on device area and bias current ratios, and linearly proportional to absolute temperature (PTAT) (Fig. 2.4).

The magnitude of this PTAT voltage is dependent on both the area ratio and current ratio of the bipolar transistor pair, but typical values produce a positive temperature coefficient of around 200uV/K. This results in a room-temperature value of approximately 50mV, necessitating amplification for any practical application.

While this PTAT voltage is as close to an ideal temperature signal as can be reasonably expected, a smart sensor also requires a temperature-invariant reference to convert this signal into the digital domain. And for a completely integrated design, this reference generator needs to be on-chip. Fortunately, the principles of on-chip reference voltage generators are well-understood, and such devices have been in common usage since the 1970s. From Eqs. 2.8 and 2.9, we have two voltages available with opposing temperature coefficients. In order to obtain a reference with zero first-order temperature dependence we simply add them together with an appropriate scale factor for ΔV_{BE} , to generate V_{REF} (Eq. 2.10). Such a reference is known as a bandgap reference, because when second-order effects are neglected the output voltage of such a reference is V_{G0} , the value V_{BE} approaches as T approaches absolute zero (Eq. 2.8, Fig. 2.4).

$$V_{REF} = V_{BE} + \alpha \Delta V_{BE}. \quad (2.10)$$

A combination of V_{BE} and ΔV_{BE} can also be used to create an enhanced temperature sensor. While a purely PTAT signal exhibits excellent linearity and repeatability, it

suffers from the presence of a large offset at room temperature. At 25°C, a 1°C temperature change induces a 0.3% change in a PTAT-based voltage or current signal. This places unnecessarily high demands on the following readout or A-D conversion circuitry if the temperature range of interest does not extend to absolute zero (-273.15°C). System resolution may be increased over a narrower range by subtracting a constant value from the PTAT signal. The zero value of this ‘intrinsically referenced’ signal therefore occurs at some point above absolute zero (Fig. 2.5), and may be used to generate an analog output that reads directly in any convenient units such as degrees Celsius or Fahrenheit [Rasmussen, 1994]. Such a signal may be generated using the same circuit techniques used in bandgap references, but subtracting V_{BE} from ΔV_{BE} rather than adding them together.

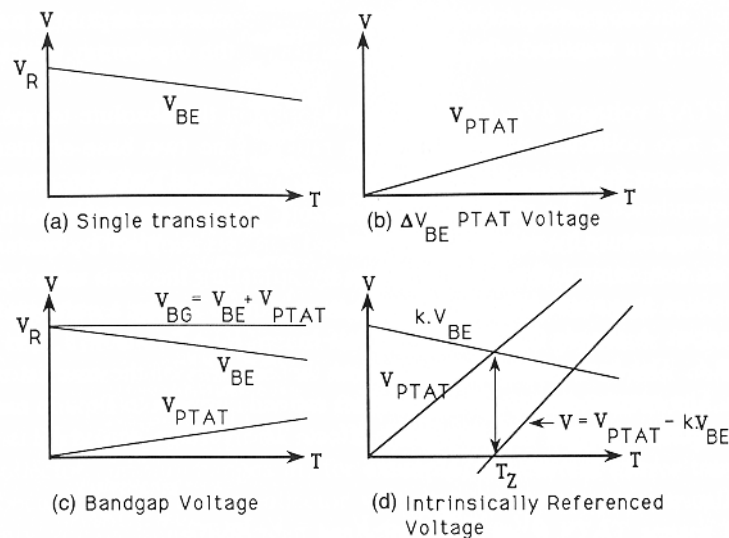


Fig. 2.5 Techniques for generating a bandgap reference and an intrinsically referenced temperature sensor [Rasmussen, 1994].

It is apparent from the above discussion that a bipolar-based temperature sensor and a bandgap reference have identical theoretical backgrounds, and similar circuit implementations. These concepts are so closely linked that for the purpose of this thesis they will be considered as one. The following section will describe the development of on-chip temperature sensors and references in the context of advancements in IC manufacturing technology.

2.2 History of Bipolar IC Temperature Sensors

Early IC manufacturing technologies produced bipolar devices operating with supplies in the tens of volts. With the progression of technology feature sizes have reduced and operating voltages have come down, and MOS devices have superseded bipolar for the vast majority of applications. This natural progression has been mirrored in the progression of published designs for bandgap references and temperature sensors in the literature.

The earliest published example of a bandgap reference designed for integrated applications was the LM109 voltage regulator [Widlar, 1971]. The bipolar-based reference generator offered advantages over previous zener-based references in the form of reduced sensitivity to manufacturing variations, reduced minimum operating voltage, and reduced noise. A simplified schematic of this voltage reference is shown in Fig. 2.6.




Fig. 2.6 Early bandgap-reference generator circuit [Widlar, 1971].

Transistors Q_1 and Q_2 are operated at a current density ratio of 10:1, thus the difference between their base-emitter voltages follows Eq. 2.9:

$$\Delta V_{BE} = \frac{kT}{q} \ln(10). \quad (2.11)$$

This PTAT voltage appears across R_3 , and neglecting base currents the voltage across R_2 becomes:

$$V_{R2} = \frac{R_2}{R_3} \Delta V_{BE}. \quad (2.12)$$

Transistor Q_3 provides the base-emitter voltage necessary for creating the reference voltage, and also buffers the PTAT voltage V_{R2} . The reference's output voltage therefore implements the bandgap technique described by Eq. 2.10:

$$V_{REF} = V_{BE3} + \frac{R_2}{R_3} \Delta V_{BE}. \quad (2.13)$$

Brokaw [1974] published an improved bandgap generator design that uses two sensing transistors and an amplifier in a negative feedback loop (Fig. 2.7). The main source of inaccuracy in the three-transistor cell of Fig. 2.6 is the assumption of zero base current in Eq. 2.12; no such assumption is required here. Additionally, this circuit is biased with a voltage source rather than the current source of Fig. 2.6, allowing greater flexibility and a simpler implementation. This circuit configuration has proved remarkably long-lived, as the same principle is still used in today's CMOS-based references.




Fig. 2.7 An improved bipolar bandgap reference generator circuit [Brokaw, 1974].

The op-amp in Fig. 2.7 monitors the voltage drop across the current-sensing resistors R supplying collector current to Q_1 and Q_2 , and aims to equalise the current flowing into each device:

$$I_{C1} = I_{C2}. \quad (2.14)$$

Under this equilibrium condition, the current density ratio between the two transistors is 8:1, and the difference between their base-emitter voltages is PTAT:

$$\Delta V_{BE} = \frac{kT}{q} \ln(8). \quad (2.15)$$

This PTAT voltage appears across R_2 , and given equal currents through both transistors, the voltage across R_1 becomes:

$$V_{R1} = 2 \frac{R_1}{R_2} \Delta V_{BE}. \quad (2.16)$$

The circuit's output is taken from the common base connection, and is of the same form as Eq. 2.10:

$$V_{REF} = V_{BE1} + 2 \frac{R_1}{R_2} \Delta V_{BE}. \quad (2.17)$$

Both of the above designs used the bandgap technique to generate a temperature-invariant reference voltage. The same circuit techniques can easily produce a temperature-dependent voltage or current. An early example [Timko, 1976] is a two-terminal device producing an output in the current domain. It achieves a very wide sensing range (-125°C to 200°C) by tying the IC substrate to an internal circuit node rather than the negative supply, thus reducing the impact of reverse-bias junction leakage at high temperatures. The output current is derived from ΔV_{BE} , and is therefore PTAT. An output signal calibrated to the Celsius scale is also possible [Meijer, 1980] using the intrinsically referenced technique discussed in Section 2.1.

2.3 Curvature Compensation for Bipolar Sensors

The most significant non-linearity present in the above bandgap circuits originates from the diode junction voltage V_{BE} . Fig. 2.4 illustrates that while it is possible to calibrate the reference to achieve a zero first-order temperature coefficient at room temperature, higher-order curvature present in V_{BE} causes the reference to deviate at lower and higher temperatures. Several methods can be used to reduce this non-linearity. If the transistor bias current is made strongly dependent on temperature, i.e.

$$I_C = kT^\alpha, \quad (2.18)$$

where the exponent α is equal to γ in Eq. 2.8, then V_{BE} will be a linear function of temperature [Rasmussen, 1994]. Fig. 2.8 illustrates the reduction in nonlinearity as the bias current is changed from $\alpha = 0$ (constant current) up to $\alpha = 3$.




Fig. 2.8 Residual curvature in V_{BE} with various bias current temperature dependencies [Rasmussen, 1994].

The main drawback of this method of curvature compensation is the difficulty in producing the power-law bias current temperature dependency. Both circuits in Figs. 2.6 and 2.7 internally used PTAT currents ($\alpha = 1$), and no practical circuit has been published for generating higher order currents. This solution is further complicated by the fact that complete linearisation will usually require a non-integer power relation of α .

One design [Meijer et al. 1982] uses the linearising effect of α in a slightly different way. Two stacks of transistors are constructed so that their base-emitter voltages add together (Fig. 2.9). One stack is supplied with a PTAT current, and the other with a temperature-invariant (CTAT) current. The base-emitter voltages within the two stacks therefore possess different curvatures (Fig. 2.10).

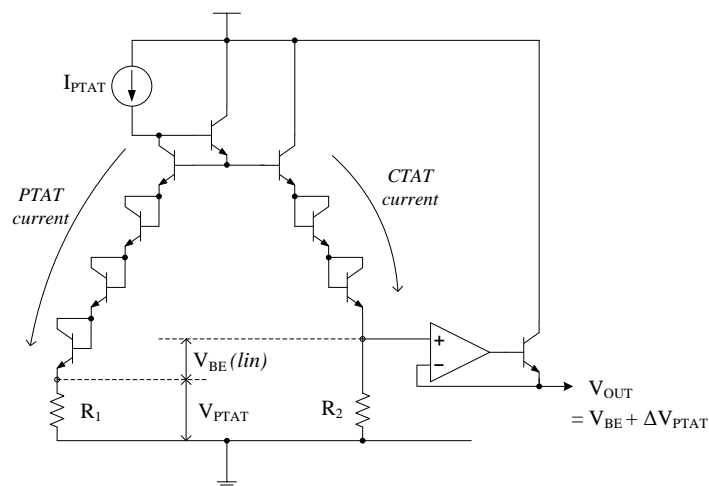


Fig. 2.9 Circuit technique to produce a linearised V_{BE} and bandgap reference [Meijer et al., 1982].




Fig. 2.10 Difference in V_{BE} curvature when biased with constant (CTAT) and linearly temperature-dependent (PTAT) currents.

Typical bipolar transistors have a temperature exponent γ in Eq. 2.8 of around 4. Therefore a transistor biased with a PTAT current source will have a V_{BE} curvature roughly 25% less than one biased with a constant current, represented by the shaded area in Fig 2.10.

The curvature correction technique implemented in Fig. 2.9 adds three copies of this shaded area to one PTAT-biased base-emitter junction, to achieve a V_{BE} with near-zero high-order curvature. The left-hand circuit leg in Fig. 2.9 contains 4 stacked transistors biased with a PTAT source. The right-hand leg contains 3 transistors, and is biased by a constant current. The top of both stacks connect to a common node, so the voltage difference between the stacks is present across the resistors R_1 and R_2 . Due to the PTAT bias current, the voltage across R_1 is also PTAT:

$$V_{R1} = R_1 k \Delta V_{BE}, \quad (2.19)$$

where k is the gain of the current mirror and PTAT generator (not shown in Fig. 2.9).

The voltage difference between the two transistor stacks is equal to one PTAT-biased V_{BE} plus 3 times the difference between the constant-biased and PTAT-biased transistors (the shaded area of Fig. 2.10). This modified V_{BE} is essentially a linearised base-emitter voltage:

$$V_{R2} - V_{R1} = V_{BE}(lin). \quad (2.20)$$

The circuit output is taken from the top of R_2 , and consists of the sum of Eqs. 2.19 and 2.20:

$$V_{R2} = V_{BE}(\text{lin}) + R_1 k \Delta V_{BE}. \quad (2.21)$$

Thus by employing stacks of base-emitter junctions, the bandgap reference of Fig. 2.9 compensates for curvature in the same way as using an integer power-law temperature-dependent bias current would (Eq. 2.18). Meijer et al. reported a 20:1 reduction in thermal nonlinearity when compared to an uncompensated bandgap reference (Fig. 2.11). Unfortunately this compensation method is not physically realisable on modern CMOS processes because the only bipolar transistors available cannot be stacked vertically as required in Fig. 2.9. The limitations of bipolar devices on CMOS processes will be discussed further in Section 2.4.

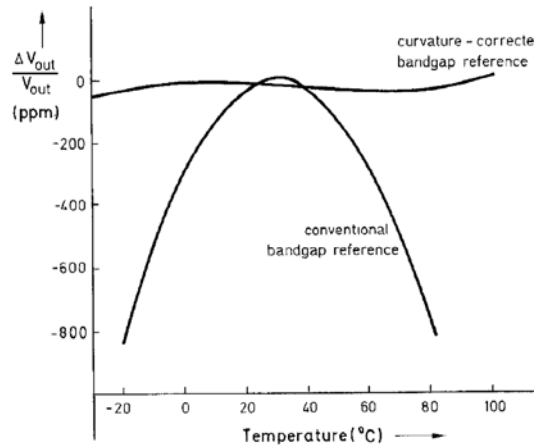


Fig. 2.11 Comparison of curvature compensated and non-compensated bandgap references [Meijer et al., 1982].

The curvature compensation technique discussed above used different temperature-dependent currents to linearise V_{BE} directly. It is also possible to correct the curvature in V_{BE} with an independently-generated non-linear voltage. Song and Gray [1983] used a ΔV_{BE} -generator circuit biased with temperature-dependent currents to produce a non-linear voltage:

$$\Delta V_{BE} = V_T \ln \left(\frac{A_2}{A_1} \right) \left(\frac{I_0 + I_T}{I_0 - I_T} \right), \quad (2.22)$$

where A_1 and A_2 are the transistor emitter areas. I_0 is a constant current and I_T is linearly proportional to temperature. The magnitude of I_T can be independently adjusted to cancel the non-linearity present in V_{BE} .

Gunawan et al. [1993] used negative feedback around a V_{BE} -generator circuit to produce a linearised V_{BE} -dependent current, which is added to ΔV_{BE} to produce a bandgap-derived reference current in the usual way:

$$I_{REF} = 2(I_{VBE} + I_{NL}) + I_{\Delta VBE}. \quad (2.23)$$

With the exception of [Song and Gray, 1983], all of the curvature-compensation techniques discussed in this section were designed for implementation on bipolar IC processes, and significantly improved the temperature insensitivity of the references. However the vast majority of modern IC processes use CMOS transistors, and the limited-function bipolar transistors available on such processes combine with the larger voltage offsets of CMOS op-amps to reduce the accuracy of an equivalent reference implemented on CMOS technology. Thus CMOS-based references require additional circuit techniques to achieve similar accuracy to the curvature-compensated references discussed above. These techniques will be described in Section 2.5.

2.4 History of CMOS IC Temperature Sensors

The sensors and references discussed up to this point were all manufactured on early bipolar processes, and therefore both the sensors and supporting circuitry were constructed with bipolar transistors. The development of self-aligned polysilicon-gate MOS transistors marked the beginning of an industry-wide transition from bipolar to CMOS processes. So naturally, it became necessary to implement the same bandgap techniques on these new processes.

The current flowing between the drain and source of a MOS device is controlled by the voltage applied to its gate terminal (Fig. 2.12). The transistor begins to conduct when the gate-source voltage V_{GS} exceeds a threshold voltage V_{TH} . Analog circuits most commonly bias MOS transistors in the saturation region, where the drain-source voltage V_{DS} is larger than the gate overdrive voltage V_{OD} :

$$V_{DS} > V_{OD}, \quad (2.24)$$

where:

$$V_{OD} = V_{GS} - V_{TH}. \quad (2.25)$$




Fig. 2.12a




Fig. 2.12b

**Simplified cross-section of an NMOS transistor (2.12a),
and equivalent schematic (2.12b).**

The drain-source current in the saturation region exhibits a square-law dependence on gate voltage, and is described by:

$$I_D = \frac{1}{2} \mu_n C_{ox} \frac{W}{L'} (V_{GS} - V_{TH})^2. \quad (2.26)$$

The term μ_n represents the mobility of charge-carrying electrons in NMOS devices on the target process, and together with the unit gate capacitance C_{ox} forms a process-dependent constant. The ratio W / L' represents the aspect ratio of the MOS device, where the effective length L' is shorter than the drawn device length L due to physical effects.

Eq. 2.26 applies when the transistor is turned on, i.e. $V_{GS} > V_{TH}$. Rather than abruptly turning off as V_{GS} drops below V_{TH} , the device enters the *subthreshold* or weak inversion region [Razavi, 2001, Chapter 2]. In this operating region the drain current exhibits an exponential dependence on the gate voltage:

$$I_D = I_0 e^{\frac{V_{GS}}{\zeta V_T}}. \quad (2.27)$$

With the exception of the constant ζ , Eq. 2.27 is identical in form to the characteristic diode equation (Eq. 2.3). It is therefore possible to extract a PTAT voltage using the same method employed by bipolar references; by operating two MOS transistors at different current densities in the subthreshold region [Kölling et al., 1990]:

$$\Delta V_{DS} = V_T \ln \left(\frac{I_{D1} (W/L)_2}{I_{D2} (W/L)_1} \right). \quad (2.28)$$

Increasing this current-density ratio is equivalent to increasing emitter area A of a bipolar transistor in Eq. 2.9.

Unfortunately, finding an MOS equivalent of the other essential component of a bandgap reference, V_{BE} , is not so easy. Circuits based on the MOS threshold voltage V_{TH} [Kölling et al., 1990] exhibit the desired negative temperature coefficient, together with substantial random mismatches and long-term drift. Circuits extracting the fundamental silicon bandgap voltage V_{G0} from MOS transistors [Vittoz, 1985] exhibit less process sensitivity, but cannot be implemented with standard CMOS transistors.

Despite the use of novel circuit techniques, MOS transistors remain fundamentally inferior to bipolars as the critical elements in temperature sensors and voltage references. The conducting drain-source channel is sensitive to mobile ions that become trapped below the gate oxide, while the conducting region of bipolar transistors is within the bulk silicon and remains largely unaffected by ion contamination. To obtain the highest possible accuracy, it is therefore necessary to use bipolar transistors as the active sensing elements for measuring temperature.

Two types of parasitic bipolar transistors exist in CMOS technologies, and can be used in bandgap and temperature-sensing circuits when certain restrictions are observed. A lateral bipolar transistor is formed between the drain and source regions of a MOS transistor placed in a well, provided the conducting drain-source channel is turned off (Fig. 2.13). This may necessitate the application of a negative gate voltage, particularly in modern low-threshold processes. Various designs for bandgap references have been published [Krummenacher and Oguey, 1989; Montané et al., 1998], proving that the lateral bipolar device is indeed capable of substituting for a general-purpose bipolar. Lateral bipolars manufactured on a modern CMOS process can achieve high gains (50 – 100) due to the reduced base width [Hastings, 2006, Chapter 8]. However the lateral bipolar exhibits several disadvantages, the most significant is its sensitivity to manufacturing process characteristics due to current flowing near the silicon surface. This, coupled with the lack of characterisation or parameter control for this device in




Fig. 2.13 Simplified cross-section of a lateral PNP transistor on a CMOS process.

modern CMOS processes means that good design with lateral bipolar transistors is difficult at best.

A second type of CMOS-compatible bipolar transistor makes use of the parasitic conduction path existing in the lateral bipolar transistor from the emitter (drain), through the base (well), and into the substrate (Fig. 2.14). This parasitic vertical bipolar reduces the collector current of a lateral bipolar, but it can also act as a transistor in its own right. The silicon substrate acts as the device's collector, hence it is known as a substrate bipolar transistor. The biggest drawback of such a device is immediately apparent - the collector is permanently tied to substrate, the most negative potential in modern P-substrate processes. Fortunately this does not prevent its use in bandgap reference generator circuits.




Fig. 2.14 Simplified cross-section of a substrate PNP transistor on a CMOS process.

While a substrate bipolar constructed on an older high-voltage CMOS process exhibits current gains of around 100, modern low-voltage process typically produce devices with gains of less than 10 [Hastings, 2006, Chapter 8]. Despite this, substrate bipolars are generally preferred over laterals due to reduced sensitivity to surface effects, and their more ideal I-V relationship [Bakker and Huijsing, 1996; Wang and Meijer, 2000]. Moreover, in some modern submicron processes the substrate bipolar is the only device containing a forward-biased diode junction that is modelled or indeed supported at all, due to the risk of minority carriers inducing destructive latch-up.

Modern CMOS processes employ a P-substrate and N-well, so the substrate bipolars manufactured on these processes are PNP devices with their collector tied to the most negative potential. This restricts the choice of circuit configurations available to the designer. Two general topologies are possible for generating V_{BE} and ΔV_{BE} using substrate bipolars, as shown in Fig. 2.15. Both circuits operate by biasing substrate

bipolars Q_1 and Q_2 at a well-defined current density ratio, and ensuring the voltages at nodes X and Y are equal:

$$\frac{I_{C1}}{A} = \frac{I_{C2}}{nA}, \quad (2.29)$$

$$V_X = V_Y. \quad (2.30)$$

Given the satisfaction of these two conditions, the difference in V_{BE} between Q_1 and Q_2 drops across R_1 , and follows the PTAT relationship of Eq. 2.9:

$$\begin{aligned} V_{R1} &= \Delta V_{BE} \\ &= V_T \ln\left(\frac{nI_{C1}}{I_{C2}}\right). \end{aligned} \quad (2.31)$$

The substrate bipolars' bias current is therefore PTAT, and can be extracted using PMOS current mirrors as illustrated in Fig. 2.15.




Fig. 2.15a




Fig. 2.15b

Current mirror (2.15a) and opamp-based (2.15b) CMOS bandgap references [Razavi, 2001].

Given a suitable amplifier, the topology of Fig. 2.15b exhibits a greater supply rejection ratio [Razavi, 2001, Chapter 11]. Furthermore, adding cascode transistors to the current mirrors of Fig. 2.15a raise the minimum supply voltage above what can be tolerated in modern low-voltage CMOS processes. Therefore Fig. 2.15b is the more common implementation of a bandgap reference in CMOS technologies [Bakker and Huijsing, 1996; Meijer et al., 2001].

2.5 Analog Techniques for Accurate CMOS Circuitry

As mentioned in Section 2.4, bipolar transistors are preferred over MOS devices as the critical sensors in bandgap references made on CMOS technologies. However the associated biasing and support circuitry must still use MOS transistors. Random threshold voltage mismatches in MOS transistors cause larger DC offsets in the current mirrors and op-amps used in Fig. 2.15. Additionally, MOS devices typically exhibit significant low-frequency flicker noise, so-called $1/f$ noise [Gray et al., 2001, Chapter 11].

Chopping

One method of eliminating both DC offsets and low-frequency $1/f$ noise is *chopper stabilisation* [Enz and Temes, 1996]. The operating principle of this technique is illustrated in Fig. 2.16. A low-frequency signal is passed through a chopper that periodically inverts the input polarity. This modulated signal is then fed into the amplifier, where noise and offsets are modelled as a voltage source at the input. The signal remains spectrally separated from offsets and low-frequency noise due to the modulating effect of the input chopper. A second synchronous chopper at the amplifier output demodulates the signal back down to the baseband, while the low-frequency noise and offsets are modulated above the frequency range of interest.




Fig. 2.16 Implementation of a chopped differential op-amp and the associated frequency spectrum at various circuit nodes.

The choppers in Fig. 2.16 are implemented as cross-connected switches (Fig. 2.17). MOS technology is ideal for implementing chopping because the op-amp input stage requires virtually zero input bias current. MOS transistors can therefore be used as switches, exhibiting zero on-state offset voltage and only a small load capacitance on the switch terminals. Depending on the expected signal swing, the switches in Fig. 2.17 may be implemented as near minimum-sized N-type or P-type devices, or a parallel combination of both. A differential op-amp is depicted in Fig. 2.16, although single-ended op-amps can just as easily implement the output chopper by swapping the output current mirror connection (Fig. 2.18).




Fig. 2.17 Implementation of the chopping block of Fig. 2.16 using MOS switches.




Fig. 2.18 Chopping implemented on a simple CMOS op-amp with single-ended output.

The op-amp of Fig. 2.18 can be directly substituted into the bandgap reference generator of Fig. 2.15b to produce a chopped reference with greatly reduced offsets and low-frequency noise [Bakker and Huijsing, 1996]. The reference generator of Fig. 2.15a is more difficult to combine with chopping as more switches are required, some of which carry the bipolars' DC bias current. Hence Fig. 2.15a is rarely used in modern CMOS designs.

Dynamic Element Matching

A technique similar in principle to chopping is known as *dynamic element matching* (DEM). It is a modulation technique that greatly reduces mismatch between several identical circuit elements. Consider a circuit requiring an accurately defined current ratio (Fig. 2.19a), a common requirement in CMOS-based PTAT and bandgap-generator circuits such as Fig 2.15. Such a current ratio is implemented with matched unit transistors (Fig. 2.19b). Even with the use of advanced layout techniques such as symmetric arrays and dummy rings, the matching of unit transistors is typically limited to around $\pm 1\%$ [Hastings, 2006, Chapter 12].




Fig. 2.19a




Fig. 2.19b

A ratioed current source (2.19a), and accurate circuit-level implementation (2.19b).

Dynamic element matching involves periodically swapping the unit current source supplying the lesser of the two matched currents (Fig. 2.20). DC mismatch between the unit sources is thereby modulated to the chopping frequency, and the average current ratio is more accurately defined than is possible with good layout alone.




Fig. 2.20 Dynamic element matching swaps the single unit source between all possible sources.

DEM is not only limited to application in current sources. Any circuit needing an accurately defined ratio created by two or more unit devices is a potential application. One possibility is the multi-bit DACs used in delta-sigma ADCs [Schreier and Temes, 2005, Chapter 6]. The DACs in this application are required to meet the linearity requirements of the overall converter, which is typically much greater than the number of bits output by the DAC itself. Dynamic element matching is almost a necessity in this case, to ensure the DAC is as linear as possible.

2.6 History of Smart Temperature Sensors

Early on-chip temperature sensors designed in the 1970s existed in a largely analog world. It was acceptable for a sensor to output a voltage or current derived from its internal sensor core. But the continual advancement of CMOS technology provided the raw materials for a digital revolution, to the point where today apparently simple devices such as fridges and heaters are designed with embedded controllers. Thus these smart devices demand similarly smart temperature sensors to reduce board-level complexity and reduce system costs.

The A-D converter of a smart temperature sensor is required to output the signal in digital form while consuming minimal power and die area. It should be at least as accurate as the analog sensor, and will ideally require no calibration. While advances in IC manufacturing technology have increased the density of logic circuitry, passive components that play a critical role in the operation of ADCs have remained highly variable over many process generations. Thus on-chip ADCs are designed to avoid the limitations of variable analog components while best utilising the capabilities of digital logic available on the target IC manufacturing process.

A simple method of obtaining a digital-compatible signal is by using a voltage-controlled or current-controlled oscillator whose frequency is determined by the temperature signal. The output frequency is then measured by an external clock, which can be very accurate when based on a crystal oscillator (Fig. 2.21). Such temperature-dependent oscillators were easy to implement on early bipolar processes, but required

external passive components due to the large area consumption and high variability of on-chip capacitors [D. van Maaren et al., 1982].

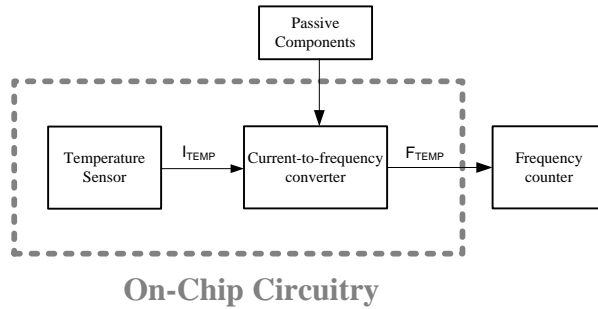


Fig. 2.21 Block diagram of early smart temperature sensor [D. van Maaren et al., 1982].

The need for external passive components is not the biggest disadvantage of the structure in Fig. 2.21, however. No matter how accurate the temperature sensor is, the sensor output must be converted to a frequency before being compared to the external clock. This means system accuracy is directly affected by variations in the current-to-frequency conversion ratio, which is in turn affected by variations in the values of passive components such as resistors and capacitors.

If an on-chip temperature-invariant reference is used, much greater accuracy can be achieved by using the external clock to measure the frequency *ratio* between the temperature-dependent and temperature-invariant frequencies [Meijer et al., 1988]. This is because variations in the current-to-frequency conversion ratio affect both the signal and the reference equally. The current-to-frequency converter can therefore use less precise on-chip passive components without impacting system accuracy (Fig. 2.22).

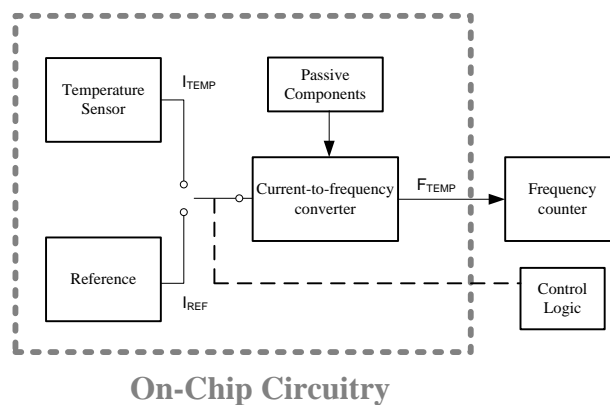


Fig. 2.22 An improved smart sensor with frequency output [Meijer et al., 1988].

However the circuit of Fig. 2.22 remains sensitive to noise near the current-controlled oscillator's switching threshold, and the accuracy is limited by the resolution of the external clock. Such a converter can be regarded as an intermediate step between an analog transducer and a smart sensor with digital output.

As IC manufacturing technology has advanced, the density and speed of digital logic has continued to increase. This progression has encouraged designers to find better ways to trade speed for accuracy. A desirable ADC design would operate at a high clock frequency with analog components of limited accuracy, yet still convert a low-frequency temperature signal with high accuracy. A delta-sigma ($\Delta\Sigma$) data converter (either A-D or D-A) is the most common implementation of a family known as *oversampling* converters, that do precisely this. They use a low-resolution quantiser to sample the input signal faster than the Nyquist rate, and a feedback loop around the quantiser 'shapes' the quantisation noise above the signal band [Schreier and Temes, 2005, Chapter 1]. The output of the quantiser is then passed through a digital filter, which removes the undesired high-frequency quantisation noise.

A block diagram of a simple first-order $\Delta\Sigma$ ADC is shown in Fig. 2.23. The quantiser is shown as a comparator, a common choice due to their simple implementation and inherent linearity. The digital filter is shown as a simple counter, which outputs the average of the quantiser bitstream. More complex high-order filters offer better suppression of the high-frequency quantisation noise at the price of increased complexity and layout area. The optimal choice of filter order and architecture is discussed further in Section 2.7.




Fig. 2.23 A simple one-bit first-order delta-sigma ADC.

The *oversampling ratio* (OSR) defines how many coarse samples of the input are taken to produce one high-resolution output value. A higher OSR will increase the conversion resolution, at the price of increased operating frequency. Typical $\Delta\Sigma$ converters employ OSRs of between 8 and 512.

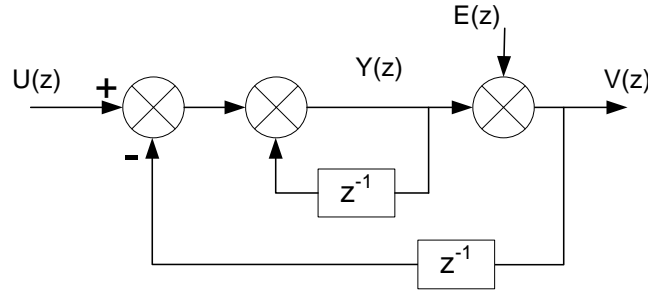


Fig. 2.24 A discrete z -domain model of the first-order $\Delta\Sigma$ modulator in Fig. 2.23.

Fig. 2.24 shows a discrete z -domain model of the first-order modulator in Fig. 2.23. The modulator output $V(z)$ is the sum of the integrator output $Y(z)$ and the additive noise source $E(z)$, representing the effect of the comparator:

$$\begin{aligned} V(z) &= Y(z) + E(z) \\ &= U(z) + E(z) + z^{-1}Y(z) - z^{-1}V(z). \end{aligned} \quad (2.32)$$

Equation 2.32 can be rearranged into the general form:

$$V(z) = STF(z) \cdot U(z) + NTF(z) \cdot E(z), \quad (2.33)$$

where:

$$STF(z) = 1, \quad (2.34)$$

and:

$$NTF(z) = 1 - z^{-1}. \quad (2.35)$$

The *signal transfer function* (STF) modifies the input signal as little as possible. In this simple modulator the STF is unity, although the STF of a more complex modulator may be unity only within the signal band. The *noise transfer function* (NTF) describes the shaping applied to the comparator error signal $E(z)$ to shift it above the signal band. A more complex $\Delta\Sigma$ modulator will implement a higher-order NTF, resulting in less in-band quantisation noise at the price of an increased high-frequency peak. Thus higher-order NTFs require higher-order digital filters.

Quantisation noise is random in nature, and is typically modelled as a white noise source with flat spectral density. Fig. 2.25 illustrates the shaping effect of the NTF on this spectrally-flat quantisation noise. The OSR and NTF order are chosen so that no significant quantisation noise exists within the input signal bandwidth, to the level of accuracy required of the converter. The magnitude and spectral characteristics of the shaped quantisation noise then determine the type of digital decimation filter required at the modulator output.

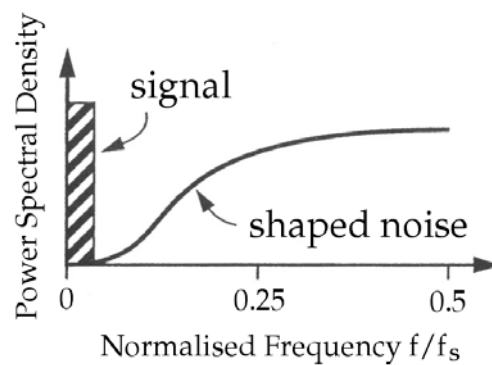


Fig. 2.25 Spectrum of the input signal and shaped quantisation noise of a typical $\Delta\Sigma$ modulator [Schreier and Temes, 2005, Chapter 1].

The trade-off between speed and resolution used by $\Delta\Sigma$ modulators makes them ideal for converting a temperature signal to the digital domain. The continual increase in digital logic density offered by subsequent IC generations has allowed the implementation of increasingly complex on-chip $\Delta\Sigma$ ADCs. A first-order $\Delta\Sigma$ ADC was implemented in a smart temperature sensor using a bipolar process and ECL logic [Riedijk and Huijsing, 1991]. The temperature signal and reference were supplied as currents to the integrator, allowing the input summing junction in Fig. 2.23 to be implemented by simply connecting the input current sources together (Fig. 2.26). The sensor and $\Delta\Sigma$ modulator occupied an area of 3mm^2 . The digital filter was not implemented on-chip due to the target processes' low logic density.

A later design by Bakker and Huijsing [1996] built a similar first-order current-input $\Delta\Sigma$ modulator on a $2\mu\text{m}$ CMOS technology. This design included the digital counter on-chip to provide an 8-bit digital output. The sensor and ADC occupied a total die area of 1.5mm^2 .

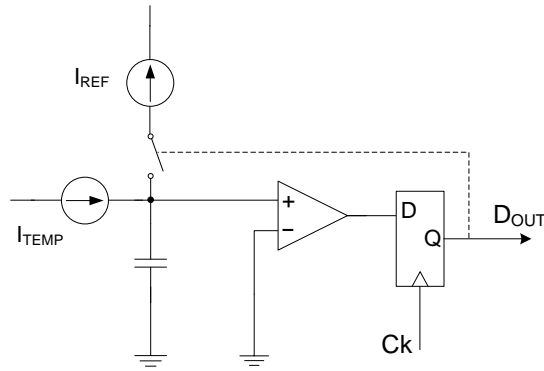


Fig. 2.26 First-order sigma-delta modulator implemented on bipolar technology [Riedijk and Huijsing, 1991].

Both of the above designs used a first-order modulator and digital filter. To produce an N-bit digital value, a digital counter requires 2^N analog samples. Greater accuracy can be obtained with fewer samples if a higher-order $\Delta\Sigma$ modulator is used, at the cost of increased circuit complexity. The first smart sensor design to use a second-order modulator was published more recently [Pertijns, Niederkorn et al., 2005]. The ADC achieves a resolution of $\pm 0.05^\circ\text{C}$, a resolution much greater than the accuracy of a typical bipolar-based integrated temperature sensor. In this situation where the analog front-end limits the sensor's overall accuracy, creating a sensor with the highest possible accuracy requires the use of analog techniques discussed earlier in this chapter. Curvature, offsets, and low-frequency noise can be reduced by methods such as bias current calibration, chopping, and dynamic element matching.

The basic bandgap reference generators discussed in Section 2.2 were self-biasing, and produced both V_{BE} and ΔV_{BE} from the same core circuit. This design [Pertijns, Niederkorn et al., 2005] separates the bias generator and the two temperature-dependent signals into three separate sub-circuits. Creating the raw V_{BE} and ΔV_{BE} signals separately provides more flexibility for the V_{BE} generator to implement a calibrated current source (Section 2.1) and the ΔV_{BE} generator to implement base current compensation (Section 3.5). Both circuits also implement chopping and dynamic element matching (Section 2.5) to reduce the offset of CMOS op-amps and current sources to negligible levels. Fig. 2.27 shows simplified schematics of the circuitry used to create V_{BE} and ΔV_{BE} .




Fig. 2.27a




Fig. 2.27b

Simplified ΔV_{BE} (2.27a) and V_{BE} (2.27b) generator circuits (bias current generator not shown) [Pertijs, Niederkorn et al., 2005].

Unlike the earlier smart sensor of Fig. 2.26, this design does not supply temperature-dependent signal and reference currents to the $\Delta\Sigma$ modulator input. Rather, the raw V_{BE} and ΔV_{BE} currents created by the circuitry in Fig. 2.27 are supplied directly to the modulator input, and the summing action required to produce a bandgap reference voltage (Eq. 2.10) is performed by the modulator itself (Fig. 2.28). This ‘implicit reference’ technique increases system accuracy by eliminating errors introduced by intermediate summing circuitry, and is discussed further in Section 3.7.

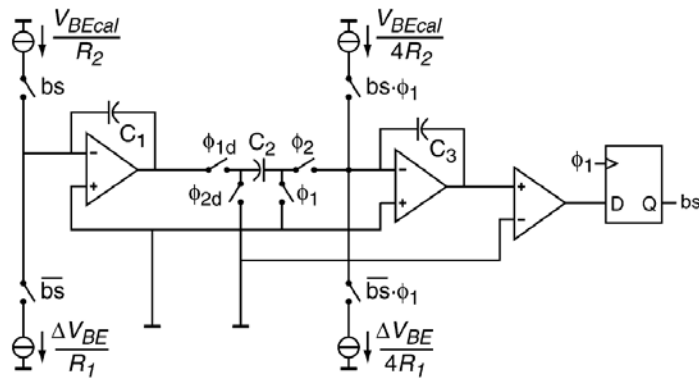


Fig. 2.28 The second-order $\Delta\Sigma$ modulator used by [Pertijs, Niederkorn et al., 2005], showing the connection to the analog bias currents of Fig. 2.27.

The second-order modulator circuitry uses a combination of continuous-time current and discrete-time switched-capacitor techniques. The modulator output is connected to an on-chip second-order symmetric decimation filter. The complete smart sensor occupies 2.5mm^2 on a $0.5\mu\text{m}$ CMOS process.

2.7 Optimal Filtering for Incremental Delta-Sigma ADCs

The first- and second-order delta-sigma ADCs discussed in Section 2.6 are operated in a way that is fundamentally different to conventional $\Delta\Sigma$ converters used in telecommunications and audio applications. The AC spectral characteristic of conventional converters is a primary performance measure while less attention is paid to DC characteristics such as linearity and offset. The ADCs used in low-frequency instrumentation and measurement applications however, do not process an uninterrupted waveform. Converters in these applications are required to accurately reproduce a near-DC signal in the digital domain. Each sample is independent from previous values, and performance measures of interest are linearity, gain, and offset.

To eliminate the influence of previous values on the present conversion, the integrators and digital filter of a $\Delta\Sigma$ ADC must be reset prior to the commencement of conversion. This is known as *incremental* or *one-shot* operation. The operation of this style of conversion is best analysed using discrete time-domain techniques, rather than the ‘linear + white noise’ approximation used in the analysis of conventional converters.

Digital filters for continuously-operating $\Delta\Sigma$ ADCs are commonly implemented using the economic *sinc* structure [Hogenuer, 1981], which possess a symmetric impulse response and are easily implemented without complex multipliers. The transfer function of a second-order sinc filter is:

$$H(z) = \frac{1}{N^2} \left(\frac{1 - z^{-N}}{1 - z^{-1}} \right)^2, \quad (2.36)$$

where N is the oversampling ratio, and the total filter length $L = (2N - 1)$. The block diagram for this second-order filter is shown in Fig. 2.29.




Fig. 2.29 Block diagram of a second-order symmetric sinc decimation filter.

The optimum trade-off between resolution and filter complexity is obtained when the digital filter is of order one greater than that of the modulator [Schreier and Temes, 2005, Chapter 3], so a second-order $\Delta\Sigma$ modulator would usually be matched with a third-order sinc filter. However this rule does not necessarily apply to incremental converters where the underlying assumption of continuous operation is invalid.

Robert and Deval [1988] used time-domain analysis to determine that a second-order non-symmetric filter is optimal for second-order incremental modulators. A non-symmetric filter is known as a *Cascade of Integrators* (CoI); the transfer function of a second-order CoI filter is:

$$H(z) = \frac{1}{L^2} (1 - z^{-1})^{-2}, \quad (2.37)$$

where L is the total filter length.

Fig. 2.30 illustrates the difference in impulse response between sinc and CoI filters. The weights for both types of filter are plotted for a total filter length $L = 100$. The impulse response of a simple first-order filter (counter) is also plotted for comparison. It is apparent that due to their non-symmetric impulse response, CoI filters cannot suppress periodic noise originating from the input signal or within the modulator itself such as modulated offsets from the chopping commonly used with CMOS op-amps (Section 2.5).




Fig. 2.30 Impulse response of first- and second-order digital filters with filter length $L = 100$.

The second-order modulator implemented by Pertijs, Niederkorn et al. [2005] required a symmetric decimation filter to average out residual mismatches modulated by the chopped op-amps used in Fig. 2.25. Based on the results given in [Robert and Deval, 1998], they chose to implement a symmetric second-order filter. However Robert and Deval did not publish an analysis of symmetric (sinc) filters, and the use of a second-order sinc filter with a second-order incremental modulator is not supported by their analysis.

A series of recent publications has thoroughly analysed the characteristics of single-loop incremental converters, with interesting results [Márkus, Silva and Temes, 2004; Márkus, 2005; Márkus, Deval et al., 2006]. It is shown that a second-order non-symmetric filter is indeed optimal for a second-order modulator, especially as the last integrator's output at the end of the conversion can be easily sampled to obtain one extra bit of resolution. But if it is necessary to remove periodic noise from the modulator output – mains-induced hum or MOS offsets modulated by chopping – then a symmetric filter is required. Márkus et al. have shown that a second-order sinc filter requires around twice the number of cycles required by an ideal filter to achieve the same resolution. Increasing the filter order to three, however, brings the required number of cycles back down to a value little more than the theoretical minimum [Márkus, Silva and Temes, 2004].

From these results, it is clear that the ADC used by Pertijs, Niederkorn et al. [2005] could be improved by implementing a third- rather than a second-order sinc filter.

2.8 Discrete Techniques for Accurate $\Delta\Sigma$ Modulators

Previous sections in this chapter have discussed techniques to reduce curvature and other sources of error in the front-end temperature sensing circuitry. A $\Delta\Sigma$ ADC is also sensitive to imperfections in the analog modulator circuitry, which can affect the accuracy of the complete smart sensor. Therefore it is important to investigate circuit techniques for the production of accurate $\Delta\Sigma$ modulators.

The modulator loop can be implemented with either continuous-time or discrete switched-capacitor (SC) circuitry. Modern implementations usually prefer switched-capacitor circuitry because of the high linearity and well-defined gains provided by matched on-chip unit capacitors. The discrete-time nature of their operation also eliminates sensitivity to clock jitter [Schreier and Temes, 2005, Chapter 6]. The only disadvantages – the need for anti-aliasing filters and the requirement for increased op-amp bandwidth – are not a significant concern in a low-bandwidth application such as temperature sensors.

Switched-Capacitor Integrators

All the first- and second-order $\Delta\Sigma$ modulators described in Section 2.7 employed integrators as the active loop filter elements. Integrators are universally preferred in $\Delta\Sigma$ modulators as they are easily and accurately implemented with switched-capacitor circuitry, and allow the creation of high-quality NTF filter notches [Adams, 1997]. A simple integrator constructed with switched-capacitor techniques is illustrated in Fig. 2.31a. The two switches are implemented with MOS transistors, and the gate voltages are controlled by non-overlapping clock phases $\Phi 1$ and $\Phi 2$ (Fig 2.31b).




Fig. 2.31a




Fig. 2.31b

Simple switched-capacitor integrator (2.31a) and switch clock phases (2.31b).

The output voltage of the integrator in Fig. 2.31a can be expressed in the time domain as:

$$V_{out}(n) = V_{out}(n-1) - \frac{C_1}{C_{int}} V_{in}(n-1), \quad (2.38)$$

which when converted to the z-domain results in the transfer function:

$$\frac{V_{out}(z)}{V_{in}(z)} = -\left(\frac{C_1}{C_{int}}\right) \frac{z^{-1}}{1-z^{-1}}. \quad (2.39)$$

From the above equations it can be seen that the integrator of Fig. 2.31a both inverts the input signal, and delays it by one clock cycle (a factor of z^{-1}). However these equations do not include second-order effects due to parasitic capacitances, particularly between node A and ground (Fig. 2.32). This unwanted capacitance is caused by the MOS switches as well as top-plate parasitics on C_1 , both of which are poorly controlled. This directly affects the integrator gain by modifying the value of C_1 in Eq. 2.39.




Fig. 2.32 Parasitic capacitances present at node A affect the accuracy of the simple SC integrator.

An improved parasitic-insensitive SC integrator is shown in Fig. 2.33 [Johns and Martin, 1997, Chapter 10]. Any parasitic capacitance present at node A is charged to the input voltage by switch S_1 , then discharged to ground by S_2 . Thus the parasitic capacitance may load the input voltage source slightly more, but the integrator accuracy is not affected. The voltage at node B alternates between ground and the op-amp's virtual ground, and unwanted capacitance on this node has no effect on circuit operation.

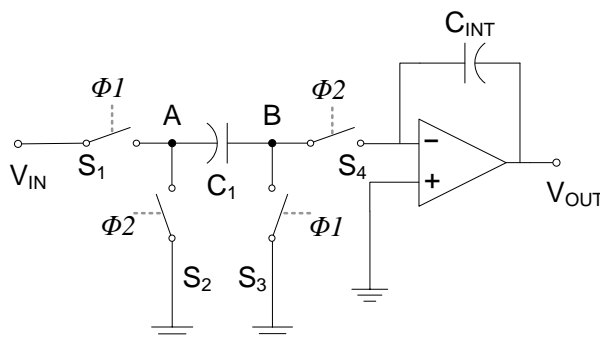


Fig. 2.33 Parasitic-insensitive SC integrator.

Besides unwanted capacitances, the MOS switches used in SC circuitry contribute to integrator inaccuracy through a mechanism known as *charge injection*. The effect of this phenomenon is to inject an uncontrolled charge onto internal nodes when the MOS switches turn off. This charge originates primarily from the channel charge present in a triode-region MOS device, although gate-drain overlap capacitance also contributes to the problem [Johns and Martin, 1997, Chapter 7]. Not only is the value of this channel charge poorly-controlled, the distribution of this charge between drain and source terminals depends on the relative impedances at each node during switch turn-off, as well as the rate of change of gate voltage during the on-off transition [Razavi, 2001, Chapter 12] (Fig. 2.34). Errors due to charge injection are thus extremely difficult to estimate through simulation.

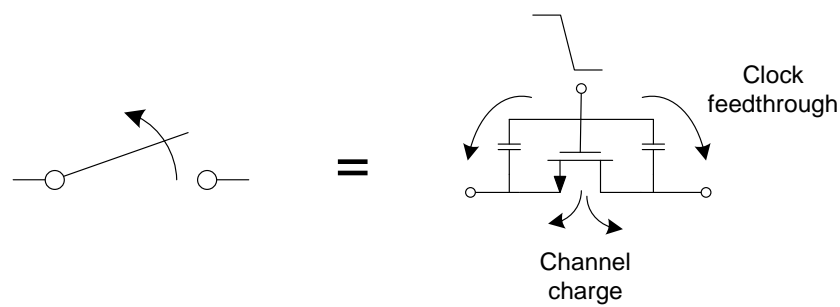


Fig. 2.34 Charge injection from a MOS switch caused by channel charge and parasitic gate capacitances.

Fortunately, charge injection errors can be reduced by a simple modification to the parasitic insensitive integrator introduced above. Shown in Fig. 2.35, the modification involves delaying the falling clock edge to switches S_1 & S_2 . The input capacitor C_1 is floating when switches S_1 & S_2 are turned off, thus channel charge contributed by these devices does not alter the charge present on C_1 . Switches S_3 & S_4 still inject charge on C_1 , but as the voltage at these switch's terminals is either circuit ground or virtual ground, the charge injected is independent of the input signal and can thus be considered a DC offset.

The integrator of Fig. 2.35a reduces the effects of charge injection to a constant DC offset. To eliminate this source of error completely requires the use of differential circuitry. As the name suggests, differential circuitry uses the difference between two voltages to represent a value. The average (common-mode) voltage carries no information, and is usually set halfway between the supply rails.

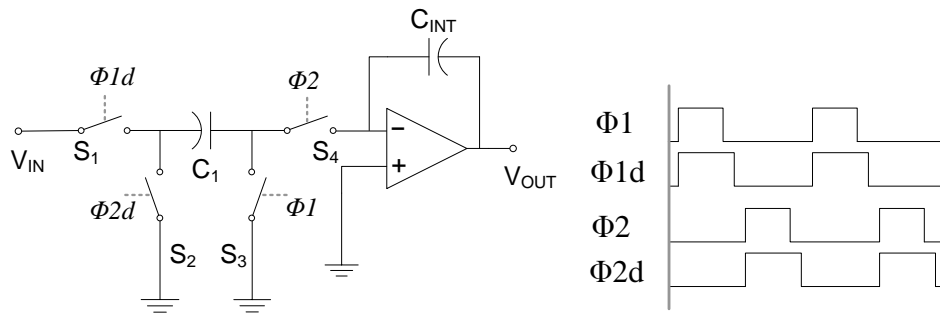


Fig. 2.35a

Fig. 2.35b

Parasitic-insensitive SC integrator modified to reduce switch charge injection (2.35a), and clock waveforms (2.35b).

Generally speaking, fully differential circuits offer increased signal swings and greater immunity to external noise sources than single-ended designs, both of which are significant advantages in a low-voltage mixed-signal IC [Johns and Martin, 1997, Chapter 6]. In the case of a precision SC integrator, a well-laid out differential version of Fig. 2.35a will exhibit virtually identical charge injection on C_1 from switches S_3 & S_4 . The resulting common-mode offset will have virtually no effect on the integrator's differential output (Fig. 2.36), allowing the integrator to achieve a high level of precision.




Fig. 2.36 Differential charge injection-insensitive parasitic-insensitive SC integrator.

Correlated Double-Sampling

Section 2.5 described analog techniques to reduce the effects of voltage offsets and low-frequency noise present in CMOS circuitry – the methods presented there were chopping and dynamic element matching. These methods are similar in that they both

modulate the offsets above the signal bandwidth, where they can be removed by filtering. In contrast, *correlated double sampling* (CDS) is a noise reduction technique that samples the unwanted offsets twice, producing a single output value with the constant offsets subtracted [Enz and Temes, 1996]. CDS is easily implemented using the switches and sampling capacitors of SC circuitry.

A SC integrator incorporating CDS is illustrated in Fig. 2.37. A single-ended version is shown for simplicity, based on the integrator of Fig. 2.35. The additional capacitor C_{OFF} stores the op-amp's input offset voltage during $\Phi 1$, and during $\Phi 2$ this stored voltage reduces the op-amp's input-referred offset to virtually zero.

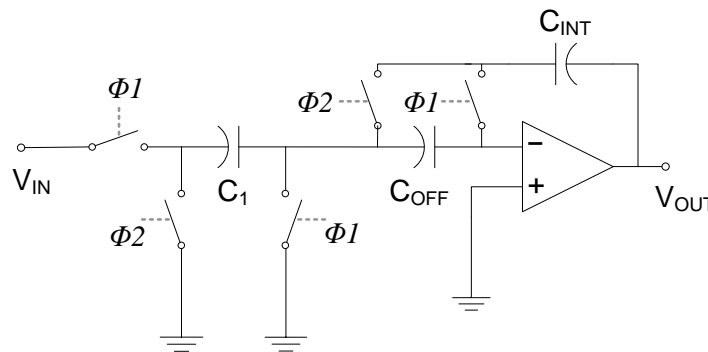


Fig. 2.37 A parasitic-insensitive SC integrator with correlated double-sampling.

The techniques described in this section can create very accurate discrete-time integrators. However, it is not usually necessary for every integrator in a $\Delta\Sigma$ modulator to employ all the techniques described here. Recall from Section 2.6 that the signal transfer function (STF) of a $\Delta\Sigma$ modulator is typically unity, so input-referred offsets and noise of the first integrator are transferred directly to the loop output. This integrator would use CDS and the parasitic-insensitive architecture discussed above. Offsets in the second integrator of a higher-order loop are attenuated by the first integrator's gain when referred to the loop input, so the additional complexity of the high-accuracy integrators discussed above would not be justified for the second and subsequent integrators in the loop.

2.9 A Case Study

This chapter has described techniques to produce accurate on-chip temperature sensors, and how they have been applied to construct smart devices providing an accurate digital output with no external support circuitry. Now we will analyse a smart sensor design with an inaccuracy of $\pm 0.1^\circ\text{C}$, to date the lowest reported for a fully integrated smart sensor [Pertijs, Makinwa et al., 2005].

The largest contributor to error in an IC bandgap reference or temperature sensor is the diode junction voltage V_{BE} . As discussed in Section 2.1, variations in process characteristics cause unpredictable variations in the gradient of V_{BE} . The temperature-voltage relationship also suffers from non-linearity due to fundamental properties of the PN junction. Both of these disadvantages are not present in a ΔV_{BE} signal, assuming sufficient precautions are taken in circuit design and layout.

The design strategy used to achieve the accuracy target of $\pm 0.1^\circ\text{C}$ was to reduce all sources of error except V_{BE} to the level of $\pm 0.01^\circ\text{C}$, and use a calibrated bias current and second-order curvature compensation to extract the highest possible accuracy from V_{BE} . The design is based on the earlier work mentioned above [Pertijs, Niederkorn et al., 2005], although meeting the desired accuracy target required a redesign of all circuit blocks.

A block diagram of the smart sensor designed by Pertijs, Makinwa et al., is shown in Fig. 2.38. The sensor front-end employs a bias current generator separate from the main sensing core. The bias generator block uses a standard bandgap reference structure to produce a PTAT bias current which is supplied to the bipolar sensor core. Chopping (Section 2.5) is employed to reduce offsets in the reference's CMOS op-amp (Section 2.4)




Fig. 2.38 Block diagram of the smart sensor of [Pertijs, Makinwa et al., 2005].




Fig. 2.39 Bipolar sensor core for an accurate temperature sensor [Pertijs, Makinwa et al., 2005].

The bipolar core is conceptually rather simple – consisting of two matched substrate bipolars, current mirrors supplying copies of the PTAT bias current, and MOS switches to direct each unit current source to either substrate bipolar (Fig. 2.39). V_{BE} is produced by directing all 6 unit sources to one substrate bipolar, and shorting the inactive bipolar’s emitter to ground. Variation in the gradient of V_{BE} is the most significant contributor to process-induced variations in the on-chip sensor. These variations can be compensated for by varying the magnitude of the V_{BE} bias current by controlling the MOS current-directing switches in Fig. 2.39. A trimming resolution of 10 bits is achieved by chopping one switch synchronously with the delta-sigma input sampling clock, thus implementing a simple delta-sigma DAC.

The sensor core can also produce ΔV_{BE} by directing one unit current source to one bipolar transistor and the remaining 5 sources to the other. Assuming good matching between the substrate bipolars, the accuracy of ΔV_{BE} depends entirely on the accuracy of this 5:1 current ratio. The switch control logic implements dynamic element matching (Section 2.5) to ensure this ratio is accurately defined.

Section 2.3 described circuit techniques for reducing the curvature present in V_{BE} . In a system incorporating an ADC, curvature in V_{BE} can be reduced by a much simpler method. Introducing a positive temperature coefficient into the bandgap voltage used as the ADC reference (Section 2.1) does not remove curvature in the reference itself. However it does reduce curvature present in the ADC output. Recall from Eq. 2.10 that the output of a bandgap reference circuit can be written as:

$$V_{REF} = V_{BE} + \alpha \Delta V_{BE}. \quad (2.10)$$

The output of a generic ADC using a PTAT temperature signal and bandgap reference is therefore:

$$D_{OUT} = \frac{V_{SIG}}{V_{REF}}$$

$$= \frac{\alpha \Delta V_{BE}}{V_{BE} + \alpha \Delta V_{BE}}. \quad (2.40)$$

Introducing a controlled linear temperature dependence into the denominator of Eq. 2.40 will cause a non-linearity in D_{OUT} that can partially cancel the non-linearity caused by curvature in V_{BE} [Rasmussen, 1994; Pertijs et al., 2001]. In other words, the reference voltage is modified to have a positive temperature dependence. This linear gradient can be created either by increasing the factor α in Eq. 2.40, or by increasing the V_{BE} bias current to reduce its negative temperature coefficient. In practice these methods are complementary, and a balance can be chosen to allow the most convenient circuit implementation.

The $\Delta\Sigma$ ADC consists of a second-order modulator and a second-order decimation filter. The modulator is implemented completely with discrete-time techniques. As discussed in Section 2.8, fully-differential SC integrators are preferred in high-accuracy low-frequency applications such as temperature measurement. Correlated double-sampling is used in the first integrator to eliminate offsets and low-frequency noise both in the front-end sensor and the integrator itself. Low-frequency chopping is also used around the entire modulator, to eliminate any residual mismatch from sampling switches (Fig. 2.40).

The raw signals V_{BE} and ΔV_{BE} connect directly to the $\Delta\Sigma$ modulator input, which implements an ‘implicit’ reference similar to the design analysed at the end of Section 2.6 [Pertijs, Niederkorn et al., 2005]. It is therefore necessary to implement the gain factor α in Eq. 2.40 directly on the first integrator. This can be done by integrating ΔV_{BE} multiple times, or altering the integrator gain by switching in more unit capacitors. A combination of two integrations and an eight-fold gain increase was used to realise a factor of $\alpha = 16$. This double-integration with 8x gain is performed only when integrating ΔV_{BE} , as determined by the comparator output during the previous integration cycle (Fig. 2.40).

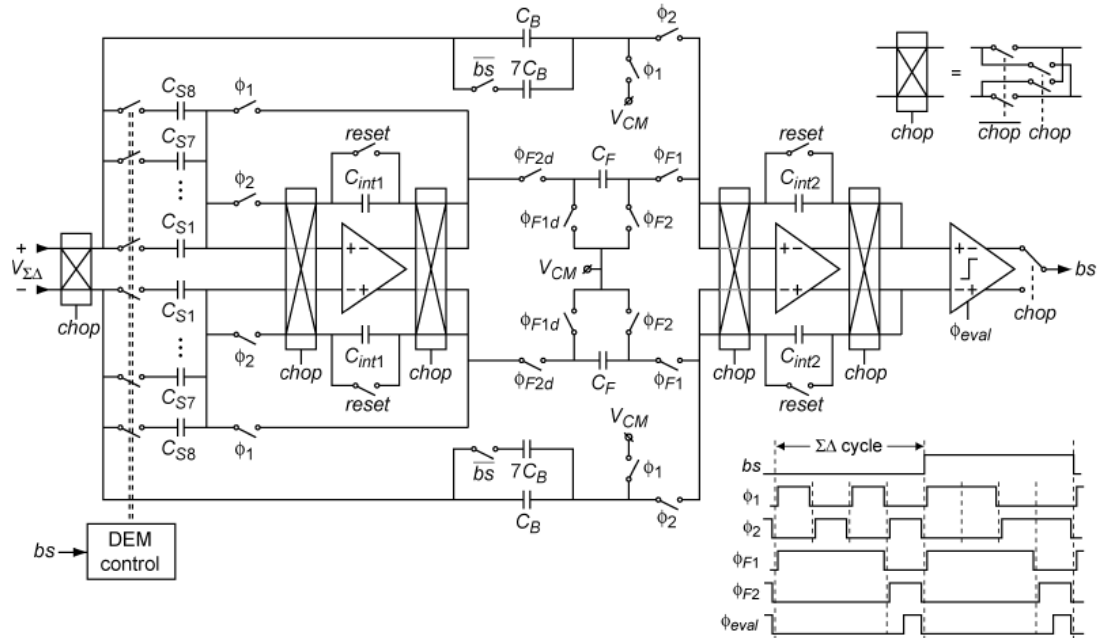


Fig. 2.40 A second-order differential switched-capacitor delta-sigma modulator [Pertjjs, Makinwa et al., 2005].

Similar to the previous design [Pertjjs, Niederkorn et al., 2005], a second-order digital filter was used. Contrary to the previous design, this filter was not implemented on-chip. An off-chip lookup table was also used to correct the small non-linearity remaining after curvature compensation as described above. The system was implemented on a $0.7\mu\text{m}$ process, and occupied a die area of 4.5mm^2 excluding the digital filter.

2.10 Process Compensation

In the advanced design studied in the previous section, several techniques were utilised to minimise many sources of inaccuracy present in a CMOS smart sensor. Despite this, variation in the characteristics of the sensing bipolar transistors still necessitated a single-point calibration – a post-manufacture step that would ideally be unnecessary. A smart sensor design was recently published that utilises batch rather than individual calibration [Aita et al., 2009]. Batch calibration takes advantage of the relatively good correlation between devices manufactured on the same wafer. This design is an optimised version of the case study sensor [Pertjjs, Makinwa et al., 2005], and achieved an accuracy of $\pm 0.25^\circ\text{C}$ over an extremely wide temperature range: -70°C to $+130^\circ\text{C}$.

While batch calibration is certainly an improvement over individual calibration, it is only an intermediate step towards the goal of a truly calibration-free sensor.

Anatomy of Bipolar Process Sensitivity

Rearranging the elementary equation Eq. 2.3, a bipolar transistor's V_{BE} is related to its saturation current I_S by the following equation:

$$V_{BE} = V_T \ln\left(\frac{I_C}{I_S}\right). \quad (2.41)$$

Stated succinctly, the goal of process compensation is to maintain the ratio I_C / I_S at a constant value, to ensure V_{BE} remains unaffected by random process-induced variations. And because the target application is a temperature sensor, this constant ratio I_C / I_S should also be maintained over the entire targeted temperature range.

The process- and temperature-dependent factors in I_S are revealed in Eq. 2.8, restated here for convenience.

$$V_{BE} = V_{G0} + V_T \ln\left(\frac{I_C T^{-\gamma}}{CA}\right). \quad (2.8)$$

The characteristic curve of V_{BE} versus temperature can be described by three parameters (Fig 2.4): intercept, slope, and curvature. The intercept is determined by the physical constant V_{G0} , which is not influenced by processing parameters. The slope of V_{BE} is determined by device layout area A and constant C in Eq. 2.8, which is highly process-dependent. The curvature present in V_{BE} is due to the term $T^{-\gamma}$, which is also influenced by processing conditions.

Were V_{BE} to exhibit significant shifts in both slope and curvature due to processing variations, the design of a process compensation system would be challenging indeed. Fortunately the spread in γ is small enough to be ignored [Pertijis et al., 2001]; experimental results confirm variations in V_{BE} curvature are insignificant relative to variations in gradient [Pertijis, Makinwa et al, 2005]. Therefore the term C in Eq. 2.8 is the main source of process-induced variations, producing a random PTAT-type variation in V_{BE} . A desirable process compensation scheme would therefore create a

substrate transistor bias current with an I_S -dependent gradient, to create an opposing PTAT-type variation in V_{BE} .

Correlated Device Characteristics

In bipolar technology, a *base pinch* resistor is formed by the narrow base region below an emitter implant identical to that used to form a bipolar transistor (Fig. 2.41). These resistors exhibit high voltage modulation and high variability ($\pm 50\%$), and are thus rarely used [Hastings, 2006, Chapter 3].




Figure 2.41a




Figure 2.41b

Cross-section of a bipolar transistor (2.41a) and base pinch resistor (2.41b) in a bipolar manufacturing process [Hastings, 2006].

The most interesting feature of these devices is that due to their construction, the large process variability is correlated to the characteristics of bipolar transistors on the same die. The degree of correlation between two variables is known as the *sample correlation coefficient*, or 'r' [Ledolter and Hogg, 2010, Chapter 1]. The value of r lies between -1 and 1, and its magnitude indicates the strength of linear correlation between two variables. Dutton and Divekar reported a linear relationship between pinched base resistance R_P and I_S with a correlation coefficient of $r = 0.83$ for a 1970s-era bipolar process [Dutton and Divekar, 1977]:

$$R_P = m \cdot I_S + c. \quad (2.42)$$

The square of the correlation coefficient is known as the *coefficient of determination*, or R^2 . This value expresses the exact proportion of variation in one variable that can be explained by the other. The above research therefore indicates a value of $R^2 = 0.69$, meaning 69% of the variation in I_S can be explained by variation in R_P .

A technological compensation circuit has been presented to reduce the effect of I_S variations on V_{BE} in bipolar technology [Amador et al., 1998], based on the data published by Dutton and Divekar. The compensation circuit subtracts a R_P -dependent current from a constant current supplying a bipolar transistor (Fig. 2.42). A smaller value of R_P increases the current I_{COMP} diverted away from the bipolar transistor, decreasing I_{BIAS} to compensate for the decrease in I_S (Eq. 2.42).




Fig. 2.42 Conceptual circuit diagram of V_{BE} process-compensation circuit [Amador et al., 1998].

Simulations showed the worst-case error in an intrinsically-referenced temperature sensor was reduced by a factor of five, using the process parameters reported by Dutton and Divekar. Unfortunately no information was presented regarding the correlation of R_P and I_S in more modern manufacturing processes, nor was a prototype device presented. Despite this lack of information, the possibility for a temperature sensor to ‘self-compensate’ is a fascinating one, and is a topic that will be pursued further in this thesis.

2.11 Summary

Smart temperature sensors require an accurate method of measuring on-chip temperature that is both compatible with standard process technologies and insensitive to variations in manufacturing parameters. The most suitable devices for this task are bipolar transistors, which can generate voltages with both positive and negative temperature coefficients. Modern CMOS processes retain the ability to create limited-

function bipolar transistors that remain the best option for temperature sensing and reference generation.

The continuous decrease in minimum feature size has enabled the integration of increasingly sophisticated on-chip A-D converters. Delta-sigma converters are ideally suited to accurate conversion of low-frequency signals, and are tolerant of analog imperfections. Converters incorporating second-order modulators far exceed the accuracy of on-chip bipolar temperature sensors, although recent research suggests the correct choice of digital filter could improve accuracy further.

Advanced smart sensor designs can achieve high accuracy and very high resolution on standard CMOS processes. Curvature and random variations in V_{BE} can be compensated for, offsets and low-frequency noise can be virtually eliminated with analog circuit techniques, and the temperature value can be converted to the digital domain very accurately despite the limited accuracy of on-chip devices. The close correlation between the characteristics of base-pinch resistors and bipolar transistors suggests that it may be possible to create a sensor with reduced sensitivity to process-induced variations.

CHAPTER 3

DESIGN METHODOLOGY

Chapter 2 presented the history of on-chip temperature sensing techniques, and how they can be used to create a fully integrated smart sensor. In order to develop a system that improves on previous research, it is first necessary to establish a framework to evaluate the strengths and weaknesses of existing designs (Section 3.1). Opportunities for improvement can then be identified within this framework (Section 3.2). The remainder of Chapter 3 will describe circuit techniques that will enable the new system to achieve the desired improvements.

3.1 Design Evaluation Criteria

The purpose of a temperature sensor is, of course, to measure the temperature of a chosen object and make that measurement available to the outside world. The task of thermally coupling the sensor to its environment cannot be performed by the sensor designer; thus when evaluating different sensor designs the sensor itself is usually considered the object whose temperature is being measured. Many conventional and unconventional techniques can be used to convey this information for further processing. In the interest of brevity, only those sensors providing an output in the electrical domain will be considered here.

Different design techniques produce sensors with different measurement characteristics. The fundamental quality of how ‘well’ a sensor works is commonly understood as how close the sensor’s output is to the actual temperature of interest. This is known as *accuracy*. Accuracy can be divided into several sub-categories. In its simplest form, the accuracy of a sensor is defined at one temperature, at one point in time. The variation of accuracy over temperature is known as *linearity*. The variation of accuracy over time is known as *repeatability* in the short term, and *drift* in the long term. The variation of accuracy between different sensors produced using the same manufacturing techniques is known as *manufacturing variation*, or *process sensitivity* when referring to sensors

constructed on an integrated circuit. Should the information be converted to the digital domain, accuracy is dependent on the *resolution* of the conversion.

The above sensor characteristics apply only to the temperature output value. Sensors implemented in the real world are subject to several other practical constraints. Any electronic device requires power to operate, and in some (remote, battery-powered) applications *power consumption* is a characteristic of interest. *Ease of implementation* refers to the amount of engineering effort required for a given sensor to be operational and successfully interfaced with the rest of a system. And of course, *cost of implementation* is an important quantity in many applications. This encompasses not only the sensor unit cost, but also the cost of implementing any support or interfacing circuitry, as well as any per-unit or per-batch calibration required by the sensor.

The above characteristics are summarised in Table 3.1. The second column indicates whether the given quantity will increase or decrease for a ‘better’ sensor design.

Table 3.1 Characteristic quantities for a generic temperature sensor, and the direction of movement that will produce a better sensor.

Quantity of Interest	Desired Direction
Accuracy	
Linearity	+
Repeatability	+
Resolution	+
Drift	-
Manufacturing variation	-
Practical considerations	
Ease of implementation	+
Cost of implementation	-
Power consumption	-

As described in Chapter 1, different sensor designs offer different trade-offs in the quantities of Table 3.1, and are thus suited for different applications. A smart integrated temperature sensor takes advantage of the capabilities of integrated circuit manufacturing techniques to produce a low-cost sensor requiring few external components. This type of sensor is therefore most suited to high-volume applications with low-to-medium accuracy requirements and a digital interface to the rest of the

system. The weaknesses of a typical smart sensor lie in the areas of linearity and manufacturing (process) variation. These characteristics are summarised in Table 3.2.

Table 3.2 The strengths and weaknesses of a typical smart integrated temperature sensor.

Quantity of Interest	'Smart' IC Sensor
Accuracy	
Linearity	Fair / Poor
Repeatability	Good
Resolution	Good
Drift	Good
Manufacturing variation	Poor
Practical considerations	
Ease of implementation	Excellent
Cost of implementation	Good
Power consumption	Good / Excellent

The advanced smart-sensor design described in Section 2.9 [Pertijis, Makinwa et al., 2005] continues to leverage the advantages of IC process technology in the form of low power, high integration, and low cost. The designers expended considerable effort to improve on the weaknesses of smart sensors. Two stages of curvature correction were used to increase linearity, and circuit techniques reduced many sources of inaccuracy so that a single-point calibration adequately compensated for the remaining process variations in V_{BE} . The resulting smart sensor had an inaccuracy of $\pm 0.1^\circ\text{C}$ and resolution of $\pm 0.01^\circ\text{C}$ over the temperature range of -55°C to $+125^\circ\text{C}$. This advanced design will be considered the 'base-line' or reference design, which further research and development work presented in this thesis will be evaluated against.

3.2 Desirable Design Characteristics

The focus of this thesis is the development of an improved smart temperature sensor. An improvement may manifest as a change in any one of the characteristics listed in Table 3.1, producing a movement in the desired direction. Particularly of interest are the weaknesses of smart sensors. Increasing the linearity and reducing sensitivity to process variations will produce a sensor capable of superseding other sensor types in many

applications; the greater manufacturing volumes will in turn lower unit costs for a mass-produced IC.

Despite its advanced design and high accuracy specification, several areas can be identified in the work of Pertijs, Makinwa et al. [2005] that present opportunities for further design innovations. Firstly, the $0.7\mu\text{m}$ manufacturing process used for fabricating the sensor is considered rather elderly by today's standards. Moving to a more modern process will primarily result in the reduction in size of digital logic, used for creating sequencing subsystems and filters for delta-sigma modulators. The use of more modern processes also introduces design challenges – these will be discussed in Section 3.3.

Secondly, the previous work used a second-order delta-sigma modulator coupled with an off-chip second-order digital filter. The research into incremental delta-sigma modulators described in Section 2.7 indicates that a third-order filter used with the same modulator will produce a similarly accurate conversion result with fewer cycles. A practical smart sensor must be as self-contained as possible, particularly such complex sub-systems as delta-sigma filters. An opportunity therefore exists to create an improved sensor by integrating a suitable third-order filter on-chip, operating with fewer cycles and occupying less area than was possible in previous work.

And finally, variations in the fundamental characteristics of the bipolar sensing transistors (Section 2.1) have historically required a multi-point calibration process. The extensive design work done by Pertijs, Makinwa et al. has eliminated random sources of error present in typical CMOS analog circuits and revealed the fundamental limits of the bipolar sensors. Variations in saturation current and emitter layout area cause variation in the gradient of V_{BE} (Fig. 2.4). This one-dimensional random variation can be corrected by a single-point calibration at room temperature, which is considerably simpler than previous designs requiring calibration at two distinct temperatures.

However, the fact remains that each device requires a calibration step. This is acceptable in high-accuracy low-volume applications, yet one of the primary advantages of smart sensors is their ability to achieve low cost through high-volume IC manufacturing. This advantage is somewhat compromised by the need for each sensor

to receive individual attention. An ideal smart sensor would require no calibration, and any post-manufacture processing steps would be common to all devices. A possible solution to this problem was described in Section 2.10, by taking advantage of the correlation between bipolar transistors and base pinch resistors constructed with the same processing steps.

Based on the above discussion, the desirable characteristics of an improved smart temperature sensor developed in this thesis can now be enumerated. This improved sensor would of course retain the advantages common to all smart sensors: **low cost** and **ease of system integration**. It would also incorporate the improvements in **linearity**, **resolution**, and **reduced calibration requirements** developed by Pertijs, Makinwa et al. And most importantly, the improved sensor would offer **lower cost**, **reduced layout area** and **more complete integration** than any previously published design through the use of a modern manufacturing process. Its A-D converter would offer similar output resolution at **reduced conversion time** by using the results of recent research into the characteristics of incremental delta-sigma modulators. And it would offer **greatly reduced sensitivity to process variations** by taking advantage of the correlation between base-pinch resistors and bipolar temperature sensing transistors. The remaining sections of this chapter will describe the methods used to implement these improvements.

3.3 High-Level System Design

Conceptually, an advanced smart temperature sensor on CMOS technology can be divided into two distinct sub-systems: the analog sensor and the ADC (Fig. 3.1). The design developed in this thesis follows the template of [Pertijs, Makinwa et al. 2005], therefore the sub-systems have similar structures (Fig. 2.38). The sensor front-end uses substrate bipolar transistors as the sensing elements, and provides V_{BE} and ΔV_{BE} signals directly to the input of the delta-sigma modulator. Traditional bandgap reference circuits are self-biasing, i.e. the bias current flowing in the bipolar transistors is derived from the PTAT voltage generated by the transistors themselves. Separating the tasks of bias current generation and temperature sensing allows easier implementation of base

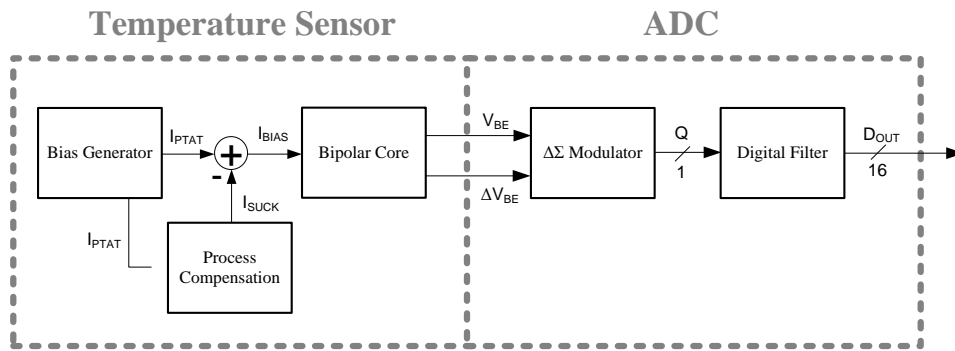


Fig. 3.1 Block diagram of the proposed smart temperature sensor.

current compensation, calibration, and interfacing with the ADC – all critical to produce a sensor with the highest possible accuracy.

The on-chip delta-sigma ADC is significantly more complex than the temperature sensor, and therefore occupies the largest proportion of layout area. It can be divided into two sub-components – a discrete-time analog modulator, and a digital filter. The filter will be implemented on-chip to achieve the highest possible system integration.

Before any detailed design work is undertaken, the first decision to be made is the choice of target IC manufacturing process. A purely digital design benefits from the reduced feature size of more modern processes without affecting functionality. Analog circuitry does not physically scale at the same rate as digital logic, while secondary considerations such as reduced supply voltage and short-channel effects complicate the migration path to a newer process. Furthermore, analog circuits typically require accurate passive components – resistors, capacitors, inductors – that are not available on the newest processes targeted at digital applications. Process extensions that provide such passive components take time to develop, thus analog designs that require these passive devices are only possible on processes a generation or two behind the leading edge.

Another issue to consider is the cost of producing a given design. As a new process enters high-volume production the cost per unit will reduce. Over time this process will be superseded and production volumes will then reduce, and newer process will offer lower unit costs due to the reduced chip area requirements when the design is migrated forward. In general the best choice is the most modern process that has entered high-

volume production, given any special limitations of the design such as analog process extensions or supply voltage.

The design of Pertijs, Makinwa et al. was implemented on a 0.7 μ m process. While processes of this generation are still available, they are seldom used for state-of-the-art designs. At the time of writing (mid-2009), 45nm digital processes are in volume production, and 32nm is in development. As described in Section 3.2, using a modern process provides the advantages of lower cost, reduced layout area, and more complete integration. These are all desirable characteristics of an advanced smart temperature sensor.

So, which process is the best to target? Of course there is no single correct answer. A good candidate would be one just behind the leading edge. Take the 130nm process node as an example. The Pentium IV 'Northwood' processors of 2003-2004 were manufactured using technology of this generation, and represented the absolute state-of-the-art at that time. IBM now offers a 130nm process with eight metal layers, and passive devices targeted at analog RF applications. While the inductors are of no interest in a low-frequency design such as a temperature sensor, the precision polysilicon resistors and linear high-density metal-insulator-metal (MIM) capacitors are. The combination of desirable passive components and high-density digital logic make this process an ideal choice for a modern smart sensor.

As a general rule, the smaller a given processes' minimum feature size is, the lower its maximum supply voltage will be. The IBM 130nm process mentioned above operates at 1.2V and 2.5V for core and I/O transistors respectively. This is significantly lower than the maximum of 5V allowed by the 0.7 μ m process of previous work. While the reduced supply voltage has no significant effect on digital logic, the operation of analog circuitry is significantly affected. This primarily manifests as reduced swings available at the output of op-amps, which are critical components of a delta-sigma modulator. Voltage headroom is also a concern in bipolar-based references and temperature sensors. The diode junctions used in such circuits have a forward operating voltage of 0.6V to 0.7V, a value that does not scale with supply voltage. The lower-threshold MOS devices produced by low-voltage processes mean that bipolar cores can be comfortably biased

at 1.2V, although this will become a significant issue when operating voltages drop below 1.0V in future processes.

The most significant disadvantage of high-density modern processes, however, is not the reduced supply voltage. Accurate IC temperature sensors rely on substrate bipolar transistors for their operation, and any degradation in the characteristics of these devices will affect the whole system. Three trends in the design of modern CMOS processes lead to a reduction in substrate bipolar transistor gain - shallow wells with retrograde doping profiles, extremely thin source/drain regions, and the use of silicidation in the source/drain regions [Hastings, 2006, Chapter 8]. In some circumstances, this can result in gains of less than one. The circuit designer has very little control over the characteristics of the manufactured devices, which are largely determined by the process itself. Fortunately gain can be significantly improved by the use of silicide block masks over the emitter region.

Forward-biased diode junctions are rarely used in modern CMOS designs. They can inject minority carriers into the substrate, inducing local voltage gradients and potentially initiating a destructive latch-up condition. For this reason their use is discouraged wherever possible. The 130nm IBM process targeted for this thesis includes modelling and layout support for a substrate bipolar transistor intended for use in bandgap applications. The transistor is not intended for general-purpose use, and is the only forward-biased junction supported by this process for any application. Pertijs et al. reported a substrate bipolar current gain of 22 in their 0.7 μ m process; the device models for the 130nm process predict a gain of around 2. This is low enough to warrant particular care in circuit design, but does not preclude accurate operation of the sensor itself.

This section has outlined the high-level structure of the proposed system, and has discussed the choice of manufacturing process suitable for a modern smart temperature sensor. The remainder of this chapter will describe the internal structure of the component blocks in Fig. 3.1, and how they implement the design criteria developed in Section 3.2.

3.4 Bipolar Core Features

As already described in Chapter 2, accurate on-chip temperature sensors use the predictable temperature characteristics of the PN diode junction to create their sensing elements. On a modern CMOS process, both substrate and lateral bipolar transistors can provide the necessary PN junction. On the target IBM 0.13 μm process however, only the substrate bipolar is modelled or supported in forward-biased applications. It is therefore the only rational choice for implementing a temperature sensor.

The integrated nature of an on-chip system allows the different sub-system blocks to be tightly integrated together. In the case of a smart sensor, this is particularly important for the connection between the sensor and ADC blocks. Connecting the delta-sigma modulator to the output of the sensing transistors as directly as possible is important to avoid errors introduced by intermediate operations such as scaling or addition. Attention must also be paid to methods of eliminating mismatch and low-frequency noise present in CMOS circuits.

As will be described in Section 3.7, the first integrator of the $\Delta\Sigma$ modulator achieves high accuracy by employing a differential switched-capacitor integrator with correlated double-sampling (Section 2.8). The ADC requires a reference voltage to convert the temperature signal to the digital domain, but it does not require the reference to explicitly exist at its input. It is enough for the positive and negative temperature coefficient signals V_{BE} and ΔV_{BE} to be alternately applied, and the feedback action of the $\Delta\Sigma$ modulator performs the addition. This implicit reference technique was published previously [Pertijs, Niederkorn et al., 2005; Pertijs, Makinwa et al. 2005] and is described further in Section 3.7.

Figure 3.2 illustrates the circuit configurations used to generate V_{BE} and ΔV_{BE} . Both circuits use a matched pair of substrate bipolar transistors as sensing elements. Both configurations can invert their output polarity without any MOS switches in series with the output signal, which can introduce thermal noise in the critical signal path. Instead the inversion is achieved by steering bias currents between the two sensor transistors. This switching action has the favourable effect of eliminating offsets in the sensing




Fig. 3.2a

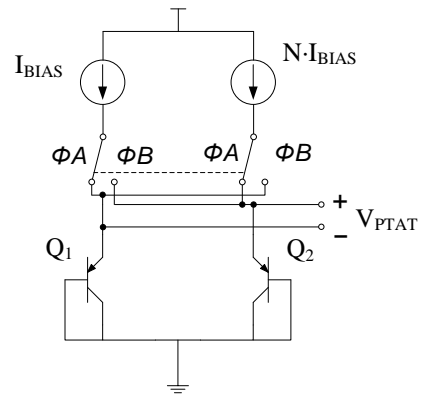


Fig. 3.2b

Conceptual circuits to generate V_{BE} (3.2a) and ΔV_{BE} (3.2b).

transistors, as any mismatch will be cancelled out by the double-sampling used in the first integrator (Section 3.7).

Figure 3.3 shows a circuit capable of implementing both Figs. 3.2a and 3.2b. Once again, a matched pair of substrate bipolars is directly coupled to the output. The bias currents of Fig. 3.2 are implemented as 5 unit current sources. Cascoding is used to increase the sources' output resistance, a technique still possible at the low supply voltage of 1.2V. V_{BE} is generated by directing all 5 current sources to one substrate bipolar and shorting the other bipolar's emitter to ground; ΔV_{BE} is generated by connecting a single current source to one transistor and the remaining 4 to the other. Control of the switches is achieved with a small block of digital logic.

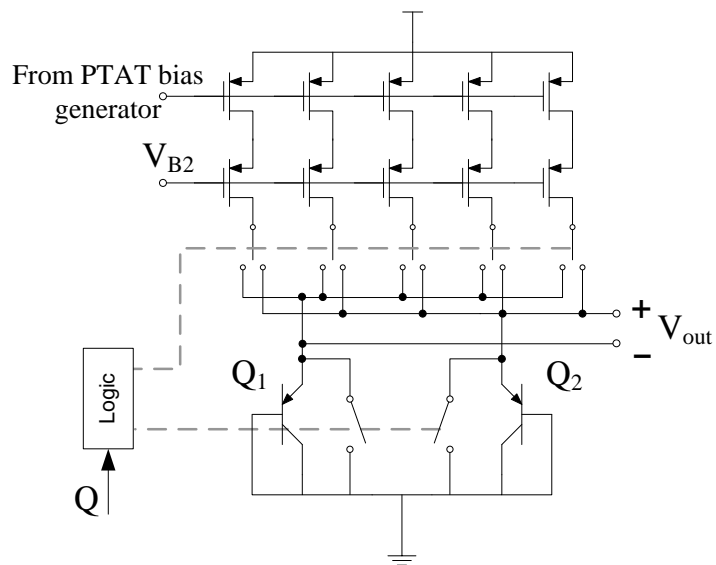


Fig. 3.3 Circuit for generating either V_{BE} or ΔV_{BE} .

The value of V_{BE} depends on the absolute value of its bias current. Calibrating this bias current is the method used by Pertijs, Makinwa et al. to compensate for I_S variations in V_{BE} . The value of ΔV_{BE} , however, is not dependent on either bipolar characteristics or the absolute value of bias current – it depends solely on the current *ratio* between the two substrate bipolars (Eq. 2.9). Because ΔV_{BE} is unaffected by the calibration step it is important for its value to be precise by design, to reduce the effect of mismatched unit current sources on the sensor’s overall accuracy. Good layout techniques can reduce mismatches in the current source transistors, but layout alone is insufficient to achieve the high accuracy targets of this system.

Dynamic element matching (DEM) was introduced in Section 2.5, and this is an ideal application of the technique. When ΔV_{BE} is being generated, the single unit source supplying the smaller current can be rotated through all 5 possible sources. This simply involves changes to the digital logic block; no additional analog circuitry is required. The additional DEM logic is not clocked when V_{BE} is generated, as doing so may result in ΔV_{BE} experiencing an uneven distribution between the unit sources.

3.5 PTAT Bias Generator Features

Separating the sensor core of Fig. 3.3 from its bias generator circuit allows the sensor to alternately supply V_{BE} and ΔV_{BE} , and allows the use of DEM to reduce mismatches in ΔV_{BE} . The design of the bias generator circuit then becomes a separate task. A constant current source is the most obvious choice of bias generator, but it is not the best choice. On-chip constant current sources are best generated from bandgap voltage references. Thus the output of such sources is affected by inaccuracies introduced by the voltage-current conversion and calculation of $(V_{BE} + \Delta V_{BE})$, as well as the inherent process variations present in V_{BE} .

Self-biasing bandgap cores (Fig. 2.15) internally use a PTAT current derived from ΔV_{BE} through the action of the feedback loop. The magnitude of this current is much less sensitive to circuit parameters and process variations than a bandgap reference-derived current, and is easily mirrored in one step to bias the sensor core. As discussed in Section 2.3, a PTAT bias current reduces the inherent second-order curvature present

in V_{BE} (Fig. 2.10). Further curvature reduction can be realised by giving the bias current a power law dependence on temperature, thus increasing the exponent α in Eq. 2.18. However such circuitry is quite difficult to implement, and the additional complexity introduces many new sources of error. Thus a simple PTAT characteristic is the optimal choice for biasing the sensor core.

A simplified diagram of the PTAT bias current generator is shown in Fig. 3.4. Based on the circuit of Fig. 2.15b, it adds cascoded current sources for increased accuracy and a small resistor to eliminate sensitivity to bipolar current gain β (explained below). The magnitude of the output bias current is determined by the core's current source ratio, and the value of R_1 . If R_1 changes with temperature, this will introduce non-PTAT variations in the output current. It is therefore important for R_1 to be as predictable and linear as possible. A 'precision' polysilicon resistor is available on the target IBM process; the high linearity specifications of this resistor make it the best choice for implementing R_1 .




Fig. 3.4 Simplified diagram of PTAT bias current generator circuit.

The equations governing the operation of a bipolar transistor (Eqs. 2.3 – 2.8) describe the relationship between V_{BE} and collector current I_C . However the collector terminal of a substrate bipolar transistor is tied to substrate, and the transistor must be biased through its emitter. In this case the I-V relationship is governed by Eq. 2.2, causing the sensor output to be sensitive to process- and temperature-induced variations in bipolar transistor current gain β .

The collector current of a substrate bipolar transistor is related to its emitter bias current through the following relation:

$$I_C = I_E \cdot \frac{\beta}{\beta + 1}. \quad (3.1)$$

It is apparent that increasing the emitter bias current I_E by a factor of $(\beta + 1) / \beta$ would cancel out this gain dependence, and restore the desired relation of Eq. 2.3.

Figure 3.4 contains a simple modification that can achieve this goal. The base current flowing out of Q_2 generates a voltage across the resistor $R_1 / 8$:

$$\begin{aligned} V_{R1/8} &= \frac{8I_{BIAS}}{\beta + 1} \cdot \frac{R_1}{8} \\ &= \frac{I_{BIAS}}{\beta + 1} \cdot R_1. \end{aligned} \quad (3.2)$$

This raises the voltage at the emitter of Q_2 , and the voltage drop across R_1 also increases due to feedback action of the op-amp:

$$V_{R1} = V_{BE2} - V_{BE1} + \frac{I_{BIAS} \cdot R_1}{\beta + 1}. \quad (3.3)$$

The current I_{BIAS} flowing through R_1 is therefore:

$$I_{BIAS} = \frac{(V_{BE2} - V_{BE1}) \cdot (\beta + 1)}{R_1 \cdot \beta}. \quad (3.4)$$

Combining Eqs. 3.1 and 3.4, the substrate bipolars' collector current is therefore:

$$I_C = \frac{(V_{BE2} - V_{BE1})}{R_1}. \quad (3.5)$$

So by introducing a β -dependent factor into the PTAT bias current, the collector current of substrate bipolars in the bias generator can be made an ideal PTAT current. And because this nearly-PTAT current is also the bias current for the bipolar sensor core (Section 3.4), the sensor transistors are also operated with a β -independent collector current.

Offsets in the bias generator circuit are inevitably introduced by the op-amp, as well as mismatch in the substrate bipolar transistors. Figure 3.5 shows how chopping is implemented in the bias generator to eliminate these error sources. The modified circuit

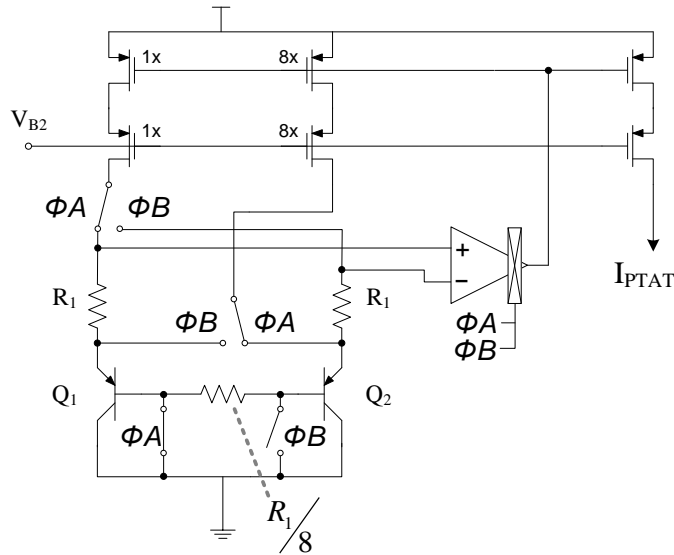


Fig. 3.5 Bias generator with chopping to reduce offset errors.

requires one additional resistor when compared to the un-chopped case of Fig. 3.4. Chopping is implemented in the op-amp output by alternating the connections of the internal current mirror (Fig. 2.18).

The circuit of Fig. 3.5 eliminates mismatch in the bipolars, resistors, and op-amp. However it does not eliminate mismatch in the 8:1 current source ratio. Errors in this ratio result in variation of the absolute value of bias current, and thus cause a PTAT-type error in V_{BE} . Pertijjs, Makinwa et al. [2005] used a single-point calibration to set the absolute value of V_{BE} 's bias current at room temperature. This was primarily to account for variations in bipolar transistor characteristics, but also compensated for gain (but not offset) errors in bias current. The aim of this work is to develop a sensor that requires no calibration procedure, and the design will investigate the possibility of utilising the correlation between pinched-base resistance and bipolar saturation current to reduce process-induced variation. Variations in the 8:1 current mirror ratio are not significantly correlated with pinched-base resistance, so the self-compensation scheme is incapable of accounting for such errors. It is therefore necessary to implement dynamic element matching similar to that used for the sensor core.

The final bias generator circuit is shown in Figure 3.6. Two levels of passgates are used to steer bias currents to either bipolar transistor, as implementing both chopping and DEM with one passgate level would require $4 \times 9 = 36$ individual gates.




Fig. 3.6 Complete PTAT bias generator with chopping and DEM.

3.6 Process Compensation Features

The chopping and DEM used in the bias generator and sensor core reduce the effects of mismatch in the analog front-end, so that the remaining error due to V_{BE} manifests as a variation of its gradient when plotted against temperature (Section 2.10). The calibration employed by Pertijs, Makinwa et al. [2005] varied the magnitude of bias current supplied to generate V_{BE} . This compensates for variations in V_{BE} gradient caused by process-induced variations in bipolar saturation current I_S (Eq. 2.41).

Section 2.10 discussed the correlation between pinched-base resistance R_P and bipolar saturation current I_S , and a process compensation circuit based on bipolar technology was described that reduces the effect of I_S variations on the gradient of V_{BE} . Unfortunately no data is available in the literature on the R_P - I_S correlation for processes more modern than that of [Dutton and Divekar, 1977]. Furthermore, base-pinch resistors are no longer modelled or supported in modern CMOS processes, and have been superseded by high-resistance polysilicon devices offering vastly superior performance characteristics.

Fortunately it is still possible to lay out a base pinch resistor in the target 0.13 μm process without gross violation of design rules. This device forms the heart of the process compensation circuit, discussed below.

The process compensation operates in a conceptually similar way to that of [Amador et al, 1998]. Illustrated in Fig. 3.7, it operates by modifying the bias current supplied to the sensor core as shown in the system block diagram of Fig. 3.1. The lower current source of Fig. 3.7 has a gain dependent on the value of the pinched base resistor R_P . This process-dependent current is subtracted from the main PTAT bias current, and the remainder flows through the transistor generating V_{BE} . Two primary parameters governing circuit operation are the current mirror gains k_1 and k_2 . When I_S takes its nominal value as determined by process parameters, the bias current supplied to the bipolar core must also assume its nominal value of 5 to maintain the operation predicted by the sensor circuit in Fig. 3.3. The mathematics behind the operation of this compensation scheme will be explored further in Section 4.4.




Fig. 3.7 Conceptual operation of process compensation circuit.

The main bias current $k_1 I_{PTAT}$ is supplied by the bias generator (Fig. 3.6), and the compensation current I_{SUCK} is also derived from I_{PTAT} . Substrate bipolar Q_1 represents the sensor core when generating V_{BE} (Fig. 3.3).

The complete compensation schematic is shown in Fig. 3.8, including simplified representations of the bias generator and sensor core circuitry (Fig. 3.1). The process-sensitive resistor R_P is grounded at one end to reduce voltage modulation effects. An op-amp ensures the voltage across resistors R_2 and R_P are equal. More current flows

through R_P when its value decreases, thus creating an R_P -dependent current mirror. The process-dependent bias current is recombined with the unmodified I_{PTAT} in the bipolar core to perform the subtraction.

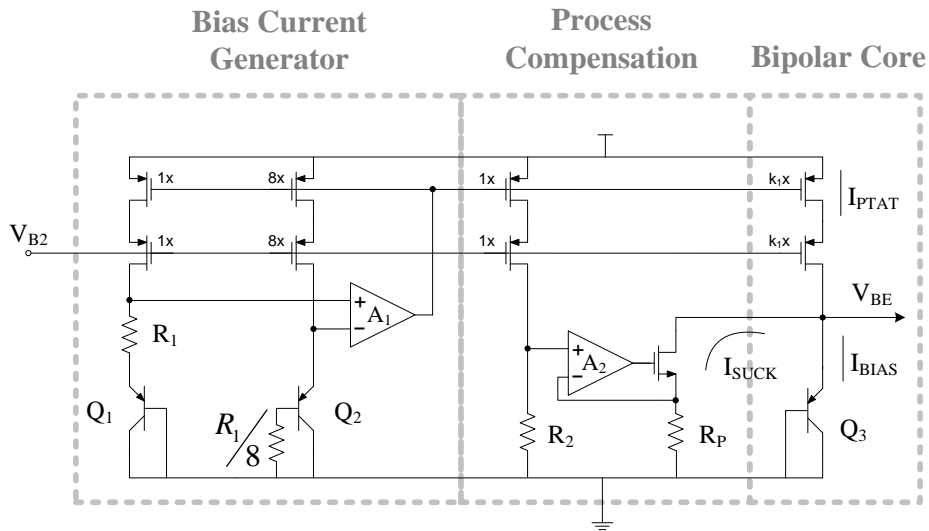


Fig. 3.8 Complete circuit diagram of process compensation scheme.

3.7 Delta-Sigma Modulator Features

‘Implicit Reference’ Input

Every A-D converter requires both an input signal and a reference; the digital output is a dimensionless quantity that represents the ratio between these two inputs. As described in Chapter 2, on-chip temperature sensors use bandgap techniques to generate V_{BE} and ΔV_{BE} . The temperature signal is taken as ΔV_{BE} due to its higher linearity and lower process variability, while the reference voltage is generated by adding V_{BE} and ΔV_{BE} together. Restating Eqs. 2.10 and 2.40 for convenience here, the ADC reference input and digital output are described thus:

$$V_{REF} = V_{BE} + \alpha \Delta V_{BE}, \quad (2.10)$$

$$\begin{aligned} D_{OUT} &= \frac{V_{SIG}}{V_{REF}} \\ &= \frac{\alpha \Delta V_{BE}}{V_{BE} + \alpha \Delta V_{BE}}. \end{aligned} \quad (2.40)$$

The simple $\Delta\Sigma$ modulator of Fig. 2.23 applies a feedback signal of either 0 or V_{REF} to its input summing node. The obvious method of generating V_{REF} in a smart sensor is to explicitly add V_{BE} and ΔV_{BE} , then apply the result to the required nodes in the modulator. This method requires additional circuitry and invariably introduces additional errors. A simple technique eliminates errors from this summing step, and allows the modulator to be connected directly to the bipolar transistors generating the raw V_{BE} and ΔV_{BE} signals.




Fig. 3.9 Summing junction at input of a generic delta-sigma modulator .

Consider the input summing junction of a delta-sigma modulator with an input of $\alpha\Delta V_{BE}$ and a reference of V_{REF} , shown in Fig. 3.9. The PTAT temperature signal $\alpha\Delta V_{BE}$ can only assume positive values, therefore the feedback signal from the comparator swings between 0 and V_{REF} as determined by the comparator output Q . The summing junction output is:

$$V_{SUM} = \alpha\Delta V_{BE}, \quad (3.7a)$$

when $Q = 0$, and

$$\begin{aligned} V_{SUM} &= \alpha\Delta V_{BE} - (V_{BE} + \alpha\Delta V_{BE}) \\ &= -V_{BE}, \end{aligned} \quad (3.7b)$$

when $Q = 1$. This summing junction can therefore be replaced with a multiplexer that selects the raw signals V_{BE} and ΔV_{BE} depending on the value of Q (Fig. 3.10). Furthermore, it is not necessary to implement a separate multiplexer as the current-steering passgates of Fig. 3.3 can implement this function. Therefore the implicit reference scheme described here simplifies the interface between the bipolar sensor core




Fig. 3.10 Multiplexer to replace the summing junction of Fig 3.9.

and the delta-sigma modulator, and does not add any further errors to the signal or reference voltages.

$\Delta\Sigma$ Modulator Loop Structure

Section 2.6 introduced the theory behind the operation of $\Delta\Sigma$ ADCs, and reviewed first- and second-order modulators used in published smart temperature sensors. A temperature signal is essentially a DC value, therefore high conversion speed is not a primary design goal. A second-order $\Delta\Sigma$ modulator can achieve a resolution much higher than a typical on-chip temperature sensor in several hundred clock cycles, and therefore represents an optimal tradeoff between conversion time and circuit complexity.

The design of a $\Delta\Sigma$ modulator involves two distinct tasks - defining the loop structure, and choosing a desirable noise transfer function (NTF). There are many possible modulator loop architectures, although practical implementations typically use one of a few common structures. Two general loop structures that can implement various filter functions are shown in Figs. 3.11 and 3.12. As in Section 2.6, the quantiser is shown as a single-bit comparator which offers inherent linearity and a simple circuit implementation. The active filter elements are integrators, preferred because of their easy implementation, high accuracy, and because they allow the creation of filters with high-quality notches.

The coefficients a_x and b_x in Figs. 3.11 and 3.12 determine the modulator's filtering characteristics, and are chosen according to the specific requirements of each application. Both the signal transfer function and noise transfer function are determined by these coefficients. The selection of these loop coefficients will be discussed in Section 4.5.




Fig. 3.11 Cascade of Integrators, FeedBack (CIFB) delta-sigma modulator structure.

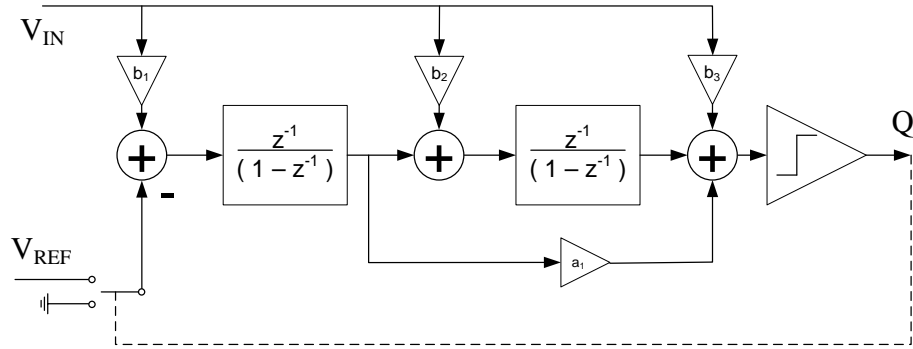


Fig. 3.12 Cascade of Integrators, FeedForward (CIFF) delta-sigma modulator structure.

Both these loop structures can implement the implicit reference technique discussed above, when certain restrictions are observed. In the case of the CIFB modulator of Fig. 3.11, if the coefficients leading to the same node are made equal, i.e.

$$\begin{aligned}
 a_1 &= b_1 \\
 a_2 &= b_2 \\
 b_3 &= 0,
 \end{aligned} \tag{3.8}$$

then the output of the multiplexer of Fig. 3.10 can be fed forward to the input of the second integrator. The CIFF modulator of Fig. 3.12 can implement the implicit-reference scheme most easily by setting $b_1 = 1$, and all other $b_x = 0$. Additional circuitry can also be added to the modulator to allow the ΔV_{BE} signal to be fed forward through b_2 and b_3 , which will allow the implementation of more complex NTFs.

The circuit-level implementation of these modulators can use either continuous-time or discrete-time techniques. Discrete-time differential switched-capacitor circuitry is most commonly used for implementing low frequency on-chip modulators, for reasons already described in Section 2.8. The op-amps used in these integrators experience capacitive loading only, and can therefore be implemented with either a one- or two-stage design. A single-stage folded-cascode topology was chosen for this work as it offers the best balance between the desired characteristics of high output swing capability, high gain, and simple implementation.

Accurate Integration

Section 2.8 has already described techniques to produce accurate discrete-time integrators. Also mentioned was the fact that due to the noise shaping effect of $\Delta\Sigma$ modulators, only the integrator closest to the input is required to achieve a high level of accuracy. A switched-capacitor integrator was presented that used correlated double-sampling (CDS) to eliminate offsets in the integrator's op-amp (Fig. 2.36).

The CDS switched-capacitor integrator used by Pertijs, Makinwa et al. [2005] as the $\Delta\Sigma$ modulator's first integrator differs from Fig. 2.36, and is shown in Fig. 3.13. During phase A, the differential op-amp output is zero, and the input + offsets are sampled on the input capacitors C_1 . At the beginning of phase B, the op-amp is configured as an integrator, and the input signal polarity is reversed. This causes the integrator output to change by a factor of:

$$\begin{aligned} \Delta V_{OUT} &= \frac{C_1}{C_{INT}} \cdot (V_{IN A} - V_{IN B}) \\ &= \frac{C_1}{C_{INT}} \cdot (V_{SIG} + V_{OS} - [-V_{SIG} + V_{OS}]) \\ &= \frac{C_1}{C_{INT}} \cdot 2V_{SIG}. \end{aligned} \tag{3.9}$$




Fig. 3.13 CDS switched-capacitor integrator [Pertijs, Makinwa et al., 2005].

Fig. 3.13 shows the input signal V_{IN} is inverted by way of a 'chopping' block (Section 2.5). Section 3.4 described how the output voltage of the sensor core is itself inverted rather than with additional switches; this is achieved by steering bias currents between two substrate bipolar transistors. The effect of this 'extended' CDS is to eliminate

offsets not only in the integrator's op-amp, but also mismatch between the two substrate bipolar transistors (Fig. 3.3).

However, there is one disadvantage to the above CDS circuit that is revealed by careful simulation. The switches of Fig. 3.13 are implemented using near minimum-sized MOS transistors, which invariably introduce parasitic capacitances (Fig. 3.14). When the op-amp output switches between zero and V_{INT} on every phase transition, parasitic capacitances on switch ΦB cause a small amount of charge to leak off C_{INT} , thereby limiting the DC gain of the integrator. The small value of on-chip capacitors and the high cycle count of the delta-sigma modulator (several hundred) mean that this leakage cannot be ignored in a high-resolution converter.




Fig. 3.14 Parasitic capacitances of a MOS switch.

An improved CDS switched-capacitor integrator is shown in Fig. 3.15, based on the single-ended design shown in Fig. 2.37. Similar to Fig. 3.13, the input signal is still sampled on C_1 during phase A. But unlike Fig. 3.13, the op-amp output remains at V_{INT} during phase A and the op-amp input offset is sampled on a third pair of capacitors C_{OFF} . This eliminates charge leakage due to large voltage transitions at the op-amp output.

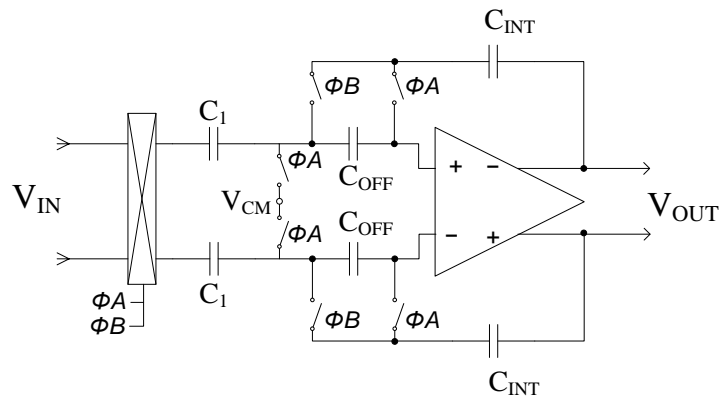


Fig. 3.15 Improved CDS integrator with negligible charge leakage.

A complete circuit-level implementation of a second-order CIFF-type modulator is shown in Fig. 3.16. The first integrator uses the CDS technique described above, and incorporates a gain-boostered folded-cascode op-amp. The second integrator does not require these advanced techniques, so employs a standard differential integrator and simple folded-cascode op-amp. The integrators are reset at the beginning of each conversion by switches across the integrating capacitors. The gains of each integrator as well as the feed-forward paths to the comparator are determined by capacitor ratios, the selection of which will be discussed in Section 4.5.

System-level chopping is also shown around the entire modulator. This eliminates residual errors due to charge injection from the numerous MOS switches, and is done at the lower speed of $(2 / \text{OSR}) C_k$ to avoid introducing further charge injection errors. The input signal, comparator, and both integrators' capacitors are synchronously inverted to maintain correct signal polarity. The modulated residual offsets can then be averaged by the decimation filter, assuming the filter implements a symmetric impulse response.

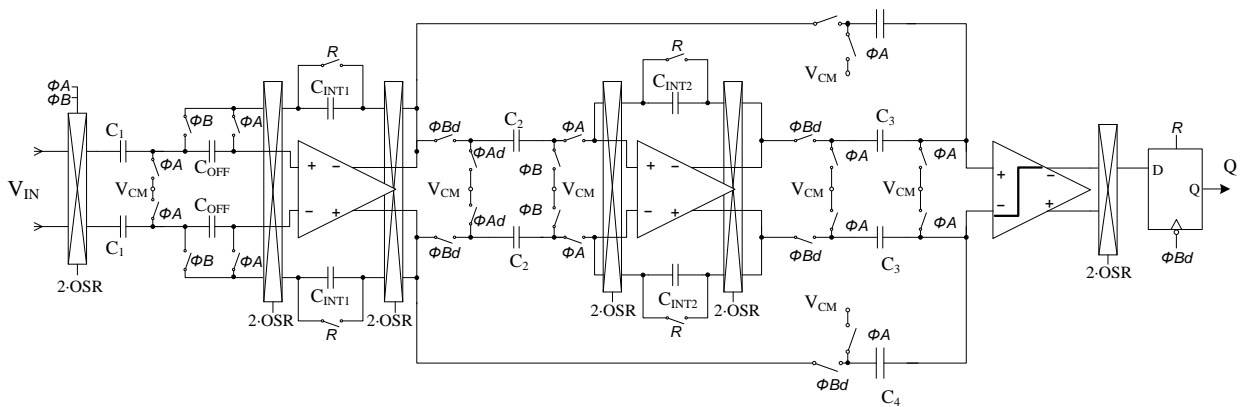


Fig. 3.16 Complete switched-capacitor modulator employing CIFF structure.

3.8 Digital Filter Features

The purpose of the digital filter is to remove the quantisation noise that has been modulated to frequencies above the signal band, while allowing the signal to pass through unmodified. As with the choice of modulator loop structure, many low-pass filter designs are possible yet only a few are commonly used. General-purpose FIR and

IIR-type filters require multiplication operations, and can occupy significant die area even on modern processes. The sinc structure introduced in Section 2.7 is easily the most common filter used in delta-sigma applications due to its acceptable performance coupled with extreme simplicity. Most importantly, it possesses the necessary symmetric impulse response to average out modulated offset residuals from the modulator's system-level chopping.

A block diagram of a third-order sinc filter employing the Hogenauer structure is shown in Fig. 3.17. The filter input Q is driven by the output of the $\Delta\Sigma$ modulator's comparator, which is a single bit in the modulator developed in Section 3.7. Arithmetic is performed within the filter by equipping each register with a summation block, denoted by the circular summation symbols. The width of the internal registers is determined by the filter order and the oversampling ratio, which are in turn determined by the required ADC accuracy (Section 4.6).




Fig. 3.17 Block diagram of a third-order sinc digital filter.

The first half of the filter operates at the full modulator clock rate, while the second half operates at a reduced speed determined by the oversampling ratio. The third-order filter produces a valid output value after $(3OSR - 2)$ input bits are received. The selection of the filter order, oversampling ratio, and internal register widths are discussed in Section 4.6.

3.9 Summary

Performance measures for evaluating temperature sensors fit into two broad categories: those that focus on the quality of output signal generated by the sensor, and those

concerned with the practical realities of implementing such a sensor in a target application. The strengths of integrated smart temperature sensors derive from their use of standard IC processes - low cost, low power consumption, and high levels of integration. The weaknesses of these devices have historically been in the areas of linearity and manufacturing-induced variations between parts.

Pertijs, Makinwa et al. [2005] presented a smart sensor significantly more accurate than previous designs. A recurring feature in the analog circuitry is the use of chopping and dynamic element matching to extract the highest possible accuracy from each subsystem. The primary source of error in this sensor is manufacturing-induced variation in V_{BE} . Several opportunities for improvement have been identified in this design, both in the sensor and ADC. A process compensation circuit has been added to the front-end sensor to utilise the correlation between pinched-base resistance and bipolar saturation current I_S . The ADC conversion time can be reduced without affecting resolution by increasing the filter order to three, and the use of a modern CMOS process allows this filter to be implemented on-chip with a minimal increase in chip area.

Circuit schematics to implement these design improvements have been presented and examined. Chapter 4 will describe in detail the many design choices required to turn these schematics into a functional system.

CHAPTER 4

DESIGN IMPLEMENTATION

Chapter 3 presented a framework for evaluating the tradeoffs involved in the design of a smart temperature sensor, and described how the design presented in this thesis can improve on previously published works within the presented framework. This design draws substantially from the smart sensor published by Pertijs, Makinwa et al. [2005], but has been significantly improved in several aspects as described in Section 3.2. Circuit diagrams were presented for each of the functional blocks within the smart sensor, and a qualitative description was given for each block. However turning a circuit diagram into a functional system involves many further design decisions and tradeoffs; presenting these decisions is the purpose of this chapter.

Besides the obvious goal of ‘making it work’, the decisions described in this chapter all intend to fulfil the design goals developed in Section 3.2. The use of a modern $0.13\mu\text{m}$ manufacturing process allows this design to pursue low power consumption, low die area, and complete on-chip integration. Incorporating design concepts published by Pertijs, Makinwa et al. allows the analog sections to achieve high accuracy levels. And utilising the correlation between base-pinch resistance and bipolar temperature sensors will significantly reduce the sensor’s sensitivity to process-induced variations.

A system-level block diagram was given in Section 3.3 showing the signal flow between each subsystem, reproduced here for convenience. Each of the blocks in Fig. 4.1 will be described in detail in the remaining sections of this chapter.

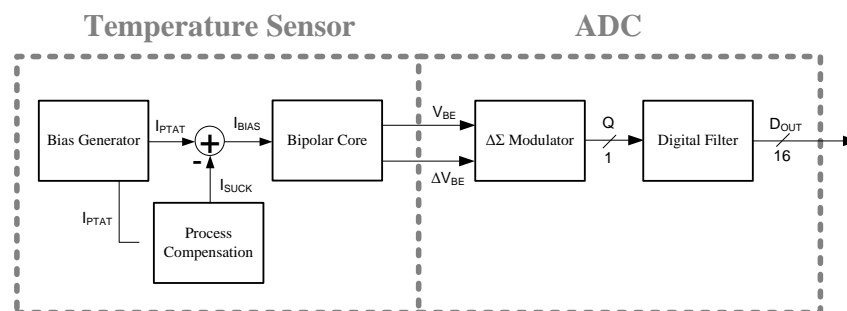


Fig. 4.1 Block diagram of the proposed smart temperature sensor.

4.1 Operational Amplifiers

The op-amp is a fundamental building block of analog and mixed-signal integrated circuits. Both the bias generator and process compensation blocks in Fig. 4.1 require op-amps with single-ended outputs, while the $\Delta\Sigma$ modulator requires op-amps with differential outputs. All these amplifiers are driving capacitive loads only, so a single-stage design will ensure a simple compensation process. The $\Delta\Sigma$ modulator is expected to require several hundred cycles to achieve the targeted resolution of 0.01°C , so a system clock frequency of around 10kHz will complete the conversion sufficiently quickly. This target bandwidth can easily be met by a low-power single-stage op-amp when driving typical on-chip capacitive loads.

The switched-capacitor integrators used in the $\Delta\Sigma$ modulator require the highest possible output swings from their amplifiers in order to accommodate the loop's internal excursions. Combined with the low system supply voltage of 1.2V, a folded-cascode architecture is most appropriate due to its wide output swing combined with wide common-mode input range and moderate power consumption. In the interests of conserving design effort, the same architecture was used for the single-ended op-amps used in the analog front-end.

Single-stage op-amps are also known as *transconductance* amplifiers – the output is taken from a high-impedance node at the junction of two current sources (Fig. 4.2a). The output stages of both single-ended and differential amplifiers are created from this basic configuration, implemented with a transistor stack as shown in Fig. 4.2b. Further design effort can be conserved by standardising the overdrive voltages of this stack over all op-amps and cascoded current sources used in the system. The bias voltages V_{B1} – V_{B4} can then be generated from one system-wide bias source.

The differential op-amps used in the switched-capacitor $\Delta\Sigma$ loop have the most demanding output swing requirements. Slightly over half the available supply (0.7V) was allocated for swing at each output leg, giving a differential swing capability of 1.4V. The remaining 0.5V was allocated as overdrives to the transistors in the current-source stack in Fig. 4.2b, taking account of the significantly lower mobility of PMOS

devices in the target process. The chosen overdrive values are shown in Fig. 4.3, along with their associated gate bias voltages $V_{B1} - V_{B4}$.




Fig. 4.2a




Fig. 4.2b

Conceptual output stage of a transconductance op-amp (4.2a) and transistor-level implementation (4.2b).




Fig. 4.3 Overdrive and bias voltages of the standard current-source stack.

PMOS Well Connection

As is common on all modern CMOS processes, the target $0.13\mu\text{m}$ process uses a P-type substrate. All PMOS devices must therefore reside in an N-well. Whenever a PMOS device's source terminal does not connect directly to V_{DD} , a design decision must be made for an appropriate connection for the N-well body terminal. Connecting the N-well to V_{DD} simplifies layout and reduces parasitic capacitances at the source terminal. This is the option chosen for the PMOS transistors in Fig 4.3, and indeed all other PMOS transistors throughout the design. The only drawback is the larger PMOS

threshold voltage than if the N-well was connected to the PMOS' source terminal; this makes no difference to the current sources in Fig. 4.3 and does not significantly affect circuit operation in other situations encountered in this design.

Single-Ended Folded-Cascode Op-amp

The bias generator block (Fig. 4.1) requires an op-amp with a single-ended output. This application does not require high output swing; in this case the common-mode input voltage presents more of a challenge. The op-amp inputs are connected to the emitters of substrate bipolar transistors (Fig. 3.6), so the common-mode input voltage will be 0.5V to 0.7V under normal operating conditions. However the start-up sequence must also be considered; upon power-up the bipolar transistors will be turned off and the op-amp's common-mode input voltage will be zero. The folded-cascode topology will continue to operate under these conditions if the input stage is constructed with PMOS devices.

The magnitude of bias current consumed by a transconductance op-amp modifies the output driving capability without significantly affecting the frequency response. The PTAT core places no significant capacitive load on the op-amp's output, so the load capacitance will be determined only by what is required to compensate the op-amp. In this situation, a lower op-amp bias current reduces both power consumption and the area required for compensation.

The final schematic for the folded cascode amplifier with PMOS input stage is shown in Fig. 4.4 along with the simulated AC output characteristics in Fig. 4.5. The amplifier achieves an open-loop gain of 49dB and unity-gain bandwidth of 14MHz with 0.3pF of compensation capacitance, while consuming only 6 μ A. One interesting advantage of this amplifier configuration derives from the decision to standardise the current-source overdrive voltages (Fig. 4.3). This amplifier drives PMOS current sources in the bias generator, so the output voltage will be approximately equal to V_{B1} . However the same bias voltage is used in the amplifier's output current mirror, so under the condition where $V_{OUT} = V_{B1}$, the amplifier in Fig. 4.4 will exhibit zero systematic offset voltage.

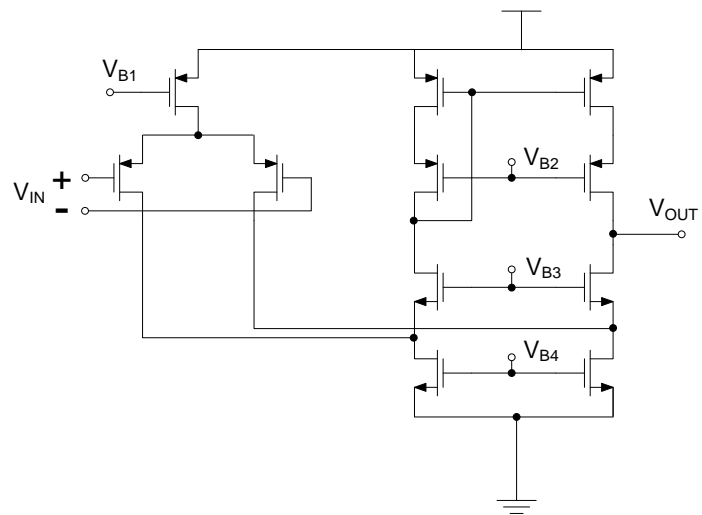


Fig. 4.4 Folded-cascode amplifier used in the bias generator circuit.

AC Response

v

Fig. 4.5 Simulated AC magnitude (top) and phase (bottom) of the amplifier in Fig. 4.4.

Differential Folded-Cascode Op-amp

To achieve high accuracy, the $\Delta\Sigma$ modulator uses fully differential switched-capacitor circuitry. So naturally, a fully differential op-amp is required. The design process is simplified by using the standardised MOS current-source stack used for the single-ended amplifier in Fig. 4.4. As the common-mode input voltage is controlled by the surrounding SC circuitry and can be biased at a convenient level, the input stage can be constructed with NMOS transistors to achieve higher gain. A slightly higher bias current of $8\mu\text{A}$ was chosen to better drive the SC loads. The final differential op-amp is shown in Fig. 4.6. The combination of differential operation and NMOS input devices allows this op-amp to achieve a slightly higher gain of 55dB.

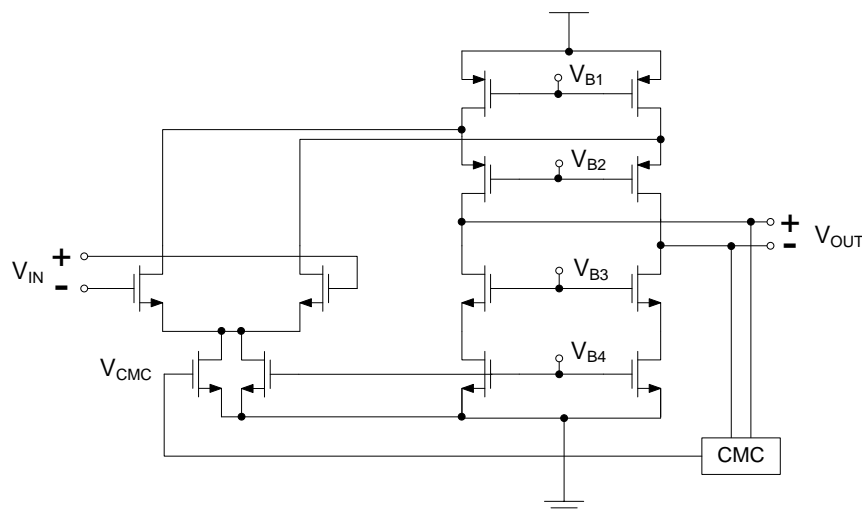


Fig. 4.6 Differential folded-cascode amplifier used in the $\Delta\Sigma$ modulator.

Common-Mode Feedback

Unlike op-amps with single-ended outputs, fully differential amplifiers require additional circuitry to control the common-mode (CM) output voltage. The op-amp's input voltage controls only the difference between its output voltages; the CM operation of the output stage is uncontrolled and resembles the opposing current sources in Fig 4.2a.

The task of detecting the average output voltage and feeding back a control signal to V_{CMC} is performed by the common-mode control block (CMC) in Fig. 4.6. This can be accomplished with either transistor-based or SC-based circuitry. Using transistors to detect the average output voltage can place further restrictions on the allowable op-amp

output swing, which is undesirable when high swing is desired. SC-based CM feedback has no such swing limitations. It does however inject small glitches on the op-amp output; this is not an issue when used in a SC integrator.

The conceptual operation of SC-based CM feedback is shown in Fig. 4.7 [Gray et al., 2001, Chapter 12]. A capacitive voltage divider C_1 across the op-amp output detects the CM output voltage, and applies it to the op-amp's CM input terminal V_{CMC} in Fig. 4.6. The op-amp itself provides the gain required to operate the CM feedback loop. As the DC voltage across C_1 is undefined, this voltage is set by a smaller pair of capacitors C_2 . The nominal feedback voltage V_{CMC} is typically not the same as the op-amp's desired output voltage, so C_2 is charged to the required $V_{OUT(CM)} - V_{CMC}$ offset.




Fig. 4.7 Conceptual operation of SC-based CM feedback for differential op-amps.

An NMOS device will act as an analog switch when supplied with sufficient gate-source voltage. Assuming V_G is limited to V_{DD} , this imposes an upper limit on the voltage that can be switched. A PMOS device exhibits complementary behavior, so is best suited to passing voltages near V_{DD} . It is therefore necessary for MOS switches experiencing wide voltage swings to be constructed with a parallel combination of N and P devices (Fig. 4.8).




Fig. 4.8 Alternative analog switch implementations using MOS devices.

In the CM control block of Fig. 4.7, the switches connecting to the op-amp's output terminals experience wide voltage swings so these are implemented using complementary devices. The CM feedback voltage V_{CMC} is well-defined, so using only NMOS devices at this node reduces the undesirable effects of charge injection. Near-minimum sized transistors are used to construct the switches: these have an on-state resistance in the tens of $k\Omega$ s which is not a concern at the targeted clock frequency of approximately 10kHz. The exact values of C_1 and C_2 in Fig. 4.7 are not critical, although smaller values are desirable to reduce the differential capacitive load on the op-amp. Making C_2 smaller than C_1 reduces output glitches when the op-amp's differential output is non-zero. The switches are driven by non-overlapping clocks ΦA and ΦB , to be discussed further in Section 4.5. The final CM control circuit is shown in Fig. 4.9.




Fig. 4.9 Final SC-based CM control block for a differential op-amp.

Gain-Boosted Differential Folded-Cascode Op-amp

Previous chapters have already discussed the need for the first integrator in a $\Delta\Sigma$ modulator loop to be as accurate as possible (Section 2.8, Section 3.7). The accuracy of SC integrators can be enhanced by increasing op-amp gain. One method of doing so which does not limit the output voltage swing nor require additional amplification stages is known as ‘gain boosting’. Shown in Fig. 4.10, a gain-boosted folded-cascode op-amp uses two auxiliary amplifiers to increase the impedance of the output current mirrors. This results in significantly higher voltage gain (albeit only at low frequencies). The op-amp in Fig. 4.10 is designed for the first $\Delta\Sigma$ integrator and is identical to Fig. 4.6 except for the addition of the auxiliary amplifiers, yet it achieves an open-loop gain of 108dB.

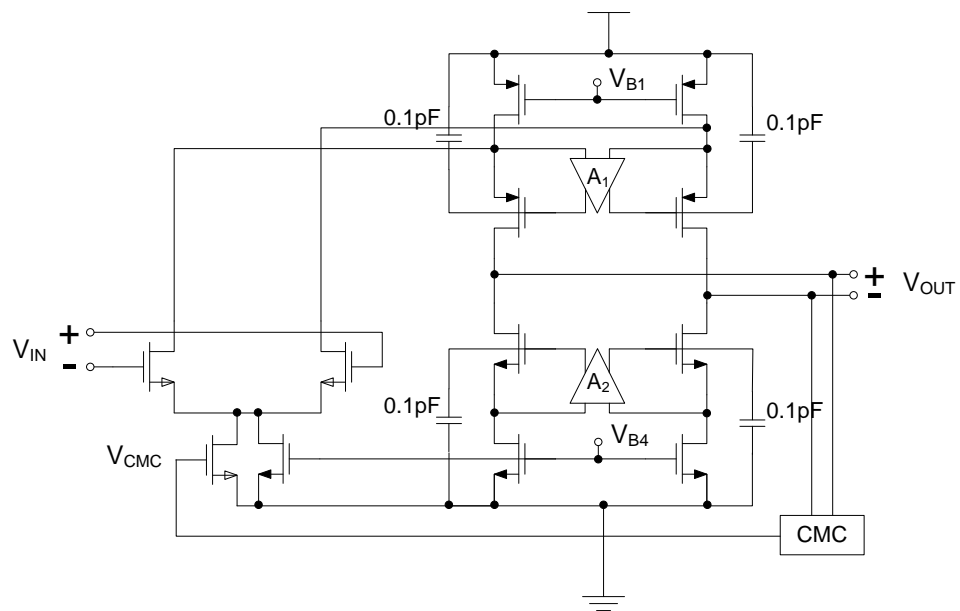


Fig. 4.10 Gain-boostered differential folded-cascode op-amp.

The auxiliary amplifiers A_1 and A_2 in Fig. 4.10 are complete differential op-amps in their own right. Their function is to maintain their differential input voltage as close as possible to zero; by doing so they significantly increase the low-frequency differential gain of the main amplifier. They do not drive any significant capacitive load and are not required to be particularly fast, they can therefore be designed to use very low bias currents. Small compensation capacitors are shown in Fig. 4.10 to ensure stable auxiliary amplifier operation.

Amplifier A_1 is supplied with a high common-mode input voltage; it can use a folded cascode structure identical to the main amplifier (Fig. 4.11a). However A_2 is supplied with a common-mode input near ground; it must therefore use a PMOS input stage (Fig. 4.11b). Both amplifiers are designed to use minimal supply current, $2\mu\text{A}$ each. Combined with the main amplifier of Fig. 4.10, the complete gain-boostered op-amp consumes a total of $12\mu\text{A}$ of supply current.

The common-mode output voltages of A_1 and A_2 should be identical to the bias voltages they replace – V_{B2} and V_{B3} respectively. This is achieved by modifying the input voltages to the CM control blocks (Fig. 4.7).




Fig. 4.11a

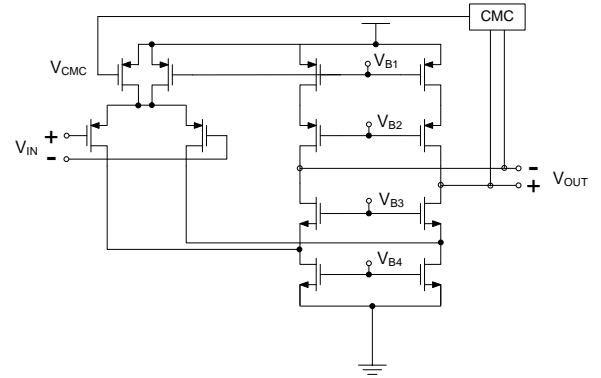


Fig. 4.11b

Auxiliary amplifiers A_1 (4.11a) and A_2 (4.11b) used in the gain-boosted op-amp of Fig. 4.10.

Amplifier Bias Voltage Generator

All the op-amps described up to this point have employed MOS transistors operating at standardised overdrive voltages to simplify both the design and layout processes (Fig. 4.3). A bias generator circuit (not to be confused with the PTAT bias current generator in Fig. 4.1) is therefore necessary to create $V_{B1} - V_{B4}$ as well as the CM feedback bias voltages. This circuit is not required to be as accurate as the PTAT bias circuit; in fact it is beneficial for the bias voltages $V_{B1} - V_{B4}$ to track variations in MOS threshold voltage. A simple bias generator circuit shown in Fig. 4.12 biases its transistors at a transconductance that is relatively independent of variations in power supply, process and temperature [Johns and Martin, 1997, Chapter 5]. It is therefore known as a constant- g_m bias circuit.




Fig. 4.12 A simple constant- g_m bias generator circuit.

The basic circuit of Fig. 4.12 requires several enhancements to be practical. Firstly, the cascade bias voltages V_{B2} and V_{B3} need to be generated. Secondly, a start-up circuit is needed to ensure the bias generator does not remain in a stable zero-current state upon

application of power. And lastly, duplicates of the bias voltages $V_{B1} - V_{B4}$ are required for supplying the SC CM-feedback blocks in order to isolate noisy SC circuitry from the amplifiers' internal nodes.

The complete constant- g_m bias generator is shown in Fig. 4.13. The bias core is modified to use cascoded current mirrors, which are biased by an additional set of transistors. The resistor value is chosen as $50k\Omega$ to set the nominal bias current through each circuit leg at $1\mu A$. The duplicate voltages for connection to the CM circuits are labelled $V_{BA} - V_{BD}$, and are generated by an independent set of transistors.

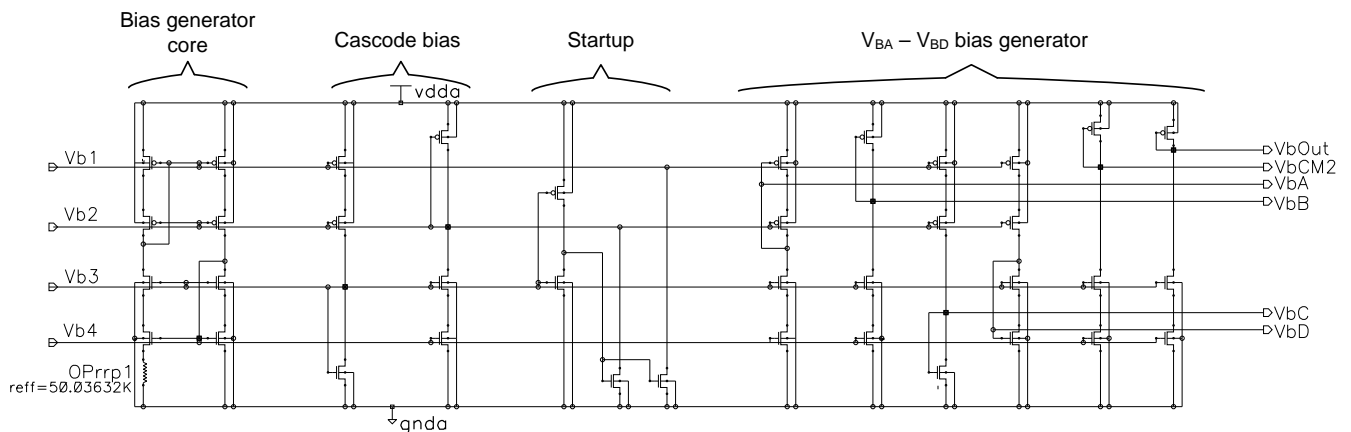


Fig. 4.13 Complete constant- g_m bias generator circuit.

4.2 PTAT Bias Current Generator

The complete circuit for the PTAT bias generator was developed in Section 3.5. It follows the circuit structure for a CMOS-based bandgap reference given in Fig. 2.15b, and uses the single-ended folded-cascode op-amp described above. Chopping (Section 2.5) is used to eliminate offsets in the op-amp and substrate bipolar transistors, and dynamic element matching is used to ensure the 8:1 bias current ratio is accurate. The complete bias generator circuit is reproduced from Section 3.5 here as Fig. 4.14.

The 8:1 current ratio was chosen more for convenience than to meet any numerical criteria; the exact magnitude of bias current is not critical as long as it is accurately reproducible between manufactured devices. A lower bias current ratio would reduce

the PTAT voltage generated by the substrate bipolars and thereby increase the relative magnitude of errors, while a larger current ratio would encounter diminishing returns due to the $\ln(I_1/I_2)$ factor in the characteristic PTAT equation (Eq. 2.9). A ratio of 8:1 is convenient to implement with symmetric layout techniques in one or two dimensions.




Fig. 4.14 Complete PTAT bias generator with chopping and DEM.

The one remaining design variable not specified in Fig. 4.14 is the value of R_1 . This resistor converts the PTAT voltage into a current that biases both the substrate bipolars in Fig. 4.14 as well as the bipolar sensor core, by way of I_{PTAT} . This bias current should be chosen to allow the substrate bipolars to operate in the most ideal region of their I-V relationship (Eq. 2.3), allowing for the fact that one substrate bipolar will be operating at 8x this current density. Based on the substrate transistor layout area and data provided by IBM, R_1 was set to 108k Ω to produce a unit current source of 0.5 μ A at room temperature. The substrate bipolars are therefore biased alternately at 0.5 μ A and 4.0 μ A.

Bias Generator Control Logic

A small block of digital logic is required to control the operation of the analog switches in Fig. 4.14. The primary task for this logic is to implement current-source DEM by swapping the unit current source supplying the 0.5 μ A bias current between all 9 possible sources. This is most simply implemented with a circular array of 9 single-bit

registers (flip-flops), as shown in Fig. 4.15. This logic is clocked from the global non-overlapping phases ΦA and ΦB . An additional startup flip-flop ensures only one register within the loop is active upon system reset. Complementary logic outputs are provided to drive the analog passgates in Fig. 4.14, as the passgates are parallel P & N MOS devices (Fig. 4.8).




Fig. 4.15 Control logic block for the PTAT bias generator circuit.

4.3 Bipolar Core

The bipolar core developed in Section 3.4 is the heart of the smart temperature sensor. It uses the PTAT current created by the bias generator circuit, and produces accurate V_{BE} and ΔV_{BE} voltages that are connected directly to the $\Delta\Sigma$ modulator input. The current sources and substrate bipolars in Fig. 4.16 are identical to those in the bias generator circuit of Fig. 4.13, so the unit current sources also supply $0.5\mu A$ each.

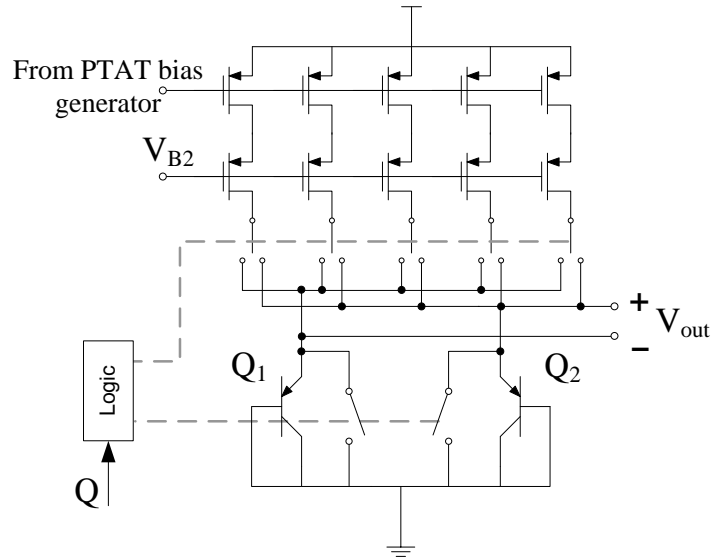


Fig. 4.16 Bipolar core generating either V_{BE} or ΔV_{BE} .

The sensor core in Fig. 4.16 uses a PTAT current ratio of 4:1. This was specified in Section 3.4, but no details were given for the reasoning behind this decision. In fact, this ratio is an important design parameter that determines the output characteristic of the temperature sensor. As stated in Section 2.9, the ADC in this smart sensor is supplied with a PTAT temperature signal and a bandgap reference voltage:

$$\begin{aligned}
 D_{OUT} &= \frac{V_{SIG}}{V_{REF}} \\
 &= \frac{\alpha \Delta V_{BE}}{V_{BE} + \alpha \Delta V_{BE}}.
 \end{aligned} \tag{2.40}$$

The digital result D_{OUT} is required to be as linear as possible. The main source of non-linearity in the sensor output is due to curvature in V_{BE} (Eq. 2.40). This curvature is influenced by the characteristics of the manufacturing process, and the temperature-dependence of the bias current (Section 2.3).

Curvature in D_{OUT} can be greatly reduced by causing V_{REF} to have a positive dependence on temperature (Section 2.9). The gradient of V_{REF} is influenced by the gradients of V_{BE} and ΔV_{BE} , as well as the factor α in Eq. 2.40. These three design variables should therefore be chosen to produce the minimum possible curvature in D_{OUT} . Table 4.1 summarises the factors affecting curvature in D_{OUT} .

Table 4.1 Summary of design variables affecting curvature in D_{OUT} .

Variable	Variables influencing variable
<i>Curvature in D_{OUT} is caused by...</i>	
Curvature in V_{BE}	- V_{BE} bias current temperature dependence - Manufacturing process characteristics
<i>Curvature in D_{OUT} can be reduced by controlling...</i>	
V_{BE} gradient	- V_{BE} bias current magnitude - Manufacturing process characteristics
ΔV_{BE} gradient	- ΔV_{BE} bias current ratio
α	- Gain applied in $\Delta\Sigma$ modulator

The gradients of both V_{BE} and ΔV_{BE} are controlled by the number of unit current sources in Fig. 4.16. Curvature in V_{BE} exhibits a complex temperature dependence which is modelled by the design kit provided by IBM, therefore the only reasonable method of optimising V_{BE} and ΔV_{BE} gradients is by simulation. A reasonable value for α of 16 was chosen based on [Pertjjs, Makinwa et al. 2005], and values for V_{BE} and ΔV_{BE} were obtained by simulation at various current levels and temperatures.

Fig. 4.17 plots V_{REF} against temperature with the ΔV_{BE} bias current ratio set to 4:1, at varying V_{BE} bias current levels. With a V_{BE} bias current of $0.5\mu A$ (one unit current source), the reference behaves like a standard bandgap reference with a near-zero first-order temperature dependence. Adding more unit current sources supplying V_{BE} produces a positive linear temperature dependence in V_{REF} , although *the curvature remains unchanged*.

The reference voltages shown in Fig. 4.17 were then used to calculate the ADC output ratio D_{OUT} , and a linear regression was performed on the resulting curves. Fig. 4.18 shows the error in $^{\circ}C$ between the actual value of D_{OUT} and its linear approximation. It is clear that the PTAT-type temperature dependence of V_{REF} at higher bias currents increases system linearity.




Fig. 4.17 Bandgap reference V_{REF} vs. temperature, produced by the bipolar core with $\alpha = 16$ and ΔV_{BE} bias ratio = 4:1.




Fig. 4.18 Deviation from linear fit in D_{OUT} when using V_{REF} from Fig. 4.17.

Finally, the maximum deviation from linearity in Fig. 4.18 is plotted against V_{BE} bias current in Fig. 4.19. The optimum choice of V_{BE} bias current is clearly $2.5\mu\text{A}$, or 5 unit current sources. By happy coincidence this is exactly the same number of unit sources required to create the 4:1 current ratio required for generating ΔV_{BE} .




Fig. 4.19 Maximum deviation from linear fit of D_{OUT} from Fig. 4.18.

Bipolar Core Control Logic

As with the bias generator, a small block of digital logic is required to control the analog passgates in the bipolar core. It also uses a circular shift register to implement DEM for the 5 unit current sources, but several additional requirements complicate the logic. Firstly, the passgates must be configured to generate either V_{BE} or ΔV_{BE} , based on the $\Delta\Sigma$ modulator output from the previous cycle (Section 3.7). Secondly, DEM is only required when ΔV_{BE} is being generated, so the circular shift register must not be clocked when V_{BE} is generated. And finally, low-frequency chopping is implemented around the entire $\Delta\Sigma$ modulator (Section 3.7). This requires the $\Delta\Sigma$ input voltage be synchronously inverted with the low-frequency chopping signal; this is achieved by inverting the output polarity sequence between phases ΦA and ΦB . The complete control block for the bipolar core is shown in Fig. 4.20.

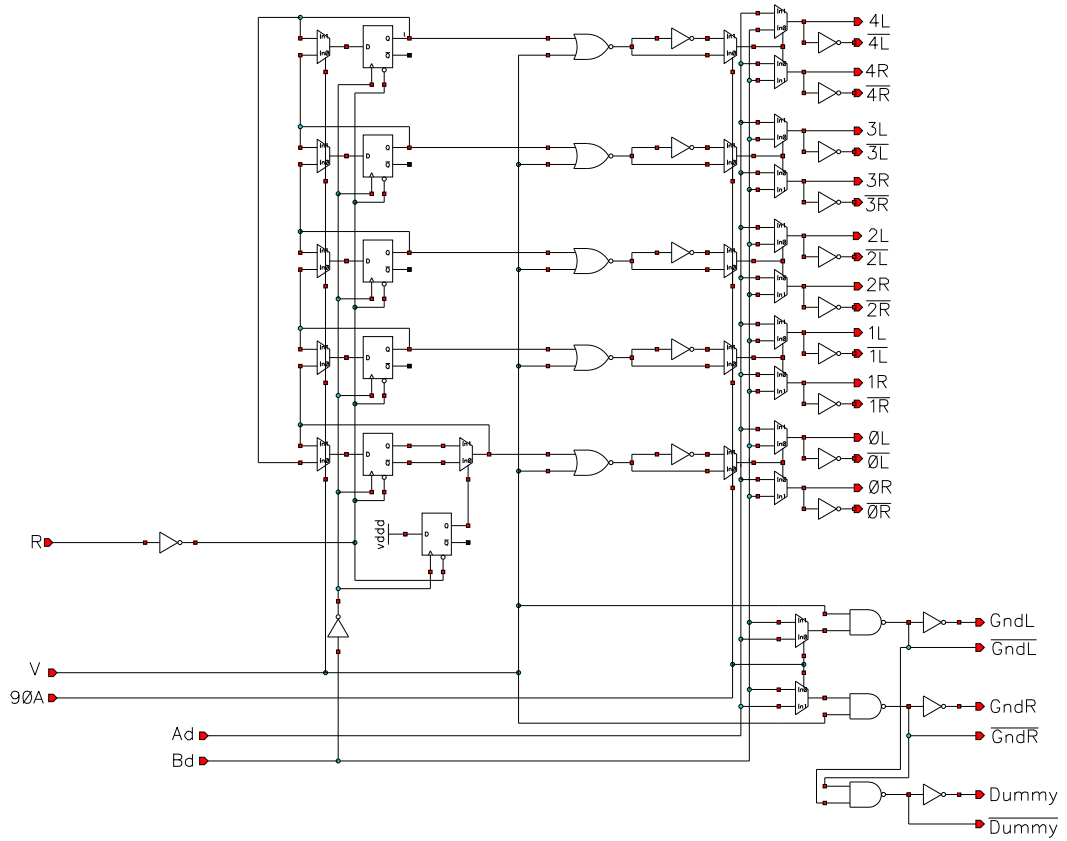


Fig. 4.20 Control logic block for the bipolar sensor core.

4.4 Process Compensation

The simulations used to derive the results presented in Section 4.3 were performed with all manufacturing parameters set to their nominal values. As shown in Table 4.1, both the curvature and the linear gradient of V_{BE} are influenced by manufacturing parameters. Fortunately the curvature in V_{BE} does not vary significantly within the same process; gradient variations in V_{BE} are the most significant contributor to random variations in the on-chip temperature sensor (Section 2.10). These gradient variations are caused by variations in bipolar saturation current I_S . A constant V_{BE} can be maintained over process variations by controlling the transistor's collector current to maintain a constant I_C / I_S ratio:

$$V_{BE} = V_T \ln \left(\frac{I_C}{I_S} \right). \quad (2.41)$$

Substrate bipolar transistors must be biased via their emitter, which introduces an additional current gain dependence into the emitter bias current:

$$I_C = I_E \cdot \frac{\beta}{\beta + 1}. \quad (3.1)$$

However Section 3.5 has already described a method to make the substrate bipolars' collector current β -insensitive, by introducing a compensating factor of $(\beta + 1) / \beta$ into the PTAT current supplied by the bias generator circuit. The collector current is therefore insensitive to process-induced changes in β , and the process compensation circuit discussed below can then seek to maintain the ratio I_E / I_S at a constant value.

Section 2.10 also discussed the correlation between pinch-base resistance R_P and bipolar saturation current I_S . Section 3.6 presented a process compensation circuit which modifies the bias current supplied to generate V_{BE} based on the value of a pinch-base resistor (Fig. 4.1), thereby cancelling gradient variations in V_{BE} . The compensation circuit operates by subtracting a process-dependent current from the main process-independent I_{PTAT} bias current, as shown in Fig. 4.21 reproduced from Section 3.6. This section will analyse the operation of this compensation circuit in detail.




Fig. 4.21 Conceptual operation of process compensation circuit.

The current I_{BIAS} supplied to the substrate bipolar transistor Q_1 in Fig. 4.21 is the difference between the two current sources, both of which are derived from I_{PTAT} created by the bias generator circuit (Fig. 4.1):

$$I_{BIAS} = I_{PTAT} \left(k_1 - \frac{k_2}{R_P} \right). \quad (4.1)$$

As stated in Section 2.10, the relationship between pinch-base resistance R_P and bipolar saturation current I_S published by Dutton and Divekar [1977] is linear:

$$R_P = m \cdot I_S + c. \quad (2.42)$$

The relationship between the bias generator current I_{PTAT} and the substrate bipolar current I_{BIAS} can therefore be expressed as:

$$I_{BIAS} = B \cdot I_{PTAT}, \quad (4.2)$$

where:

$$B = \left(k_1 - \frac{k_2}{(mI_S + c)} \right). \quad (4.3)$$

The curvature analysis presented in Section 4.3 concluded that the optimal V_{BE} bias current was $2.5\mu A$; this was obtained with no process compensation circuitry and with the target processes' parameters at their nominal values. This corresponds to a current gain between the bias generator and bipolar core of 5:

$$I_{BIAS} = 5 \cdot I_{PTAT}. \quad (4.4)$$

It can therefore be concluded that the factors k_1 and k_2 in Eq. 4.3 should be chosen so that when R_P assumes its nominal value, the factor B in Eq. 4.2 is also equal to 5.

As mentioned at the beginning of this section, complete cancellation of process-induced variations in V_{BE} requires maintaining the ratio I_{BIAS} / I_S at a constant value. Proceeding on the assumption that circuit-level techniques make I_{PTAT} process-insensitive, the ratio B / I_S must therefore be maintained at a constant value. This ideal relationship between I_S and B is represented by the straight line through the origin in Fig. 4.22. The curved line in Fig. 4.22 results from the actual relationship implemented by the process compensation circuit, as described by Eq. 4.3. The factors k_1 and k_2 are chosen to ensure the closest possible correspondence to the desired straight line over the expected 3σ variation in I_S .

The process compensation circuit developed in Section 3.6 is reproduced in Fig. 4.23. Current from the bias generator is impressed across the polysilicon resistor R_2 , producing a PTAT voltage at the input of op-amp A_2 . Negative feedback ensures the voltage across R_P is also PTAT, therefore the current sunk by the NMOS transistor is:

$$I_{SUCK} = \frac{R_2}{R_P} I_{PTAT}. \quad (4.5)$$




Fig. 4.22 The desired relationship between I_S and B , and the actual relationship implemented by the circuit of Fig. 4.20.




Fig. 4.23 Complete circuit diagram of process compensation scheme.

Resistor R_2 therefore determines k_2 in Eq. 4.1. This process-dependent NMOS current sink is connected to the I_{PTAT} current source at the emitter of Q_3 . The gain k_1 is determined by the number of unit current sources supplying I_{PTAT} to Q_3 . This will be some number larger than 5, the optimal value without the process compensation in place.

As well as exhibiting sensitivity to process parameters, the pinch-base resistor R_P will inevitably exhibit sensitivity to temperature. A linear temperature relationship in R_P will cause the substrate bipolar bias current I_{BIAS} to deviate from ideal PTAT behaviour, and introduce additional curvature into V_{BE} . A practical pinched-base resistor is also likely

to exhibit higher-order temperature dependency, making the task of correcting curvature by calibrating bias currents difficult indeed (Section 4.3).

As the pinch-base resistor is unsupported by the target 0.13 μm process, no information was available at design time regarding the temperature-dependence of R_P . Simultaneously linearising V_{BE} and optimising the process compensation system requires detailed information on the characteristics of the target processes' substrate bipolar transistors and pinched-base resistors.

Implementable Compensation Circuit

One significant problem remains when implementing the compensation scheme described in Fig. 4.23. The optimal values for the current gain factors k_1 and k_2 depend on the exact correlation between R_P and I_S on the target manufacturing process. Not only is this data unavailable, but data on the expected value of the pinch-base resistor itself is not available. The values of k_1 and k_2 in Fig. 4.23 depend on a resistor value and a current source gain, making post-manufacture adjustments impossible. An alternative process compensation scheme is therefore required for the design developed in this thesis.

One method of achieving the aim of freely-adjustable compensation parameters is by moving the bias current I_{PTAT} off-chip, performing the required amplification with external circuitry, then re-inserting the bias currents at the emitter of Q_3 in Fig. 4.23. To reduce the effects of leakage in the external circuitry, current mirrors were used to amplify I_{PTAT} before being fed off-chip, and attenuate the resulting currents as they re-enter the chip. The implemented process compensation scheme shown in Fig. 4.24 is designed to be highly flexible and can even be completely disabled, returning the V_{BE} bias current to the simple 5x relation in Eq. 4.4.

The bias current I_{PTAT} is impressed across an on-chip resistor to produce an accurate PTAT voltage. This is sampled by an off-chip instrumentation amplifier which converts the PTAT voltage to a PTAT current with variable gain, and feeds the result back into the chip as I_{HIGH} . This PTAT current is sunk to ground via the on-chip process-dependent resistor R_P , the voltage across which is used by a second external amplifier to

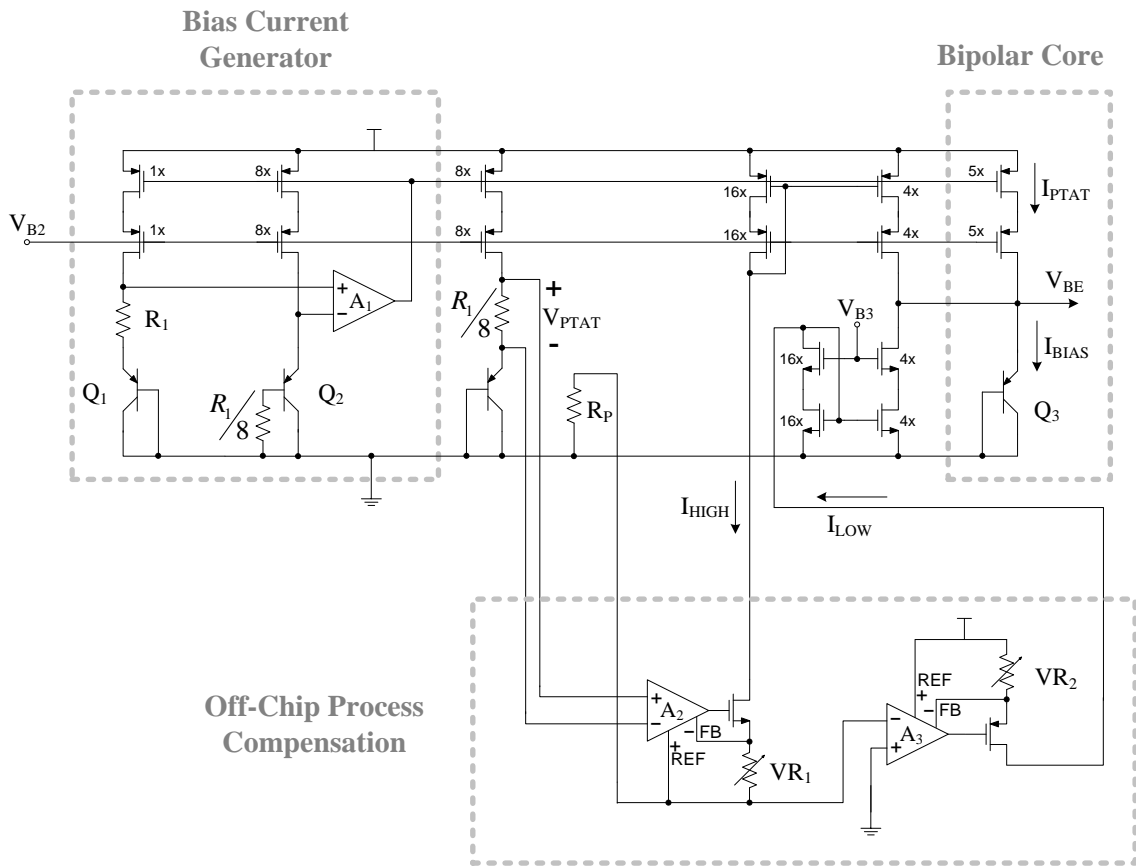


Fig. 4.24 Fully-adjustable process compensation circuit realised with off-chip amplifiers.

create a controllable process-dependent current fed back to the chip as I_{LOW} . The compensation currents I_{LOW} and I_{HIGH} are attenuated by current mirrors, and then added to the constant on-chip $2.5\mu A$ current supplied to V_{BE} .

As the bias core is already supplied with the optimal bias current without the process compensation in place, the goal of the external circuitry in Fig. 4.24 is to provide zero compensation with R_P at its nominal value, or more specifically, set $I_{LOW} = I_{HIGH}$. When external resistor VR_2 is set to the nominal value of R_P , the gain of external amplifier A_3 will be unity and I_{LOW} will therefore equal I_{HIGH} when R_P assumes its nominal value. The value of VR_1 determines the ‘gain’ of the compensation circuit, in other words the sensitivity of the compensation currents to changes in R_P .

A complete schematic diagram of the actual implemented analog front-end (bias generator, process compensation interface, and bipolar core) is shown in Fig. 4.25. The attenuating current mirrors connecting the external compensation circuitry to the bipolar core use a gain ratio of $0.25x$.

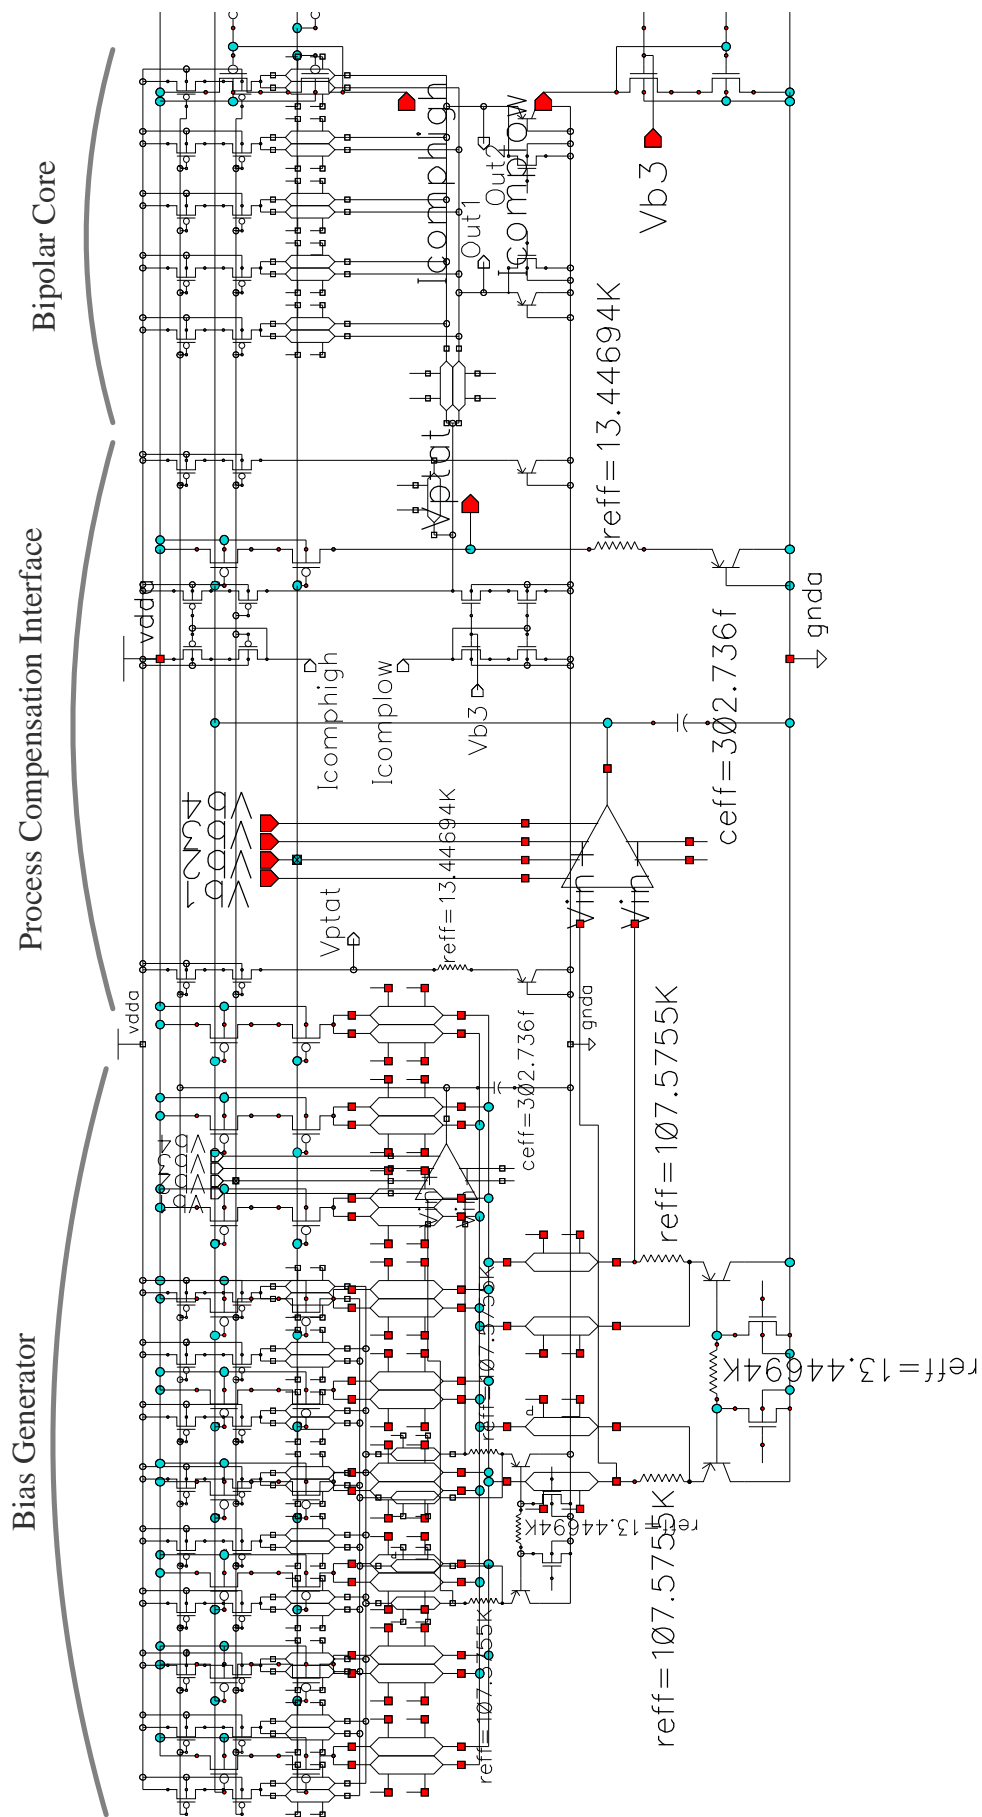


Fig. 4.25 Complete schematic of the on-chip analog front-end – bias generator, process compensation interface, and bipolar core.

Several additional features are worth mentioning here. The on-chip resistor generating V_{PTAT} is connected to ground via a dummy substrate bipolar transistor, the purpose of which is two-fold. Firstly, this raises the current mirror's output voltage closer to the voltage within the bias generator, which increases the accuracy of the mirror and therefore V_{PTAT} . Secondly, V_{BE} can now be directly measured with external equipment allowing further insight into the relationship between R_P and V_{BE} . Another dummy substrate bipolar is used to sink the compensation currents I_{LOW} and I_{HIGH} when the bipolar core is generating ΔV_{BE} , to avoid excessive switching transients.

4.5 $\Delta\Sigma$ Modulator

Designing a $\Delta\Sigma$ modulator involves choosing both a loop structure and noise transfer function (NTF). Primary considerations during the design process include the conversion time (clock cycles) and output resolution of the resulting ADC. Modulator design choices also affect loop voltage swings and chip layout area, and should therefore also be considered with respect to the design characteristics developed in Section 3.2. The overall smart sensor accuracy target is 0.1°C . The most significant contributor to inaccuracy is variation V_{BE} (Section 2.9). To ensure the ADC does not significantly contribute to this error, an ADC resolution of 0.01°C is targeted, which corresponds to a resolution of 16 bits. Section 3.7 described a design for a second-order CIFF modulator. This section will analyse the detailed decisions required to implement this modulator.

Loop Structure

Two common $\Delta\Sigma$ modulator loop structures were presented in Section 3.7 – *Cascade of Integrators, FeedBack* (CIFB), and *Cascade of Integrators, FeedForward* (CIFF). Both loop structures can implement a wide variety of NTFs. Because the NTF rather than the loop structure determines the modulator's filtering characteristics, the optimal loop structure choice is determined by circuit-level considerations.

The implicit reference technique (Section 3.7) connects the bipolar sensor core directly to the $\Delta\Sigma$ modulator input terminals. It is therefore important to minimise the capacitive load placed on these critical analog signals. The CIFB and CIFF structures are reproduced in Fig. 4.26, with the implicit reference depicted as a multiplexer at the modulator input. These structures place different capacitive loads on their input signals. The input voltage is sampled by two capacitive networks in the CIFB modulator, while the CIFF modulator samples the input with one network and achieves stability by feeding the first integrator's output forward to the comparator. The CIFF loop architecture is therefore the better choice due to the reduced capacitive loading on the bipolar core.




Fig. 4.26a




Fig. 4.26b

Second-order CIFB (4.26a) and CIFF (4.26b) $\Delta\Sigma$ loop architectures with ‘implicit reference’ front-end.

Noise Transfer Function

Quantisation noise is injected into the $\Delta\Sigma$ loop by the comparator. The NTF determines the magnitude of quantisation noise within the loop, which in turn affects loop stability and the number of cycles required to achieve a given resolution. The simplest NTF for a second-order loop is the differential function:

$$NTF(z) = (1 - z^{-1})^2. \quad (4.6)$$

The high-frequency peak in Eq. 4.6 can be limited by introducing appropriately-chosen poles into the equation. A commonly-used rule of thumb [Schreier and Temes, 2005,

Chapter 4] is to limit the maximum gain of the NTF to a suitable value, such as 1.5. The NTF is therefore implemented in the form:

$$NTF(z) = \frac{(1 - z^{-1})^2}{D(z)}, \quad (4.7)$$

where $D(z)$ can implement the poles of any high-pass filter, such as Butterworth. Using information contained in [Schreier and Temes, 2005, Chapter 8], Eq. 4.7 resolves with the help of MATLAB to:

$$NTF(z) = \frac{(1 - z^{-1})^2}{1 - 1.225z^{-1} + 0.4415z^{-2}}. \quad (4.8)$$

Section 2.8 mentioned that the first integrator within a $\Delta\Sigma$ loop is the most sensitive to analog imperfections. A switched-capacitor integrator uses an input sampling capacitor, which becomes an additional noise source due to sampled white noise. Large sampling capacitors may therefore be necessary to reduce this noise. Any modification that reduces the first integrator's gain will correspondingly reduce the value of integration capacitance required, thereby saving layout area.

Limiting the peak NTF magnitude as in Eq. 4.8 reduces signal swings at the output of the loop's integrators. On the other hand, limiting the peak NTF magnitude also increases the proportion of noise present in the signal band, which reduces the achievable ADC resolution within a given conversion time. Because the temperature input is a near-DC signal and conversion time is not a significant issue, the NTF of Eq. 4.8 was chosen for this design. Thus the modulator will require slightly longer to perform the AD conversion, but layout area will be reduced.

Capacitor Values

The complete CIFF modulator circuit diagram developed in Section 3.7 is reproduced here in Fig. 4.27. To implement the desired NTF in Eq. 4.8, the gains of both integrators and the comparator input paths must be set. Towards this aim, a z-domain representation of the CIFF modulator is shown in Fig. 4.28, including the gains of each switched-capacitor stage.

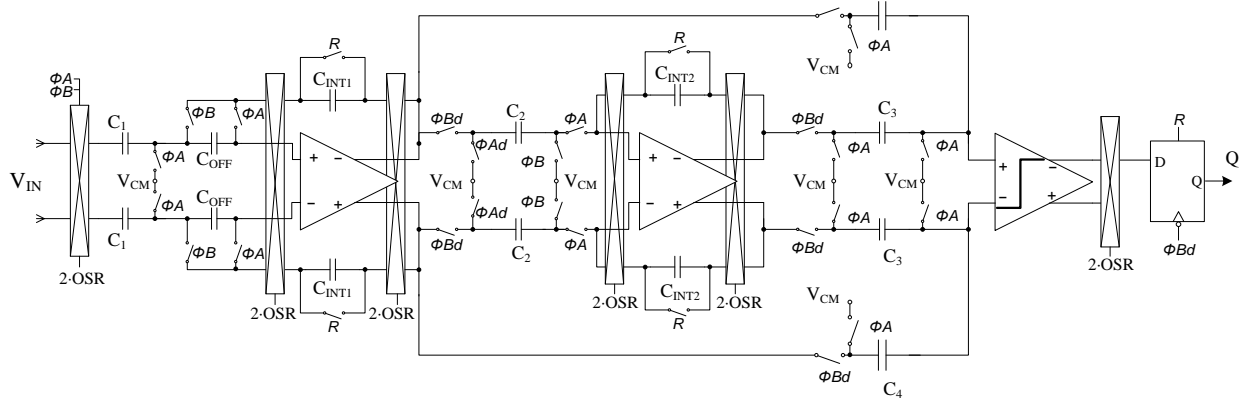


Fig. 4.27 Complete switched-capacitor modulator employing CIFF structure.

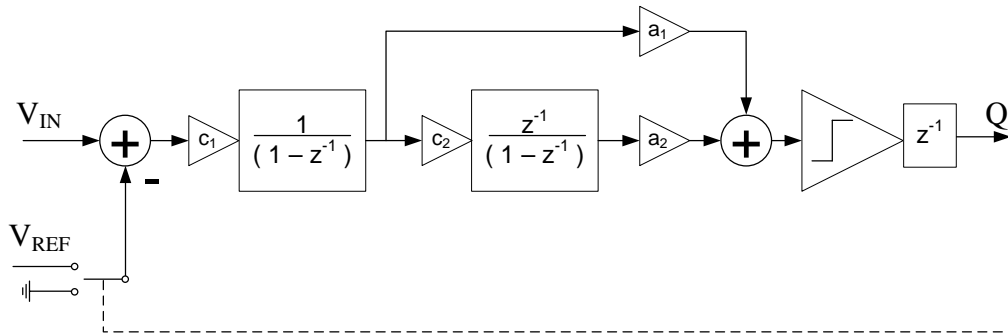


Fig. 4.28 A z-domain model of the CIFF integrator in Fig. 4.27.

The NTF of the loop in Fig. 4.28 can be written as:

$$NTF(z) = \frac{1}{1 + H(z)}, \quad (4.9)$$

where $H(z)$ represents the gain of the forward path from input summing node to loop output:

$$H(z) = \left(\frac{a_2 c_2 z^{-1}}{(1 - z^{-1})} + a_1 \right) \cdot \frac{c_1 z^{-1}}{(1 - z^{-1})}. \quad (4.10)$$

Substituting $H(z)$ into Eq. 4.9 and simplifying leads to the following expression for NTF(z):

$$NTF(z) = \frac{(1 - z^{-1})^2}{1 + (a_1 c_1 - 2)z^{-1} + (1 - a_1 c_1 + a_2 c_1 c_2)z^{-2}}. \quad (4.11)$$

Substituting values from Eq. 4.8 into Eq. 4.11 reveals the desired gain values:

Table 4.2 Preliminary $\Delta\Sigma$ loop gain values.

Variable	Value
a_1	0.775
a_2	0.2165
c_1	1
c_2	1

An important step in the design of an incremental $\Delta\Sigma$ modulator is a time-domain simulation, to ascertain the peak ADC error and to discover the maximum excursion expected at the output of each integrator. The gain of each integrator can then be scaled to ensure it will not exceed the value determined by limited op-amp output swing. Scaling down the loop's internal swings has an effect on the ADC output error, so this procedure may be an iterative one [Márkus, Silva and Temes, 2004].

The z-domain model in Fig. 4.28 was simulated in MATLAB. The conversion cycle count and integrator gains were adjusted to simultaneously meet the accuracy target (16-bit accuracy) and integrator swing limitations ($\pm 0.7V$), (Section 4.1). Both criteria were met with a conversion count of 538 cycles, and the integrator gains as listed in Table 4.3. The maximum predicted integrator swing with these parameters was $\pm 0.56V$ for both integrators in the loop.

Table 4.3 Final $\Delta\Sigma$ loop gain values.

Variable	Value
a_1	0.775
a_2	0.433
c_1	0.3
c_2	0.5

The gain factors above are determined by capacitor ratios rather than capacitor values. The only capacitor in the loop with a minimum value requirement is the first integrator's sampling capacitor C_1 . Smaller values for C_1 will increase the amount of sampled white noise seen by the first integrator, and this noise is added directly to the input signal due to the lack of noise shaping at the loop input.

The conversion length of 538 cycles was determined by simulations considering quantisation noise only, and circuit-level effects were not considered. To avoid increasing the total conversion cycle count, C_1 should be chosen so that sampled thermal noise has a negligible effect on the overall converter accuracy. Fortunately the high oversampling ratio reduces the proportion of sampled white noise within the signal band, thus easing the requirements on C_1 .

The double-sampling employed by the first integrator effectively doubles the input signal magnitude. A full-scale input signal therefore corresponds to $2x V_{REF}$ (Eq. 2.40), which is approximately 2.5V. The total noise power associated with a 16-bit accurate conversion on this signal is therefore:

$$\begin{aligned}\sigma^2_{TOTAL} &= \frac{2.5^2}{(2^{16})^2 \cdot 12} & (4.12) \\ &= 121 \times 10^{-12} \text{ V}^2.\end{aligned}$$

To avoid further increasing the required cycle count, $\frac{1}{3}$ of the total conversion error was allocated to sampled thermal noise. Based on a thermal noise power of $\sigma^2_{THERM} = 40 \times 10^{-12} \text{ V}^2$ and assuming insignificant contributions from other capacitors, a minimum value for C_1 of 1.3pF was calculated, although this was increased to 2.4pF in the final design. This large safety margin was chosen primarily to allow a unit capacitor size of $20\mu\text{m} \times 20\mu\text{m}$, the minimum size specified by IBM for good inter-device matching on the target process. The layout area cost of 0.006mm^2 (or 3% of the total device area) was considered an acceptable price to pay for this reduced thermal noise and enhanced unit matching.

All capacitors within the $\Delta\Sigma$ modulator were constructed with appropriate unit capacitors to ensure accurate reproduction of the ratios in Table 4.3. The ratio $a_1:a_2$ at the comparator input was approximated with unit capacitors in the ratio of 7:4. This 2.2% error will cause a slight change in the loop NTF, but will have negligible effect on the final converter resolution. The implemented values for the capacitors in the circuit of Fig. 4.27 are shown in Table 4.4.

Table 4.4 Final values for $\Delta\Sigma$ loop capacitances in Fig. 4.26

Capacitor	Value
C_1	1x 2.4pF
C_{1INT}	3x 2.4pF
C_2	1x 1.0pF
C_{2INT}	2x 1.0pF
C_3	4x 0.8pF
C_4	7x 0.8pF
C_{OFF}	1x 0.8pF

Comparator

In addition to the differential folded-cascode op-amps developed in Section 4.1, the $\Delta\Sigma$ modulator in Fig. 4.27 also requires a comparator. This comparator has the same relaxed speed requirements as the rest of the system, and the noise shaping effect of the loop means that it is not required to be particularly accurate either.

Comparator designs employing positive feedback latches are simple to implement, reasonably sensitive, and operate in a clocked manner that integrates well with the rest of the discrete-time loop. Fig. 4.29 illustrates the operating principle of a latching comparator. Assume that the outputs of both inverters are initialised at half the supply voltage. Any differential input signal will tend to increase one inverter's input while decreasing the other. The gain provided by the inverters themselves will increase this difference until the gates latch at logic high and low.




Fig. 4.29 Operating principle of a latching comparator.

The schematic for the comparator implemented in this design is shown in Fig. 4.30. The latch transitions from reset to evaluation phases on the falling edge of the \overline{Eval} input. A simple preamplifier reduces capacitive coupling between the latch and its inputs, isolating the large internal voltage transitions and reducing so-called ‘kickback’ into the source circuitry. An RS latch maintains the comparator’s logic outputs during the reset phase.




Fig. 4.30 Clocked comparator implemented in the $\Delta\Sigma$ modulator.

The current draw of this comparator depends on both the common-mode input voltage, and the state of the clock input \overline{Eval} . A PMOS-based preamplifier was chosen to fit the 0.55V common-mode input voltage supplied by the modulator’s second integrator. Under this condition the circuit in Fig. 4.30 draws a supply current of around 5 μ A.

ΔV_{BE} Gain Factor

Due to the unequal temperature gradients of the fundamental quantities V_{BE} and ΔV_{BE} created in the bipolar core, a multiplicative factor α is required in the expression for the bandgap reference voltage V_{REF} (Eq. 2.40). Yet the implicit reference technique discussed in Section 3.7 required the raw signals V_{BE} and ΔV_{BE} to be directly connected to the modulator input. Perhaps the only method of implementing this gain α without introducing additional circuitry and error sources is to use the modulator’s first integrator. This can be achieved by switching in additional sampling capacitors to increase integrator gain, or by performing multiple integrations.

The sensor linearity calculations in Section 4.3 used a value of $\alpha = 16$. Pertijs, Makinwa et al. [2005] achieved this gain with a combination of an eight-fold integrator gain increase and two discrete integrations. The layout area required by the 7 additional input sampling capacitors is not an insignificant issue when low area consumption is desired. This design therefore implements α with a sequence of 16 cycles on the first integrator when ΔV_{BE} is being integrated. This will result in an approximately 8x increase in conversion time, but will save significant layout area.

Because the factor α appears directly in the sensor's output equation (Eq. 2.40), it is essential for α to be implemented accurately. Pertijs, Makinwa et al. [2005] used DEM between the 8 unit sampling capacitors to eliminate gain mismatch. In this design the accuracy of α is dependent solely on the accuracy of the first integrator. Fortunately the improved CDS integrator developed in Section 3.7 achieves very low leakage, and simulations predict sufficiently accurate operation.

The complete modulator schematic from Fig. 4.27 is reproduced here in Fig. 4.31, showing the modified clock connections to the first integrator. When the bipolar core is generating V_{BE} the first integrator operates synchronously with the rest of the modulator, and clock phases ΦA and ΦB follow global clocks $\Phi AIIA$ and $\Phi AII B$. When the bipolar core is generating ΔV_{BE} the first integrator is clocked 15 times via ΦA and ΦB , and the rest of the modulator operates synchronously on the 16th cycle. These clock phases are shown in Fig. 4.32.

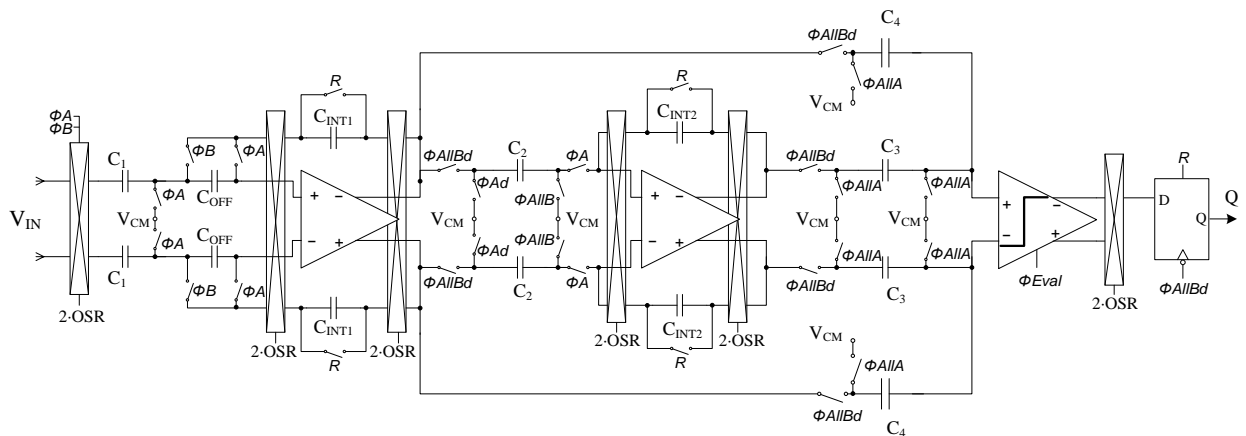


Fig. 4.31 Complete switched-capacitor modulator with multi-clocking first integrator.

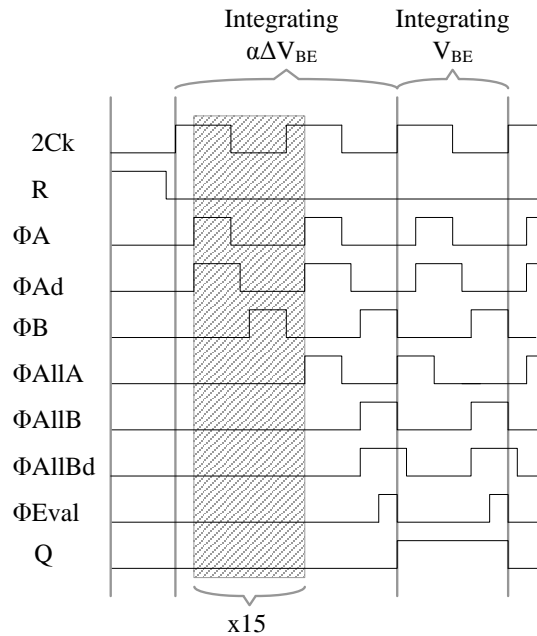


Fig. 4.32 Details of the modulator clock phases used in Fig. 4.31.

ΔΣ Modulator Control Logic

A block of digital logic is required to generate the clock phases illustrated in Fig. 4.32. Like all switched-capacitor circuits, this logic is based on the concept of two non-overlapping clock phases. The charge-injection-insensitive integrators described in Section 2.8 also require ‘falling-edge delayed’ copies of the two clock phases. The basic non-overlapping clock generator used in this design is shown in Fig. 4.33. The delay between phases is created by using a string of ‘slow response’ inverters – standard logic inverters with long MOS devices placed at their inputs to create a RC time constant. The delay between the ΦA and ΦAd falling edges is equal to the delay between ΦAd falling and ΦB rising, and is approximately 2.4ns.

The control block also needs to create the 16x integration on the first integrator, depending on the comparator’s output from the previous cycle. This is implemented with a 4-bit counter and some glue logic to ensure correct operation. The comparator is triggered half-way through the last clock phase ΦAllB via the \overline{Eval} signal. This is achieved by dividing the master clock input by two, and triggering the comparator on the second half of the master clock. Halving the input frequency also guarantees the clock phases ΦA and ΦB are equal in length. The complete modulator control block is shown in Fig. 4.34. The non-overlapping clock generator in Fig. 4.33 is represented by the block labelled ‘NOCKGen’ in Fig. 4.34.




Fig. 4.33 Non-overlapping clock generator with delayed falling edges.

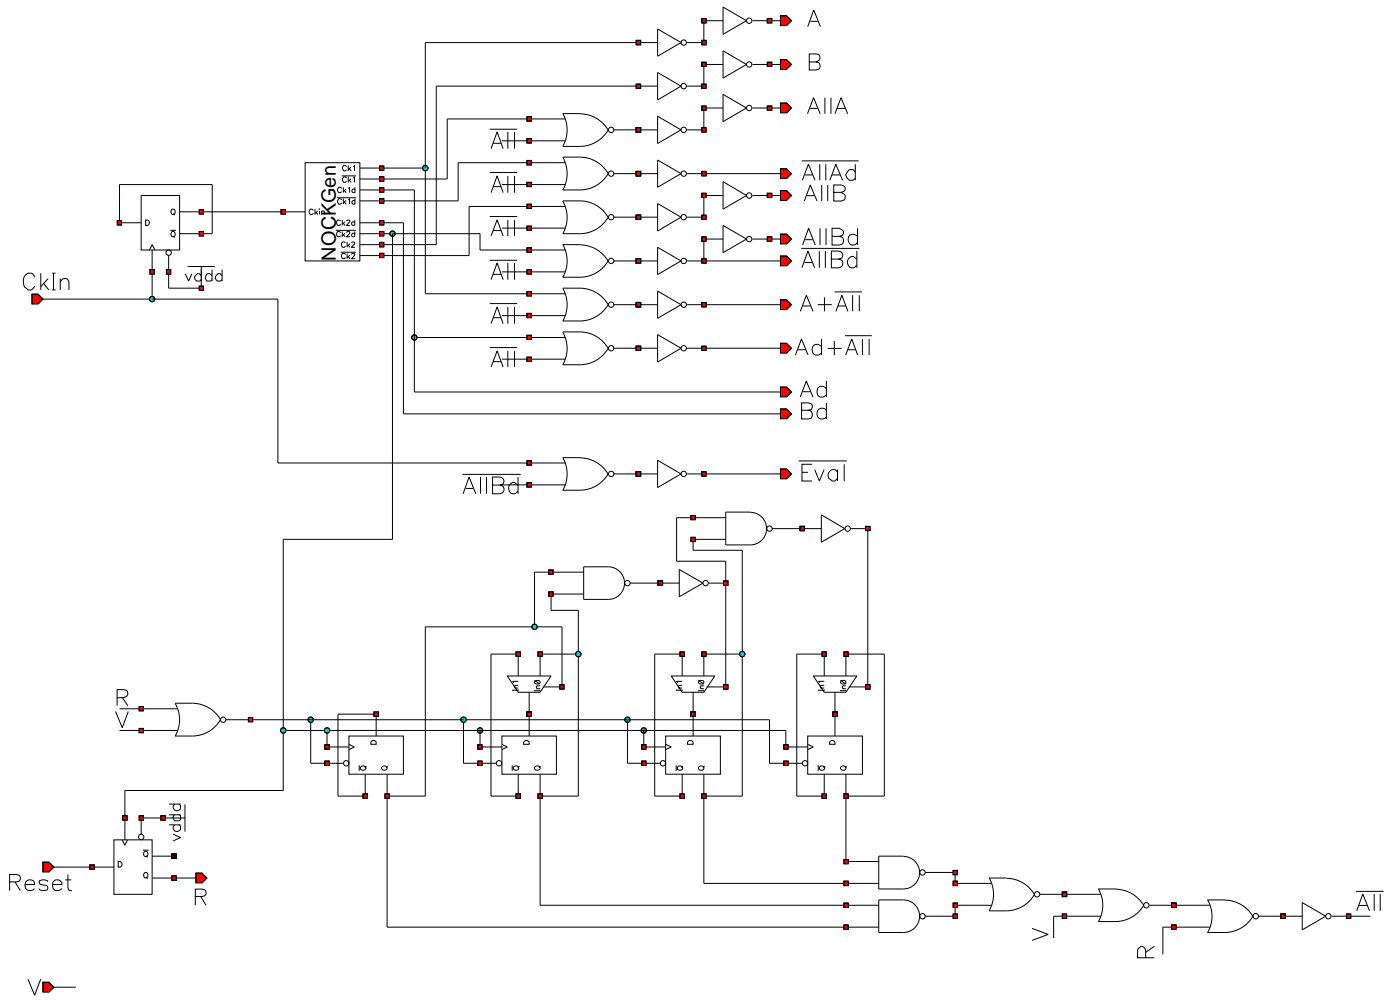


Fig. 4.34 Complete $\Delta\Sigma$ modulator control block.

Although the logic in Fig. 4.34 directly controls only the $\Delta\Sigma$ modulator, both the analog front-end and the digital filter derive their clock signals from the ΦA and ΦB signals generated by the modulator control. The clock input to the logic block in Fig. 4.34 is therefore the system master clock. This is intended to be driven by an off-chip oscillator.

Clock Frequency Analysis

The maximum device operating frequency is determined by the slowest component in the system, which in this design is the switched-capacitor integrators used in the $\Delta\Sigma$ modulator. The speed of these integrators is limited by the current available to charge their load capacitances every integration cycle. The first integrator sees the largest load of 7.2pF (Table 4.5). The current available to charge this capacitance is the 2 μ A bias current flowing through each amplifier output leg. The largest input transition on the first integrator is V_{BE} , or roughly 0.7V. The charge drawn from the op-amp output is:

$$\begin{aligned}\Delta Q &= \Delta V_{OUT} \cdot C_{1INT} \\ &= \Delta V_{IN} \cdot \frac{C_1}{C_{1INT}} \cdot C_{1INT}.\end{aligned}\quad (4.13)$$

The time required for the op-amp to supply this charge is therefore:

$$T = \frac{\Delta V_{IN} \cdot C_1}{I_{BIAS}}.\quad (4.14)$$

Using the values stated in the previous paragraph, Eq. 4.14 indicates a slewing time of 2.5 μ s for integrator 1, which sets the absolute maximum clock frequency at 200kHz. Extra settling time is required to allow for linear amplifier settling and for deviations from the ideal relation of Eq. 4.14. A conservative system clock frequency under these conditions would therefore be 20kHz.

4.6 Digital Decimation Filter

Section 2.7 discussed the choice of decimation filter for incremental $\Delta\Sigma$ modulators. Previous publications [Márkus, Silva and Temes, 2004] concluded that the optimal choice for a second-order modulator requiring a symmetric impulse was a third-order

sinc filter. To confirm this design choice, the second-order modulator in Fig. 4.28 using the NTF of Eq. 4.8 was simulated in MATLAB, and the output bitstream was filtered with second- and third-order sinc filters. The resulting conversion error was converted to an effective number of bits (ENOB), and plotted against the number of cycles required to fill the decimation filter (Fig. 4.35).

A second-order sinc filter has a total impulse response length of $(2 \cdot \text{OSR} - 1)$, while a third-order filter is $(3 \cdot \text{OSR} - 2)$ in length. Because the filter is reset at the beginning of each conversion along with the modulator's integrators, the modulator must completely fill the filter buffers before a valid output can be produced. The horizontal range of 300 to 800 modulator cycles in Fig. 4.35 therefore corresponds to an OSR range of 150 to 400 for the second-order filter, and 100 to 267 for the third-order filter.




Fig. 4.35 Plot of conversion resolution achieved by 2nd and 3rd-order symmetric sinc filters, versus number of modulator cycles to produce the result.

It can be seen that the symmetric third-order filter achieves the target resolution of 16 bits with around 500 cycles, while the second-order filter fails to reach this accuracy with 800 cycles. Thus a symmetric third-order filter is clearly the better choice for decimating the output of a second-order $\Delta\Sigma$ modulator where a symmetric impulse response is required. The block diagram of such a filter is reproduced in Fig. 4.36 from Section 3.8.




Fig. 4.36 Block diagram of a third-order sinc digital filter.

The total conversion length of 538 cycles developed in Section 4.5 corresponds to an OSR of 180 cycles. Registers $R_1 - R_3$ in Fig. 4.36 operate at the full modulator clock frequency. Every 180 cycles, registers R_4 and R_5 are clocked, and R_3 is reset. A valid output is produced on the third slow clock.

The minimum widths of the registers in Fig. 4.36 depend on the filter order and oversampling ratio. Assuming the registers are designed to wrap-around on overflow, the maximum output value of a 3rd-order filter at an OSR of 180 is 180^3 , and can be represented by $3\log_2(180) = 23$ bits. However the first register can be significantly smaller. R_1 is a simple counter with 538 single-bit inputs, and therefore requires only 10 bits to represent its maximum value. Although the complete $\Delta\Sigma$ ADC is targeting 16-bit accuracy, truncating the filter's output to 16 bits will reduce accuracy by a further bit. Therefore the entire 23 bits is loaded into a final shift register upon completion of the conversion, and is clocked off-chip on a serial line.

Designing the clock sequencing logic to ensure numerical accuracy and adequate register set-up times was perhaps the most difficult task of the entire design process. An 8-bit counter is used to create the slow clock ($Ck / 180$) activating R_4 and R_5 . Additional flip-flops manage the final output register clock and system-wide reset signal. The 8-bit counter is also used to generate a signal that inverts every 90 clock cycles (90A and 90B) which is used to control the low-frequency modulator chopping signal $2 \cdot OSR$ in Fig. 4.30. The complete digital filter including control logic is shown in Fig. 4.37.

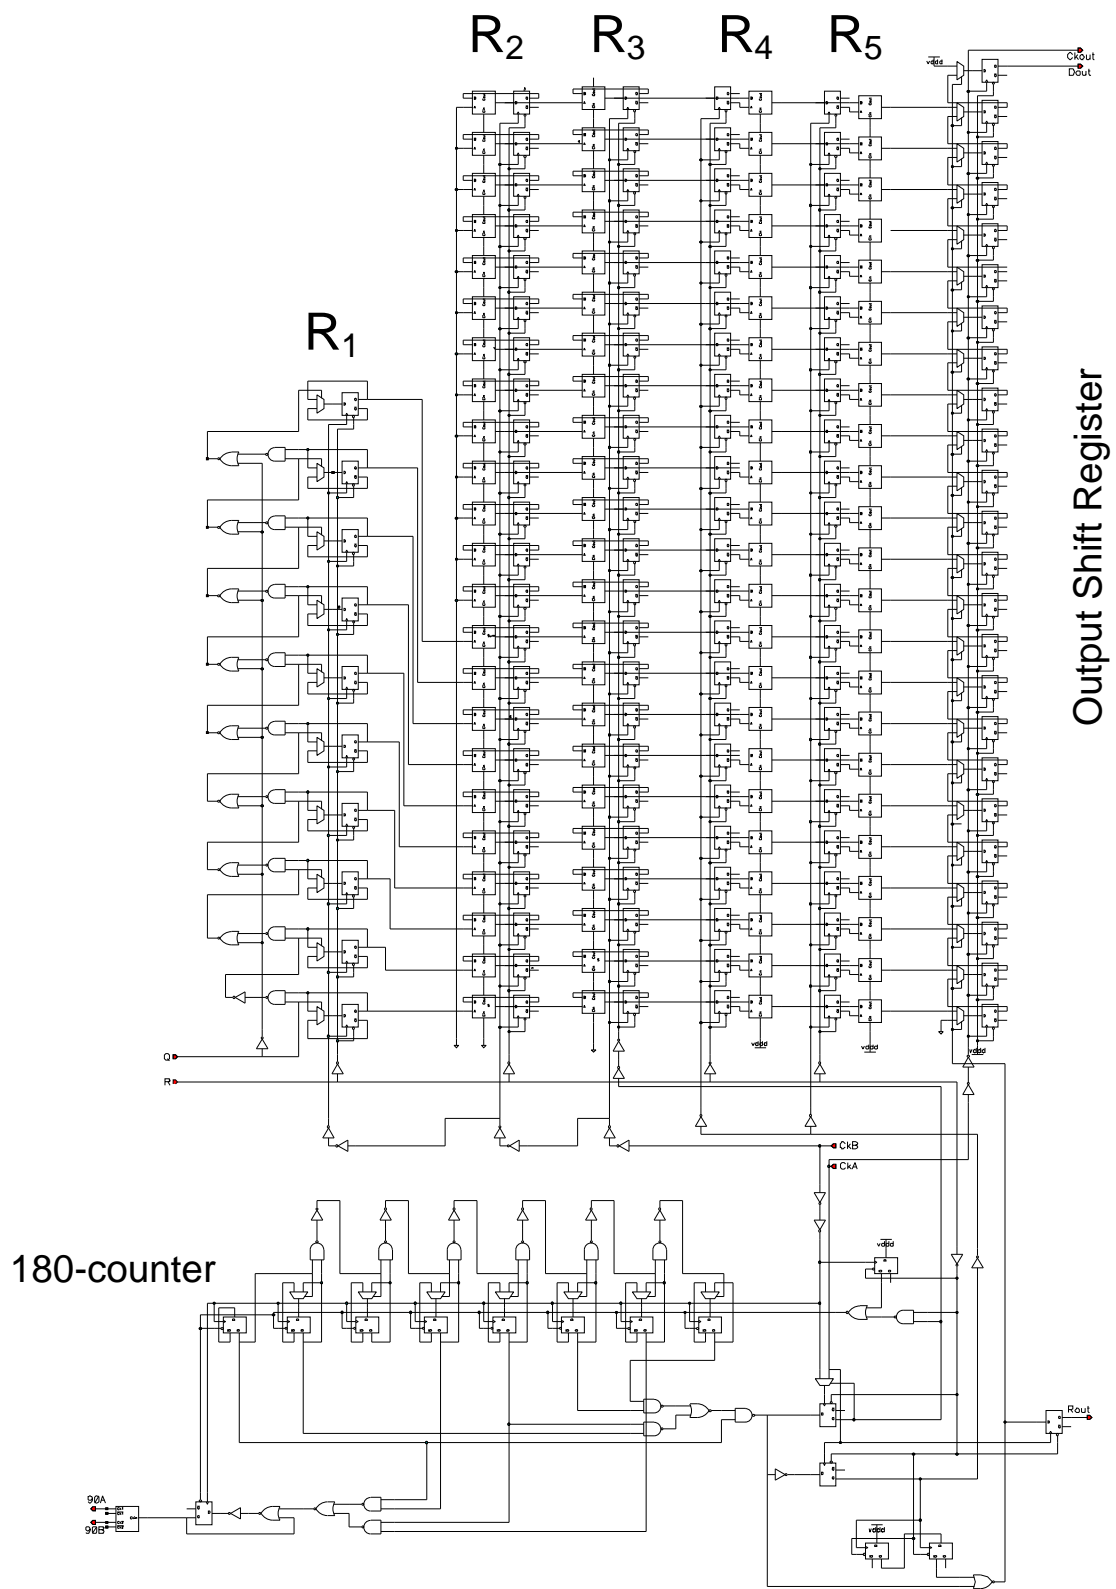


Fig. 4.37 Complete circuit diagram for third-order sinc decimation filter.

4.7 Layout

The final design task is to implement the circuitry developed in previous chapters on the target IBM 0.13 μm manufacturing process. Perhaps the most important consideration when laying out a precision mixed-signal design such as this is the need to separate sensitive analog circuitry from noisy digital logic, both physically and electrically. Independent V_{DD} and Gnd planes were used for the analog and digital sections, and were connected to separate bondpads. Digital logic blocks were surrounded by high-resistance substrate regions to minimise substrate crosstalk. The analog front-end was placed at the opposite end of the layout to the digital filter. Attention was paid to symmetric layout of sensitive analog devices – particularly the substrate bipolars in the bias generator and bipolar core circuits, the ratioed current sources, and the differential circuitry of the $\Delta\Sigma$ modulator.

As no support is provided by IBM for the pinch-base resistors critical for the operation of the process compensation circuitry (Section 4.4), it was necessary to lay these devices out manually. Two versions of R_{P} were created – type 1 that pinches the entire N-well with active-P and violates several design rules, and type 2 that commits fewer design rule violations but does not pinch the entire N-well layout area. A third unpinched N-well resistor was created for comparative purposes. Fig. 4.38 illustrates the layout of these three devices.




Fig. 4.38 Layout of N-well pinched-base resistor test devices.

Chip I/O

The smart sensor requires little in the way of external interfacing. It is a fully integrated smart sensor designed to interface directly with common digital logic devices such as microcontrollers. Its operation is entirely controlled by logic-level clock and reset inputs, and provides a serial digital output. The raw $\Delta\Sigma$ modulator output is also provided to confirm correct operation of the digital filter.

However, more external connections are necessary to check the analog front-end is operating correctly, and to provide an interface with the external process-compensation circuitry (Fig. 4.24). Particularly important for correct analog functionality are the amplifier bias voltages $V_{B1} - V_{B4}$ generated in Fig. 4.13. Table 4.5 lists the digital and analog I/O implemented on the test device.

Table 4.5 Complete list of input and output ports on the implemented design.

Name	Direction	Description
Digital		
2Ck	I	System clock input. Internally divided by 2.
R	I	Reset control. Device must be reset after power-up.
R _{CK}	O	Raw (unfiltered) clock out.
R _D	O	Raw data out. Changes on R _{CK} rising edge.
F _{CK}	O	Filtered clock out.
F _D	O	Filtered data out. Changes on F _{CK} rising edge.
Analog		
V _{B1}	O	Amp bias voltage 1
V _{B2}	O	Amp bias voltage 2
V _{B3}	O	Amp bias voltage 3
V _{B4}	O	Amp bias voltage 4
R ₁	O	N-well resistor
R ₂	O	Pinched-base resistor (type 1)
R ₃	O	Pinched-base resistor (type 2)
V _{PTAT+}	O	Positive end of PTAT-generating resistor.
V _{PTAT-}	O	Negative end of PTAT-generating resistor. Also V _{BE} .
I _{HIGH}	I	Process compensation input. Tie to V _{DD} to disable.
I _{LOW}	I	Process compensation input. Tie to ground to disable.

Power and ESD Considerations

The internal analog and digital power buses are brought out to separate pins on the device package: V_{DDA} and $GndA$, V_{DDD} and $GndD$. These should be directly connected on the application circuit board. IBM specifies a supply voltage of 1.20V for core transistors, and this is the voltage used in design work and simulations throughout this chapter. To ensure stable operation the power pins should be connected to a well-regulated and bypassed source. The total supply current consumed by the smart sensor is simulated as $52\mu A$.

The low-frequency signals used to interface with the test device present no significant challenge when selecting an electrostatic discharge (ESD) protection scheme. The IBM design kit contains many detailed rules to ensure robust ESD performance for production parts, but modifying the layout to meet all these rules would require a significant and unnecessary design effort. All that is needed is to minimise ESD damage due to careful and limited handling during sensor characterisation. All input and output pins were protected by IBM-supplied RC clamps to substrate (analog ground). The analog and digital power planes were joined through anti-parallel ESD diodes.

The complete layout was submitted through the MOSIS academic program and a batch of 30 devices was received, packaged in 28-pin ceramic DIP packages. Fig. 4.39 shows the external connections to the 28-pin DIP package. The final smart sensor layout and die photo are shown in Fig. 4.40a and 4.40b respectively. The active circuitry occupies an area of $0.48\text{mm} \times 0.43\text{mm} = 0.21\text{mm}^2$.




Fig. 4.39 External connections to sensor on a 28-pin DIP package.

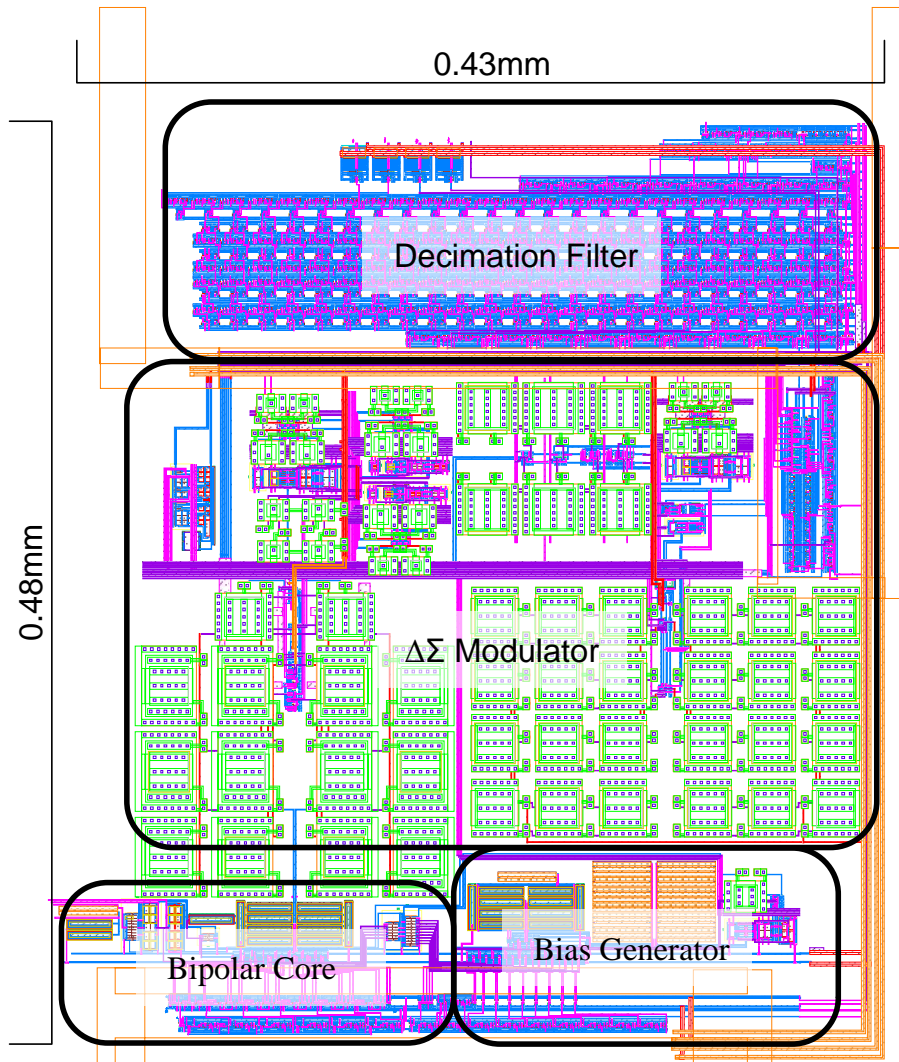


Fig. 4.40a Complete layout of active circuitry for the smart temperature sensor.




Fig. 4.40b Chip photo of active die area.

4.8 Software Tools

The smart sensor developed in this thesis was implemented on a full-custom IC as described in the preceding sections of Chapter 4. The primary design software employed for this task was the Virtuoso Analog Design Environment from Cadence. Circuit schematics were entered at the device level with the Virtuoso Schematic Editor, and tested with the Spectre simulator. A range of simulation techniques were employed on the different sub-circuit blocks, particularly DC sweeps on the front-end sensor and transient analyses on the switched-capacitor amplifiers, control logic, and decimation filter. A final transient analysis was performed on the complete system, taking around 2 weeks of CPU time and generating 40GB of waveform data.

The design process was supported by the design kit provided by IBM for the targeted 0.13 μm process. The design kit included device models based on BSIM4.4 SPICE, and programmable layout cells (P-cells) ensured the performance of manufactured devices exhibited good correspondence to simulation.

Full-custom layout was performed in the Virtuoso Layout Editor. Digital sections were constructed using standard cells drawn by hand. All routing was done manually. A range of checks were performed on the layout with the Assura tool: design rule (DRC), local and global pattern density, and floating gate. Layout-vs-schematic (LVS) checks were performed upon completion of each circuit sub-cell. IBM requires the designer to manually fill several upper metal layers to meet minimum pattern density requirements. This was achieved by adding a DRC-acceptable pattern of squares to unused regions of the die.

CHAPTER 5

EXPERIMENTAL SETUP & RESULTS

This chapter presents the work undertaken to experimentally verify the smart sensor design developed in the preceding chapters. The first task is to confirm correct operation of the various circuit blocks, and that the sensor output responds to variations in die temperature. Given satisfactory sensor operation, the exact sensitivity of the device to temperature should be quantified. Several critical electrical parameters are generated by the analog front-end and should also be recorded: V_{BE} , V_{PTAT} , and R_P . And furthermore, because the design is intended to investigate the effect of process-induced variations in all of the above parameters, as many individual sensors as possible should be characterised to obtain information on the correlation between recorded parameters.

5.1 Testing Methodology

The testing process can be divided into three distinct sequential phases – functional verification, sensor output characterisation, and process-compensation characterisation. Each test stage is dependent upon the successful completion of preceding tests. The first stage of testing will confirm the internal bias voltages are correct, that the ADC produces a reasonable output at ambient temperature, and that the decimation filter is producing a numerically correct result. The second stage of testing will sweep the test devices over a range of temperatures and gather data from several test points: ADC output, substrate bipolar V_{BE} , and pinch-base resistance R_P . The final stage will analyse the analog data gathered during the temperature sweeps and attempt to configure the process-compensation circuitry to reduce inter-device variation.

Configuring the process compensation circuitry for optimal operation requires advance knowledge of the temperature sensitivity of substrate bipolar transistors and pinch-base resistors, as well as the correlation between these two parameters when manufacturing process parameters change. At design time, the only information available from the IBM design kit is of the temperature characteristics of the substrate bipolar. Therefore a

primary goal of testing is to gather data on the temperature behaviour of the pinch-base resistor, and to test as many individual devices as possible to gather data on the R_P - V_{BE} correlation. During this process the off-chip compensation circuitry will be disabled.

The analog sensor front-end uses circuit techniques to eliminate all sources of error except for thermal gradient variations in V_{BE} . An externally-connected substrate bipolar transistor allows V_{BE} to be measured directly, along with the pinch-base resistor and the sensor's ADC output. This allows the R_P - V_{BE} process correlation to be determined directly, while simultaneously confirming the correlation between process-induced changes in the thermal responses of V_{BE} and sensor output. The process compensation scheme is based on the assumption that V_{BE} thermal gradient variations are the most significant contributor to sensor output variations. Any random variation in sensor output that is not correlated with V_{BE} will reduce the effectiveness of the process compensation scheme. The data gathered here allows the V_{BE} -sensor output correlation to be analysed.

The raw data gathered during this first phase of testing can then be analysed to determine the feasibility of the proposed process compensation scheme. Effective compensation requires a strong correlation between R_P and V_{BE} – determined entirely by the manufacturing process – and a strong correlation between V_{BE} and sensor output – determined mostly by the sensor electrical design. Assuming the data gathered here confirms these correlations, the process compensation circuitry can then be calibrated appropriately. The remainder of this chapter will describe the techniques used to implement the methodology described above.

5.2 Electrical Testbed

Data Capture

The smart sensor is designed to interface with a microprocessor or other programmable logic device. Its operation is controlled by two logic-level inputs (clock and reset), and it provides a 23-bit output over a serial line synchronised to the internal filter clock. Due to the quantity of data expected to be generated during testing, the logical destination for the sensor's digital output values is a PC. One possible method of transferring the

sample values is by using a microprocessor to read the serial data, then transferring it to the PC via RS-232. However it was expected that development of the required firmware would take a significant amount of time, so an alternative method was sought.

A *digital storage oscilloscope* (DSO) was available with a RS232-based PC interface. The associated software was capable of exporting waveform data to a *comma-separated values* (CSV) file, which is easily imported into a spreadsheet. The only development work required to implement this data flow is a method of extracting digital data from the raw voltages seen by the scope. This is easily done with a spreadsheet-based macro (VBScript), and involves significantly less development time than custom microprocessor firmware. This DSO-based transfer method was therefore adopted for the smart sensor testbed.

Sensor Operational Control

If the *device under test* (DUT) is supplied with a continuous clock signal while reset is de-asserted, the internal ADC will produce a continuous sequence of output values. For testing purposes it is desirable for the DUT to produce one temperature reading on command, then stop. This requires re-asserting the reset signal upon completion of the ADC conversion. Conversion completion can be detected by the presence of data on the filtered ADC output port F_D .

A circuit to control this single-conversion operation is shown in Fig. 5.1. The DUT is held in reset until a pushbutton is activated, causing the ADC to begin the conversion sequence. A short time delay is implemented after data is detected on the filtered output pin, to allow time for the 23-bit result to be clocked out. The control circuit then re-asserts the reset signal.

Two clock/data output pairs are provided for connection to the DSO – the 23-bit filtered result, and a combination of the raw and filtered bitstreams. The combined output is necessary as the DSO has only two input channels, and checking the filter operation requires capturing both the filter's input and output for the same conversion.

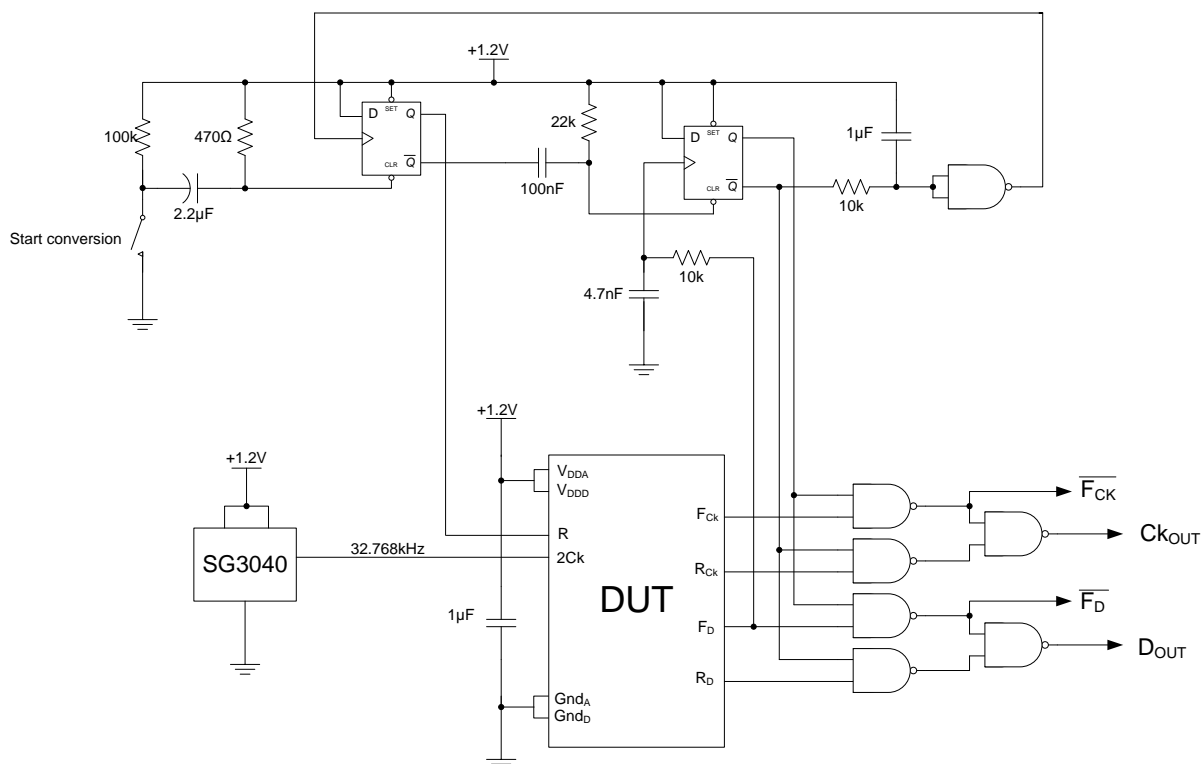


Fig. 5.1 DUT clock and reset control circuit for single-conversion operation.

A continuous clock signal is necessary for the ADC’s differential amplifiers to maintain the correct CM output voltage, even when the system is held in reset. This continuous clock signal is provided by the device labelled SG3040 in Fig 5.1 – a standard 32.768kHz oscillator module manufactured by Epson Toyocom. This clock is divided by the DUT’s clock generation circuitry (Fig. 4.34) to produce an operating frequency of 16kHz, well within the conservative limit of 20kHz established in Section 4.5. As the logic gates interface directly with the DUT, they must be capable of operating at the chip supply voltage of 1.2V. The logic gates were therefore selected from the low-voltage 74LV CMOS logic family.

Precision Current Source

Pinch-base resistors are no longer supported on modern IC manufacturing processes. One primary reason is that they exhibit significant voltage-modulation effects. That is, their resistance varies with the applied voltage. Assessing the temperature- and process-sensitivity of these resistors requires accurate resistance measurements, however. It can therefore be expected that simply measuring the value of R_p with a digital multimeter (DMM) set to measure resistance will generate non-repeatable results.

A much more repeatable way to measure on-chip R_P resistance is by supplying a constant bias current, and measuring the resulting voltage drop. This is done by using an op-amp to equalise the voltage drop across biasing and sensing resistors (Fig. 5.2). The current source can be calibrated by adjusting any of the three resistors shown. The DMM used to measure the resulting R_P voltage drop is an HP 974A, which has an accurate 500mV range with very low input bias current. The current source magnitude was set at 40 μ A to maintain the expected voltage drop within this range.




Fig. 5.2 Precision current source for measurement of on-chip pinched-base resistors.

The op-amp used in Fig. 5.2 is not required to be unusually fast or accurate, therefore the general-purpose JFET-input TL072 was chosen for its sufficiently low input bias current specification (65pA typical). The large threshold of the discrete MOS used as the output buffer requires a gate control voltage below ground, making dual power supplies a necessity.

External Process-Compensation

Section 4.4 described an external process-compensation circuit that can be freely calibrated on-the-fly. This is necessary because the required compensation currents are unknown at design-time, as discussed in Section 5.1. The off-chip section of the process-compensation circuit is reproduced in Fig. 5.3. This circuit is not required during the initial testing phase, therefore the compensation feedback ports I_{HIGH} and I_{LOW} are tied to chip V_{DD} and Gnd respectively.

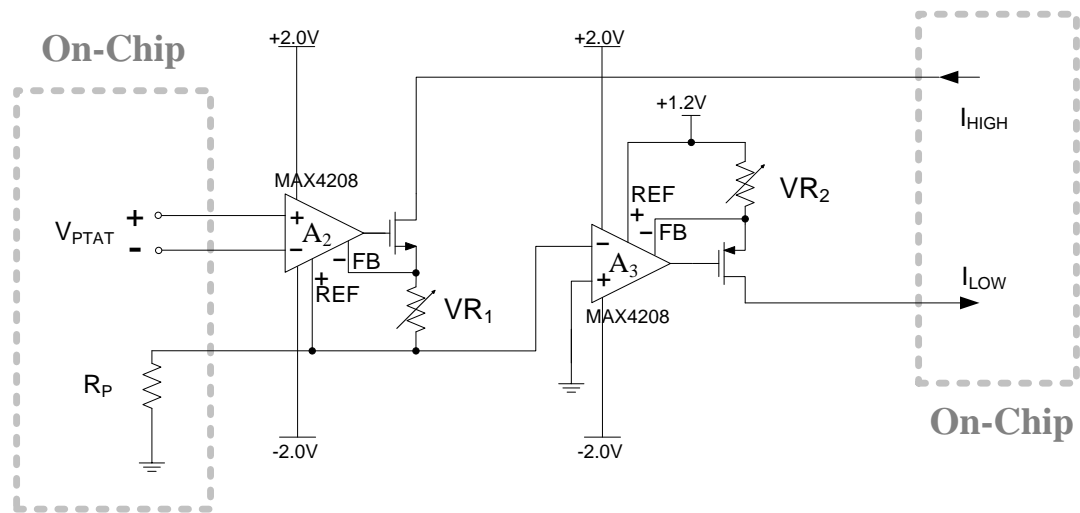


Fig. 5.3 External process-compensation circuitry.

Unlike the current source in Fig 5.2, the op-amps in the process-compensation circuitry are required to be highly accurate. The PTAT voltage used to derive the compensation currents has a value of 53mV at room temperature, and the on-chip resistor creating this voltage is biased by a current of 4 μ A. The op-amps therefore require both very low offset voltage and very low input bias current specifications. Additionally, both op-amps in Fig. 5.3 are configured to provide level-shifting as well as amplification in an effort to reduce circuit complexity. This requires instrumentation amplifiers with floating outputs and fully-differential feedback (so-called ‘differential-differential’ amplifiers). A commonly-available op-amp meeting all of the above requirements is the MAX4208. It uses a chopped MOS input stage to achieve a typical input offset voltage of $\pm 3\mu$ V and input bias current of 1pA.

The discrete N and P MOS devices used as output buffers in Fig. 5.3 have large threshold voltages. Therefore like the current source in Fig. 5.2, the process compensation circuit also requires dual power supplies. However the MAX4208 op-amps have a limited maximum supply voltage, requiring a separate lower dual supply. A voltage of ± 2.0 V provides sufficient op-amp output swing while remaining within the MAX4208 supply limits.

Testbed Power Supplies

The DUT requires a supply voltage of 1.2V, with an expected current draw of around 50 μ A. A regulator is necessary to ensure this voltage is stable and reproducible. The process compensation circuitry requires dual ± 2.0 V supplies with a draw of a few mA, and the precision current source requires dual supplies at around ± 4.5 V with similar current levels. Because the exact magnitude of the dual supplies is not critical, a pair of voltage regulators was eliminated by supplying the higher ± 4.5 V rails from AA alkaline batteries. The batteries provide clean low-noise power, and reduce cable clutter by eliminating a separate power supply. The total supply current drawn by the testbed will not significantly drain the batteries over several hours of testing.

The lower supply voltages are drawn from the ± 4.5 V rails through suitable regulators. The ± 2.0 V rails are created by complementary LM317 and LM337 adjustable regulators. The 1.2V DUT supply is below the range of common adjustable regulators, so a low-power fixed output TPS71712 was used. The complete testbed power supply schematic is shown in Fig. 5.4.




Fig. 5.4 Testbed power supply circuitry.

The circuit blocks described above were constructed on a mixture of veroboard and custom-milled PCBs. The complete electrical testbed is shown in Fig. 5.5




Fig. 5.5 Implemented circuitry for the electrical testbed.

5.3 Thermal Testbed

Characterising a temperature sensor obviously involves recording its response over a range of temperatures. The operating temperature range for the target 0.13 μ m manufacturing process is specified as -55°C to +125°C, and the smart sensor will be tested over as much of this range as possible.

Two temperature-controlled environments were available covering a range from 0°C to well above the maximum operating temperature: a Contherm CAT1150CP incubator, and a Contherm CAT8150 oven. The absolute thermal accuracy of these enclosures is not important as a separate calibrated thermal probe will be used. Rather, the most important performance measure is the ability to hold the enclosure at a fixed temperature set-point. These appliances use PID control to achieve a temporal accuracy of $\pm 0.2^\circ\text{C}$, which is sufficiently accurate for the testbed application. The DUT will be placed inside the temperature-controlled cabinet and connected to the electrical testbed by ribbon cables. To reduce risk of mechanical damage when swapping between chips, a zero-insertion force (ZIF) socket was used to interface with the DUT.

The thermal testbed should be capable of measuring the DUT temperature as accurately as possible – a resolution of at least 0.01°C is desired to match the expected ADC accuracy. This requires an accurate thermometer placed in close proximity to the DUT. A platinum PT100 sensor was available, along with a Grant Squirrel 2020 datalogger to read out the sensor value. This sensor was calibrated at 0.0°C with an ice bath, and at 100.0°C in boiling water. The datalogger is equipped with a 24-bit $\Delta\Sigma$ ADC and can be configured to display the temperature measurement with more than 5 decimal places of accuracy. However a significant amount of random noise on this signal makes accurate temperature measurement from the on-board display impossible.

Accuracy can be increased by using the logging functionality to average successive sensor readings, then transferring the stored data to a PC through a USB connection. The sampling rate was set to the maximum value of 10Hz, and each logged value consisted of an average of 10 samples to give a logging rate of 1Hz. This averaging reduces the random error somewhat, and a resolution of 0.01°C can be obtained by visual estimation from the logged data (Fig. 5.6). A secondary advantage of plotting the measured temperature data is that any subtle trends in the data become clearly visible, indicating the enclosure temperature has not yet stabilised.




Fig. 5.6 Example of temperature data captured from datalogger.

The example data in Fig. 5.6 was chosen specifically to illustrate a phenomenon exhibited by the datalogger – the measured temperature shows sudden step changes that are not caused by actual temperature variations. These steps occur at unpredictable intervals, and have a magnitude of around 0.05°C. The reference thermometer and by extension the entire test setup is therefore limited to an accuracy of $\pm 0.025^\circ\text{C}$.

To obtain accurate data, it is necessary to ensure that the reference thermometer and DUT are at exactly the same temperature. One method of doing so is to thermally attach both to an aluminium block. The reference thermometer was attached by drilling a hole of diameter slightly larger than the probe and applying thermal paste. Attaching the IC package proved somewhat more challenging as the ceramic cover over the chip cavity protruded above the package face. This was solved by using layers of adhesive thermal pads to connect the ceramic package to the aluminium block (Fig. 5.7).




Fig. 5.7 Method used to attach DUT and reference thermometer to isothermal aluminium block.

Unfortunately, this isothermal setup did not perform as well as expected. The ZIF socket is exposed to cabinet temperature, and conducts any temperature changes through to the DUT via the package pins. Therefore the DUT temperature is an unknown value somewhere between the block and chamber temperatures, significantly reducing the accuracy of the test setup.

An alternative setup was then considered where the DUT and thermometer were placed in close proximity, and enclosed by thermally insulating material. The DUT and thermometer stabilise at the surrounding ambient temperature, which is a low-pass

filtered version of the cabinet temperature. Eliminating the large thermal mass of the aluminium block means the complete setup stabilises at the cabinet set-point in around 30 minutes, rather than several hours for the setup in Fig. 5.7.

The thermal insulation is not required to completely isolate the DUT from the cabinet; rather it is used to attenuate short-term fluctuations in cabinet temperature. A convenient material meeting this modest insulation requirement over the entire test range is corrugated cardboard. An enclosure was therefore constructed using this material. A complete schematic of the entire thermal and electrical test setup is shown in Fig. 5.8.

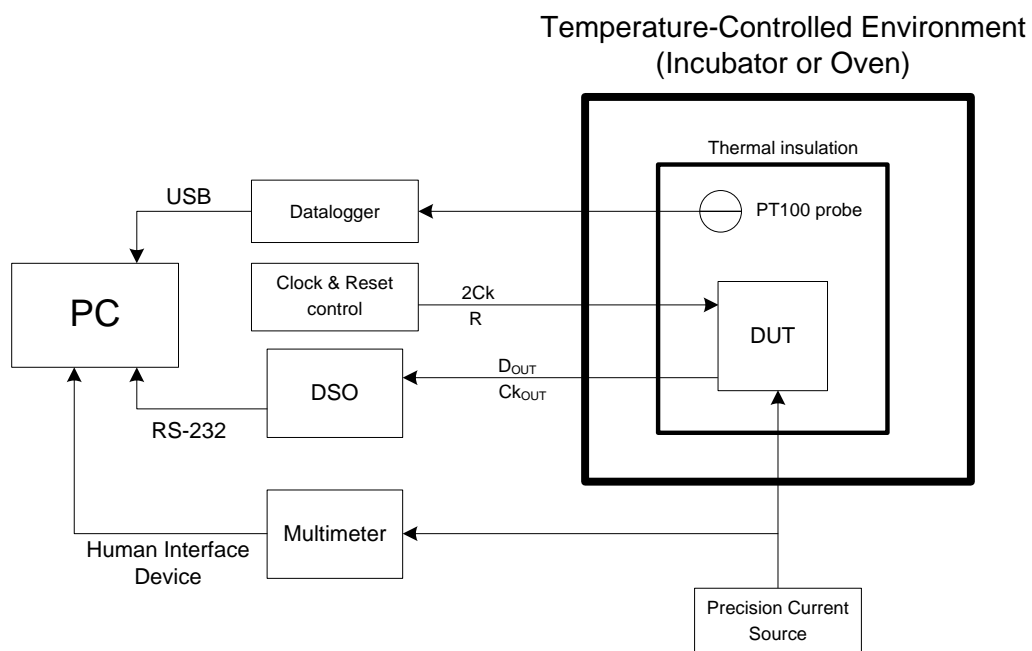


Fig. 5.8 Schematic of complete thermal and electrical testbed.

The equipment used in the thermal testbed is pictured in Fig. 5.9. The temperature-controlled environment in this case is the incubator; the oven is not visible in this picture.




Fig. 5.9 Test equipment used in the thermal testbed.

5.4 Functional Verification

Confirming that the smart sensor operates as intended is a significant first step in the testing process. Besides the final digital output value, several test points were configured to allow the operation of internal circuit blocks to be independently verified (Section 4.7).

A test device was initially powered up without clock or data inputs, and the DC operation of the front-end analog circuitry was confirmed. The bias voltages $V_{B1} - V_{B4}$ and the temperature-dependent V_{BE} and V_{PTAT} voltages were all within acceptable tolerances.

The sensor was then clocked for one ADC conversion cycle, and both the raw and filtered bitstreams were captured by the DSO. An example of the filtered 23-bit output waveform is shown in Fig. 5.10. The waveforms are shown inverted, due to the inherent inversion of the NAND gates used to buffer the chip output in Fig. 5.1.




Fig. 5.10 Typical filtered output bitstream captured on DSO.

Although Fig. 5.10 shows the complete 23-bit filter output waveform, less than 23 data bits are observable. This is because the clock transitions are clustered together in groups of 2 or 3 and separated by longer periods of inactivity. This phenomenon can be traced back to the operation of the $\Delta\Sigma$ modulator, specifically the 16-count integration performed by the modulator's first integrator (Section 4.5).

The decimation filter's output register is loaded with data at the end of the ADC conversion and produces the waveform in Fig. 5.10 as the $\Delta\Sigma$ modulator begins its next conversion cycle. The filter is necessarily clocked synchronously with the $\Delta\Sigma$ modulator, and the output register creating the waveforms in Fig. 5.10 also shares this clock. Therefore when the modulator is integrating ΔV_{BE} the 16 individual integrations delay the comparator's output, and the clock to the filter's output register is likewise delayed. The time between each clock pulse in Fig. 5.10 thus depends on the bitstream generated by the $\Delta\Sigma$ modulator as it begins its next conversion cycle.

The raw $\Delta\Sigma$ modulator output bitstream (single-bit) is shown for a typical room-temperature conversion in Fig. 5.11. This bitstream is an exact replica of the modulator comparator's output Q in Fig. 4.30, and is the same bitstream used by the decimation

filter to produce the final ADC output value. The Y-axis in Fig. 5.11 represents the weight given to each sample by the symmetric third-order filter. The final ADC output value can be calculated by multiplying the 511-byte bitstream by the corresponding weights in the third-order sinc filter impulse response. A full-scale input to the $\Delta\Sigma$ modulator would produce a continuous stream of '1s' at the comparator output Q, which would be multiplied by the filter impulse response to produce a filter output value of $180^3 = 5,832,000$. This value is therefore the full-scale output value of the $\Delta\Sigma$ ADC.




Fig. 5.11 Raw $\Delta\Sigma$ modulator bitstream multiplied by 3rd-order sinc filter impulse.

The final filtered output is shown in Fig. 5.12. This is a 23-bit digital number clocked off-chip serially (Fig. 5.10). This 23-bit value represents the filter's final output in standard unsigned binary format. Each bit is therefore weighted by its corresponding power of 2, to produce the same numerical value of 2,628,647 represented by the unfiltered bitstream in Fig. 5.11. This corresponds to 45.07% of the ADC's full-scale output. The exact correspondence between the summation of weighted bits in Figs. 5.11 and 5.12 confirm the correct numerical operation of the decimation filter.

The $\Delta\Sigma$ ADC output value in Fig. 5.12 provides information on the magnitude of the temperature-dependent ΔV_{BE} input in relation to the implicitly-generated reference V_{REF} (Section 3.7). However it provides no information on the exact value of temperature




Fig. 5.12 Filtered ADC output bitstream with binary weighting.

measured by the sensor front-end. The numerical correspondence between ADC output and temperature is determined by many parameters including the temperature-dependent magnitude and curvature of V_{BE} and ΔV_{BE} , and the magnitude and curvature of V_{BE} bias current (Table 4.1). As these critical parameters are substantially dependent on the characteristics of devices manufactured on the target process, it was not surprising that the room-temperature ADC value of the test device differed by several °C from that predicted by simulation. This has no effect on the analysis regime described in Chapter 6, as all error measurements are taken from a linear regression on individual or aggregate sensor data.

Upon measuring the current drawn by the first test device, it was discovered with some consternation that the device supply current was 4.3mA when the value expected from simulation was 52µA. However the chip was functioning as intended, so this problem most likely originates outside the active circuitry. The anomalous current was drawn through the chip’s analog supply pins, and the magnitude of this current was reasonably constant between tested devices ($\pm 5\%$). One strong possible cause is the ESD protection placed across all I/O pins and the analog and digital supply rail. Unfortunately the exact cause of this unusually high bias current cannot be determined from information available from the available I/O test points.

Although functional operation of the sensor is not affected by this anomalous current draw, the additional on-chip power dissipation is a concern for accurate temperature

measurement. The 28-pin ceramic package has a specified thermal resistance of 34°C/W in still air. Therefore the 5mW drawn by the DUT will result in a steady-state temperature error of 0.18°C. This significant error source can be avoided by obtaining a temperature reading directly after powering up the sensor, and removing power as soon as the output value is captured by the DSO.

5.5 Numerical Results

Having confirmed correct sensor operation at room temperature, data can now be gathered on the sensor's temperature-dependent behavior. The thermal testbed (Section 5.3) was used to sweep several test devices over the range of 0°C to 120°C, and the electrical testbed (Section 5.2) was used to record the ADC output, the analog voltage V_{BE} , and the pinch-base resistors R_1 , R_2 , and R_3 . As described in Section 5.1, gathering data from many test devices will allow the correlation between pinch-base resistance and sensor characteristics to be determined.

The time required for testing is largely determined by the thermal settling time after exchanging the DUT. This settling time increases as the chamber temperature moves away from ambient temperature, and over 1 hour was required at the extremes of the test range. This restricted the number of test devices that could be characterised in the time available. In order to complete the testing process in a reasonable amount of time, 6 test devices were characterised.

The observant reader will notice that 0°C data points are missing for two test devices. This is because the devices themselves ceased to operate correctly at low temperatures. The likely cause of this is a marginal metal path, via, or contact opening at low temperature. Interestingly, another device not used for data collection exhibited the same behavior at room temperature. The raw ADC data gathered from the 6 test devices is shown in Fig. 5.13 It is immediately apparent that the sensor output is a near-linear function of temperature. This confirms the correct operation of the $\Delta\Sigma$ modulator, the only circuit block not explicitly tested in Section 5.4. Moreover, it confirms the correct operation of the entire smart sensor as a whole.




Fig. 5.13 Output of ADC versus temperature.

Fig. 5.14 shows the raw V_{BE} voltage produced by the analog front-end. The negative linear temperature dependence is clearly visible. Higher-order curvature is also visible, particularly at higher temperatures.




Fig. 5.14 Substrate bipolar voltage V_{BE} versus temperature.

The resistance of R_1 (simple N-well resistor) is shown in Fig. 5.15. At higher temperatures it exhibits a reasonably linear dependence on temperature, while at lower temperatures significant nonlinearities are present.




Fig. 5.15 N-well resistor R_1 versus temperature.

Base-pinch resistors R_2 and R_3 are constructed slightly differently. The entire N-well of R_2 is pinched, while R_3 is only partially pinched to better comply with the target processes' design rules. However in testing it was determined that both resistors exhibited virtually identical behavior. Therefore Fig. 5.16 presents only the data of R_2 versus temperature, as R_3 would be graphically indistinguishable. Interestingly, the base pinch resistors exhibit less random variation than the un-pinched N-well resistor above.




Fig. 5.16 Base-pinch resistor R_2 versus temperature.

CHAPTER 6

ANALYSIS & DISCUSSION

This chapter analyses the results obtained in Chapter 5, from 6 test devices over the temperature range of 0°C to 120°C. The process-compensation circuitry was disabled during the data gathering process in order to obtain information on the correlation between the unmodified sensor output and the raw value of R_p . It is hoped the data obtained through this process can be used to calibrate the process-compensation circuitry and thereby reduce process-induced variations in sensor output. The sensor output is first analysed in Section 6.1 for linearity. A detailed analysis is performed in Section 6.2 of possible causes of linearity errors seen in the sensor output. The secondary data gathered from the analog front-end is analysed in Section 6.3 to determine the correlation between R_p and the raw substrate bipolar voltage V_{BE} . Section 6.4 analyses the overall correlation between R_p and sensor output that is necessary to achieve the desired process compensation. And Section 6.5 discusses the performance of the prototype temperature sensor in the context of the design goals developed in Chapter 3 and previously published smart sensor designs.

6.1 Sensor Output Linearity

The analog front-end uses a PTAT temperature signal, so the ideal sensor output is a straight line with zero output at -273.15°C. The V_{BE} and ΔV_{BE} bias currents were chosen to cancel second-order curvature in the sensor output. Simulations using the IBM-supplied device models predict that the residual deviation from a linear fit would be 0.06°C (Fig. 4.19).

A linear regression was applied to the experimental data from 6 test devices presented in Section 5.5, and Fig. 6.1 shows the residual error in °C from the average linear fit. Several features from Fig. 6.1 are worth discussing. The large peak error of nearly 5°C is contributed by one test device (#5) which exhibits significantly different output characteristics to the other 5 devices. Given the limited data available, it is difficult to




Fig. 6.1 Deviation from average linear fit for the 6 tested smart sensors.

determine if this difference is due to normal or abnormal variations in process parameters, or even a yield issue.

The remaining 5 devices display a much more linear output response with lower spread. However these devices all exhibit a significant deviation from linearity at the highest temperature point of 120°C. This is most likely caused by reverse-biased junction leakage at high temperature affecting the operation of the low-power analog front-end. If the temperature range is restricted to avoid the deviation at 120°C and device #5 is excluded from the regression analysis, the remaining 5 linear devices are accurate to 1°C over the range of 0°C to 100°C (Fig. 6.2).

Fig. 6.3 shows a plot of the gradient / intercept data pairs resulting from individual linear regressions on each sensor’s temperature curve. An ideal PTAT sensor output with zero ADC output at 0K would produce linear approximations with zero intercepts. Although the actual intercept values in Fig. 6.3 deviate from this ideal, the intercepts are very closely correlated to the corresponding linear gradients. The high R^2 value of 0.995 [Ledolter and Hogg, 2010, Chapter 8] indicates that variations in ADC output characteristic involve only one degree of freedom, and no further information is gained by maintaining separate gradients and intercepts. The correlation analysis in Sections 6.3 and 6.4 of this chapter will therefore use a linear approximation of the sensors’ ADC output gradients with the intercept set to zero.




Fig. 6.2 Deviation from average linear fit of 5 smart sensors.

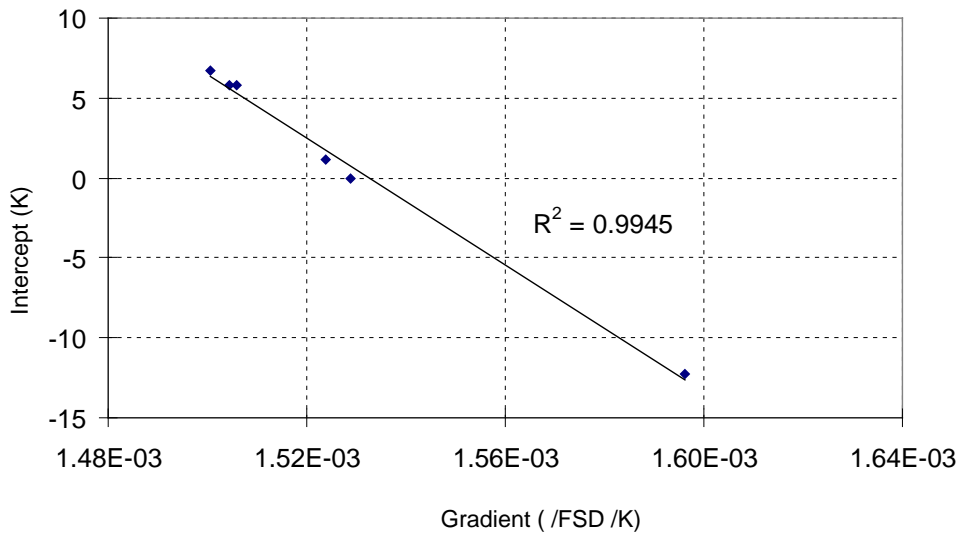


Fig. 6.3 Correlation between gradient and intercept of linear approximations of individual device output curves.

6.2 Sources of Inaccuracy

Despite excluding one tested device, the maximum deviation from an average linear fit in Fig. 6.2 was still 1°C over the limited temperature range of 0°C to 100°C – far greater than the 0.06°C predicted by simulation. This section will investigate the factors contributing to this large inaccuracy.

High-Temperature Junction Leakage

The five devices plotted in Fig. 6.2 all exhibit a significant positive output deviation at 120°C. This deviation takes the form of a step change of 1°C to 2°C, and is not consistent with residual high-order curvature. As mentioned previously in Section 6.1, this deviation is most likely to be caused by increased junction leakage at high temperature affecting the analog front-end. Interestingly, device #5 – which was excluded from Fig. 6.2 due to its significantly differing gradient – exhibits no visible step change at 120°C, although it may be obscured by the large second-order upward curvature in this device's output. This high-temperature error step can be avoided by restricting the temperature range used for analysis to 0°C – 100°C. Doing so reduces the maximum error in Fig 6.2 from 3°C to 1°C. Subsequent analyses in this section will use this restricted temperature range.

Process Parameter Variations

The largest deviation from the average linear fit in Fig. 6.1 is contributed by one device which exhibits substantially different output characteristics from the other five. Even among the five linear devices re-plotted in Fig. 6.2, significant gradient variations are visible at lower temperatures. Process parameter variations are the one factor not taken into account during simulation. To determine the error contribution from inter-device variations, the measured device output curves were plotted against their individual linear regressions rather than the group average linear regression. It can be seen that the peak error from linear fit reduces from 3.4°C in Fig. 6.1 to 1°C in Fig. 6.4 over the temperature range of 0°C to 100°C. The remaining 1°C error is still 17 times greater than the value predicted by simulation, and must be accounted for by other means.




Fig. 6.4 Deviation from individual linear fit for the 6 tested smart sensors from 0°C to 100°C.

Residual Curvature

The amount of curvature present in V_{BE} is influenced by the temperature-dependency of the bias current I_{PTAT} as well as the characteristics of the substrate bipolar transistors themselves. While the bias current generator produces a precise PTAT voltage (Section 3.5), this is converted into a PTAT current by the polysilicon resistor R_1 in Fig. 3.4. Curvature in the sensor's ADC output is therefore influenced by the temperature-dependent behaviour of both substrate bipolar transistors and polysilicon resistors.

The bipolar bias current magnitude was chosen in Section 4.3 to cancel second-order curvature predicted by simulations. However the individual sensor error curves in Fig. 6.4 clearly show that second-order curvature is a significant contributor to linearity errors in all tested devices. It is therefore apparent that the IBM-provided device models are not sufficiently accurate to support high-accuracy temperature sensor design.

To determine the error contribution from second-order curvature, second-order polynomial regression was applied to the measured device output curves. The resulting errors plotted in Fig. 6.5 are much more random in nature, and the peak error is reduced




Fig. 6.5 Deviation from individual 2nd-order polynomial fit for the 6 tested smart sensors from 0°C to 100°C.

to 0.43°C. Uncompensated second-order curvature therefore causes an error of around 0.5°C in each sensor’s final output.

Measurement Error

Several factors can contribute to error when measuring the DUT temperature. Section 5.3 showed how the reference thermometer output from the Grant datalogger exhibited unpredictable steps of around 0.05 °C. Another possible error source is incomplete thermal equalisation between the reference thermometer and DUT. Although reasonable care was taken during the data gathering process, the possibility of settling errors of up to 0.05°C cannot be discounted.

Another significant source of measurement error is directly due to variations in the sensor’s ADC output. Successive readings taken from a DUT at thermal equilibrium exhibit random variations of around 0.5°C. This effect was more visible at low temperatures, and reduced to negligible levels at the upper end of the temperature range. Fig. 6.6a shows a typical sequence of ADC readings obtained from device #3 at room temperature. The median temperature of 13 data points is 23.3°C, yet two readings peak at 23.8°C. The maximum difference between readings in Fig. 6.6a is 0.66°C. Fig. 6.6b

shows a series of readings obtained from the same device at 100°C, plotted with the same vertical scale. Random variation is greatly reduced at high temperature, and the maximum difference between 7 successive readings is now 0.03°C.




Fig. 6.6a Sequential ADC readings obtained from device #3 at room temperature.




Fig. 6.6b Sequential ADC readings obtained from device #3 at 100°C.

This trend of reduced random variation at high temperature is also visible upon reviewing the polynomial-fit errors plotted in Fig. 6.5. The largest random errors are at the lower end of the temperature range. At 100°C, all 6 devices are within 0.04°C of their respective second-order approximations. Interestingly, it was observed during

testing that once the devices had been exposed to high temperature, the random variations in ADC output remained at a much lower level when re-tested at room temperature.

When determining the actual ADC output value at a particular temperature, occasional large deviations such as shown in Fig. 6.6a were manually discarded. The final reading was then produced by averaging the remaining data points in the sequence. However in some cases the ADC output sequence contained many such large steps, to the point where outliers could not be isolated by visual inspection. The final ADC value was then produced by averaging a data sequence with a difference of 0.5°C between the highest and lowest points. Therefore the error contribution from ADC output uncertainty may be up to 0.25°C.

The largest contributor to residual error in Fig. 6.5 is device #1 with a peak deviation of 0.43°C. The other 5 devices exhibit similar deviations of around 0.15°C. The reason for the higher errors from device #1 is unknown, but the residual error shown by devices #2 to #6 can be completely accounted for by the uncertainty in ADC output as discussed above.

Although the exact cause of this unpredictable and temperature-dependent random variation is unknown, a possible explanation can be found in the test device's external connections. The global analog bias voltages $V_{B1} - V_{B4}$ were connected to external package pins to allow the operation of the bias voltage generator to be confirmed and corrected if necessary (Section 4.1). These four voltages are generated by transistors biased with around 2μA of current, and are therefore extremely sensitive to stray charge entering through the external pins. They are also directly next to the unfiltered digital bitstream output from the $\Delta\Sigma$ modulator (Fig. 4.38). This unfortunate layout decision means that any stray current leaking from the high-level digital output to the sensitive analog bias voltages will significantly affect the operation of all on-chip op-amps, used in the bias current generator and the $\Delta\Sigma$ modulator. This explanation is consistent with the observed temperature dependence of random variations, as a high-temperature bake will drive off moisture in any external surface contamination and reduce leakage into the bias voltage pins.

Summary of Error Contributions

The various sources of error in the smart sensor itself and the test procedure presented above can now be summarised. The total deviation of 5°C seen in Fig. 6.1 can be completely accounted for, with the exception of the higher random errors seen in device #1 (Fig. 6.5). Table 6.1 presents each error source discussed in this section, along with an estimated error contribution for each source.

Table 6.1 Summary of error sources in the measured smart sensor response.

Total deviation from average linear fit	5°C
Error Source	Estimated Error Contribution
Process parameter variations	4°C
High-temperature junction leakage	2°C
Residual second-order nonlinearity	0.5°C
Measurement error	
- Datalogger error	0.05°C
- Thermal settling error	0.05°C
- ADC random variation	0.25°C

6.3 Pinch-Base Correlation with V_{BE}

The process-compensation scheme implemented on the smart sensor developed in this thesis is fundamentally based on the correlation between pinch-base resistance R_P and bipolar saturation current I_S (Section 2.10). Given a constant temperature and bias current, the base-emitter voltage V_{BE} across a substrate bipolar will vary with changes in I_S . Therefore the first step in characterising the process-compensation system is to investigate the correlation between pinch-base resistance and the on-chip V_{BE} produced by the analog sensor front-end.

Analysing the correlation between two quantities requires reducing both those quantities to single-dimensional variables. When plotted on an X-Y plane, a collinear correlation between the two variables is then visible as a straight-line characteristic. It is therefore necessary to reduce the R_P and V_{BE} curves plotted in Section 5.5 into single-dimensional variables.

V_{BE} Characterisation

Fig. 6.7 plots the error when the V_{BE} curves were subtracted from an average linear fit. As expected, significant second-order curvature is present. Significant inter-device variations are also apparent.




Fig. 6.7 Deviation from average linear fit of V_{BE} generated by the 6 test devices.

Earlier in Section 6.1, it was found that the gradients and intercepts of the sensors' output approximations were very closely correlated ($R^2 = 0.995$). Maintaining separate gradients and intercepts (two dimensions or degrees of freedom) was therefore redundant when one could be predicted by the other. The gradients and intercepts from the linear approximations of V_{BE} are plotted against each other in Fig. 6.8. Although a straight-line relationship is still in evidence, the lower R^2 value of 0.77 indicates that gradient or intercept cannot be derived entirely from the other. Thus the V_{BE} curves cannot be represented by a one-dimensional variable.




Fig. 6.8 Correlation between gradient and intercept of linear approximations of individual V_{BE} curves.

The large second-order curvature visible in Fig. 6.7 suggests a second-order polynomial regression is more appropriate for representing V_{BE} . This regression takes the form:

$$V_{BE} = aT^2 + bT + c, \quad (6.1)$$

and thus incorporates 3 degrees of freedom. It was found that the coefficients a , b , and c in Eq. 6.1 were very closely correlated, with an average R^2 value of 0.994 between the three coefficients. Although the polynomial regression is a much better approximation of V_{BE} , it can be seen from the high inter-coefficient correlation that the third degree of freedom is unnecessary. The correlation analyses performed later in this chapter will therefore use the gradient / intercept pairs of the linear approximation plotted in Fig. 6.8.

R_p Characterisation

The temperature-dependency of pinch-base resistors R_1 and R_2 presented in Section 5.5 exhibit a complex shape that does not easily lend itself to the approximations used to characterise V_{BE} and the ADC output. Above 60°C both resistors are approximately linear with temperature, while below 60°C significant curvature is visible. Both types of resistor show no significant variation in gradient between devices; process-induced variations manifest most significantly as variations in curve offset (Figs. 5.15 and 5.16).

The process-compensation circuitry described in Section 4.4 uses the value of R_P to control the substrate bipolar bias current I_{BIAS} . Any change in R_P with temperature will affect the temperature-dependency of I_{BIAS} , thereby introducing additional curvature into V_{BE} . An ideal base-pinch resistor would therefore exhibit a flat temperature characteristic while still remaining sensitive to variations in process parameters.

It can be seen from Fig. 5.16 that the base-pinch resistors implemented on the target process are far from this ideal behaviour. The magnitude of R_2 varies by 23% of its room temperature value over the tested range of 0°C to 120°C, while process-induced variations are only 3% of the room temperature value. This will present a considerable challenge when linearising the sensor output, but in this section only the resistors' process sensitivity is of interest. It is therefore necessary to isolate the process-induced variation in R_P from the large temperature sensitivity into a single degree of freedom.

Towards this aim, it is noted that both R_1 and R_2 exhibit roughly linear temperature-dependence above 60°C. Linear regression was used to determine the resistance of R_1 , R_2 , and R_3 at exactly 80°C using the measured data from 60°C to 100°C. This provides a measure of the offset of each resistor's response curve, which is the parameter most sensitive to process variations.

Correlation Analysis

As described in Section 4.7, three well resistors were designed to investigate the effect of base-pinch resistor layout. Resistor R_1 is a standard N-well resistor included for comparative purposes. R_2 is an ideal pinch-base resistor where the entire width of N-well is pinched with P-active, thereby violating several IBM design rules. R_3 does not pinch the entire width of N-well, thus causing fewer design rule variations. The values of these three resistors at 80°C were individually correlated against the gradients and intercepts of V_{BE} 's linear approximations on each of the 6 test devices (Fig. 6.9).

All possible combinations of R_P and V_{BE} in Fig. 6.9 were analysed, but only the pairs that produced a result of significance are labelled. The un-pinned resistor R_1 is not correlated in any significant way to V_{BE} . R_2 and R_3 both exhibit some degree of correlation, but the more ideal construction of R_2 gives noticeably better results.

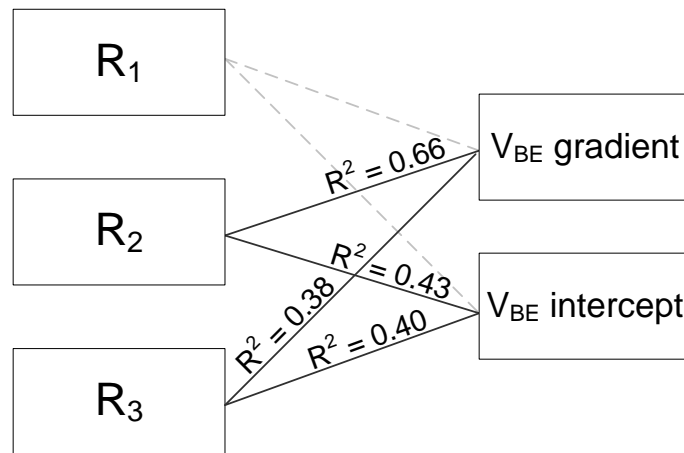


Fig. 6.9 Diagram showing possible correlations between on-chip resistors and V_{BE} characteristics.

Based on the analysis in Section 2.10, it is expected that process variations will primarily affect the gradient of V_{BE} . And indeed, the highest R^2 value of 0.66 is between R_2 and V_{BE} gradient. This relationship is plotted in Fig. 6.10.




Fig. 6.10 Correlation between value of R_2 @ 80°C and V_{BE} 's gradient.

The correlation plotted in Fig. 6.10 is similar in strength to the relationship between R_p and I_s reported by Dutton and Divekar [1977] with $R^2 = 0.69$. Obviously, many IC process generations have passed between then and now. In fact, the measured correlation of $R^2 = 0.66$ in Fig. 6.10 is quite significant when it is considered that I_s was not measured directly. The measured values of V_{BE} obtained from the 6 test devices were also influenced by variations in bias current magnitude, which depends on the reproducibility of the polysilicon resistors used in the bias current generator circuit (Fig. 3.4) as well as the effectiveness of the implemented chopping and DEM.

It can therefore be concluded that pinch-base resistors exhibit a significant correlation to the characteristics of the substrate bipolar transistors manufactured on the targeted $0.13\mu\text{m}$ CMOS process.

6.4 Pinch-Base Correlation with Sensor Output

The smart sensor developed in this thesis uses circuit-level techniques to eliminate many sources of error and noise in the analog circuitry. The most significant contributor to variations in sensor output characteristic is expected to be temperature gradient variations in V_{BE} . The process-compensation circuitry requires an end-to-end correlation between pinch-base resistance and sensor output characteristic. Section 6.3 has demonstrated that a link exists between pinch-base resistance and the characteristics of V_{BE} . This section will therefore investigate the remaining link between V_{BE} and the sensor output.

V_{BE} is used to create the ADC reference voltage (Eq. 2.10), and therefore directly affects the sensor's output (Eq. 2.40). It was determined in Section 6.1 that inter-device variation in ADC output characteristics could be represented by a single degree of freedom: the gradient of the linear approximation of each device's output curve. However both the gradient and intercept of V_{BE} 's linear approximations in Section 6.3 were necessary to represent inter-device variation. Fig. 6.11 shows the possible relationships between V_{BE} and the sensor output, along with the most significant correlations from Section 6.3.

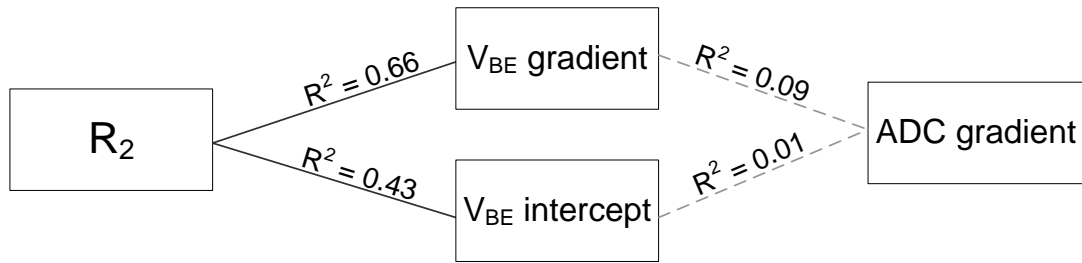


Fig. 6.11 Diagram showing possible correlations between on-chip resistor R_2 , V_{BE} , and ADC output characteristic.

It can be seen from the R^2 values in Fig. 6.11 that the gradient and intercepts of V_{BE} are not correlated in any significant way to the ADC gradient. This surprising result appears to be in contradiction to the physical reality of the sensor itself. Further investigation reveals that if a linear regression is performed on V_{BE} with a fixed intercept, the resulting single-dimensional gradient exhibits a somewhat more significant correlation to ADC gradient. Plotted in Fig. 6.12, this relation has an R^2 value of 0.55.




Fig. 6.12 Correlation between V_{BE} single-dimensional gradient representation and ADC output gradient.

The correlation in Fig. 6.12 does not appear convincing given the close physical relationship between V_{BE} and the ADC output. Particularly, the single outlying data point at the top-right of the graph appears to be distorting the linear fit. The remaining 5 data points in Fig. 6.12 exhibit a linear trend opposite in slope to the best-fit line. This suggests that another factor is influencing the ADC output besides the gradient or offset of V_{BE} .

The most significant correlations discovered in Sections 6.3 and 6.4 are shown diagrammatically in Fig. 6.13. While a chain of correlated variables exists between R_p and the ADC output, the direct correlation necessary for correct operation of the process compensation circuitry is $R^2 = 0.035$. This miserable result eliminates the possibility of successfully configuring the external process-compensation circuitry to reduce gradient variations in sensor output gradient. As mentioned above, this is due to the lack of a strong correlation between V_{BE} and the ADC output.

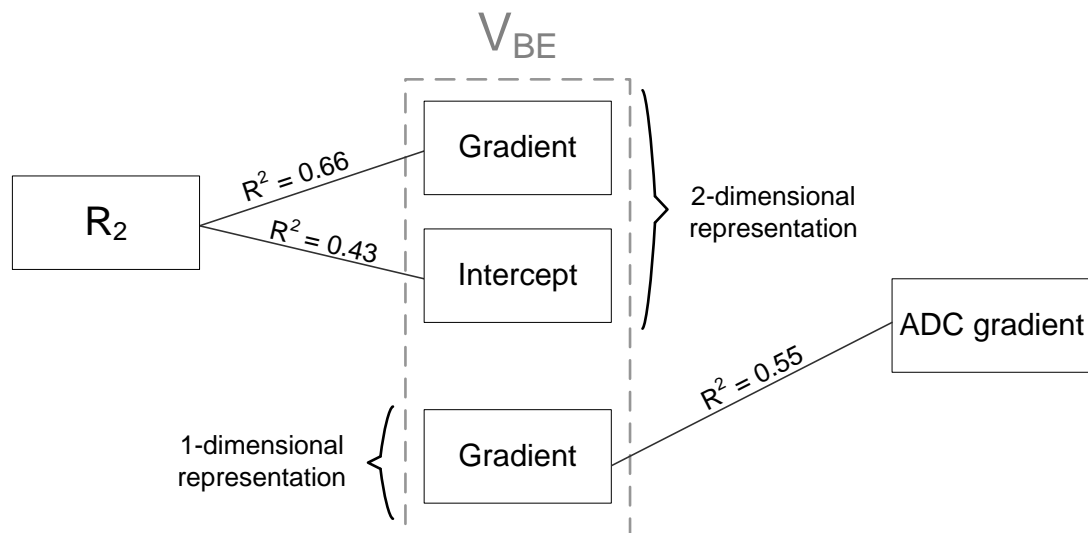


Fig. 6.13 Diagram showing significant correlations between on-chip resistor R_2 , V_{BE} , and ADC output characteristic.

6.5 Discussions

Having analysed the experimental data gathered in Chapter 5, it is now time to evaluate the developed smart sensor against the design characteristics discussed in Chapter 3. The design criteria listed in Section 3.2 can be divided into three distinct categories. The sensor should build on the advantages common to all smart IC sensors: low support circuitry requirements, low power, and low cost. It should incorporate advanced analog techniques to achieve a level of accuracy higher than past smart sensor designs. And it should implement a system to reduce the errors caused by process variations to reduce the need to calibrate the manufactured devices. These three design goals are now individually examined.

Evaluation of Typical Smart Sensor Features

The fundamental advantage of smart temperature sensors over discrete sensors arises from their use of standard IC manufacturing technology to integrate the signal conditioning and ADC onto the same device as the thermal sensor itself. Smart sensors therefore offer the advantages of high reliability, low power consumption, low cost, and easy integration into microprocessor-based control systems. The developed smart sensor matches or exceeds the specifications of existing smart sensors in all these aspects.

The prototype sensor is completely self-contained, from the first bias voltage generator to the last digital output shift register. Its only electrical inputs are a clock and reset command, and the only essential output is a serial clock/data bitstream. While it does not implement an addressable bus protocol such as I²C, the output data is easily captured with edge-triggered digital logic. In fact, this simple logic interface is actually easier to implement than a complete bus protocol in many cases.

The physically manufactured test devices require 4.3mA of supply current due to unexpected parasitic conduction paths, most likely within ESD protection devices. The simulated current draw is 52 μ A, which is less than the value of 75 μ A reported by Pertijs, Makinwa et al. [2005], and significantly less than the 200 μ A drawn by the mass-produced TCN75 described in Chapter 1.

The die area required by on-chip circuitry is a significant contributor to the per-unit cost of a mass-produced IC. Several methods have been employed to reduce the prototype sensor's layout area to 0.21mm², excluding bondpads. Pertijs, Makinwa et al. reported a die area of 4.5mm² including bondpads and test circuitry, but excluding the $\Delta\Sigma$ decimation filter. The prototype smart sensor developed for this thesis therefore requires approximately 10 times less die area to perform the same sensor function. This substantial area reduction was achieved through two methods. A modern 0.13 μ m manufacturing process was employed, which reduced the size of all circuitry but had the most impact on the digital decimation filter and sequencing logic. Secondly, the area required by capacitors in the $\Delta\Sigma$ modulator was significantly reduced by optimising capacitor values at the expense of conversion time (Section 4.5). The on-chip capacitors

used in the $\Delta\Sigma$ modulator occupied 0.06mm^2 in the implemented device, which is 29% of the total active circuitry area. Without taking these measures to reduce capacitor values, the area occupied by capacitors alone could have easily been 5 times larger, resulting in a total die area of 0.45mm^2 .

Evaluation of High-Accuracy Analog Techniques

The smart sensor design published by Pertijs, Makinwa et al. [2005] is taken as the base-line for the sensor developed in this thesis, as it is the most accurate smart temperature sensor design published to date. It achieves an uncalibrated error of $\pm 0.3^\circ\text{C}$ over the temperature range of -55°C to 125°C . Single-point calibration reduces this error to $\pm 0.1^\circ\text{C}$.

The 6 prototype sensors tested in Chapter 5 displayed a maximum deviation of 5°C from an average linear fit over the range of 0°C to 120°C . The most significant contributor to this error is inter-device variation caused by process parameter shifts (Table 6.1). As the circuitry used in the prototype sensor is essentially a refined and optimised version of that published by Pertijs, Makinwa et al., the cause of this large error is likely to be found elsewhere. The fact that the gradient and intercepts of the individual ADC output linear regressions display a very strong correlation (Fig. 6.3) indicates that the cause of this error is an unexpected systematic parameter variation, rather than random noise or measurement error.

A likely cause of the large inter-device error lies in the substrate bipolar transistors implemented on the target process. Pertijs, Makinwa et al. designed their smart sensor on a $0.7\mu\text{m}$ CMOS process, and reported a current gain of $\beta = 22$ for the substrate bipolars implemented therein. The design kit for the $0.13\mu\text{m}$ IBM process used to produce this prototype sensor predicts a substrate bipolar current gain of $\beta = 2$.

The value of β is significant because the equation used to predict the temperature-dependency of a base-emitter junction (Eq. 2.3) is a function of the transistor's collector current, while substrate bipolar transistors must be biased through their emitter. Although the bias current generator is designed to make the collector current independent of β (Section 3.5), in reality perfect cancellation will not be achieved if β

mismatch exists between the transistors in the bias generator and the bipolar core (Fig. 3.1). The compensation factor $(\beta + 1) / \beta$ in Eq. 3.4 is derived from the substrate bipolars in the bias generator circuit, which will not completely compensate base current in the bipolar core if β mismatch exists. And given the extremely low value of β in the target process, small mismatches will result in significant collector current errors.

Evaluation of Process Compensation

Variations in manufacturing process parameters cause variations in the temperature-dependence of the substrate bipolars' base-emitter junctions. Two circuits were presented in Section 4.4 that compensate for this variation by using the value of a pinch-base resistor to adjust the magnitude of bias current supplied to the bipolar transistors. The prototype smart sensor implemented a current-mirror interface to allow external process-compensation circuitry to modify the operation of the bipolar sensor core (Fig. 4.23), although this interface was disabled during testing.

The data gathered from the 6 test devices was analysed in Sections 6.3 and 6.4. It was found that the fundamental correlation between R_P and V_{BE} necessary for correct operation of the process compensation system does indeed manifest on the target manufacturing process. However the necessary assumption that V_{BE} is the most significant contributor to sensor output variations was not valid on the prototype device, so the complete process-compensation system could not be experimentally verified. Nonetheless, the pinch-base compensation system developed in this thesis remains a promising technique for increasing the uncalibrated accuracy of smart integrated temperature sensors.

CHAPTER 7

CONCLUSIONS

The experimental results obtained from the prototype smart sensor have been analysed and discussed in Chapter 6 with respect to the design methodology developed in Chapter 3. In this final chapter, research outcomes are evaluated against the Objectives established in Section 1.1. Key findings are summarised in Section 7.2 as recommendations for creating high-performance smart sensors in modern CMOS processes. And Section 7.3 identifies significant opportunities for future research work.

7.1 Summary of Research

Objective 1

Research Objective 1 was to review literature on the subject of integrated smart sensors and develop a design methodology improving upon the shortcomings identified therein. This objective has been satisfied by the contents of Chapters 2 & 3 of this thesis. The literature review incorporated various methods of sensing temperature with devices manufactured on standard IC processes (Section 2.1), the history of temperature sensors on bipolar and CMOS processes (Sections 2.2 & 2.4 respectively), and methods of increasing on-chip sensor accuracy through linearisation and offset reduction (Sections 2.3 & 2.5). On-chip ADC designs progressed towards the favourable characteristics of $\Delta\Sigma$ data converters (Section 2.6), and optimal techniques for implementing the decimation filter and modulator circuitry were examined in Sections 2.7 & 2.8. A smart sensor design was examined in detail in Section 2.9 that incorporated many of the design features discussed in previous sections, and achieves an accuracy level of $\pm 0.1^\circ\text{C}$. The sparse literature available on the subject of the correlation between base-pinch resistors and bipolar transistors was reviewed in Section 2.10.

Chapter 3 opened with a review of the strengths and weaknesses of existing smart sensor designs. The primary advantages of smart temperature sensors originate from their integrated nature: low cost, low power consumption, high integration, and high

reliability. The most significant shortcoming is the low accuracy of the on-chip thermal sensor. A suitable design methodology therefore involved continuing to build on the advantages inherent to smart sensors, while addressing the accuracy shortcomings by incorporating advanced analog techniques (Section 2.9) and compensating for process-induced variations with base-pinch resistors (Section 2.10). Chapter 3 went on to detail circuit-level schematics that enable a prototype design to meet the specified design goals.

Objective 2

Research Objective 2 involved designing, fabricating, and testing a prototype smart sensor using the methodology developed in Objective 1. Chapter 4 described in detail the many design decisions required to implement the prototype circuitry presented in Chapter 3. An experimental testbed suitable for testing the sensor was described in Chapter 5, along with the data gathered from 6 test devices over the temperature range of 0°C to 120°C. The specific circuit-level improvements incorporated into the prototype smart sensor were:

- Uncalibrated bias current accuracy was improved by incorporating DEM into the bias generator circuit. Bias current magnitude then depended solely on the value of one polysilicon resistor.
- Loading on the sensitive $\Delta\Sigma$ modulator input signals was reduced by employing a CIFF modulator loop structure.
- A high-accuracy differential SC integrator at the $\Delta\Sigma$ modulator's input reduced capacitor layout area by implementing the ΔV_{BE} gain factor as a series of 16 integrations.
- ADC accuracy was increased without affecting conversion time by implementing a 3rd-order on-chip decimation filter.
- Circuitry was developed to reduce process-induced variations in V_{BE} gradient by modifying substrate bipolar bias current based on the value of a base-pinch resistor.

Objective 3

Research Objective 3 was to evaluate the prototype smart sensor against its simulated performance, and against the performance of previous smart sensor designs. An extensive analysis of experimental data gathered from the 6 test devices was performed in Chapter 6. It was concluded in Section 6.5 that the prototype sensor easily surpassed previous designs in terms of die area and power consumption. The primary contributing factors to this advantage were the use of a modern 0.13 μm CMOS manufacturing process, and reduction of the $\Delta\Sigma$ modulator's integrating capacitor values at the expense of conversion time.

Although the prototype device was fully functional, its output inaccuracy of 5 $^{\circ}\text{C}$ fell far short of the targeted $\pm 0.1^{\circ}\text{C}$. The most significant contributor to this error was inter-device variation, which was not anticipated during the design process as the analog front-end circuitry was explicitly designed to minimise variation between devices. The most likely cause of the inter-device variation was incomplete base-current compensation between bias generator and bipolar core substrate transistors, compounded by the extremely low substrate bipolar β in the target IC process.

Data gathered from base-pinch resistors and substrate bipolar transistors indicated that a significant correlation between R_P and V_{BE} does exist. However the aforementioned random variation in sensor output characteristics swamped the correlation between substrate bipolar V_{BE} and sensor output, therefore rendering ineffective any attempt to reduce inter-device variation with the process-compensation circuitry.

7.2 Recommendations for High-Performance Smart Sensor Designs

From the discussion above and the more detailed analyses in Chapter 6, the following recommendations can be made for IC designers wishing to create a state-of-the-art smart temperature sensor:

- The accuracy of on-chip thermal sensors is fundamentally limited by the gain of substrate bipolar transistors created on the target manufacturing process. A

substrate bipolar β of around 5 to 10 should be considered a minimum for accurate operation

- Layout area, particularly in the $\Delta\Sigma$ decimation filter, can be reduced by selecting a modern CMOS process with small minimum feature size. However these more advanced processes typically offer poor substrate bipolar performance. Select the most advanced process that meets the minimum figure for β specified above.
- At low values of β , matching between substrate bipolar transistors used in the bias current generator and bipolar sensor core is critical. Use common-centroid layout techniques to ensure all substrate bipolars in the analog front-end match as closely as possible.
- Vendor-supplied process development kits cannot be relied upon to achieve maximum accuracy. Sensor linearity depends on the temperature-dependence of polysilicon resistors and substrate V_{BE} (both gradient and curvature). The parameters of these devices must be characterised as accurately as possible.
- Process-compensation is possible using base-pinch resistors created with standard CMOS layers. Laying out such a resistor is likely to violate several design rules on a modern CMOS process – do the best you can. The temperature-dependence of this resistor must be characterised to the same accuracy as the polysilicon resistors and substrate bipolars. As pinch resistors exhibit significant voltage modulation, this characterisation must be performed at the application's bias current temperature-dependence and magnitude.
- Bias current ratios within the bipolar sensor core must be chosen to minimise curvature in the final ADC output. This is only possible with accurate models of base-pinch & polysilicon resistance and substrate bipolar V_{BE} temperature-dependence.

7.3 Opportunities for Further Research

Research investigating the role of base-pinch resistors in reducing the effect of process parameter variations in substrate bipolar-based thermal sensors is surprisingly rare in published literature. Data gathered from the prototype sensor developed in this thesis revealed a significant correlation between R_P and substrate bipolar V_{BE} , although other limitations prevented the full significance of this correlation from being realised. This presents an opportunity of significance for a researcher in the field of smart sensors to follow the design presented in this thesis and the recommendations in Section 7.2 to establish a functional correlation between R_P and sensor output gradient, thereby producing a smart IC-based temperature sensor with an uncalibrated accuracy exceeding any published or commercially-available design to date. The ultimate smart thermometer.

REFERENCES

- R. W. Adams, "The Design of High-Order Single-Bit $\Sigma\Delta$ ADCs," in *Delta-Sigma Data Converters*, S. R. Norsworthy, R. Schreier and G. C. Temes, Eds. Piscataway, NJ: IEEE Press, 1997, pp. 165-192.
- A. L. Aita, M. Pertijs, K. Makinwa, and J. H. Huijsing, "A CMOS smart temperature sensor with a batch-calibrated inaccuracy of $\pm 0.25^\circ\text{C}$ (3σ) from -70°C to 130°C ," in *IEEE Int. Solid-State Circuits Conf. 2009*, pp. 342-343.
- R. Amador, A. Polanco, H. Hernández, E. González and A. Nagy, "Technological compensation circuit for accurate temperature sensor," *Sensors and Actuators A: Physical*, vol. 69, no. 2, pp. 172-177, Aug. 1998.
- A. Bakker and J. H. Huijsing, "Micropower CMOS temperature sensor with digital output," *IEEE J. Solid-State Circuits*, vol. 31, no. 7, pp. 933-937, July 1996.
- A. P. Brokaw, "A simple three-terminal IC bandgap reference," *IEEE J. Solid-State Circuits*, vol. 9, no. 6, pp. 388-393, Dec. 1974.
- J. R. Carstens, *Electrical Sensors and Transducers*. NJ, NY: Regents / Prentice Hall, 1993.
- S. Dimitrijević, *Understanding Semiconductor Devices*. NY: Oxford University Press, 2000.
- R. W. Dutton and D. A. Divekar, "Bipolar models for statistical IC design," in: *Process and Device Modeling for Integrated Circuit Design*, F. Van De Wiele et al. The Netherlands: Addison-Wesley, 1977, pp. 461-518.
- J. J. Ebers and J. L. Moll, "Large-Signal Behavior of Junction Transistors", in *Proc. IRE*, vol. 42, no. 12, pp. 1761-1772, Dec. 1954.

C. C. Enz and G. C. Temes, "Circuit Techniques for Reducing the Effects of Op-Amp Imperfections: Autozeroing, Correlated Double Sampling, and Chopper Stabilization," in *Proc. IEEE*, vol. 84, no. 11, pp. 1584-1614, Nov. 1996.

J. Fraden, *Handbook of Modern Sensors*. NY, USA: Springer, 2004.

P. R. Gray, P. J. Hurst, S. H. Lewis, and R. G. Meyer, *Analysis and Design of Analog Integrated Circuits*. NY, USA: John Wiley & Sons, 2001.

M. Gunawan, G. Meijer, J. Fonderie and J. H. Huijsing, "A curvature-corrected low-voltage bandgap reference," *IEEE J. Solid-State Circuits*, vol. 28, no. 6, pp. 667-670, June 1993.

A. Hastings, *The Art of Analog Layout*. NJ, USA: Prentice Hall, 2006.

E. B. Hogenauer, "An economical class of digital filters for decimation and interpolation," *IEEE Trans. Acoust. Speech, Signal Processing*, vol. 29, no. 2, pp. 155-162, April 1981.

D. Johns and K. Martin, *Analog Integrated Circuit Design*. USA: John Wiley & Sons, 1997.

A. Kölling, F. Bak, P. Bergveld, and E. Seevinck, "Design of a CMOS Temperature Sensor with Current Output", *Sensors and Actuators: A, Physical*, vol. 22, no. 1-3, pp. 645-649, June 1989.

P. Krummenacher and H. Oguey, "Smart temperature sensor in CMOS technology," *Sensors and Actuators A: Physical*, vol. 22, no. 1-3, pp. 636-638, June 1989.

J. Ledolter and R. V. Hogg, *Applied Statistics for Engineers and Physical Scientists*. NJ, NY: Prentice Hall, 2010.

J. Márkus, "Higher-order incremental delta-sigma analog-to-digital converters," Ph.D. dissertation, Budapest University of Technology and Economics, Budapest, Hungary, 2005. [Online]. Available: <http://www.mit.bme.hu/projects/delsig01/index.html>

J. Márkus, P. Deval, V. Quiquempoix, J. Silva and G. C. Temes, "Incremental Delta-Sigma Structures for DC Measurement: an Overview," in *IEEE Custom Integrated Circuits Conf. 2006*, pp. 41-48.

J. Márkus, J. Silva and G. C. Temes, "Theory and applications of incremental converters," in *IEEE Trans. Circuits and Systems I*, vol. 51, no. 4, pp. 678-690, April 2004.

G. Meijer, "An IC temperature transducer with an intrinsic reference," *IEEE J. Solid-State Circuits*, vol. 15, no. 3, pp. 370-373, June 1980.

G. Meijer, A. Boomkamp, and R. J. Duguesnoy, "An accurate biomedical temperature transducer with on-chip microcomputer interfacing," *IEEE J. Solid-State Circuits*, vol. 23, no. 6, pp. 1405-1410, Dec. 1988.

G. C. M. Meijer, P. C. Schmale, and K. van Zalinge, "A New Curvature-Corrected Bandgap Reference", *IEEE J. Solid-State Circuits*, vol. 17, no. 6, pp. 1139-1143, Dec. 1982.

G. Meijer, G. Wang and F. Fruett, "Temperature Sensors and Voltage References Implemented in CMOS Technology," *IEEE J. Sensors*, vol. 1, no. 3, pp. 225-234, Oct. 2001.

Microchip Technology Inc. (2010, Feb.). *TCN75 Datasheet* [Online]. Available: <http://www.microchip.com/wwwproducts/Devices.aspx?dDocName=en010752>

E. Montané, S. A. Bota and J. Samitier, "A compact temperature sensor for a μm^2 CMOS technology using lateral p-n-p transistors," *Microelectronics Journal*, vol. 29, no. 4-5, pp. 277-281, April-May 1998.

M. A. P. Pertijs, A. Bakker and J. H. Huijsing, "A High-Accuracy Temperature Sensor with Second-Order Curvature Correction and Digital Bus Interface," in *Int. Symp. Circuits and Systems 2001*, vol. 1, pp. 368-371.

M. A. P. Pertijs, K. A. A. Makinwa and J. H. Huijsing, "A CMOS Smart Temperature Sensor With a 3σ Inaccuracy of $\pm 0.1^\circ\text{C}$ From -50°C to 120°C ," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2805- 2815, Dec. 2005.

M. A. P. Pertijs, A. Niederkorn, X. Ma, B. McKillop, A. Bakker and J. H. Huijsing, "A CMOS Smart Temperature Sensor With a 3σ Inaccuracy of $\pm 0.5^\circ\text{C}$ From -50°C to 120°C ," *IEEE J. Solid-State Circuits*, vol. 40, no. 2, pp. 454-461, Feb. 2005.

W. Rasmussen, "Thermal sensors," in *Sensor Technology and Devices*, L. Ristic. MA, USA: Artech House, 1994, pp. 287-291.

B. Razavi, *Design of Analog CMOS Integrated Circuits*. NY, USA: McGraw-Hill, 2001.

F. R. Riedijk and J. H. Huijsing, "An integrated absolute temperature sensor with digital output," in *Int. Conf. Solid-State Sensors and Actuators Tech. Dig. 1991*, pp. 479-482.

J. Robert and P. Deval, "A second-order high-resolution incremental A/D converter with offset and charge injection compensation," *IEEE J. Solid-State Circuits*, vol. 23, no. 3, pp. 736-741, June 1988.

R. Schreier and G. C. Temes, *Understanding Delta-Sigma Data Converters*. Piscataway, NJ: IEEE Press, 2005.

B. Song, and P. R. Gray, "A Precision Curvature-Compensated CMOS Bandgap Reference", *IEEE J. Solid-State Circuits*, vol. 18, no. 6, pp. 634-643, Dec. 1983.

M. P. Timko, "A two-terminal IC temperature transducer," *IEEE J. Solid-State Circuits*, vol. 11, no. 6, pp. 784-788, Dec. 1976.

Y. P. Tsividis, "Accurate analysis of temperature effects in I_C/V_{BE} characteristics with application to bandgap reference sources," *IEEE J. Solid-State Circuits*, vol. 15, no. 6, pp. 1076-1084, Dec. 1980.

D. van Maaren, J. Klijn, and G. C. M. Meijer, "An Integrated Micropower Low-Voltage Temperature-Controlled Oscillator", *IEEE J. Solid-State Circuits*, vol. 17, no. 6, pp. 1197-1201, Dec. 1982.

E. A. Vittoz, "The Design of High-Performance Analog Circuits on Digital CMOS Chips", *IEEE J. Solid-State Circuits*, vol. 20, no. 3, pp. 657-665, June 1985.

G. Wang and G. Meijer, "The temperature characteristics of bipolar transistors fabricated in CMOS technology," *Sensors and Actuators A: Physical*, vol. 87, no. 1-2, pp. 81-89, Dec. 2000.

R. J. Widlar, "New developments in IC voltage regulators," *IEEE J. Solid-State Circuits*, vol. 6, no. 1, pp 2-7, Feb. 1971.

APPENDICES

Conference Proceedings I

R. P. Fisk and S. M. Hasan, "Analysis of Internally-Generated Noise in Bandgap References," in *Proc. Electronics New Zealand Conf.*, Christchurch, New Zealand, Nov. 2006, pp. 18-23.

Conference Proceedings II

R. P. Fisk and S. M. Hasan, "Incremental Delta-Sigma Modulators for Temperature Sensing Applications," in *Proc. Int. Conf. Mechatronics and Machine Vision in Practice*, Auckland, New Zealand, Dec. 2008, pp. 63-67.

Conference Proceedings III

R. P. Fisk and S. M. Hasan, "Low-Cost Temperature Sensor on a Modern Submicron CMOS Process," in *Proc. Electronics New Zealand Conf.*, Otago, New Zealand, 2009, pp. 43-48.